

SLApp Helpdesk — User Manual

For technicians and administrators

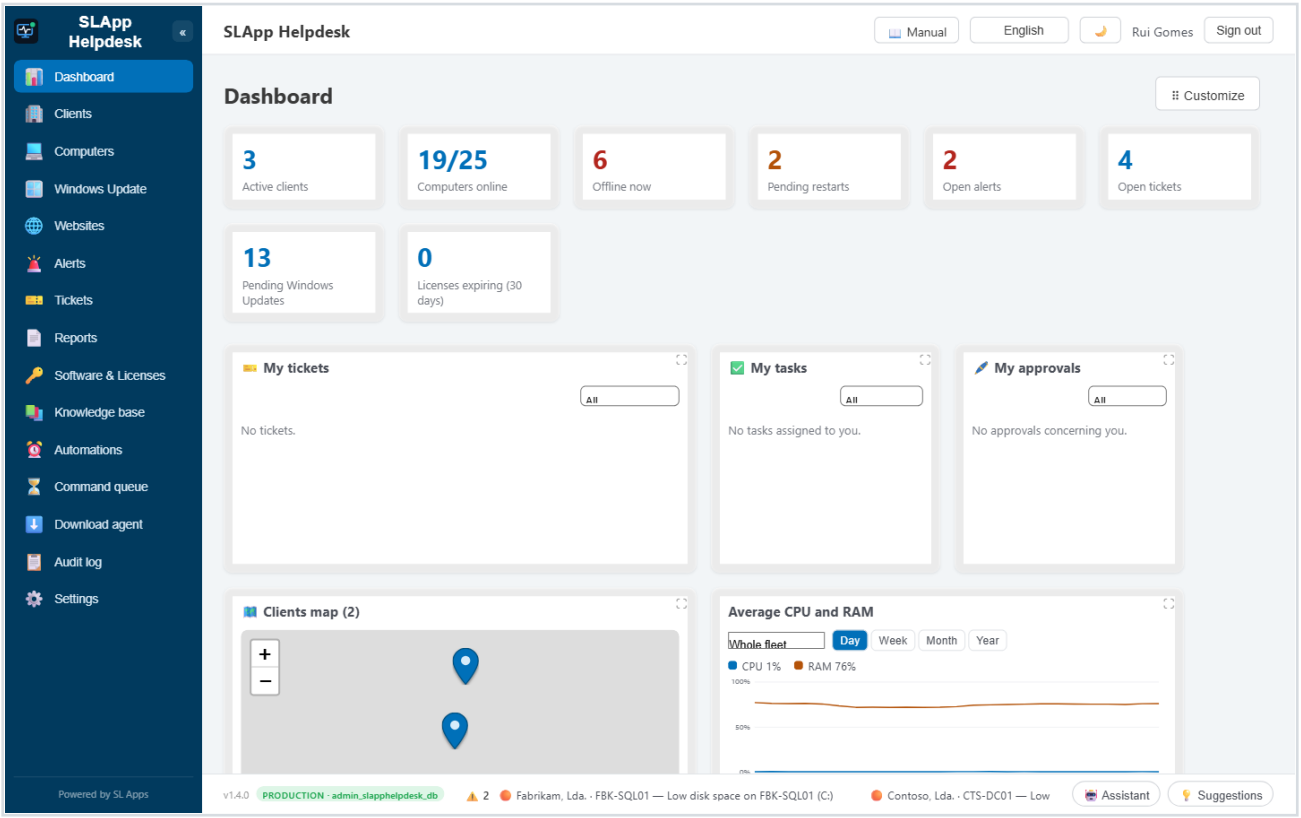
Contents

1. Getting around SLApp Helpdesk 1
2. Signing in and user approval 6
3. Dashboard 11
4. Clients 15
5. Computers 32
6. Alerts 51
7. Tickets 61
8. Staff changes 69
9. Agenda 73
10. Reports 79
11. Knowledge Base 88
12. Automations 91
13. Maintenances 94
14. Windows Update and patch policies 103
15. Vulnerabilities 107
16. Software packs (administrators only) 112
17. Command Queue 116
18. Infrastructure overview 119
19. Hyper-V 120
20. Active Directory 127
21. Exchange 132
22. SQL Server 135
23. Storage / iSCSI 139
24. Failover Cluster 144
25. DHCP 147
26. IIS 151
27. Backups 154
28. Third-party backups 193

29. Syslog 196
30. Network configs 201
31. Printers 205
32. Lifecycle 208
33. Websites 213
34. Domains 215
35. Microsoft 365 219
36. Download Agent and installing on computers 227
37. Software & Licenses (operator only) 233
38. Audit Log (administrators, or by explicit permission) 241
39. Transactions (administrators, or by explicit permission) 246
40. Settings (administrators only) 250
41. Help tools: AI Assistant, Suggestions and this manual 268
42. Getting help 273

1. Getting around SLApp Helpdesk

SLApp Helpdesk is a web application for IT support companies that combines remote monitoring and management (RMM) with professional services automation (PSA). On the RMM side it keeps an eye on your clients' computers through a small agent installed on each machine: hardware and software inventory, Windows Update status, outdated software with actively exploited vulnerabilities, the hardware lifecycle, alerts, remote commands and scripts, and, on Windows, scheduled backups of folders and SQL Server databases, a daily health check of domain controllers and of third-party backup software, and configuration backups of network devices. The server itself, with no agent involved, also watches your clients' websites, Internet domains and Microsoft 365. On the PSA side it gives you the day-to-day tools of a support business: clients and contacts, tickets, on-site visits, a knowledge base, reports and automations. This section explains the vocabulary the app uses and how the screen is organised, so the module chapters that follow make sense.



The main layout: sidebar on the left, header at the top, the active module in the middle and the footer bar at the bottom.

Key concepts

- Tenant (your company): the IT support company that owns this installation of SLApp Helpdesk. Everything you see (clients, computers, tickets, users, settings) belongs to your company and is invisible to any other company using the same service.
- Client: a company you provide support to. Contacts, tickets and assets are always attached to a client. Computers normally are too, but a computer whose agent has just been installed can arrive without a client: you then move it to the right one from the Computers module (Move).
- Contact: a person who works at a client. Contacts can be given access to the Client Portal, where they open and follow their own tickets. A contact is not a technician.
- User (technician): a member of your own staff who signs in to SLApp Helpdesk with a Microsoft work or school account. Users can be administrators (they see Settings and the Audit log) or regular technicians.

- Agent: the small program installed on each monitored computer (Windows, Linux or macOS). It reports inventory and status to the server and carries out the commands you send from the app. A computer only appears in the app once its agent has checked in.

The sidebar

The dark column on the left lists the modules of the application. Click a module to open it in the main area; the active one is highlighted. Clicking the SLApp Helpdesk logo at the top always returns you to the Dashboard. The first entries are modules of their own: Dashboard, Clients, Computers, Alerts, Tickets, Agenda, Reports and Knowledge base. The other modules are gathered in four groups: Automation & updates (Automations, Maintenance, Windows Update, Vulnerabilities, Software packs and Command queue), Infrastructure (Overview, Hyper-V, Active Directory, Exchange, SQL Server, Storage / iSCSI, Failover Cluster, DHCP, IIS, Backups, Third-party backups, Syslog, Network configs, Printers and Lifecycle), Cloud & web (Websites, Domains and Microsoft 365) and Administration (Download agent, Software & Licenses, Audit log and Transactions). Settings comes last. In this manual a module that sits in a group is written as Group › Module, for example Infrastructure › Backups.

- A group has no page of its own: clicking it opens one of its modules, the one of that group you used last since you opened the app (its first module the first time). While one of its modules is open, the group is highlighted and a submenu beneath it lists the modules of the group, with the open one highlighted; click another one there to switch. Clicking the group while you are already in one of its modules does nothing.
- When Settings is active, a submenu in the same style appears beneath it with its groups: Company, License, General, Users and Mails. You switch between groups in this submenu rather than with tabs inside the page.
- Settings, Audit log, Software packs and Backups are only shown to administrators (Transactions too, and only with the Contracts module of the Complete plan), and Software & Licenses only to the operator company (SL Apps itself); customer companies never see it. A group whose modules are all hidden from you is not shown, and a group left with a single module shows that module directly, as a normal entry: a technician without admin rights at a customer company, for example, sees Download agent on its own instead of an Administration group.
- The « button at the top of the sidebar collapses it to a narrow strip of icons (its tooltip reads Collapse); the same button becomes » (Expand) to bring the labels back. Collapsing gives wide tables more room; while collapsed, hovering an icon shows its name, whether it is a module, a group or an entry of a submenu.
- The Powered by SL Apps line at the bottom is only visible while the sidebar is expanded.

On a phone or a narrow window (up to about 900 pixels wide) the sidebar is hidden: the menu button at the left of the header (its tooltip reads Open menu) slides it in over the page, always with its labels, and tapping outside it closes it again. Tapping a module opens it and closes the menu; tapping a group or Settings opens it but keeps the menu open, so you can then pick the module, or the Settings group, you want in its submenu. While your company's license is suspended, the app stays on Settings › License and every entry of the sidebar is shown but disabled.

On a phone (a screen up to about 640 pixels wide) each row of tabs is a single button showing the current tab, for example Overview. Tap it to open the list of all the tabs, with the same icons and counters, and tap one to switch to it; tapping outside the list or pressing Escape closes it. On a detail page such as a computer or a ticket, the main action (Remote access on a computer) stays visible and the other actions of the header are under More, with the same names and confirmations.

The header



The header: Manual link, language switcher, theme toggle, your name and Sign out.

On the left, the top bar shows your company's Logo (or, without one, its Icon) and its name, as set in Settings › Company; SLApp Helpdesk is shown while they load, or when the company has neither, and a change saved in Settings › Company shows at once. On the right, a small set of controls is available from every module:

- Manual opens this user manual as a PDF in a new browser tab (tooltip: Open the user manual (PDF) in the current language). The PDF matches the language currently selected in the app (English, Portuguese or French).
- The language switcher shows a flag, the name of the current language and a small arrow (tooltip: Language / Idioma / Langue). Click it and choose English, Português or Français; the current language is highlighted in the list. Your choice is remembered in this browser; the app starts in English until you pick something else, regardless of your browser's own language. If you also choose a language in My profile › Options (or simply use the switcher while signed in), it is saved on your account and follows you to any browser, computer and the mobile app. Since version 1.14.81 the messages that come from the server, such as the reason an action was refused, and the alert messages follow this language too.
- Moon / sun switches between the light and dark theme (the tooltip reads Switch to dark theme or Switch to light theme). The theme is also remembered in this browser. Changing it while signed in also saves it on your account (see Options below).
- Your name, as reported by your Microsoft account, with your photo (or your initials). Click it to open My profile: there you add or change your photo (PNG, JPEG, WebP or GIF, up to 2 MB) and fill in your job title, department, phone and mobile, then Save profile. The Account card shows what comes from Microsoft (name and e-mail, updated at every sign-in), your access (Administrator or User) and your first and last access. Every user sees and edits only their own profile. Under Remote access the Account card also shows whether you Can connect without asking for authorization and Can turn off session recording; only an administrator can change them (for anyone else the boxes are greyed out with Only an administrator can change this).
- Sign out ends the session and returns you to the login page.

1. Click the language button in the header (flag plus language name).
2. Pick English, Português or Français from the dropdown.
3. The interface changes immediately; from now on the Manual link also opens the PDF in that language.

My profile: Options

Below your photo and contact details, My profile has an Options card with personal preferences. They are saved on your account, not in the browser, so they follow you to any browser, computer and the SLApp Helpdesk mobile app. Only you see them: an administrator who opens your profile with Edit profile does not see the Options card. Change what you need and click Save options; a notice confirms Options saved (or explains which field was not accepted). Language and theme are applied at once.

- Language: English, Português or Français. The app opens in this language every time you sign in, wherever you sign in. The language switcher in the header changes the same setting. If your account has no language saved yet, the app saves the one it is showing, Portuguese or French, when you open it (since version 1.14.81), so that your other browsers and the mobile app open in it too; a language already saved is never replaced, and an app showing English saves nothing.
- Theme: Light, Dark or Same as the system. With Same as the system the app follows the light or

dark mode of your computer or phone and switches by itself when that changes. The moon / sun button in the header changes the same setting.

- Start page: the module that opens after you sign in. The list shows only the modules you can actually open (according to your access and your company's plan), grouped as in the sidebar; Default (Dashboard) keeps the usual behaviour. If the module you chose is no longer available to you, the app simply opens the Dashboard.
- Email notifications, four boxes: A ticket is assigned to me (on by default; when someone else assigns you a ticket, never when you assign it to yourself), A client replies on a ticket of mine (on by default; when the client answers from the client portal on a ticket assigned to you, or confirms there an on-site visit you are a technician of), Critical alerts (off by default; the alerts of the highest severity in your company, as soon as they are raised, only those of the modules you have access to (see Who receives the alert emails in the Alerts section), without duplicating an email the alert rule already sends you) and Daily summary (off by default; every morning at about 08:00, Lisbon time: your open tickets, the ones with the SLA overdue and the ones assigned to you yesterday, sent once a day). The emails are written in your chosen language, or in the company's email language if you have not chosen one.
- Signature: up to 1000 characters (a counter shows how many you have used), for example your name, job title and phone. It is added at the end of your public replies to clients, never to internal notes.

> **i** When you write a public reply on a ticket and have a signature, the reply box shows an Include my signature box (ticked by default) with a faded preview of the signature. Untick it to send that one reply without it.

Installing the Android app

Until SLApp Helpdesk is published on Google Play, the Android app is installed from a download. Open My profile in a browser on your computer: the Mobile app card shows a Download APK button with the version and size of the current app, a QR code of the same download link and the installation steps. The card does not appear inside the mobile app itself. If it reads that the Android app is not available for download yet, no build has been published to your installation.

1. On the phone, open the download link: scan the QR code with the phone camera, or open My profile on the phone and tap Download APK.
2. When Android asks, allow installing apps from this source (the browser you used for the download). Android asks this only the first time.
3. Open the downloaded file and tap Install, then open SLApp Helpdesk and sign in with your Microsoft account.

Updates: when a newer version is published, the app shows a bar at the top reading New version X of the app is available, with an Update button and a ×. Update downloads the new APK through the phone browser; open the file when the download finishes and the new version installs over the old one, keeping you signed in. × hides the bar until the next version. The app checks again when you return to it, at most every 6 hours.

> **i** The app will later also be published on Google Play. On iPhone and iPad it is coming to the App Store; until then, open SLApp Helpdesk in Safari and add it to the Home Screen.

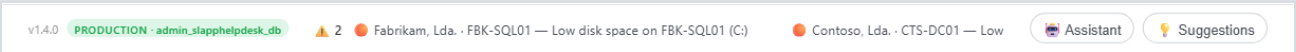
My profile: My vault

My profile also has a My vault card with your own personal password vault - the same vault you have in

the client portal, in the browser extension and in the agent's Password manager window (folders, entries, sharing, import), opened with the same master password. Click Open my vault and it opens inside the page; Open in a new window opens it in a window of its own, and Close puts it away. It is the vault of the contact that is you: an active contact with your e-mail, with portal access, in a client of your own company (for example your company's own client). If there is none, the card says so and takes you to Clients to create it; if there are several, choose which one. The vault session lasts 2 hours and opens only the vault; when it ends, open it again.

> **i** Nobody else, technicians and administrators included, can read your vault: it is end-to-end encrypted with your master password. Each opening is recorded in the Audit log. The card is not shown inside the mobile app.

The footer bar



The footer: version and environment on the left, the open-alerts ticker in the middle, Assistant and Suggestions on the right.

The bar at the bottom of the window stays visible whatever module you are in. From left to right:

- The application version (for example v1.2.3) followed by an environment badge reading Production or Local dev and, after a dot, the name of the database the app is connected to. This is the first thing to quote when you report a problem.
- The alerts ticker. When your company has open alerts, a warning counter shows how many there are and the alerts scroll continuously across the middle of the bar in the form client · computer — message (an alert that is not about a computer, such as a domain alert, leaves the computer out). Hardware lifecycle alerts are not scrolled one by one: they are summed up in a single item at the end, for example 12 hardware lifecycle alerts, while the counter still counts each of them. A red dot marks a critical alert and an orange dot any other severity. The list refreshes every minute and disappears entirely when nothing is open. Hovering it shows, for example, 3 open alerts — open the Alerts module, and clicking anywhere on the ticker jumps straight to the Alerts module.
- Assistant opens and closes the AI Assistant panel in the bottom-right corner of the window. You can ask it questions in plain language about your environment; the conversation is kept while you move between modules, and earlier conversations can be reopened from its History button (‹ Back returns to the chat). Inside the panel, × (Close panel) only hides it and keeps the current conversation, whereas Close conversation ends it and starts a fresh one; the maximize button maximises the panel and the restore button restores its normal size. The footer button is highlighted while the panel is open.
- Suggestions opens a small form where any signed-in user can send an idea to the SL Apps product team: choose a Type (Change suggestion or New feature), fill in a Subject and a Description, then click Send suggestion (it stays disabled until both fields are filled in; Cancel discards the form). A confirmation reads Suggestion sent and the team replies to your email if they need more details.

> **i** The AI Assistant and the suggestion form are described in more detail later in this manual.

Notices at the bottom right

Whenever you click a button that saves, deletes, sends, runs or starts something, a short notice appears in the bottom-right corner, just above the footer bar, saying what happened — for example Network scan started, Restart sent to PC-01 or Maintenance "Monthly cleanup" saved. Actions that finish later say they started and, where the page keeps track of them, show a second notice when they end (Network scan finished — 12 devices found); a command sent to a computer that is offline says it was queued for when the computer connects. A successful action shows a green tick and its notice

closes by itself after a few seconds; a red notice means the action failed and gives the reason returned by the server — it stays on screen longer. Since version 1.14.81 that reason is in the language of the app, except a text that another service (Microsoft, Stripe, the AI provider…) or the computer itself returns, which is shown as it came; when the server cannot be reached, takes too long to answer or is not available for a moment, the notice says so in plain words, for example No connection to the server. Try again in a moment. Pointing at a notice keeps it open, and × closes it. Up to four notices are shown at a time, the newest at the bottom. The notice of a new chat message from a computer also appears here: click it to open that conversation.

Right-click menus

Right-click any row of a list or table — in every module — to open a menu with the same actions that row already offers: its row buttons, its Actions menu, or what clicking it does (Open, Edit, Delete…). It is the fast way to act on a row without hunting for its buttons.

- Where a row has an identifier — a name, IP address, email, ticket number — the menu also offers Copy <thing>.
- The menu follows the same rules as the buttons: entries that are hidden or disabled there (for example because the computer is offline, or you are not an administrator) are hidden or disabled here too.
- Anything destructive (Delete, Remove, Revoke…) only opens the same confirmation as the row's own button; nothing is deleted straight from the menu.
- Use the arrow keys to move through the menu and Esc to close it. Right-clicking inside a text field keeps the browser's normal copy/paste menu.

> **i** A few rows whose own delete button acts immediately (for example attachments or custom fields) do not offer Delete in the menu, so nothing is ever removed without the same safeguards.

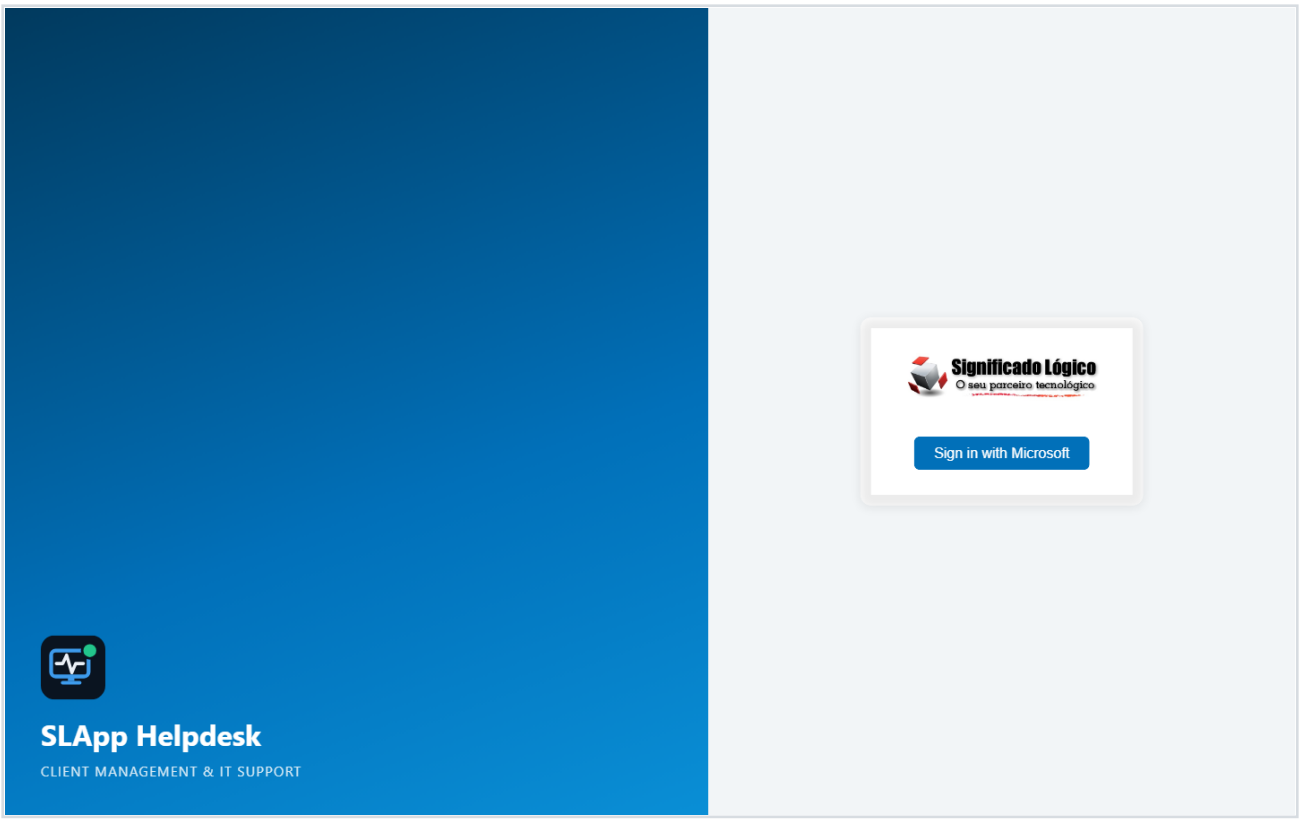
Moving between modules

Many screens link to each other so you rarely need to go through the sidebar. Clicking a computer's name wherever it appears (in an alert, on the Dashboard) opens the Computers module with that computer's detail sheet already open. Create ticket on an alert switches to Tickets with a new ticket pre-filled from the alert. The Dashboard cards for your tickets, tasks and approvals open the corresponding ticket directly, on the relevant tab. In every list, the toolbar of the tab or section you are in (its title, buttons, totals, filters and status buttons) and the table header stay fixed right under the page header while you scroll, so the controls and column names remain at hand on long pages (on very short screens the toolbar scrolls normally to leave room for the list).

> **i** Switching modules does not lose your work in the AI Assistant panel, but it does close any open form: finish or cancel what you are editing before clicking another module in the sidebar.

2. Signing in and user approval

SLApp Helpdesk has no passwords of its own. You sign in with your Microsoft work or school account, and the application decides what you are allowed to see based on the account it finds (or creates) for you. This section explains the login screen, what happens the very first time someone from your company signs in, why a new colleague may be held in a waiting state, and how an administrator approves, blocks or removes users.



The login screen: one button, Sign in with Microsoft.

The login screen

When you open the application without an active session you land on a page with the SLApp Helpdesk logo, the tagline Client management & IT support and a single button: Sign in with Microsoft. Clicking it sends you to the Microsoft login page; once Microsoft confirms who you are, you are sent straight back into the application. If your browser already has a Microsoft session (for example because you use Outlook on the web), the round trip is usually silent and you will not be asked for a password.

1. Open the application address in your browser.
2. Click Sign in with Microsoft.
3. Choose or type your work account on the Microsoft page and complete any multi-factor prompt your organisation requires.
4. You are returned to the application. Your name appears in the top-right corner next to the Sign out button, together with the Manual link, the language selector and the light/dark theme toggle.

> **i** Use the Microsoft account your company gave you, not a personal one. Personal Microsoft accounts (outlook.com, hotmail.com and similar) and guest accounts (an account invited into another organisation's Microsoft directory) never create a workspace. A personal account only gets in where a workspace already lists its domain as authorised, and a free email domain such as gmail.com or outlook.com can no longer be added to that list.

What happens the first time you sign in

The application looks at the email address of your Microsoft account and works out which company workspace you belong to. Several situations are possible:

- Your email domain matches the list Domains authorized to sign in of an existing workspace: an account is created for you in that workspace. If nobody there is an administrator yet (you are the very first person), you become an administrator immediately and can start working at once. Otherwise your account is created with the status Pending and you must wait for an administrator to approve

you.

- Your email address was pre-registered when the workspace was set up (the primary contact of a provisioned client receives a welcome email inviting them to sign in): your Microsoft account is linked to that pre-created administrator account and you are in right away.
- No workspace authorises your domain yet and you signed in with a work or school account of your own organisation: a workspace is created on the spot for your company with a free trial — 1 user and 5 computers, on the plan and for the number of days the service operator set for trials (the Complete plan and 30 days by default). You become its administrator at once, you receive a welcome email and the operator is notified automatically. If your company is not a customer of the operator yet, the new workspace is named after your email domain.
- Your company is already a customer of the service operator (its customer record lists your domain) but has no workspace yet: the workspace is created from that record instead of a new one — its name, tax ID, address and domains are copied — and you become its administrator, even when the record has other contacts. You are added to its contacts, or your contact is reactivated if it had been removed, as an ordinary contact without any special permission.
- No workspace can be created with a personal Microsoft account, a guest account or a free email domain; when your company's customer record already has a workspace (ask an administrator of that workspace to add your domain to its authorised domains), is waiting for a licence the operator is setting up, has been deactivated or lists another domain that a workspace already uses; or when several customer records or several workspaces list your domain (it is ambiguous). In those cases contact the operator; the screen shows Your email domain is not authorized to sign in to this application.
- The workspace has a user limit (a free trial allows a single user) and it is already full — for example because a colleague signed in first: you see This plan only allows 1 user(s) — contact an administrator to upgrade.

On every later sign-in the application simply recognises your Microsoft identity, refreshes your name and email from Microsoft, records the time as your last access and re-checks two things: that your workspace is still active and that your email domain is still authorised.

The waiting and blocked screens

If the application cannot let you in, you do not see the login page again; instead you see a short full-page notice with a Sign out button. Read the title and the sentence underneath it to know what to do:

- Your account is awaiting approval (hourglass icon), with the text Your account is awaiting approval by an administrator. Your account exists but an administrator of your company has not yet approved it. Ask them to do so under Settings › Users, then click Sign out and sign in again.
- No access (padlock icon), with the text Your account does not have access to the application. Contact an administrator. Your account was rejected or deactivated by an administrator.
- No access with Your email domain is not authorized to sign in to this application. Either your workspace no longer lists your email domain as authorised (see below), or you are signing in for the first time and no workspace authorises your domain or could be created for it — the reasons are listed under What happens the first time you sign in.
- Error loading session: This account no longer has access to the application. Your company's workspace has been suspended, typically because its licence expired or was disabled. It comes back automatically once the licence is active again.
- Error loading session: This plan only allows 1 user(s) — contact an administrator to upgrade. The workspace has reached the number of users its plan allows; an administrator must free a seat (deactivate someone) or upgrade the plan before you can be added.

- Error loading session: Multiple pending accounts match this email — contact an administrator. A rare set-up inconsistency (two pre-registered accounts with the same email); an administrator has to clean it up.
- Error loading session: … with any other text, for example Could not connect to the database. A technical problem, usually the database being unreachable; try again in a moment.
- Multi-factor authentication required (lock): your company requires multi-factor authentication and your sign-in did not use it, or Microsoft did not say how you signed in. Click Sign in again and complete the second step (Authenticator app, code, security key…), or Sign out. If Microsoft never asks for a second step, ask your Microsoft 365 administrator to enable MFA for your account. See Requiring multi-factor authentication in the Settings chapter.

> **i** These checks run on every request, not only at login. If an administrator rejects, deactivates or deletes your account while you are working, the application signs you out automatically at the next action instead of silently failing.

> **i** If a notice says Too many requests in a short time. Please wait a moment and try again., the application received too many requests from you, or from your network, in a short time (a script or a browser tab stuck reloading, for example). Wait a few seconds and carry on; nothing was lost.

Approving and managing users (administrators)

Only administrators see the Settings entry in the sidebar. Open Settings › Users; the first tab, Users, lists every account in your workspace — everyone who has ever signed in, plus any account pre-registered during set-up that has not signed in yet (the second tab, Groups, is about technician teams and is covered elsewhere). Accounts in the Pending state are always listed first so they are impossible to miss; the rest follow in name order. The table has the columns Name, Email, Status, Admin, Active and Last access, followed by a column of action buttons.

- Status shows one of three badges: Approved, Pending or Rejected.
- Admin shows Yes or No. Administrators can change settings, manage users and see the Audit log; everyone else is a regular technician.
- Active shows Yes or No. A deactivated user keeps their account but is shown the No access screen when they try to sign in.
- Last access is the date and time of the person's most recent request to the application; it shows — for an account that has never signed in.
- Approve and Reject appear only while the account is Pending.
- Reauthorize appears only on a Rejected account and sets it back to Approved.
- Make admin / Remove admin toggles administrator rights. Edit profile opens that user's profile in a window (photo, job title, department, phone and mobile); the change is recorded in the audit log.
- Deactivate / Reactivate toggles the Active flag without losing anything.
- Delete removes the account permanently after a confirmation dialog titled Delete user.

1. Ask the new colleague to sign in once with their Microsoft account. Their account is created with the status Pending and they see the awaiting-approval screen.
2. As an administrator, open Settings › Users. The pending account is at the top of the list.
3. Check the Name and Email columns to make sure it really is your colleague, then click Approve. Click Reject instead if the request is unexpected.

4. If the person should also manage settings and users, click Make admin.
5. Tell the colleague to click Sign out on the waiting screen and sign in again; they now enter the application.

> The application does not email administrators when a new account is waiting. If someone tells you they are stuck on the awaiting-approval screen, Settings › Users is the place to look.

> Prefer Deactivate over Delete when a technician leaves. Deleting is final: the person can no longer sign in and, although tickets, remote sessions and actions they worked on are kept, those records stop showing their name. You cannot delete your own account (the page shows You cannot delete your own account if you try), but nothing stops you from removing your own admin rights, so make sure another administrator exists before you do.

Authorised email domains (Settings › Company)

Under Settings › Company, below the company details, logo and icon, there is a field called Domains authorized to sign in. It is the list of email domains whose users may sign in to your workspace; it is also what lets the application recognise that a new colleague belongs to your company rather than somewhere else. The help text under the field sums up the rule: only emails from these domains can sign in; empty means no restriction; a new user from an authorised domain waits for approval in Settings › Users.

1. Open Settings › Company (administrators only).
2. In the Domains authorized to sign in field, type a domain such as yourcompany.com in the box (the placeholder suggests the format) and press Enter, type a comma, click Add, or simply click elsewhere — any of these adds it. You can paste a full email address or a web address: the part before the @, any http:// or https:// prefix and any trailing path are dropped automatically, and the domain is stored in lower case.
3. Each domain appears as a small chip prefixed with @. Click the × on a chip to remove it. A chip with a padlock and no × is a locked domain, which cannot be removed (see below); hovering it shows Company domain set when the company was created; it cannot be removed, and the same line appears below the chips.
4. Click Save at the bottom of the page. Saved next to the button confirms the change.

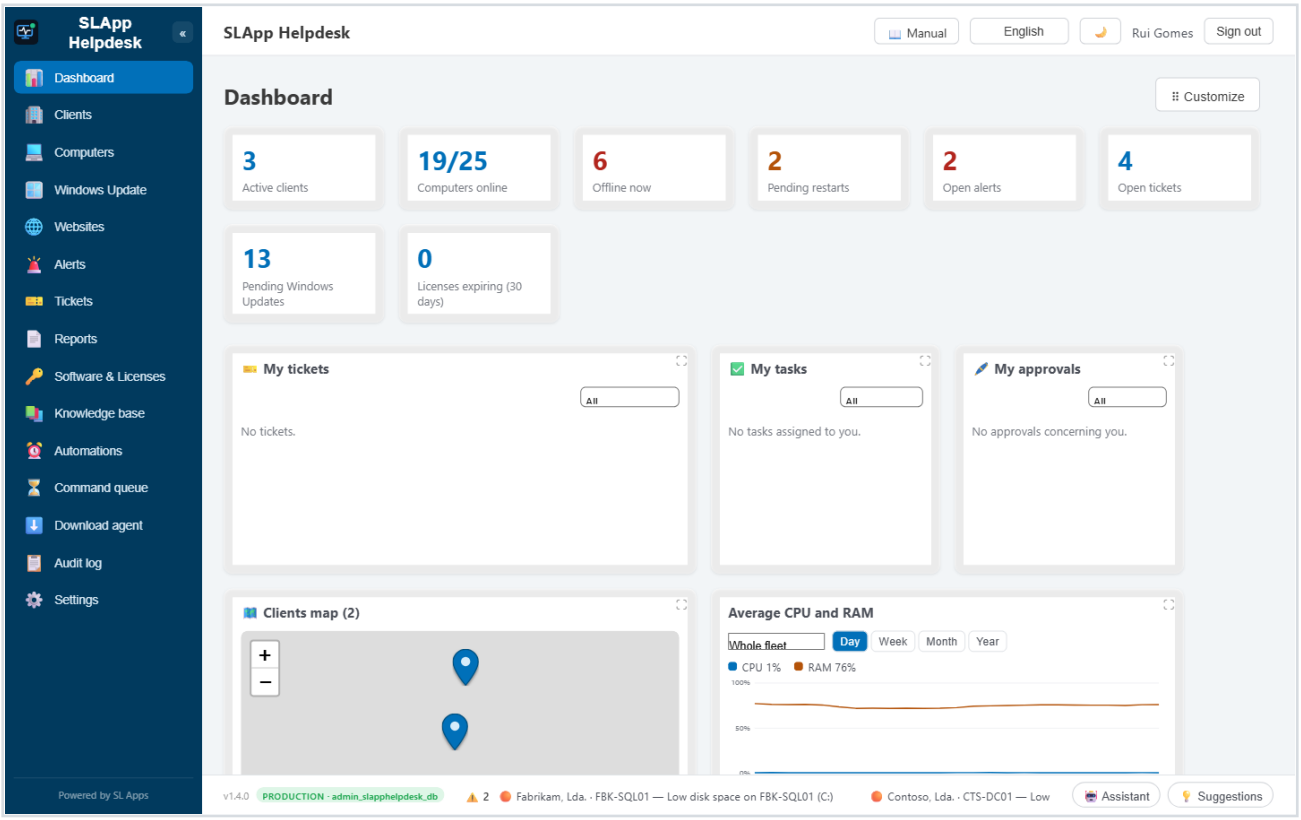
> Some domains are refused when you add them, and then nothing is saved: a free email domain such as gmail.com (the message above Save says … is a free e-mail domain and cannot be an authorized domain.), or a domain another company already uses, as one of its authorised domains or as the domain of another customer of the service operator (… cannot be an authorized domain: another company already uses it. Contact support if it belongs to your company.). Domains already in the list are not checked again, and any domain can be removed except a locked one.

> **i** Locked domains: the domains your workspace was created with — the ones in the list when it was created, from a license or from a self-service sign-up — are locked, except one that another company also authorizes. They show a padlock and always stay in the list, so a new colleague from those domains is always matched to your workspace; nobody in your company can remove them, not even an administrator, and a list with a locked domain cannot be emptied either: the application refuses to save it. Other domains are added and removed as before. Workspaces created before version 1.14.85 get the same lock once, automatically: the domains of their list that are also on their company's customer record with the service operator — nothing is added to the list. To change a locked domain, contact the service operator.

> **i** Keep at least one domain in the list for a workspace provisioned for your company (a locked domain does that for you). With the list empty, a colleague signing in for the first time cannot be matched to your workspace and may end up creating a separate one of their own or being refused. Existing users are only affected if the list is non-empty and their domain is missing from it: from the next request they are signed out with the not-authorised message. Each user can have one session in a browser and one in the mobile app at the same time (customer companies): opening the app in a second browser or on a second phone shows "Account opened on another device" on the first one, with Use here (take the session back) and Sign out. Tabs and pop-up windows of the same browser count as one session.

3. Dashboard

The Dashboard is the first screen you land on after signing in. It gives you a one-glance picture of the whole fleet you support: how many computers are online, what needs attention (alerts, pending restarts, Windows Updates), how the tickets are going, and a personal panel with the tickets, tasks and approval requests that concern you. Everything on it is read-only, but most items are clickable shortcuts into the Computers and Tickets modules.



The Dashboard: indicator tiles on top, cards on the grid below.

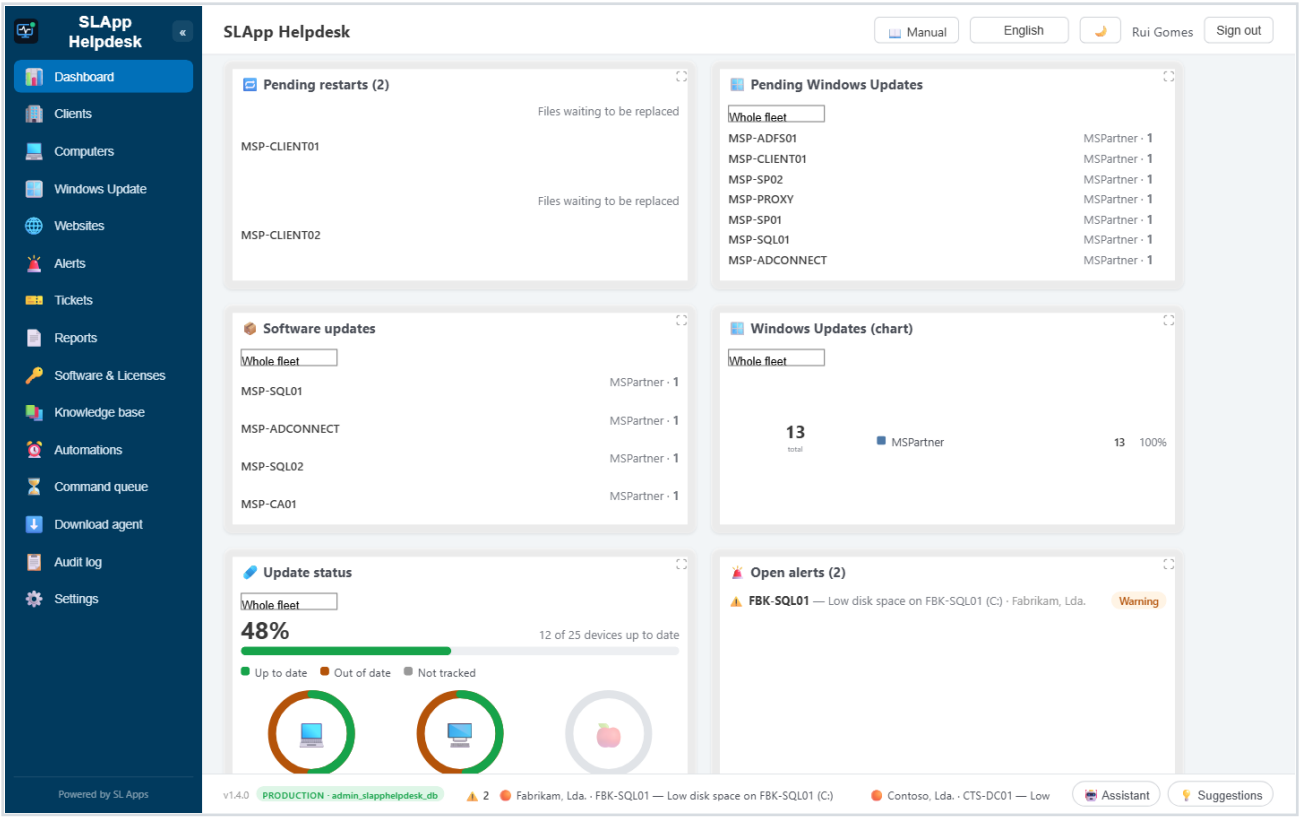
Indicator tiles

The top row shows eight counters. A computer counts as online when its agent has reported in the last two minutes.

- Active clients: number of active clients in your organisation.
- Computers online: online computers out of the total active computers (for example 48/51).
- Offline now: computers not reporting right now. Shown in red when above zero, green when everything is online.
- Pending restarts: computers whose agent reports that a restart is required. Amber when above zero.
- Open alerts: unresolved alerts, counting computer alerts, alerts from monitored network devices and the client alerts (domains, Microsoft 365, vulnerabilities, hardware lifecycle, network configurations, Active Directory and third-party backups). Website alerts are not counted here, nor listed in the Open alerts card. Red when above zero.
- Open tickets: tickets whose status is not Resolved or Closed.
- Pending Windows Updates: total number of Windows Updates waiting to be installed across the fleet (updates, not computers).
- Licenses expiring (30 days): active licenses that expire within the next 30 days. This counter is only meaningful for the operator organisation; other organisations always see 0.

Cards

Below the tiles, the cards sit on a grid of squares. Several cards have a client selector that starts on Whole fleet; pick a client to narrow the card to that client's computers or tickets. Lists scroll inside the card when they are longer than the card. Every computer name on the Dashboard is a link: click it to open that computer's detail page.



Further down the Dashboard: pending restarts, pending Windows Updates, software updates and the Windows Updates chart.

- My tickets: tickets assigned to you or to a group you belong to (up to 100, most recently updated first), each line showing the ticket number, subject and status. A ticket with a timer running comes

first and shows, in green in place of its status, the time running (to the second, since the earliest running timer) and who is timing it, for example 0:12:34 · Rui Gomes. The status filter also offers Counting time, which keeps only those tickets. The filter at the top offers All, Open, In progress, Pending client, Resolved and Closed. Click a line to open the ticket.

- My tasks: ticket tasks assigned to you, open ones first. Filter between All, To do and Completed; completed tasks appear struck through. Click a task to open its ticket on the tasks tab.
- My approvals: approval requests where you are an approver, pending ones first, each with a coloured badge (Pending, Approved, Rejected, Cancelled). Click a request to open its ticket on the approvals tab.
- Clients map (n): an OpenStreetMap view with one marker per client location that has coordinates; the marker uses the client's icon if one is set, otherwise a blue pin. Click a marker to see the client name, location, address and number of computers. If no location has coordinates, the card tells you to add them in the client's detail page, Locations tab.
- Average CPU and RAM: two lines, CPU and RAM, averaged over the selected scope. Choose Whole fleet or a client, and a period with the Day, Week, Month or Year pills (last 24h by hour, last 7 days and last 30 days by day, last year by month). The legend shows the latest values; move the mouse over the chart to read the exact values for a point in time. With too little data you see Not enough metrics in this period yet.
- Online by client: one stacked bar per client, green for online and red for offline, with the count shown as online/total (hover the bar for the online and offline figures). Computers with no client assigned are grouped under — no client —. The client selector reduces the card to a single client.
- Offline the longest: up to ten offline computers with how long ago each one last reported (for example 3 h ago or 12 min ago), or never connected for agents that never checked in. When nobody is offline the card says Everything is online.
- Agents by version: one bar per installed agent version with the number of computers on it; computers whose version is not known yet are grouped under ?. Handy for spotting machines that did not update.
- Tickets by status: one bar per status that has tickets (Open, In progress, Pending client, Resolved, Closed), sorted by count, with a client selector.
- Pending restarts (n): every computer waiting for a restart, with the reason reported by the agent (or the client name when there is no reason).
- Pending Windows Updates: up to 50 computers with updates waiting to be installed, most updates first, showing the client and the number of pending updates per machine, with a client selector.
- Windows Updates (chart): a donut of the same pending updates, with the total in the centre and a legend showing the count and percentage of each slice. On Whole fleet the slices are clients; after choosing a client the slices become that client's computers. Hover a slice or a legend line to highlight it.
- Update status: the percentage of devices with no pending Windows Updates, the X of Y devices up to date count, a progress bar, and one ring per category: Windows PC, Windows Server, Mac and Linux. Green is Up to date, amber is Out of date, a grey ring means no devices in that category. If you have Mac or Linux computers, their rings show Not tracked because the agent does not report patch status for them yet; they are left out of the percentage. A client selector narrows the card to one client.
- Software updates: up to 50 computers that have installed software with a newer version available, with the count per machine and a client selector.
- Open alerts (n): up to 30 unresolved alerts, critical ones first, each with an icon for the type (disk, memory, CPU, network, offline, and one for each type of client alert), the computer, the message,

the client and a Critical or Warning badge. Alerts from monitored network devices are prefixed with the device name. A vulnerability or hardware lifecycle alert links to its computer like a computer alert; the other client alerts show the item they are about (a domain, a Microsoft 365 user or license, a network device, a domain controller or a backup job) in place of the computer, as plain text, or the client's name when the alert concerns the client as a whole. Hardware lifecycle alerts always come after all the others, whatever their severity.

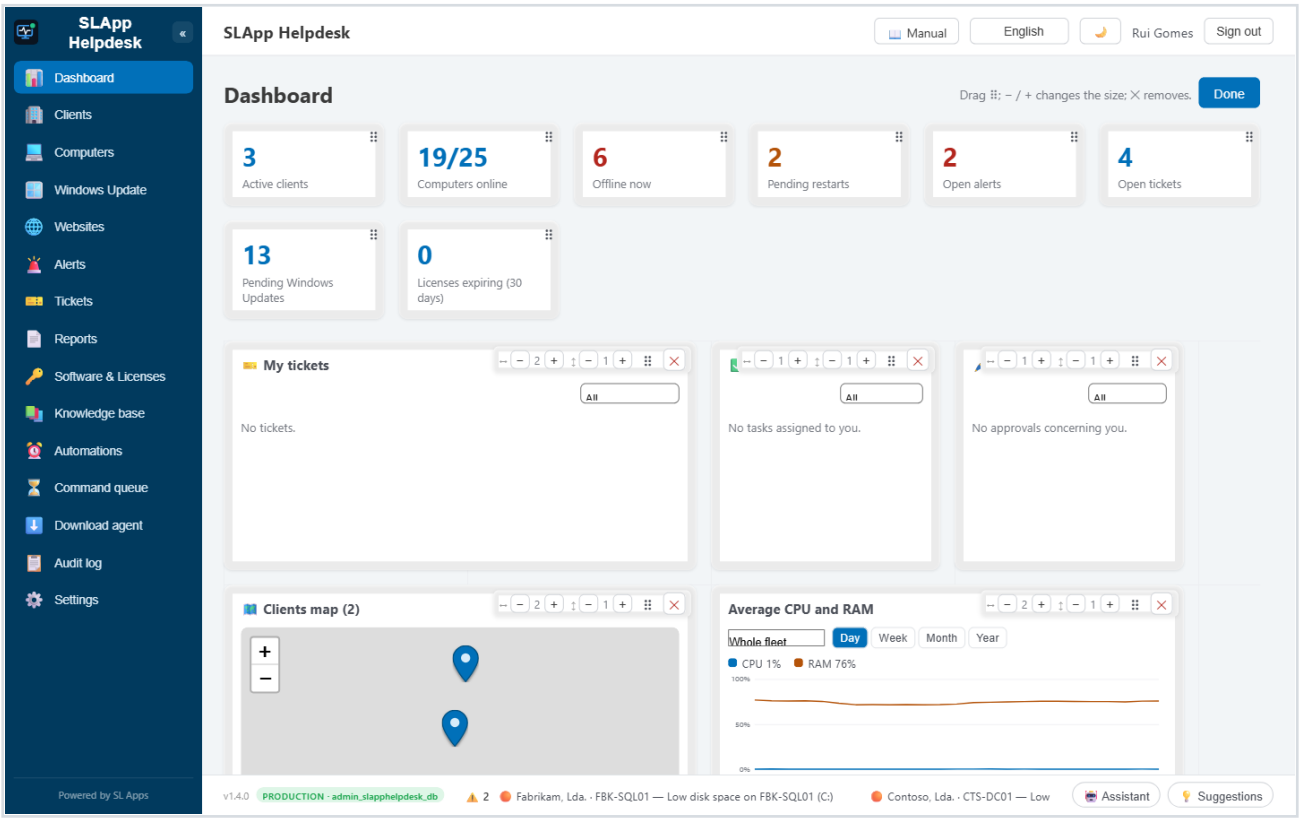
> ℹ Hardware lifecycle alerts (an operating system out of support, a warranty about to expire) stay open for months and can be dozens at once. So that they never push an operational alert out of sight, the Open alerts card lists them after every other alert and the footer ticker sums them up in a single item (for example 12 hardware lifecycle alerts); the Open alerts tile and the card's title still count each of them.

Maximizing a card

1. Click the maximize icon in the card's top-right corner (tooltip Maximize). The icon is not shown while Customize mode is on.
2. The same card opens in a large window, with all its selectors still working.
3. Click × Close, or click outside the window, to go back.

Customizing the layout

Click Customize in the page header to rearrange the Dashboard. The grid of squares becomes visible, each card gets a small toolbar, and the header shows the reminder Drag the grip; - / + changes the size; × removes. Click Done when you are finished.



Customize mode: grid, per-card toolbar and the Add card… list.

- Width and height: next to the horizontal arrow (width) and vertical arrow (height) markers, the - and + buttons change how many squares the card spans, and the current number is shown between them. Width goes from 1 up to the number of columns that fit on your screen; height goes from 1 to 6.
- Grip handle: drag the card anywhere on the grid. A dashed ghost shows where it will land; if that

spot is taken, the card slides down to the first free space in that column and the other cards stay put.

- ×: removes the card from the Dashboard.
- Add card…: this drop-down only appears while at least one card is removed. Pick a card to bring it back; it is placed at the bottom-left of the grid so it never lands on top of another card.
- Indicator tiles: in Customize mode each tile also gets a grip handle so you can drag it to a new position in the top row.

1. Click Customize.
2. Drag cards by their grip handle, resize them with - / +, and remove the ones you do not need with ×.
3. Use Add card… to restore a card you removed earlier.
4. Click Done to leave Customize mode. The layout is kept automatically.

> **i** Your layout, tile order and removed cards are stored in the browser you are using, not in your account. On another computer or browser you start again from the default layout. There is no reset button; to get the default back, clear the site data for this app in your browser.

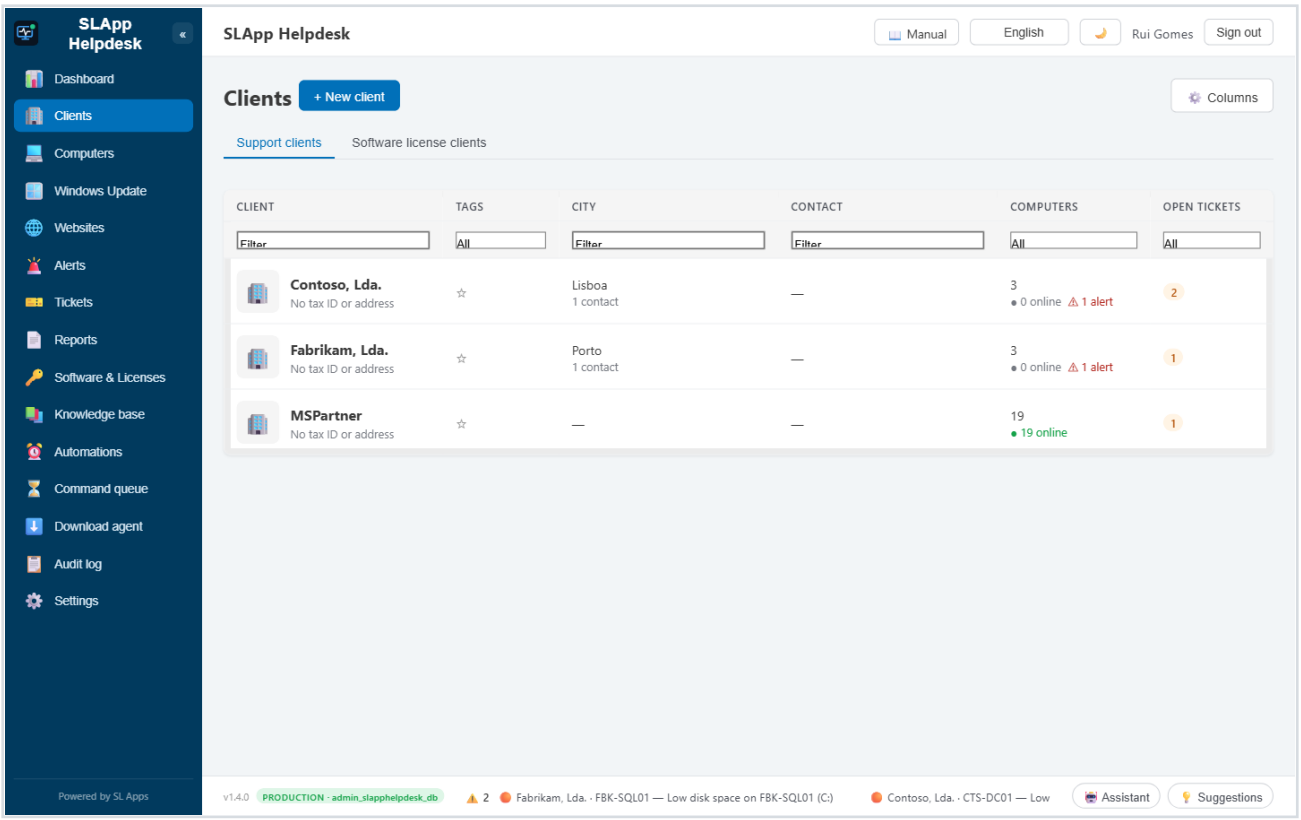
> **i** The Dashboard refreshes itself every minute: the thin bar in the page header counts down to the next refresh (hover it to read Time until the next automatic refresh), and the refresh button next to it (Refresh now) reloads the figures at once. Both are hidden while you customize the layout.

> **i** Cards added in future versions of the app show up automatically: each one takes its default position, or the first free spot below it if that position is already occupied by one of your cards, so you never miss a new one.

4. Clients

The Clients module is the home of every company you support. Each client record holds its identification, contacts, physical locations, computers, SNMP devices, alerts, tickets, contracts, a password vault and licenses, the websites and domains you monitor for it, plus an optional connection to the client's Microsoft 365. Almost everything else in the application (tickets, computers, portal logins) hangs off a client, so this is usually the first screen you fill in for a new customer.

The client list



The Clients list with the count of computers, online machines, alerts and open tickets per client.

Opening Clients shows one row per active client, sorted by name, with a + New client button at the top right. Click anywhere on a row to open that client's sheet. There is no search box: the row of filters under the column headers narrows the list instead, a click on a header sorts by that column, and Columns chooses the columns and their order (Tax ID, Address, Postal code, SLA, Notes and Created can be added there).

- Client: the client's icon (or a building placeholder), the name and, on a second line, the Tax ID, address and SLA name when they exist (otherwise "No tax ID or address"). A Self-service badge marks clients created automatically when someone signed in with an unrecognised Work/School Microsoft account; a ConsultisMed badge marks clients created by a ConsultisMed clinic activating its license.
- City: the city, plus "N contacts" when the client has contacts.
- Contact: the client's general phone and email.
- Computers: the number of active computers, how many are online (a machine counts as online if its agent reported in the last two minutes) and, when there are any, the number of open alerts.
- Open tickets: tickets that are not resolved or closed.

Creating and editing a client

The New client and Edit client forms are the same. Only the name is mandatory (Save stays disabled until it is filled in); everything else can be completed later.

- Name *, Tax ID, Phone, Email, Address, Postal code, City: the identification shown on the list and on the Data tab. The address, postal code and city also feed an automatic "Sede" (head office) location, see Locations below.
- Email domains: type a domain and press Enter, a comma or Add (leaving the field also adds what you typed); each one becomes a removable chip. Whatever you type is normalised (lowercased, and a pasted email or URL is reduced to its domain). Any Microsoft account with an email in one of these domains automatically gets access to the Portal, and the domains also drive the Microsoft 365 user

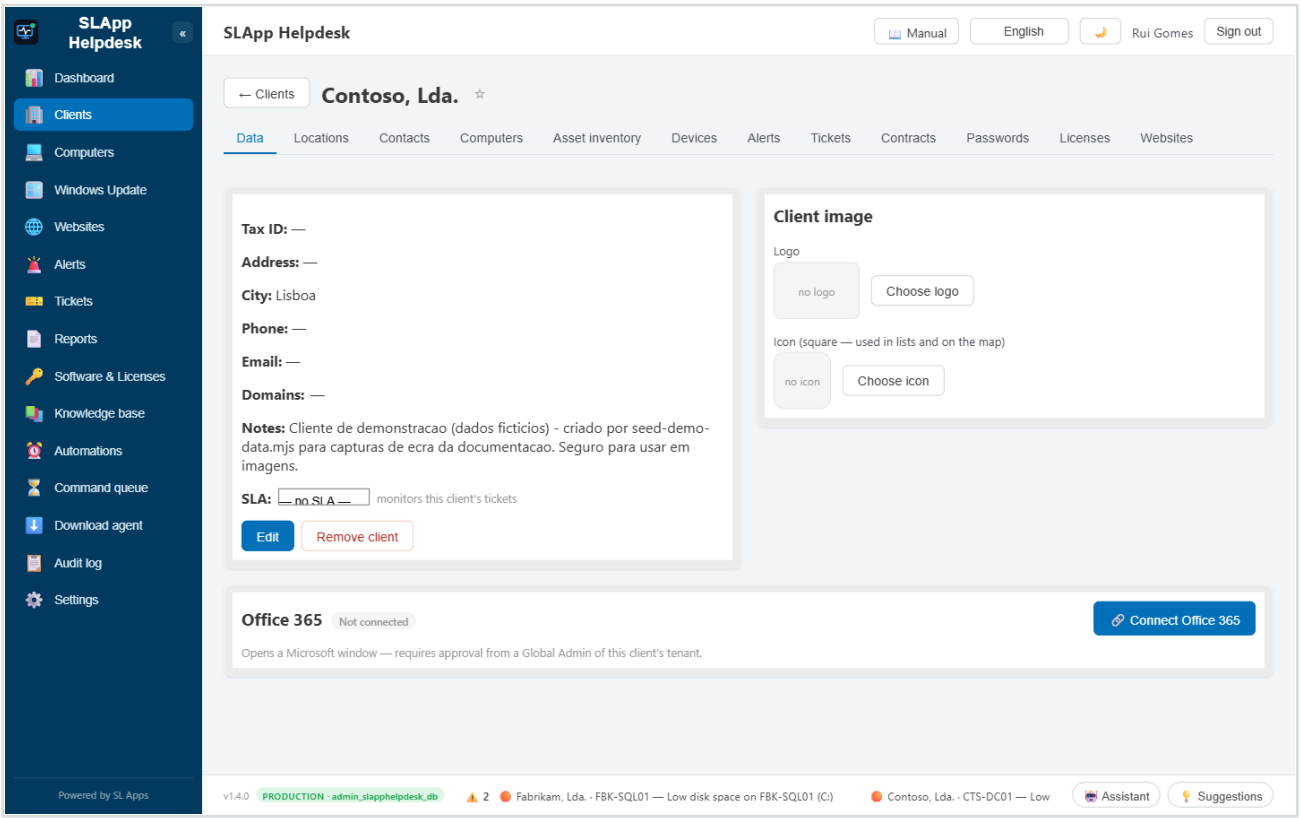
import filter and are proposed by Suggest in Cloud & web › Domains.

- Only the primary contact can open tickets in the Portal: when ticked, only the contact marked as Primary contact can open new tickets from the Portal.
- Notes: free text.

1. Click + New client.
2. Fill in at least Name *; add the tax ID, address and contact details you have.
3. Add the client's email domains (one at a time, Enter adds each).
4. Tick Only the primary contact can open tickets in the Portal if the client wants a single point of contact for support requests.
5. Click Save. The new client appears in the list; open it to add contacts, the logo and the rest.

> The logo, icon, SLA and Microsoft 365 connection are not in this form. They live on the Data tab of the client sheet and save immediately when changed.

The client sheet and the Data tab



A client sheet: the header with the logo and name, the row of tabs and the Data tab open.

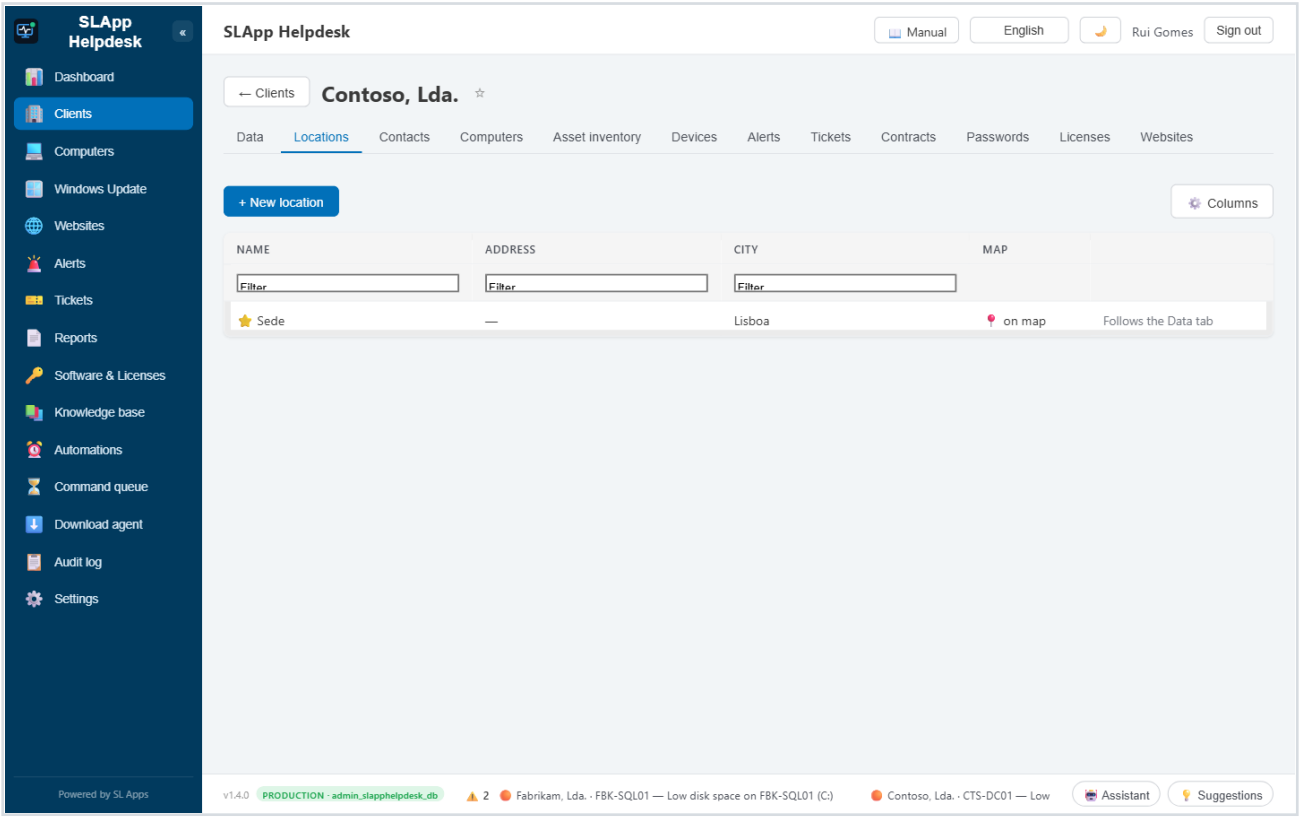
The sheet header shows a ‹ Clients button (the browser's back button also closes the sheet), the client's logo (or its icon when there is no logo) and its name, followed by the tabs: Data, Locations, Contacts, Staff changes, Computers, Asset inventory, Devices, Alerts, Tickets, Contracts, Passwords, Licenses, Websites, Domains and Microsoft 365.

The Data tab shows the identification fields read-only (Tax ID, Address, City with the postal code, Phone, Email) with an Edit button, the domain chips, the notes and an SLA drop-down ("— no SLA —" or one of the active SLAs from Settings). Choosing an SLA saves straight away and that SLA then monitors this client's tickets. Next to it, the Client image card has two uploads: Logo (wide, shown in the sheet header) and Icon (square — used in lists and on the map). Choose logo / Choose icon picks a file (the button reads Replace logo / Replace icon once there is one) and Remove clears it; images are

resized in the browser and saved at once, with no Save button. The Microsoft 365 card sits below (see the dedicated section).

Locations

A client can have several physical locations (head office, stores, branches). Each one is placed on the Dashboard map, so keep the addresses accurate. The table lists the primary location first and shows Name (a star marks the primary location), Address, City (with the postal code), Map ("on map" when coordinates exist, otherwise "no coordinates") and the Edit and Remove buttons (removal is confirmed first).

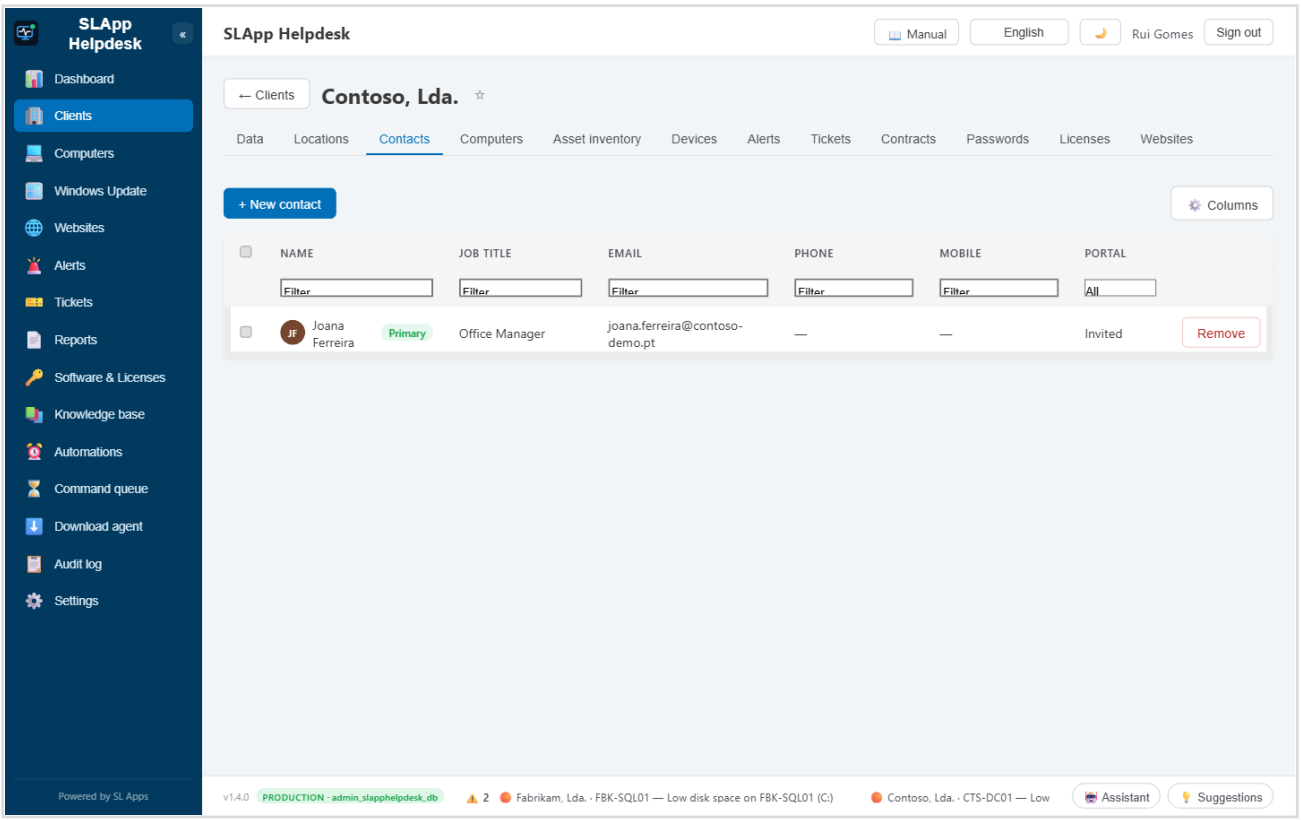


The Locations tab of a client.

One row is special: "Sede" is created automatically from the address on the Data tab when the client is first saved and is refreshed every time the client is edited, so it cannot be edited or removed here; the row simply says "Follows the Data tab".

1. Click + New location.
2. Enter the Location name (for example "Head office" or "Amora store"), the Address, Postal code and City.
3. Tick Primary location if this is the main site.
4. Click Save. On save, the map position is computed from the address; if the address cannot be resolved the location is still saved, just without coordinates, and you can edit it later to try again.

Contacts



The Contacts tab: primary contact badge, portal status and bulk selection.

Contacts are the people you deal with at the client, and also the accounts allowed into the client Portal. The table shows an avatar (a photo synced from Microsoft 365, or coloured initials), Name (with a Primary badge for the primary contact, or a Staff changes badge for another contact allowed to request staff changes), Job title, Email, Phone, Mobile and Portal. The Portal column reads "Linked" once the contact is tied to a Microsoft account — which happens immediately for contacts imported from Microsoft 365, or when a Portal sign-in is matched to the contact — "Invited" while portal access is allowed but no Microsoft account is tied yet, and a dash when portal access is off. The Columns button lets you choose which columns to show; Billing and Buys packages are available there and start hidden, with for the contacts that have Can see the billing in the portal or Can buy credit packages in the portal. Click a row to edit it, or use Remove at the end of the row (confirmed first).

1. Click + New contact.
2. Fill in Name *, Job title, Phone and Mobile.
3. Enter the Email (used to link the Portal login); it must match the Microsoft account the person will sign in with.
4. Leave Portal access ticked if the person should be able to use the Portal; untick it for people who must not see tickets.
5. Tick Primary contact for the main point of contact (relevant when "Only the primary contact can open tickets in the Portal" is on).
6. Tick Can request staff changes to let the person ask for new employees and leavers in the Portal (see the Staff changes chapter). The primary contact always can: for one, the box is ticked and cannot be changed.
7. Tick Can see the billing in the portal to let the person see the client's billing in the Portal: the invoices of its contracts with their PDFs, and the credit packages bought with their invoices. These contacts are also the ones who get the invoice PDFs by email (see the Transactions chapter).
8. Tick Can buy credit packages in the portal to let the person request or buy credit packages for the client's credit contracts (see Packages requested in the portal below); each request is made in that

contact's name. The person also needs Portal access.

9. Click Save.

> The two billing boxes do not depend on Primary contact, and the primary contact needs them too. Without Can see the billing in the portal, it no longer sees the Invoices card in My company; without Can buy credit packages in the portal, it still sees its credits but cannot get a package (when packages are on sale, the card tells it to ask you for that permission); with neither, it does not see the package purchases either. A contact that is not the primary contact and has either box ticked gets a Billing tab in the Portal, when your plan includes the Contracts module, with what it is allowed to see. When version 1.14.45 was installed, both boxes were ticked once for the contacts that were primary at the time, so nothing they saw disappeared. New contacts, including those imported from Microsoft 365 or created on their first Portal sign-in, have both unticked, and making a contact primary later does not tick them. A Portal session opened from the agent never has these permissions, and View as shows exactly what the contact viewed can see. Each change is recorded in the Audit log on the Contact.

> Administrators editing an existing contact also see Can use "View as" in the portal. It lets that contact preview the client portal as any other contact of every client of your organisation, which is meant for your own staff who help clients through the portal. New contacts never have it; contacts that do carry a View as badge next to their name, and removing it takes effect at once. Every use is recorded in the Audit log on the Contact, and only an administrator can change the email or the Portal access of a contact that has it.

> To remove several contacts at once, tick their checkboxes and click Remove selected (N). Removal is confirmed first and is logged in the Audit log. The bulk buttons act only on the selected rows you can see: the header checkbox ticks or unticks only the rows visible under the current column filters, and a contact ticked earlier that a filter now hides is neither counted in (N) nor removed or moved.

Move to another client. When a person changes company (for example between two clients of the same group), move the contact instead of removing it and creating it again: the photo, the contact's own vault, its language, its Microsoft link and its View as permission go with it. Tick the contacts and click Move to another client (N), next to Remove selected (N), or right-click a row and choose Move to another client for that contact alone.

1. In the Move N contacts to another client window, type in Search client… and pick the Target client. The list shows the active clients of your organisation except the current one.
2. Under What happens, read the impact worked out for the contacts you picked; only the lines that apply are shown. Emails that are not in the email domains of the target client (when it has any) are listed below it.
3. Click Move. A toast confirms "N contacts moved to <client>" and the list reloads. Cancel closes the window without changes.

What changes for each moved contact:

- It joins the target client as a regular contact: never primary and without Can request staff changes, Can see the billing in the portal or Can buy credit packages in the portal (they were about the billing of the previous client; under What happens, the window counts the contacts that lose them). The primary contact of the target client is not changed.
- Computers of the previous client no longer have it as their contact, and assets of the previous client are no longer assigned to it: the assignment ends today in the Assignment history, and that period stays visible on the asset of the previous client.
- Its monthly report deliveries of the previous client are removed (recorded in the Audit log), and its pending requests to connect the previous client's Microsoft 365 are dropped.

- Vault: it loses access to the folders shared with it by people who stay at the previous client (pending invitations are deleted, active shares are revoked and the folder is marked so that its passwords are changed). Company folders it owns stop being offered to everyone in the company, though the people who already have them keep them. If the target client has company folders, they are offered to it.
- Portal: from its next request it sees the target client; a portal session opened from the agent ends.
- Each move is recorded in the Audit log on the Contact, as moved, with the previous and the new client.

The history stays with the previous client, untouched: tickets (with the same contact and requester), comments, ratings, visits, approvals, staff changes, time entries, KB articles and reports. An existing ticket of the previous client can still be edited and saved with that contact, but the contact can no longer be chosen anew on a ticket, or as the owner of an asset, of the previous client. If, meanwhile, one of the selected contacts was removed or moved by someone else, nothing is moved and a message asks you to refresh the list.

> ℹ Microsoft 365: a contact imported from the previous client's Microsoft 365 keeps its Microsoft link, so a new user import in the previous client updates it where it is now and never duplicates it. A contact without that link would be created again in the previous client if someone imports the same person from its Microsoft 365; the window warns about it.

Staff changes

The Staff changes tab lists the client's requests for a new employee or for an employee leaving, newest first: the ones its contacts sent from the Portal and the ones a technician recorded here with + New employee or + Employee leaving. A click on a request opens its ticket on the Staff change tab, where the Microsoft 365 and SLApp steps are carried out one at a time. The tab, the request forms and the steps are described in the Staff changes chapter.

Computers and folders

The Computers tab lists the computers linked to this client, already grouped, with the same tools as the global Computers module. At the top, the "Folders of this client:" bar shows the existing folders as chips; click + New folder to add one (a small window asks for the name), or click the x on a chip to delete it after a confirmation. Deleting is harmless: the computers using it keep the folder name, it just disappears from the drop-down.

- Group by: Folder (the default) or Location; the choice is remembered per client. The next drop-down filters to a single folder or location (All folders / All locations), and the search box ("Search (name, IP, OS, user…)") looks through every visible field of the computer.
- Each group heading is collapsible and your collapsed groups are remembered. Computers without a folder or location always appear last, under "No folder" / "No location".
- Columns: Name (click to open the computer sheet), OS, Local IP, Last user, Agent (with an "outdated" badge when an older agent is installed), Software ("N to update") and Windows Update ("N pending") — both badges are clickable and open that computer's Software tab — Alerts, Status (Online/Offline).
- Row buttons: remote access, Actions (opens the computer's remote actions), Move (to another client or folder) and the pencil button, Edit name/folder/notes, which also lets you set Contact (who uses this PC) and Location (where this machine is).
- Bulk selection: tick several rows (or Select all) and a bar appears with "N selected". Select all ticks or unticks only the computers visible under the current filter and search, and the bar counts and acts only on the selected computers that are visible: pick the destination client, optionally tick Set folder

and choose a folder, then click Move N to client; Actions on N opens the bulk actions window for those machines; Cancel selection clears the choice. The client drop-down starts at "— no client —": moving with that value leaves the computers unassigned to any client (and clears their contact), so pick the client first.

> ℹ If the tab says "No computers linked", install the agent from Administration › Download agent in the sidebar, then link the machine to this client with Move in the Computers module.

Asset inventory, Devices and Alerts

Asset inventory lists this client's equipment records; assets generated automatically from computers have their name, type and serial number fed by the agent, while status, owner, warranty and custom fields are free to edit. Assets also carry a purchase price and supplier, and keep an event history: log a breakdown, a repair with its cost, a maintenance or a note — events can be added and deleted here. The available event types are configured in Settings › General › Assets.

Open an asset (from this tab, or from the Asset sub-tab of a computer) to see its page. In Contact (owner) you choose which of this client's contacts has the equipment; contacts of other clients are not offered, and the server refuses a contact that does not belong to the asset's client. On the page of an existing asset, below the custom fields, the Assignment history table keeps who had it and when, newest first: Contact, From, To (the current holder shows a Current badge) and Recorded by (who made the change). The name is copied into the history, so it still reads the same if the contact is later renamed or deactivated. The Client Portal shows the same table, read-only, on the page of the asset; only the staff app can change who an asset is assigned to.

When you change the contact, a date field appears under it: Assigned on when you pick a person, or Returned on when you clear the dropdown to take the asset back. Leave it empty to record today, or enter an earlier date if the hand-over already happened. The date cannot be in the future, nor earlier than the last change already in the history. It also appears when you create an asset with a contact already chosen.

Assets that already had a contact when the history was introduced got a first line dated with the day the asset was created in SLApp, because the real hand-over date was never recorded. That date is shown with a ~ in front (hover it to read the explanation) and is only approximate.

An asset generated automatically from a computer follows the computer. If the computer is moved to another client, the asset moves with it and loses its contact, since that person belongs to the previous client, and that period ends in the history. The history only lists people of the client the asset belongs to now, so the new client never sees the previous client's contacts; if the asset ever returns to the first client, those earlier lines show again.

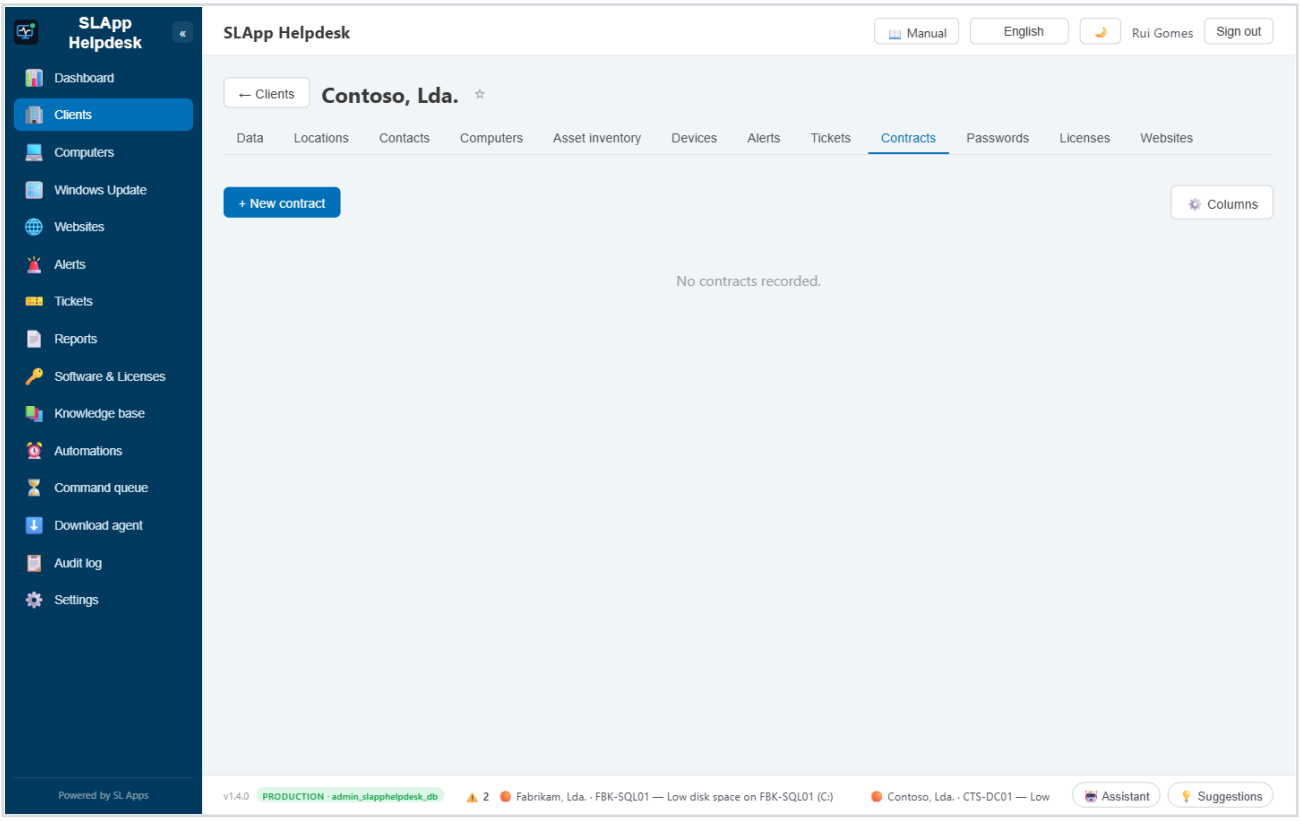
Devices gathers the SNMP devices (NAS, routers, printers and other equipment) polled by this client's agents. Each device is configured on the sheet of the computer that polls it, so this tab is read-mostly: every card shows the device name, Online/Offline, an Inactive badge when the device is switched off in its configuration, the open alert count when there is one, the model, a clickable chip with the polling computer (opens that computer's sheet), then the address and port, the type and when it was last updated. Poll now asks the agent to poll immediately instead of waiting for the interval (unavailable on inactive devices). The last polling error, if any, is shown in red; when the device did not answer it says what to check — for example that SNMP must be enabled on the device and that the community, SNMP v2c and the UDP port must be right (see When a device shows No SNMP response in the Computers section for the full checklist). A device never polled reads "No data yet — the agent will poll shortly." When data exists you see Temperature, Firmware (with an update badge when newer firmware is available), Uptime, CPU, RAM and System status, plus a Show disks, volumes, RAID and UPS (for NAS), Show supplies and detail (for printers) or Show interfaces and detail link; a printer's card also

shows a bar for each supply (see the Printers section). While a device is offline none of these readings is shown; its last known state is kept, so a reboot or a dropped interface during the outage is still reported when the device answers again. The list refreshes itself every 15 seconds. The Generate alerts checkbox at the top of the tab is the same company-wide switch for every SNMP device except printers (administrators only; see Turning alert generation on or off per module in the Alerts section).

Alerts shows this client's computer and device alerts with the same look as the Alerts module: Computer, Message, Severity (Critical or Warning), Opened at and Status. It also lists the alerts of the websites assigned to this client (down, slow response, SSL certificate and hosting expiry), with the site name at the start of the message, and the alerts of the server's own checks for this client: its domains, its Microsoft 365, vulnerabilities, hardware lifecycle, network configuration backups, Active Directory and third-party backups. Open alerts are shown by default with an "N open" badge; tick Include resolved to see history. Resolve closes an alert by hand: any technician can resolve a computer, device or website alert, but on the alerts of the server's own checks only administrators can. The list refreshes every 30 seconds.

Tickets and Contracts

Tickets is a read-only list of this client's tickets (#, Subject, Status, Priority, Created at); to work on them, go to the Tickets module.



The Contracts tab.

Contracts records the agreements you have with the client: Name, Type, Value (excl. VAT), Periodicity, Hours, Start and End, with an Expired badge once the end date has passed and a restart badge next to the name of a contract with automatic billing (hover it to read Automatic billing). Click a row to edit or use Remove (confirmed first). In the form you fill in Name *, Type (the active contract types configured in Settings › General; an old type that has since been deactivated stays selectable), Value (€, excl. VAT), Periodicity (One-time, Monthly or Yearly), Automatic billing (see Automatic billing of a contract below), Start, End (optional), the Hours bank fields when the contract includes support hours, the Travel and kilometres fields (see Travel and kilometres of on-site visits below) and Notes.

Hours bank of a contract

A contract can include a number of support hours: its hours bank. The billable time your technicians log on the client's tickets counts against it, period after period; the hours used beyond those included are extra hours, billed at the contract's own price, and the administrators are warned by email when the hours are running low and when they run out. The Hours bank part of the contract form has four fields, and a contract whose hours are left empty simply has no hours bank:

- Hours included per month, Hours included per year or Hours included (in total): the label follows the Periodicity. A Monthly contract includes these hours every calendar month, a Yearly one every contract year, and a One-time contract (or one without a periodicity) once, for its whole duration. Empty or 0 means no hours bank.
- Price of an extra hour (€): what each hour beyond those included is billed at, which gives the Extra value in Reports › Hours & billing. Leave it empty when extra hours are agreed case by case: they are still counted, only without a value.
- Warn at (% used): the share of the hours included at which the administrators receive a warning email (80 when left empty). Another email follows when the hours run out, at 100%; see Hours bank warning emails below.
- Round each entry up to: No rounding, or 5, 10, 15, 30 or 60 min. Every time entry that counts against this contract is rounded up on its own, so with 15 min an entry of 20 minutes counts as 30. Only the hours bank and the billing use the rounded time; the ticket keeps the time exactly as logged.

Only billable time counts (the Billable box of each entry on a ticket's Time tab), and only finished entries: a timer that is still running counts once it is stopped. An entry belongs to the day it started, the day its timer was started or it was added by hand. A contract is in force on the days between its Start and End dates (either can be empty; a one-time contract, or one without a periodicity, that has no Start date is in force from the day it was created), and only the time of those days counts against it. Within them, the periods follow the Periodicity, with dates in UTC: Monthly counts each calendar month; Yearly counts contract years from the anniversary of the Start date (calendar years when there is no Start date); One-time, or no periodicity, counts a single period from the Start date, or from the day the contract was created, up to the End date, open-ended when there is no End.

Which contract an entry counts against is set on each ticket, on its Time tab: automatically, the client's contract with an hours bank in force on the day of the entry (when several are, the one that started most recently), or a contract chosen by hand for that ticket. Billable time that no contract with an hours bank covers is billed apart, as time outside any hours bank (see Billable time and the hours bank in the Tickets chapter).

The Hours column shows each contract's hours bank in the current period: a small bar and the hours used out of those included, for example 7.5 h / 10 h, followed in red by the extra hours (+2 h) once more than the hours included has been used. The bar is green below the contract's warning percentage, amber from it and red at 100%; hover it to read the exact percentage used and the warning level. A contract whose dates do not cover today shows only its hours, in grey (for example 10 h included), and a contract without an hours bank shows a dash.

Click the Hours cell, or right-click the row and choose See the hours used, to open the Hours bank window of that contract (a click anywhere else on the row still opens the contract form). The window opens on the current period or, for a contract whose dates do not cover today, on its last period when it has ended and on its first one when it has not started yet. For a monthly or yearly contract, the buttons either side of the period (Previous period and Next period) move from one period to the next, never before the contract starts nor into the future; a one-time contract has a single period. Four figures sum up the period, Included, Used, Remaining and Extra hours (the latter with its value when the contract has a price for an extra hour), above the usage bar, and the table lists the time entries counted in it:

Date, Ticket, Technician, Time counted and Notes. An entry that was rounded up has an up arrow next to its time; hover the time to read how much was logged and how much is counted. Click an entry to open its ticket on the Time tab.

Hours bank warning emails

Whenever the time of a ticket changes (a timer stopped, time added, edited, deleted or marked as billable or not, the ticket's contract changed, or the ticket closed while timers were running), the app checks the contract that ticket's time counts against today. When the hours used in its current period reach the contract's Warn at (% used), every active administrator of your company with an email address receives an email titled Hours bank at 80% (with the contract's own percentage), naming the client and the contract, the hours used out of those included, the period and, when there are any, the extra hours so far (with their value when the contract has a price for an extra hour). When the hours used reach 100%, the email reads Hours bank at 100% and says that the hours included have run out and that further time is billed as extra hours.

- Each warning goes out once per contract and period: while the hours keep going up, a level already reported is not sent again, and every new period starts with no warnings sent.
- Once the hours have run out, only the 100% email is sent. When a single change takes a contract from below its warning level straight past 100%, the administrators get the 100% email alone; the warning-level email never follows it.
- Changing the hours included (a top-up, for example), Warn at (% used) or Round each entry up to clears the warnings already sent for the current period, so the contract warns again the next time its time changes, if it is still at or past a level. Only a change of the value in effect counts: typing 80 into an empty Warn at (% used), which already stands for 80, or emptying it again, changes nothing.

> **i** The hours bank is always worked out from the time entries and the contracts as they are now. Changing an entry's minutes or its Billable box, choosing another contract on a ticket, or editing a contract's hours, dates or rounding updates the Hours column, the Hours bank window and Reports › Hours & billing at once, also for periods that have already ended; removing a contract moves its time to another contract in force on those days, or outside any hours bank.

Travel and kilometres of on-site visits

When a technician finishes an on-site visit (see the Agenda chapter), SLApp logs on the visit's ticket the time spent at the client and, as a separate entry, the travel time, and keeps the kilometres driven. The Travel and kilometres part of the contract form says how the contract bills them. It only matters for a contract with an hours bank: the travel of a visit that no contract with an hours bank covers is billed like any other time outside an hours bank, and its kilometres are listed without a value.

- Bill travel time: ticked by default. Unticked, the travel time counted against this contract is never billed: it does not count against the hours and stays out of Reports › Hours & billing, like time that is not billable.
- Travel counts against the hours included: ticked by default, and only available while Bill travel time is ticked. Ticked, travel counts against the hours bank like any other time, with the contract's rounding, its warnings and the client report. Unticked, it stays out of the hours used and is billed apart, at the contract's Price of an extra hour (with no value when the contract has none).
- Price per km (€): the price of each kilometre. The kilometres of a finished visit count against the same contract as the visit's time, whether or not its travel time is billed; with the field empty they are listed but not billed.

In the Hours bank window, travel counted in the hours has a car icon before its time; hover it to read Travel time of an on-site visit. Travel billed apart and the kilometres, with their value, appear in the

Travel and Kilometres columns of Reports › Hours & billing.

Automatic billing of a contract

Tick Automatic billing, under the Type, Value and Periodicity fields of the contract form, and the app creates a transaction awaiting payment at the start of each period of the contract: its value plus VAT, at the rate set in Settings › General › Connectors › Payments (23% by default). The hint under the box shows that rate to administrators. The transactions are listed in Administration › Transactions (see Contract transactions in the Transactions chapter), where they are marked as paid or cancelled and where their invoice PDF is attached, and the client finds them in the portal. Automatic billing is part of the Contracts module of the Complete plan.

- Only administrators can tick or untick the box. A technician sees it greyed out with the contract's current state and the hint Only an administrator can switch automatic billing on or off., and can still edit the rest of the contract.
- It needs a Value greater than 0 and a Periodicity, and a One-time contract also needs a Start. Until then the box is greyed out with a hint saying what is missing; the server refuses it too, also when the value of a contract that is already billed automatically is cleared.
- Monthly: one transaction per calendar month, created on the 1st, or on the contract's Start when it starts later in that month, always for the full value, never prorated. Yearly: one per contract year, on the anniversary of the Start (on 1 January when there is no Start). One-time: a single transaction, on or after the Start.
- Nothing is ever billed backwards and no period is billed ahead: ticking the box bills from the period in progress that day onwards, and the form then says Automatic billing since <date> - earlier periods are never billed. Unticking it stops new transactions and keeps the ones already created; ticking it again later bills from the period in progress at that moment.
- Periods the contract does not cover, before its Start or after its End, get no transaction. A cancelled transaction is never created again for the same period. When the periodicity or the Start of a contract changes, a new period that would overlap one already billed is skipped, so the same days are never billed twice.

On a monthly contract with an hours bank, the transaction of each month also carries the previous month's extras of that contract, the same figures as in Reports › Hours & billing: Extra hours (the hours used beyond those included, at the Price of an extra hour), Travel billed separately and Kilometres (at the Price per km), each only when it has a value and only for months from the one in which automatic billing was turned on. A month's extras are only ever billed once. When such a contract ends, the month after its last one gets a transaction with those extras alone, if it had any. Yearly, one-time and credit contracts are billed for their value only: the packages of a credit contract are already transactions of their own.

The app checks the contracts every hour, the first time a few minutes after the server starts. When a check creates transactions, a single email lists them all, with the client, the contract, the period, the amount excl. VAT, the VAT and the total: it goes to the addresses in Send the package requests to (Settings › General › Connectors › Payments) or, when that field is empty, to the administrators. Each transaction created is recorded in the Audit log.

Credit contracts

Instead of an hours bank, a contract can work with credits: in its form, set Model to Credits (Hours bank is the default). The client buys credit packages, and every billable time entry on its tickets debits credits: the Credits per hour of the entry's intervention type times the time counted with the contract's rounding, worked out per entry to 0.01 credit (half a hundredth rounds up). Credit contracts are part of the Contracts module of the Complete plan. Create the intervention types and the packages first, in

Settings › General › Clients (see Intervention types and credit packages in the Settings chapter): until there is at least one type, the form says so.

- Model: Hours bank or Credits. A credit contract has no hours included (choosing Credits clears them) and needs a Start, today by default. The model can only change while the contract has no time and no packages yet; after that it is locked, and the way to change it is to end the contract and create a new one.
- Low balance warning below: the balance under which the administrators are warned by email; leave it empty for no low-balance warning (the warning for a negative balance is always sent). Tick Also warn the client to send the warnings to the client's primary contacts as well. See Credit balance warnings below.
- Default intervention type *, required, and Intervention type for travel, optional (without it travel uses the default type). A time entry saved without a type, such as a timer stopped when its ticket was closed, gets the contract's default type (the travel type for travel) automatically, at the rate that type has at that moment.
- The rounding of each entry stays in the form, and the travel options become one drop-down, Travel: Travel debits credits (the travel time debits credits like the work), Bill travel separately (with the field Hourly price of travel billed separately, which values that travel in Reports › Hours & billing) or Free. Changing the rounding or the travel option recalculates every debit of the contract, past ones included, as the form recalls.
- Clients can buy packages in the portal, ticked by default, and Packages for sale in the portal: All active packages, or Only these packages and a tick per package. See Packages bought in the portal below.
- Only administrators can change the model, the warning threshold, Also warn the client, the intervention types and the portal options; the other technicians see them read-only.

The balance is never a counter that the app keeps: it is worked out every time from the packages (purchases and positive adjustments) and the debits (time entries and negative adjustments), so editing, deleting or retyping an entry, or voiding a package, shows at once. Each package is valid for its validity in months from its purchase date, or never expires; the credits that expire first are spent first, even from a package bought later. The window shows Valid until, the last day a package can be used. When the debits go past the credits, the balance becomes negative, In debt, and the debt is taken from the next package bought. After the contract's End the client keeps using its credits for as long as it has valid packages with credits left. A contract whose balance is not zero cannot be removed: void or adjust its credits first.

On the Contracts tab, the Hours column of a credit contract shows its balance with its state, OK, Low or Negative, in words and in colour. Click it, or right-click the row and choose Credits and movements, to open the Credits window of the contract. Four figures sum it up, Balance, In debt, Next expiry and Used this month; Packages lists every package with its Package, Date, Valid until, Credits, Remaining and Status (Active, Used up, Expired, Not yet valid for one dated in the future, or Voided, struck through); and Movements, between the From and To dates, lists Date, Type (Package bought, Adjustment, Time debited or Expired), Description (the package, the adjustment and its reason, or the ticket number and subject: click it to open the ticket on its Time tab), Intervention type, Time, the Credits added or taken and the Balance after it. A time entry with no intervention type is marked No intervention type — debited 0.

1. Buy package (administrators; in the Credits window or on the row's right-click menu): pick a Package, whose credits, price and validity are shown under it, or Custom package with its Name, credits, price and validity. The Purchase date is today by default; it can be earlier, but not before the contract's start, or up to 31 days ahead (the package is then Not yet valid until that day). Add

Notes if you want and save. A double click never records the package twice.

2. Adjust credits (administrators): enter the Credits (+ or -) and the Reason, which is required, and the date. Positive adds credits, like a package, and can have an expiry date (leave it empty for credits that never expire); negative removes them, like time debited.
3. Void (administrators), on a package or an adjustment: give the reason and confirm. Packages and adjustments are never deleted: a voided one stays in this window, struck through, and leaves the balance and the portal.

Packages requested in the portal. Once an administrator ticks Sell credit packages in the client portal in Settings › General › Connectors › Payments, the contacts of a client with a credit contract who have Can buy credit packages in the portal (see Contacts above) can get packages in the portal (see Buying a credit package in the Client Portal Guide). How the client pays is set in the same card, under How clients pay for credit packages: Invoice (no online payment), the default, or Online payment (Multibanco, MB WAY, card). With Invoice, the portal only records the request and no payment provider is needed: the client clicks Request package, the administrators get an email to issue the invoice, the client gets one confirming the request, and the request waits as Awaiting payment, marked Invoice, in the Portal orders table of the contract's Credits window. On the Contracts tab, the Hours column of the contract shows N awaiting payment, and Settings › General › Clients › Credit packages lists, above the packages, the Requests awaiting payment of every client, each with Open to go straight to its Credits window. A request awaiting payment never expires, and a contract can have at most 3 at a time.

1. Mark as paid (administrators), once the client has paid the invoice: in the Mark the request as paid window, enter the Invoice number and a Note if you want (both optional) and click Mark as paid. The package is added at once with the credits and the price of the request, its validity counted from today; it is marked via portal, the request shows Paid with Marked as paid by, the name and the date, and the client gets a receipt by email. A double click never adds the package twice. If the contract was removed or is no longer a credit contract, no package is added and a message says why.
2. Cancel request (administrators), for a request that will not be paid: enter the Reason, which is required, and click Cancel request. No credits are added and the client gets an email saying the request was cancelled, without the reason; the request stays in the table as Cancelled, with who cancelled it, when and why. A paid request cannot be cancelled: void its package instead.

The email of each request goes to the administrators unless Send the package requests to, in the same Payments card, lists other addresses. Administration › Transactions (administrators only) lists every package of every client in one place, requested in the portal or registered by the team, with Mark as paid and Cancel on the requests awaiting payment.

With Online payment, the client pays in the portal by Multibanco or MB WAY through your Eupago account or by card through your Stripe account: the money goes to your account, and the package's price is the total the client pays. A package is only added once the payment is confirmed. In the Credits window it is marked via portal, with its Payment reference, the state of the Payment and who it was Requested by, and the Portal orders table lists every order made in the portal: waiting for payment, expired, failed, or paid but in need of attention. A Multibanco reference can be paid for 3 days, an MB WAY request for 5 minutes and a card payment page for 24 hours; a payment that still arrives after its order expired adds the package all the same. Changing the way clients pay only affects new requests: the ones already made keep theirs.

> ℹ With Online payment, a payment received for an order the team had cancelled, or for a contract that was removed or is no longer a credit contract, adds no package: the administrators get an email and the order shows why. On a cancelled order, Create the package now adds it; otherwise refund the client in Licenses › Payments. Refunding a portal payment there voids its package, while voiding a portal package here does not refund the client. Voiding the package of a request paid by invoice does not issue a credit note either: handle it in your invoicing.

Credit balance warnings

Whenever the time of a ticket changes, a package is bought or voided, the credits are adjusted or the contract is edited, and every 6 hours for the credits that expire on their own, the app checks the balance of the client's credit contracts. There are two levels: Low, when the balance drops below the contract's Low balance warning below, and Negative, when it drops below zero. The approved administrators get the email, in the language set for emails, with the client, the contract, the balance, the threshold and the next expiry; with Also warn the client ticked, the client's primary contacts get one too, in the client's notification language, with a link to the portal and no prices nor technicians' names.

- Each level is sent once while the balance stays in it, and is armed again when the balance comes back: to the threshold or above for Low, to zero or above for Negative. A package that leaves the balance still low sends nothing new; one that lifts it above the threshold means the next drop warns again.
- When the balance jumps straight below zero, only the Negative warning goes out, never the Low one after it.
- Changing the threshold, the model, the dates, the rounding or the travel option of the contract starts the warnings again from scratch.

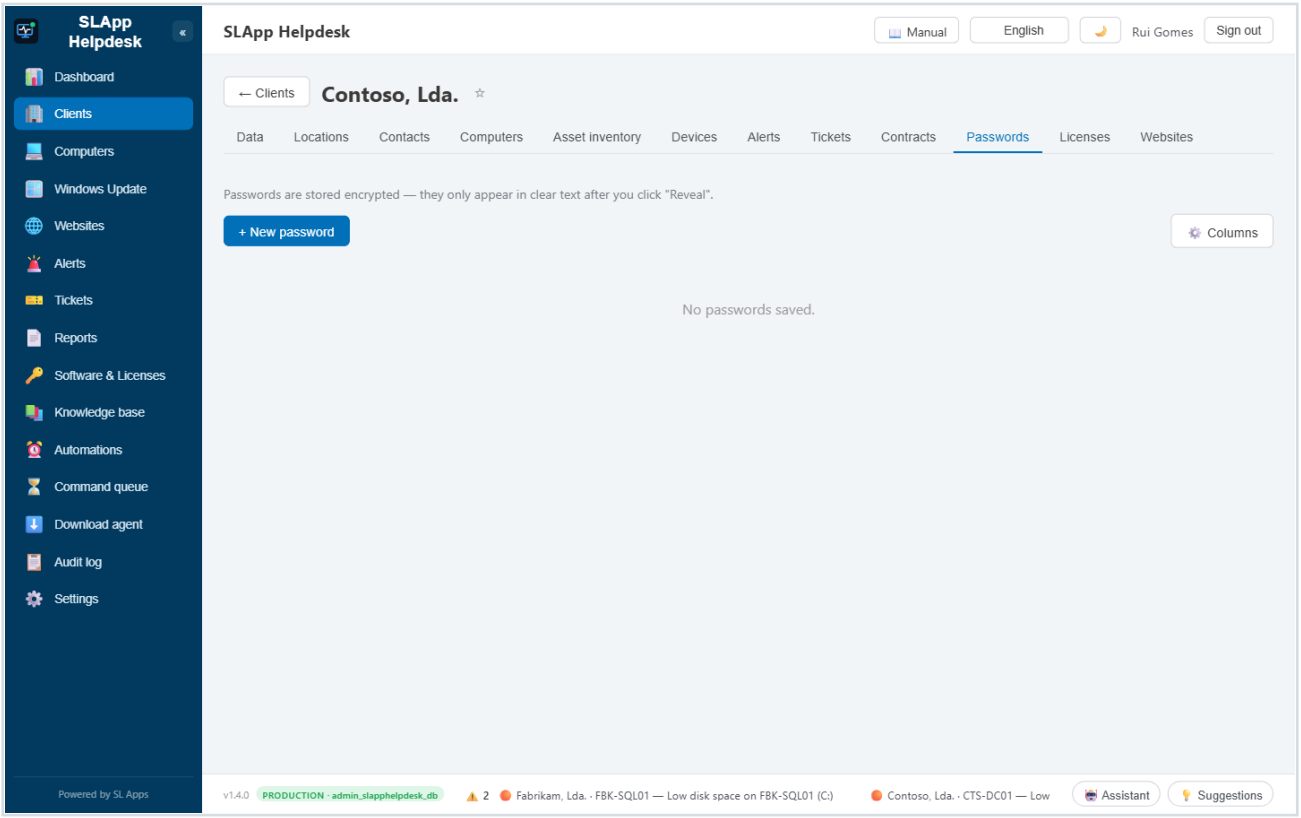
Transactions of the client (administrators only)

Below the list of contracts, administrators see a Transactions card with every transaction of this client: the credit packages, requested in the portal or registered by the team, and the transactions created by the automatic billing of its contracts. It is the same list as Administration › Transactions, kept to this one client, with the same rows, columns, statuses, totals, detail windows and actions (see the Transactions chapter). Technicians who are not administrators do not see the card, and neither do organisations whose plan does not include the Contracts module, exactly as with the Administration page. Amounts include VAT.

- There is no Client filter and no Client column: every row belongs to this client. The other filters stay, Period, Type, Status, Method / origin and Search; the Period starts at Last 12 months, and Clear filters goes back to it.
- Licence subscription invoices are never listed here, not even in the service operator's organisation: they stay in Administration › Transactions. The Type filter therefore offers All types, Credit packages and Contracts.
- The actions are the same: on a contract transaction, its details, Mark as paid, Cancel and the invoice PDF (attach, view, replace or remove); on a credit package, Mark as paid and Cancel on a request awaiting payment, Open for its contract's Credits window, and the invoice PDF. After each action the card is refreshed, and so is the list of contracts above it, with its credit balances. Buying a package or adjusting the credits from the list of contracts refreshes the card too.
- Export CSV downloads the rows shown, in a file whose name includes the client's name. The columns you choose for this card are saved apart from those of Administration › Transactions, so changing one never changes the other.

Passwords vault

Passwords are stored encrypted and only appear in clear text after you click Reveal; the list itself never contains the secret. Columns: Title, Username, Computer (or a dash for client-wide entries), Password (masked), Updated, and the Reveal/Hide, Copy and Remove buttons. Copy puts the password on the clipboard without displaying it; click the title, username or computer cell to edit an entry.



The Passwords tab: the client's encrypted vault.

1. Click + New password.
2. Give it a Title * (for example "Local admin" or "Router") and the Username.
3. Type the Password * or click Generate for a strong 16-character one. The strength indicator checks for at least 12 characters, lowercase, UPPERCASE, a digit and a symbol.
4. If the client has computers, choose one in Computer (optional) or leave "— Client-wide —".
5. Add Notes if useful and click Save.

> i When editing, the password field shows "New password (leave blank to keep the current one)": saving with it empty keeps the stored secret. Passwords attached to a computer also appear on that computer's Passwords tab.

Licenses and payments

What you see here depends on who you are. Every tenant gets Computer software: all license keys stored on this client's computers (Windows, SQL Server, Office and so on), with Computer, Product, Edition, Key and Files columns. Keys marked "auto" were detected by the agent; rows highlighted in red and marked "no longer used" are keys the agent stopped detecting on that machine. Reveal key shows a key on demand, then Copy and Hide; to edit a key or attach files, open the computer sheet › Licenses.

Only the operator tenant (the company running the platform) also sees the sub-tab bar with SL Apps licenses, used to assign licenses of its own products. The table shows Product, Endpoint, Key, Seats, Users, Total, Activation and Expiration, with a Payment button on each row; click a row to edit it (the

product cannot be changed afterwards). + Assign license asks for Product *, Endpoint (domain of the application that uses this license), Key (type one or click Generate), No. of seats and Price per seat (€), No. of users and Price per user (€), Activation, Expiration and Periodicity (Monthly, Semiannual or Yearly); prices are pre-filled from the product and the form shows the resulting Total per period. Provision ConsultisMed clinic creates the clinic tenant in ConsultisMed for this client: the form asks for the Clinic app domain (required) and an Email domain (for staff login; defaults to the client domains), then Provision creates a demo license of 1 user + 1 seat if the client has none and returns the ConsultisMed tenant id and the Kiosk code for the reception PCs. Running it again on a client already provisioned simply shows the existing tenant, and the license row then carries a "ConsultisMed tenant linked · Kiosk code" note. Below the table, the Payments section lists payment requests (Date, Description, Method, Value, Reference, Status). Payment on a license row (pre-filled with the product and total), or + Generate ad-hoc payment, creates a Multibanco reference or an MB Way request (Description, Method, Value (€), and MB Way mobile number for MB Way) that you then pass on to the client. Check refreshes a pending payment; Refund (administrators only, on paid payments) asks for a Reason * and, for Multibanco, a Destination IBAN *.

Microsoft 365 connection

At the bottom of the Data tab, the Microsoft 365 card shows the state of the connection to the client's Microsoft 365 as a badge — Connected, Needs reconnect, Failing, Old connection or Not connected; hover it for the reasons, which are also listed under the card — with the tenant's name and default domain. Connecting lets SLApp read the client's users, licenses, MFA and mailboxes every 6 hours, and lets you act on its users from the Microsoft 365 tab, import the tenant's users as contacts (with their photos) and collect BitLocker recovery keys. Connect Microsoft 365, Reconnect and Disconnect are for administrators; other technicians see Ask an administrator to connect, and once the client is connected they can import users and work with the BitLocker keys too.

1. On the Data tab, click Connect Microsoft 365. A Microsoft window opens and the button reads "Awaiting approval…". If nothing opens, allow pop-ups for this site and try again.
2. Have a Global Administrator of the client's Microsoft 365 sign in and accept the requested permissions. SLApp then creates its own app, SLApp Helpdesk - Microsoft 365, in the client's tenant; the window lists each step and closes by itself when they are all OK (with a warning or a failure it stays open, so that you can read it). The card switches to Connected, with who connected it and when, and the date the App secret expires.
3. Click Import users. When the client has email domains, the list opens filtered to Client domains (the email domains from the client form) with those users pre-selected; switch to All domains to see guests and external accounts. If the client has no domains, the Client domains filter is disabled and the list opens on All domains with every user pre-selected. The counter shows how many are visible out of the total.
4. Adjust the checkboxes and click Import N contacts. Imported people appear in the Contacts tab with name, job title, email, phones and photo.

> **i** Importing is safe to repeat: users already imported are refreshed from Microsoft 365 (name, job title, email, phones and photo), which also means any manual edits you made to those fields on an imported contact are overwritten by the next import. A Microsoft 365 user who already exists as a contact of the client without the Microsoft link (created by hand, by a staff request or invited to the portal) is matched by email or UPN and linked to that contact instead of being added a second time; the contact keeps its portal access and other settings.

Each step of the connection, Reconnect (which repairs a connection that Needs reconnect), Disconnect and the sync are described in the Microsoft 365 chapter. A client connected with the earlier Connect

Office 365 button shows Old connection: Import users and the BitLocker keys keep working, but nothing is synced until an administrator clicks Reconnect, which upgrades it.

Once connected, a BitLocker recovery keys card appears. Import from Microsoft 365 fetches the keys stored in the client's Intune/Entra ID (the app SLApp creates in the client's tenant has the BitLocker read permission; an old connection only has it if that permission was consented in the tenant) and stores them encrypted; keys already held are skipped. Each key is matched to one of the client's computers by the Entra device ID the agent reports, falling back to the device name. Matched rows are highlighted in green with a link to the computer, the others read "unmatched" (a machine without the agent or not managed). Re-link to computers redoes that matching without re-reading the keys, which is useful after agents have reported their device IDs. The table shows Device, Computer, Drive, Key (masked, with Reveal/Hide) and Imported; Remove (confirmed first) deletes a key from this database only and does not touch the client's tenant.

> **i** Disconnect (administrators, confirmed first) makes SLApp forget the connection and its secret, which it also revokes in the client's tenant when it can, and stops syncing the client; contacts and BitLocker keys already imported are kept. The app stays in the client's tenant: a card under the Microsoft 365 card says how to delete it in the Microsoft Entra admin center, with a Copy App ID button.

Websites

The Websites tab lists the monitored websites assigned to this client — status, SSL, the hosting subscription’s setup and expiration dates, and today’s visits. Sites are assigned from the Websites module: one by one in the edit form, or in bulk with the Assign to client button in the list’s selection bar.

Domains

The Domains tab lists the client's monitored domains with their registration expiry, registrar, SPF, DKIM and DMARC status, last check and open alerts, with the same filters and sorting as Cloud & web › Domains. It is read-only: a click opens a domain's detail window, and domains are added, changed and checked in the Domains module (see the Domains chapter).

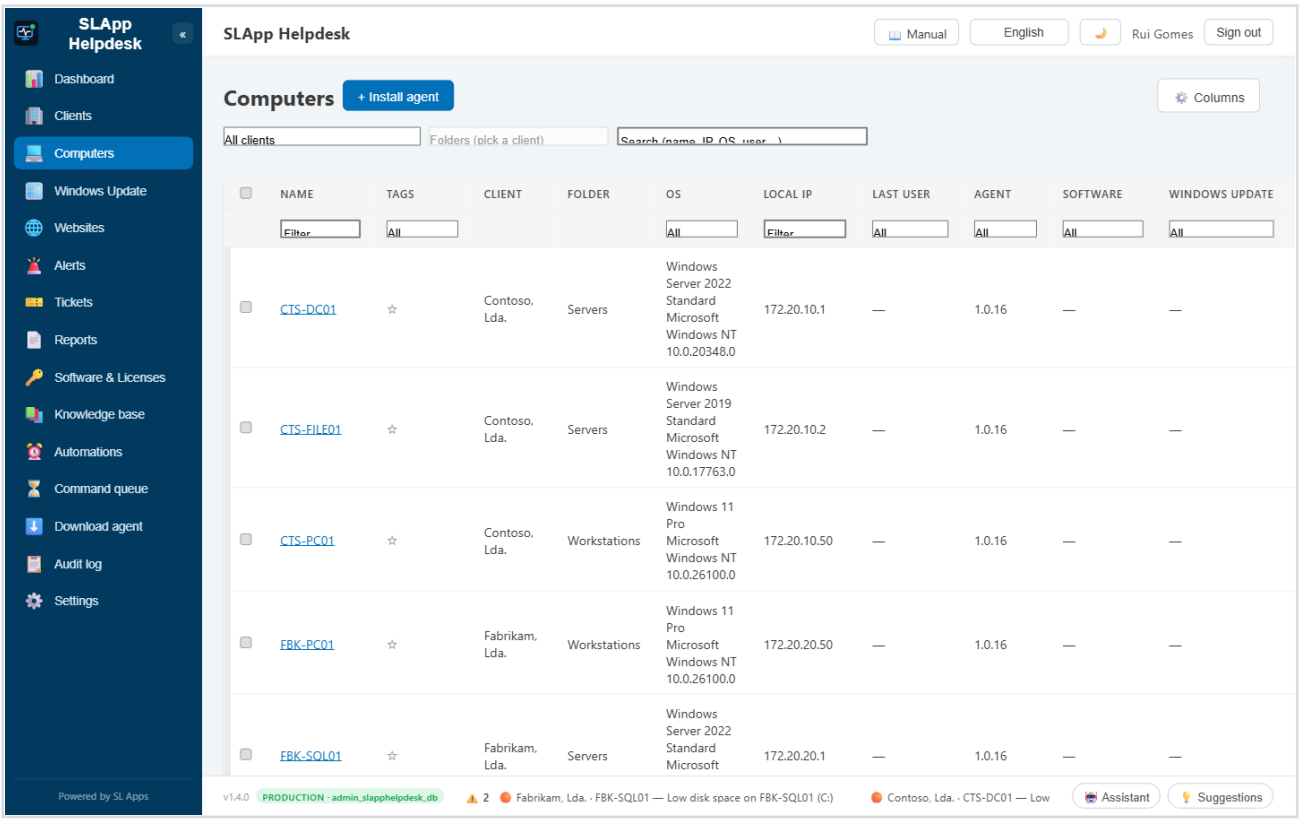
Microsoft 365

The Microsoft 365 tab shows the client's Microsoft 365: the state of the connection and its last sync, the users, licenses, administrators and domains read from the tenant, the open alerts, Sync now and the user actions (reset a password, block a sign-in, sign out everywhere, licenses, automatic replies, new user). It is the same page as in Cloud & web › Microsoft 365, described in the Microsoft 365 chapter. A client that is not connected shows Not connected to Microsoft 365, with how to connect it on the Data tab.

5. Computers

The Computers module is the inventory of every machine that has the agent installed. The list gives you the health of the whole fleet at a glance; opening a computer takes you to its detail page, where you can see everything the agent collects (hardware, disks, operating system, security, software, events), keep your own notes and passwords, scan the local network, monitor SNMP equipment, and run commands or connect remotely.

The computers list



The computers list with the client/folder filters, the search box and the per-column filter row.

The list refreshes on its own every 20 seconds. At the top you have + Install agent (opens the installation wizard), a client dropdown (All clients, — no client — for machines not yet assigned, or one client), a folder dropdown (only enabled after you pick a client, because folders belong to a client) and a search box that matches any text in the row (name, IP, OS, user and so on). Under the column headers there is a second filter row; clicking a header sorts by that column (Client and Folder are not sortable).

- Name — click it to open the detail page. Filter by free text.
- Client and Folder — where the machine is filed. Use the dropdowns at the top to filter.
- OS — operating system and version; the filter is a dropdown of the versions present.
- Local IP — free-text filter (partial match works).
- Last user — Windows account last seen signed in; dropdown filter.
- Agent — agent version, with a dropdown filter of the versions present. An outdated badge appears when a newer version is published (hover to see which one).
- Software — "N to update" when the agent found newer versions; clicking the badge opens the Software & Windows Update tab. Filter: With pending / Without pending.
- Windows Update — "N pending" updates; same click behaviour and filter.
- Alerts — number of open alerts (red badge). Filter: With alerts / Without alerts.
- Security — Protected, At risk or Unknown (Unknown just means the agent has not sent a security snapshot yet). Filter: Protected / At risk / Unknown.
- Status — Online / Offline plus a restart badge when a restart is pending (hover for the reason). Filter: Online, Offline or Restart pending.
- Devices — a dish icon with N when the computer polls SNMP devices; red if any of them is offline (hover for the count). Clicking it opens the detail page. Filter: With devices / Without devices.
- Actions — a per-row menu (see below).

> "Online" reflects the agent's service channel. A machine can be online with nobody signed in — that matters for remote actions run as the signed-in user and for the native remote access session.

Selecting computers, bulk actions and the row menu

Tick the checkbox on one or more rows (the header checkbox selects every row currently visible after filtering). A bar appears with N selected, a client dropdown, a Set folder checkbox with a folder picker, and the buttons Move N to client, Actions on N and Clear selection.

1. Filter or search so that only the machines you want are visible, then tick them (or tick the header checkbox).
2. In the selection bar, choose the destination client. Leave it on — no client — only if you really want to unassign them.
3. Optionally tick Set folder and pick a folder of that client (leave the folder empty to clear it).
4. Click Move N to client. A Confirm bulk move dialog summarises the move; click Move to apply or Cancel to back out.

> The client dropdown in the selection bar starts on — no client —. Always read the confirmation dialog: moving a selection to "no client" removes those computers from their current client.

Actions on N opens a dialog with the same tools as the Remote actions tab, applied to every selected computer in parallel: Restart, Shut down, Force sync, Force update, Install SLApp Vault extension and Remove SLApp Vault extension (these two also switch the Password manager entry of the agent's tray menu on and off; Windows only — a non-Windows machine in the selection shows up as failed with a reason, it is never silently skipped), a Saved script (optional) picker, Interpreter (CMD or PowerShell), Run as (Signed-in user or System (elevated/admin)), a Command / script box and Run on N. Restart, Shut down and Force update first ask you to Confirm (or Cancel). A results table with Computer, Result and Detail fills in as each machine replies — Executed, Waiting (no channel right now — the command stays queued and runs when the machine connects), Failed or Running… — with a summary line above it.

The Actions menu on each row gathers: a Remote access section with SL App (native) (disabled when offline), Remote Desktop (through the agent) — the RDP tunnel described under Remote access below, greyed out with the reason under it while the computer is offline or when you may not open tunnels, and not offered on a phone or tablet — and any third-party tools the agent detected (TeamViewer, AnyDesk, Remote Desktop Connection, Chrome Remote Desktop, with their ID when known); Remote actions (opens the detail page on that tab); Move (opens the Move dialog, which asks you to Confirm after you click Move); and Remove, which always asks you to confirm first (Cancel / Remove). Removing a computer also uninstalls its agent from the machine — see Removing a computer and its agent, below.

Removing a computer and its agent

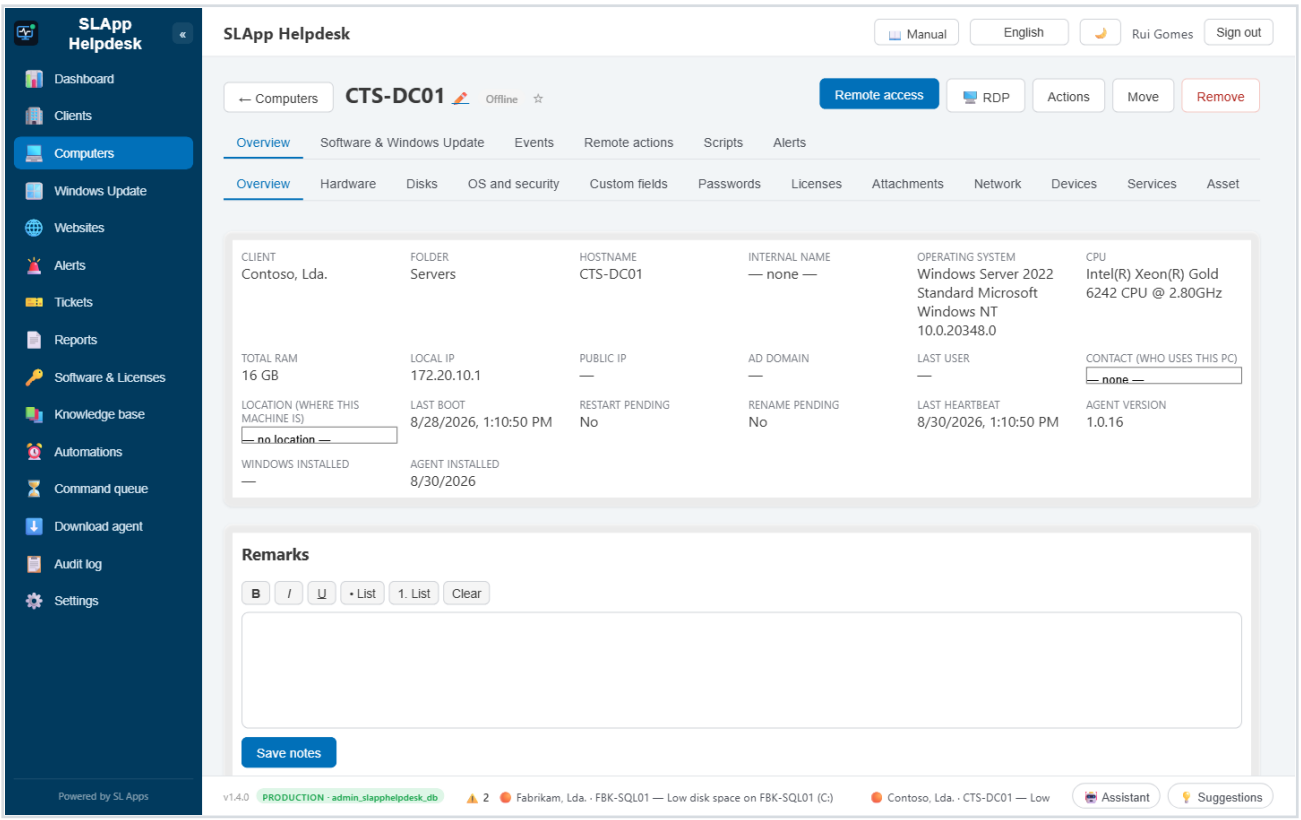
Remove — in the row's Actions menu, in the right-click menu of the row, or in the header of the computer's detail page — takes the computer out of the Computers list and, with agent 1.0.48 or later on Windows and 1.0.14 or later on Linux and macOS, uninstalls the agent from the machine, so the machine does not register itself again as a new computer. The confirmation window says so before you confirm. There is no undo. If the server cannot take the removal, the window shows the reason and nothing is changed. When the computer has backup jobs, an administrator also sees a line in the window with the number of jobs (for example “This computer has 2 backup jobs. They will be paused when the computer is removed (their run history stays).”, or “This computer has 1 backup job. It will be

paused when the computer is removed (its run history stays).”). The jobs are paused, not deleted: their run history stays and the files already written to the destinations are untouched. See the Backups section.

- What is removed from the machine: the agent's service, its program folder, its tray icon (the menu-bar icon on Linux and macOS) and its saved identity (the credential it uses to talk to the server); on Windows also its entry in the list of installed apps, its Start Menu shortcut, the firewall rule of the syslog receiver and the password manager's data of each user, and from agent 1.2.0 the data that was still waiting on the computer to be sent to the server (usage samples, Syslog lines and watched events, see What the agent does on the computer in the Download Agent and installing on computers section), which is lost. Nobody on the computer is asked or told. Software installed through software packs or scripts is not uninstalled, and neither is the SLApp Vault browser extension: if it is switched on for the computer, untick Install the SLApp Vault extension on this computer (or use Remove SLApp Vault extension) before you remove the computer, otherwise Chrome and Edge keep the extension.
- When: a machine that is on and connected uninstalls the agent in a minute or two. The agent waits for two heartbeats in a row to be sure, so a high Heartbeat interval (Settings › General › Monitoring) makes it longer. If the computer is off or has no network, the agent uninstalls itself a minute or two after it next reaches the server. A computer removed before this behaviour existed (before app 1.9.123) left no mark: if its agent was off at that time, it registers again as a new computer when it next connects, as it used to.
- Older agents: agents older than 1.0.48 (Windows) or 1.0.14 (Linux and macOS) cannot uninstall themselves. The server simply stops accepting them, so the computer does not come back, but the old agent stays on the machine doing nothing useful until it updates itself (agents check for a new version every 6 hours); the new version then uninstalls itself a minute or two later.
- Installing the agent again on the same machine always creates a new computer; a removed computer is never brought back. Use the installer or command from Download agent (or the + Install agent wizard in Computers): if you chose a client there, the new computer appears under it; if you chose no client, it appears without one and you assign it with Move. Removing a computer frees its license seat at once, and the new computer takes one again.
- What does not follow to the new computer: it starts with no history. Software, alerts, events, charts, custom fields, notes and the asset generated from the computer stay with the removed record, which is no longer listed. Tickets and passwords stay with the client but remain linked to the removed computer, not to the new one. Alerts that were open when you removed it stay open: resolve them in Alerts. What you linked to the old computer one by one — alert rules, maintenance schedules, software packs, tags, SNMP devices, Syslog settings — does not follow either: link it again on the new computer.

> **i** To manage a removed machine again, install the agent on it again with an installer or command from Download agent (do not rename the installer file: the token is in its name). Install after the old agent has gone: if the machine was off when you removed it, switch it on, let it connect and wait a couple of minutes. If you install earlier, on Windows the agent registers the computer again as a new one; on Linux and macOS the new agent uninstalls itself once more, so wait until it has gone and run the install command again. The backup jobs of a removed computer stay paused on it and do not move to the new one: in Backups, open each job, choose the new computer in Computer (agent), click Save, then open the job's Actions menu and choose Resume.

The detail page and the Overview tab



The computer detail page: header with status badges and connection buttons, the Overview sub-tabs and the usage charts.

The header shows ‹ Computers, the computer name with a pencil button to rename it (this changes only the name in the app, never the Windows hostname), the Online/Offline badge, a restart pending badge when applicable and, while online, two dots for the agent's channels: service (elevated commands and tunnels) and session (the signed-in user's session, used by remote access). On the right: Remote access (its menu also opens the RDP tunnel, Remote Desktop (through the agent)), Actions, Move and Remove (Remove asks you to confirm, and uninstalls the agent from the machine). The four main tabs are Overview, Software & Windows Update, Events and Remote actions; Overview has its own sub-tabs: Overview, Hardware, Disks, OS and security, Custom fields, Passwords, Licenses, Attachments, Network, Devices and Asset. Your browser's Back button closes the detail page.

- Identity card — Client, Folder, Hostname, Operating system, CPU, Total RAM, Local IP, Public IP, AD domain, Last user, Last boot, Restart pending, Last heartbeat and Agent version (with an outdated badge when a newer one is published).
- Contact (who uses this PC) and Location (where this machine is) — dropdowns that link the machine to one of the client's contacts and locations. They only work once the computer is assigned to a client (otherwise you see "assign a client first"). The agent keeps the contact up to date automatically from the user signed in to the machine; tick Locked next to the contact when you have set it by hand and do not want the agent to overwrite it.
- Remarks — a rich-text editor for your own notes about the machine. Open alerts — listed here when there are any.
- Usage — CPU, RAM and Disk (system) over Hour, Day, Week or Month. Network — Sent and Received traffic, with an All adapters filter when more than one adapter is up. Both have Real time (refreshes every 5 s while you stay on the page) and Maximize.

> ℹ Last user comes from the agent, which reads the account signed in at the machine even though the agent service itself runs as SYSTEM, and it resolves Microsoft 365 / Entra ID sign-ins to the user's full name and address. When that user matches one of the client's contacts, the computer is linked to that contact automatically — unless the contact is Locked.

> ℹ Charts are interactive: scroll the mouse wheel or use the + / - buttons to zoom, drag left or right to pan, and double-click (or the reset button) to get back to the full period. Gaps longer than five minutes (machine asleep or off, agent not running) are left blank rather than joined with a line. A Windows agent from 1.2.0 that only lost its connection to the server keeps its samples on disk and sends them once it is back, so such a gap fills in afterwards (up to about 20 MB or 7 days of samples; see What the agent does on the computer in the Download Agent and installing on computers section). Real time asks the agent for a sample through the user session, so it only advances when someone is signed in. The usage data (CPU, RAM, disk and network) is kept for 30 days by default; an administrator can shorten that (down to 7 days) in Settings › General › Retention.

The Hardware lifecycle card

Under the identity card, the Hardware lifecycle card shows the computer as the Lifecycle module sees it (see the Lifecycle section). Age gives the age in years, marked ~ when it is approximate, with where it comes from and the date it counts from. Replace by gives the date the computer is due for replacement, in red once it has passed, and after how many years. Operating system shows the product, its support chip (Supported, Ending soon, Unsupported or Unknown), its end-of-support date and, when the support ends soon or has ended, the version to move to. Warranty shows its chip, its end date and its source, why it can or cannot be looked up and the result of the last lookup. Recommendation shows Replace, Upgrade OS, Renew warranty or OK with the next step and its date, and Open lifecycle alerts appears when there are any. Hover the values for more detail. Check warranty asks Dell, Lenovo or HP for the warranty at once and Edit dates opens the purchase date, purchase price and warranty end date of the computer; each button appears only when it can be used. A virtual machine has no age, replacement date or warranty, only the support of its operating system.

The Active Directory & third-party backups card

On a Windows computer the Overview also has the Active Directory & third-party backups card. It shows the read-only check that the agent runs on the computer once a day (see The daily check in the Active Directory section): Last check, Next check and Check in progress while a check waits for its result (hover it when the computer is offline), with a Check now button that runs the check at once, at most once every 5 minutes. A computer that has not been checked yet says so, and a red line appears when checks returned no readable result.

- Active Directory: Role is what the check detected, Domain controller, Domain member, Not in a domain or Not detected yet, with a PDC emulator badge for the PDC emulator of its domain. Monitoring as a domain controller offers Automatic detection, Always monitor as a domain controller and Never monitor as a domain controller (see Domain controllers: detection and marking in the Active Directory section). For a domain controller, Domain controller health shows its health, its domain and its open alerts, followed by one chip per part: dcdiag, Replication, Services, Time with the offset from the PDC emulator, and SYSVOL/NETLOGON; hover a chip for what is wrong. The replication partners, services, FSMO roles and privileged groups are in Infrastructure › Active Directory.
- Third-party backups: Backup software found lists the supported products the check found, in orange when one could not be read. Custom events (N) opens the custom events of this computer. The jobs of the computer are listed with Job, Product, Last result, Last success, Expected and Alerts, plus Last run, Duration and Message through Columns, and can be sorted and filtered as in the module.

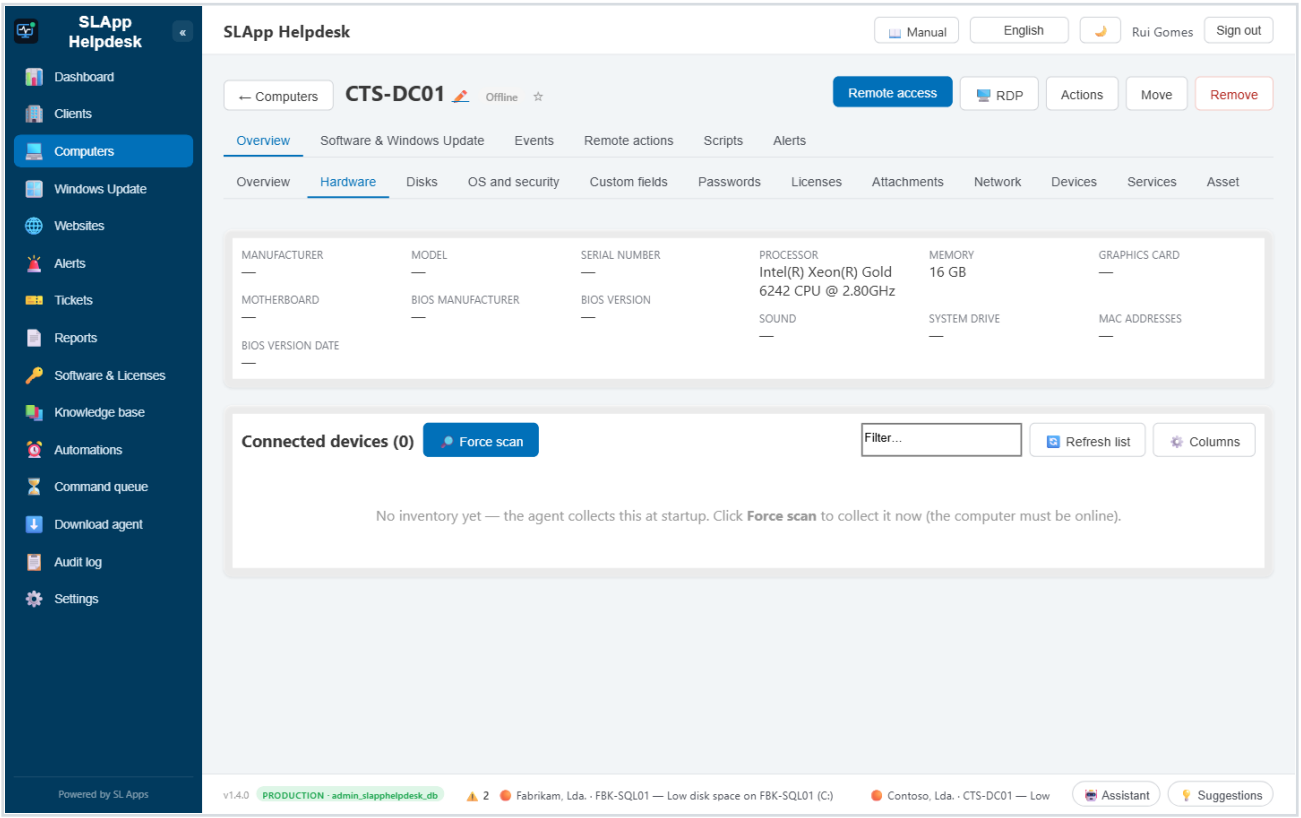
Click a job to open its window, or right-click it to check now, change its expected frequency, open the custom events, copy its name or, for administrators, forget it (see the Third-party backups section). With no job the card says No backup jobs found on this computer.

- Notes under the jobs say when the backup part of the check keeps failing, when Synology Active Backup for Business is installed (its results stay on the NAS and cannot be read here) and when Acronis is installed but no backup events were found (Acronis writes to the Windows event log only when that option is switched on).

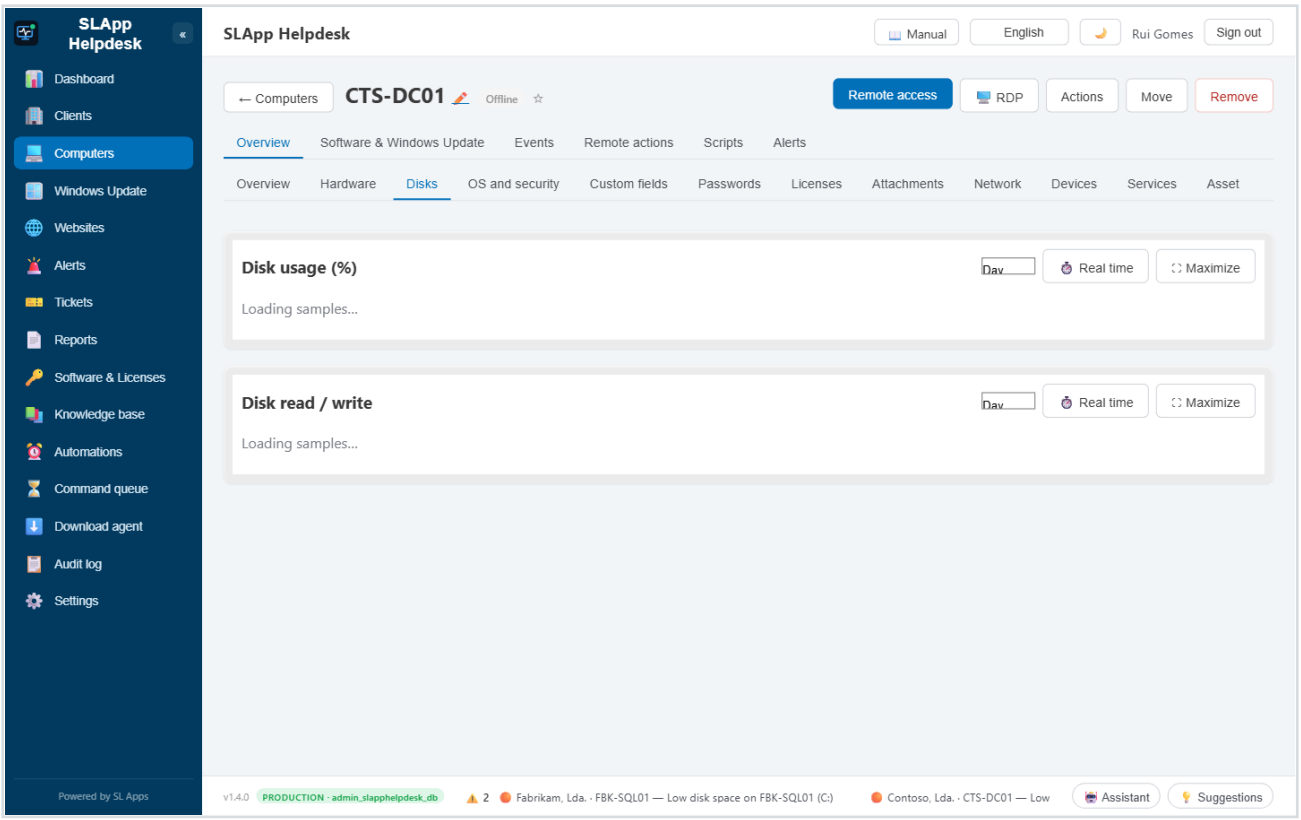
> **i** Choosing Never monitor as a domain controller for a computer that is monitored as one asks you to confirm, then closes its Active Directory alerts; the backup part of the daily check keeps running. Linux and macOS computers do not have this card.

Hardware, Disks and OS and security

Hardware lists Manufacturer, Model, Serial number, Motherboard, BIOS manufacturer, BIOS version, BIOS version date, Processor, Memory, Graphics card, Sound, System drive and MAC addresses, followed by the Connected devices card (USB, ports, printers, Bluetooth…) grouped by device class with Name, Manufacturer and Status, plus a filter box. The agent collects devices at startup; Force scan asks for a fresh collection right now (the computer must be online) and Refresh list re-reads what is stored.

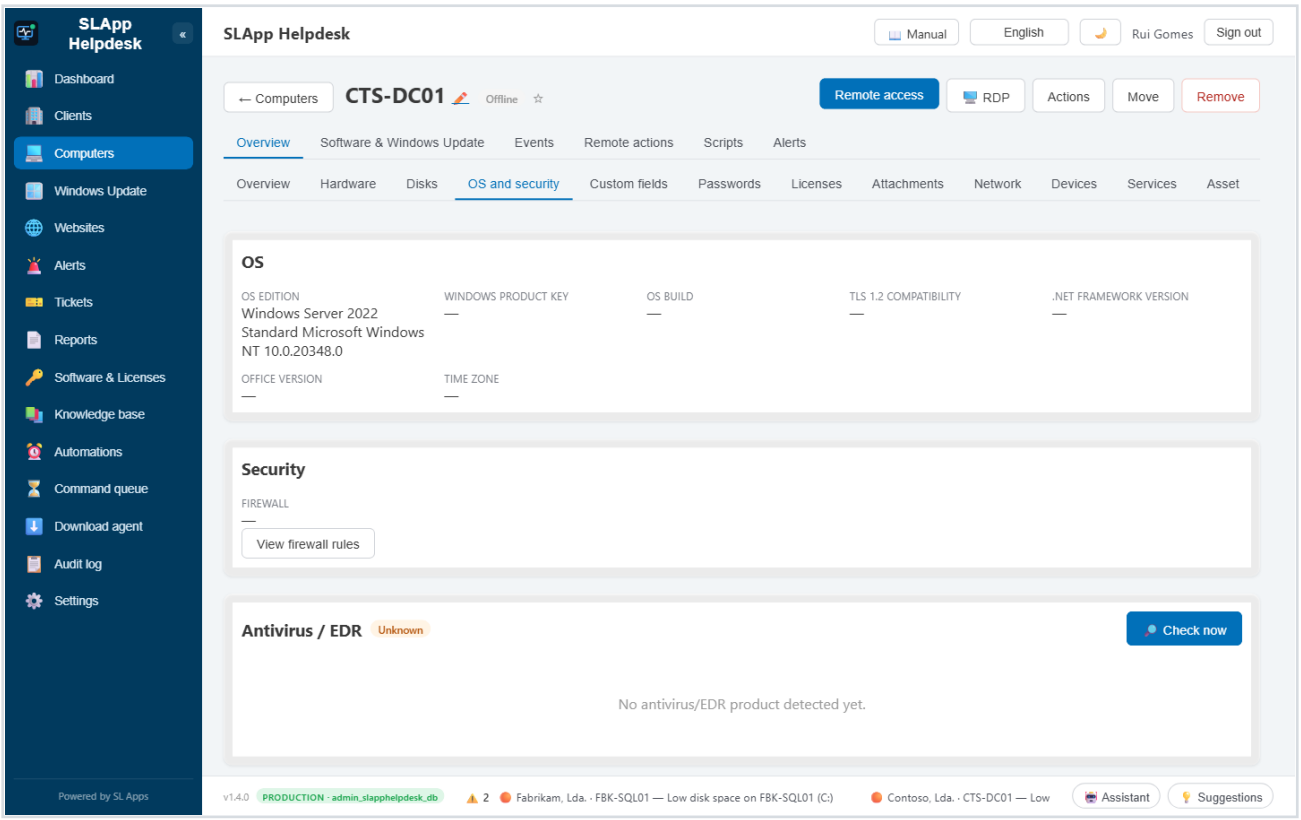


The Hardware tab.



The Disks sub-tab: logical drives with BitLocker status and trend, physical disks, and the per-disk charts.

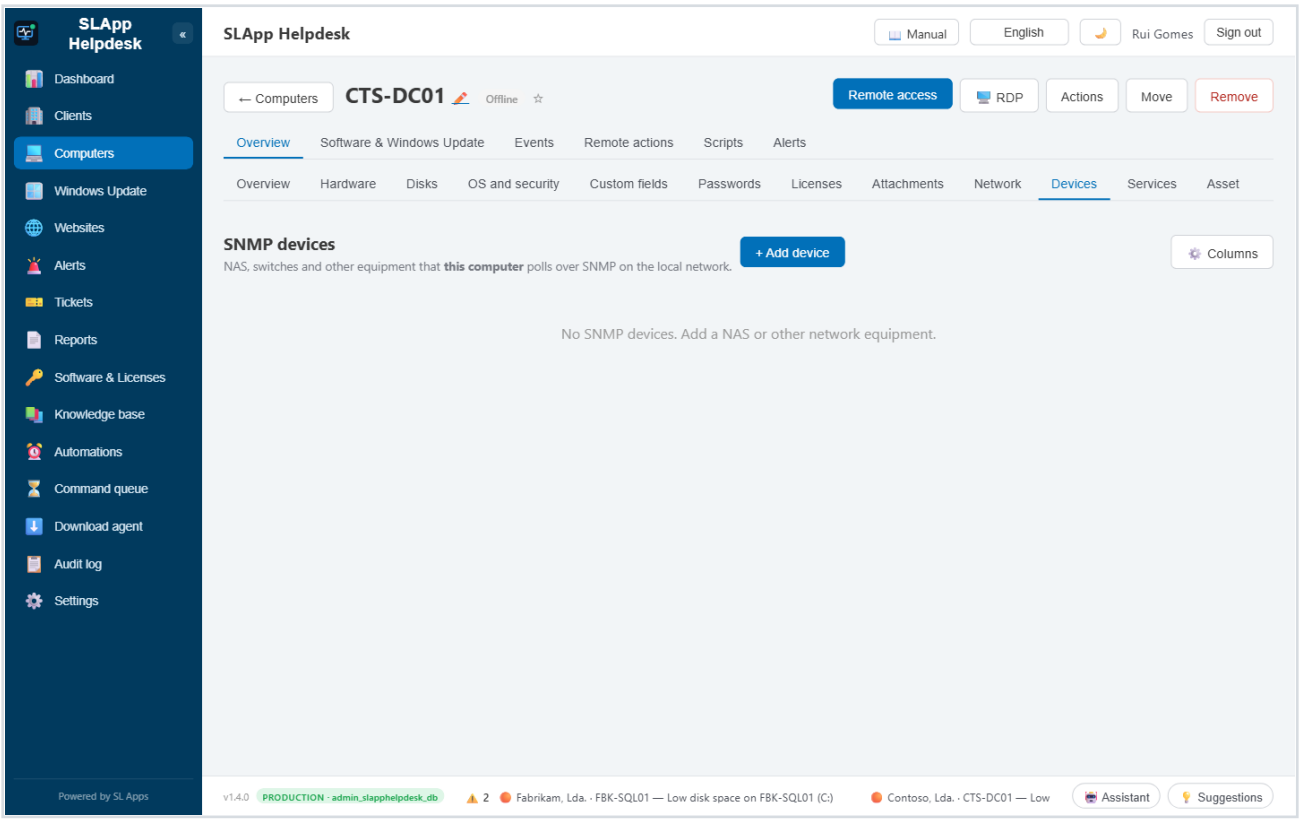
- Disks table — Drive, Model, Serial number, BitLocker (On/Off; when recovery keys were imported from Intune you get a Reveal key button, then Copy and Hide), Total, Free, Trend (a mini chart with up/down/flat arrows and the change in percentage points over the selected period) and Usage.
- Physical disks — Index, Type, Model, Serial number and Drive.
- Disk usage (%) and Disk read / write charts — one line per drive (solid = read, dashed = write), with an All disks filter, period, Real time and Maximize. Read/write data requires agent 0.8.40 or later.



OS and security: OS details, firewall rules on demand and the Antivirus / EDR check.

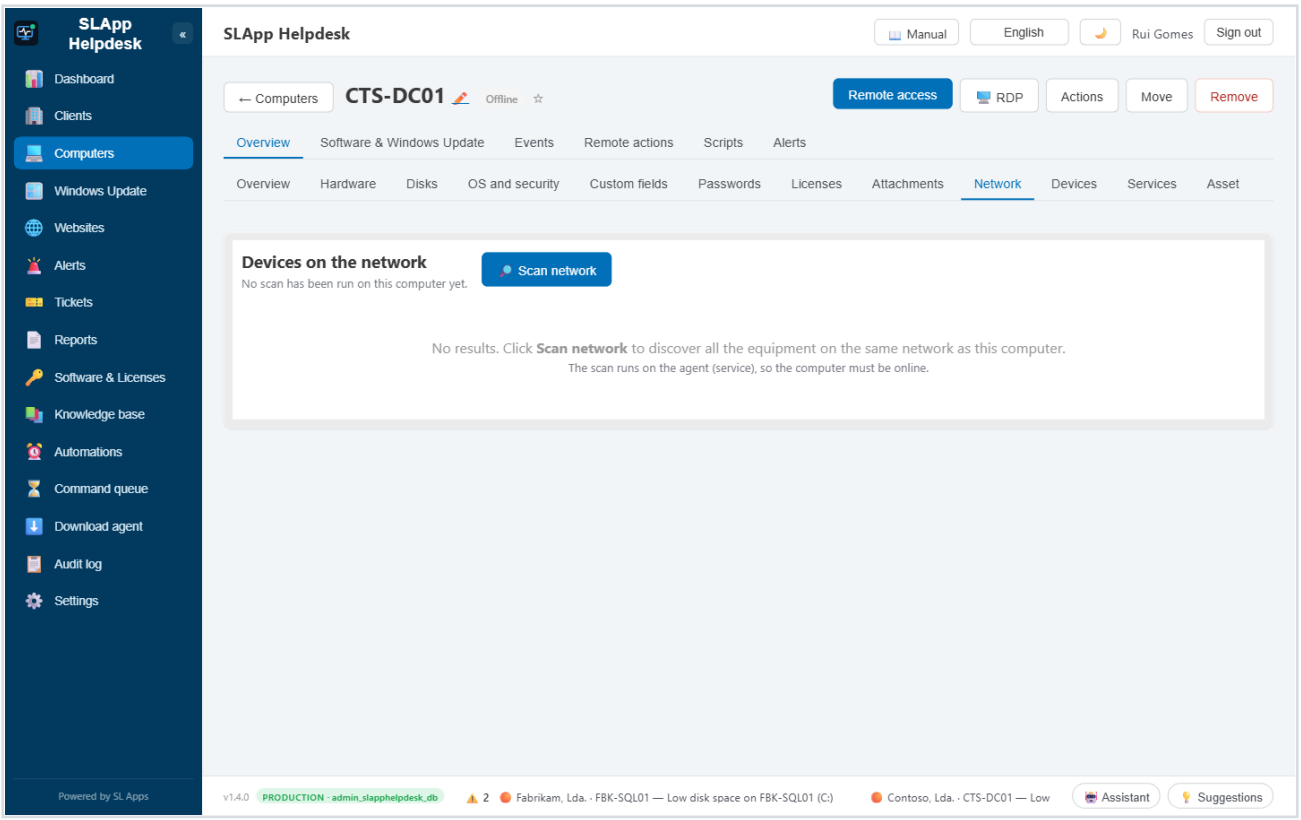
OS and security has three cards. OS shows OS edition, Windows product key (masked; click the eye icon to reveal), OS build, TLS 1.2 compatibility, .NET Framework version, Office version and Time zone. Security shows the Firewall state and a View firewall rules button: the rules are never collected automatically, so the first click triggers a scan on the machine and shows the table (Name, Direction, Action, Enabled, Profile, Protocol, Local port, Remote port) with a filter row including a Port… box; Update re-scans and Close collapses the section. Antivirus / EDR shows the Protected / At risk / Unknown badge, a Check now button that asks the agent for a fresh security snapshot, and the detected products with Type (Antivirus or EDR), Enabled/Disabled and Up to date/Out of date. On a Windows machine, a fourth card, SLApp Vault extension, pushes the SLApp Vault personal password manager's browser extension to Chrome or Edge with a single checkbox, Install the SLApp Vault extension on this computer — nothing to install by hand, since it installs by Windows enterprise policy the next time the agent connects. The same checkbox also switches on the personal password manager on the computer itself: while it is ticked, the agent's tray menu offers Password manager (see Download Agent and installing on computers), and unticking it removes that entry again. The same toggle is also available in bulk from the computers list (see above).

Custom fields, Passwords, Attachments, Network and SNMP devices



The Devices tab: network devices monitored by this computer over SNMP.

- Custom fields — free Key / Value pairs you type yourself (for example a TeamViewer ID). + New field opens the New custom field window (Key and Value); Edit in the row's menu, or a double-click on the row, opens the same window filled in to change the value (the key of an existing field stays as it is); Remove deletes the field after a confirmation.
- Passwords — the same encrypted vault as the client, filtered to this computer: + New password, then a table with Title, User, Password and Updated. Each row has Reveal (it becomes Hide once shown), Copy and Remove (asks to confirm); click the Title or User cell to edit the entry.
- Licenses — product keys detected by the agent (marked auto) or added by hand with + New license, stored encrypted (Reveal key, then Copy and Hide), with files attached per license. Attachments — Upload file stores any document (invoice, warranty…) listed with File name, Size and Created; click the file name to download, Remove to delete. Asset — the asset record linked to this computer.



The Network sub-tab after a scan: adapters, discovered devices and the Connect buttons.

Network starts with the Network adapters card (Name, Status, IP, Speed). Below it, Devices on the network is an on-demand scan run by the agent's service: click Scan network (the computer must be online; it takes one to two minutes) and the table lists every device found with IP (badges this PC and gateway), Name, Type, Manufacturer, MAC, Ports / services and Latency — all sortable and filterable, with a general filter box as well. The pencil button next to a name opens Edit device, where Label (own name), Type and Notes are saved against the MAC so they survive later scans. The Connect column offers one-click connections: HTTP, HTTPS, HTTP:8080, HTTPS:8443 and Proxmox open the device's web page in a new tab through the agent (so it works from outside the client's network), RDP downloads a .rdp file, SSH/FTP/Telnet/RTSP hand over to the app registered on your PC, and Share copies the \\ip path for Explorer. The web connections (HTTP, HTTPS, HTTP:8080, HTTPS:8443 and Proxmox) are tunnels: a technician whose tunnel permission was removed in Settings › Users sees them greyed out with You are not allowed to open RDP/HTTP tunnels. Ask an administrator (Settings › Users)., and a failure to open one is always shown as a notice. The RDP button of a device row only downloads a .rdp file and is not affected.

Devices is where you configure SNMP devices that this computer polls on its local network — Synology NAS, QNAP NAS, Router / Firewall (for example DrayTek), Generic (SNMP) or Printer (see the Printers section). The tab is a table with one row per device and the columns Name, Host (address and port), Type, Status (Online or Offline, plus Inactive when the device is switched off), Alerts (open alerts), Model, Last polled (updated x ago), Last error, Temperature, Firmware (with an update badge when a newer firmware exists), Uptime, CPU, RAM and System. Click a header to sort and use the row under it to filter. The Columns button lets you choose which columns to show and reorder them; SNMP version and Interval are available there and start hidden. Every row has Poll now, Edit and Remove (asks to confirm). Click a device name to expand its detail (hover it to read Show disks, volumes, RAID and UPS, Show supplies and detail for printers, or Show interfaces and detail for routers), which shows Active VPN tunnels, Interfaces, Volumes, RAID / Storage Pools, Disks and UPS, or the supplies of a printer. Under Interfaces, Choose which ones alert + internal name lets you exclude ports from the network alert and give them a friendly name. A long Last error is cut short in the table; hover the cell to read all of it. The Generate alerts checkbox next to Columns turns the alerts of all your company's

SNMP devices except printers (which have theirs in Printers) on or off; only administrators can change it (see Turning alert generation on or off per module in the Alerts section).

1. Open the Devices sub-tab and click + Add device.
2. Fill in Label and choose the Type. The Type list starts on Synology NAS; choose QNAP NAS for a QNAP, Router / Firewall for a router or firewall, Printer for a network printer (its form asks for supply thresholds instead of the ones below) and Generic (SNMP) for any other equipment, such as a server's iDRAC or a switch.
3. Fill in SNMP version (only v2c is offered), Host / IP, Port (161 by default), Community and Interval (seconds) (300 by default); keep Active on Yes. Host / IP and Community are required: Save stays disabled until both are filled in. If you type 162 in Port, a warning note appears under the field: 162 is the port devices use to send traps, while SLApp asks its questions on 161. The note never blocks saving, because a device can listen for queries on another port.
4. Leave Raise alerts for this device ticked and adjust the thresholds: Volume — warning (%) (85 by default), Volume — critical (%) (95), Temperature — warning (°C) (60) and CPU — warning (%) (off by default); tick Interface/VPN that was up and goes down (routers) and Firmware update available as needed (both start ticked) and Email the team when an alert opens if you want emails (starts unticked). Critical health (offline, degraded RAID, failed disk) always alerts; empty thresholds are off.
5. Click Save. The agent polls within about a minute — use Poll now if you do not want to wait for the interval.

> **i** Spaces typed or pasted before or after the Host / IP or the Community are removed when you save, so a stray space from copy and paste cannot turn a correct community into a wrong one. A community saved earlier with a stray space is cleaned the same way when the agent reads it, with no need to edit the device. When you edit a device, leave Community empty to keep the stored one; Host / IP is still required. Changing Host / IP, Port or Type starts the device's comparison from scratch, so replacing a piece of equipment does not raise a false reboot or interface-down alert.

When a device shows No SNMP response

A device that does not reply in time is asked a second time (agent 1.0.43 or later) before it is reported offline, so one lost request does not raise an alert. If it still stays silent, Status turns Offline, the Last error column says what to check and, with Raise alerts for this device on, a Critical alert opens in the Alerts module, its message starting with No SNMP response: followed by that same text. The usual text is "no reply (timeout). Check SNMP is enabled on the device, the community, v2c and UDP 161 access.", where the port is the one set on that device. When that port is 162, the timeout text ends with one more sentence, "The configured port is 162, the port for traps; SNMP queries use 161.", in Last error as well as in the alert message. When the device itself is up but nothing listens on that UDP port, the text is "SNMP port unreachable: the device is up but nothing listens on UDP 161. Check SNMP is enabled on the device and the port." Work through these checks in this order, from the most likely cause.

- SNMP is enabled on the device. Check that the SNMP agent is switched on in the device's own settings.
- The Community matches exactly, capitals included. SNMP v2c never answers a wrong community, so a mistake here looks exactly like a timeout. Edit the device and type the community again.
- The Port is the SNMP agent's own port, 161 unless someone changed it on the device. Do not use 162: that is the port devices use to send traps (alerts), not the one they answer queries on. The device form warns you when Port is 162, and a timeout on that port reminds you of it.

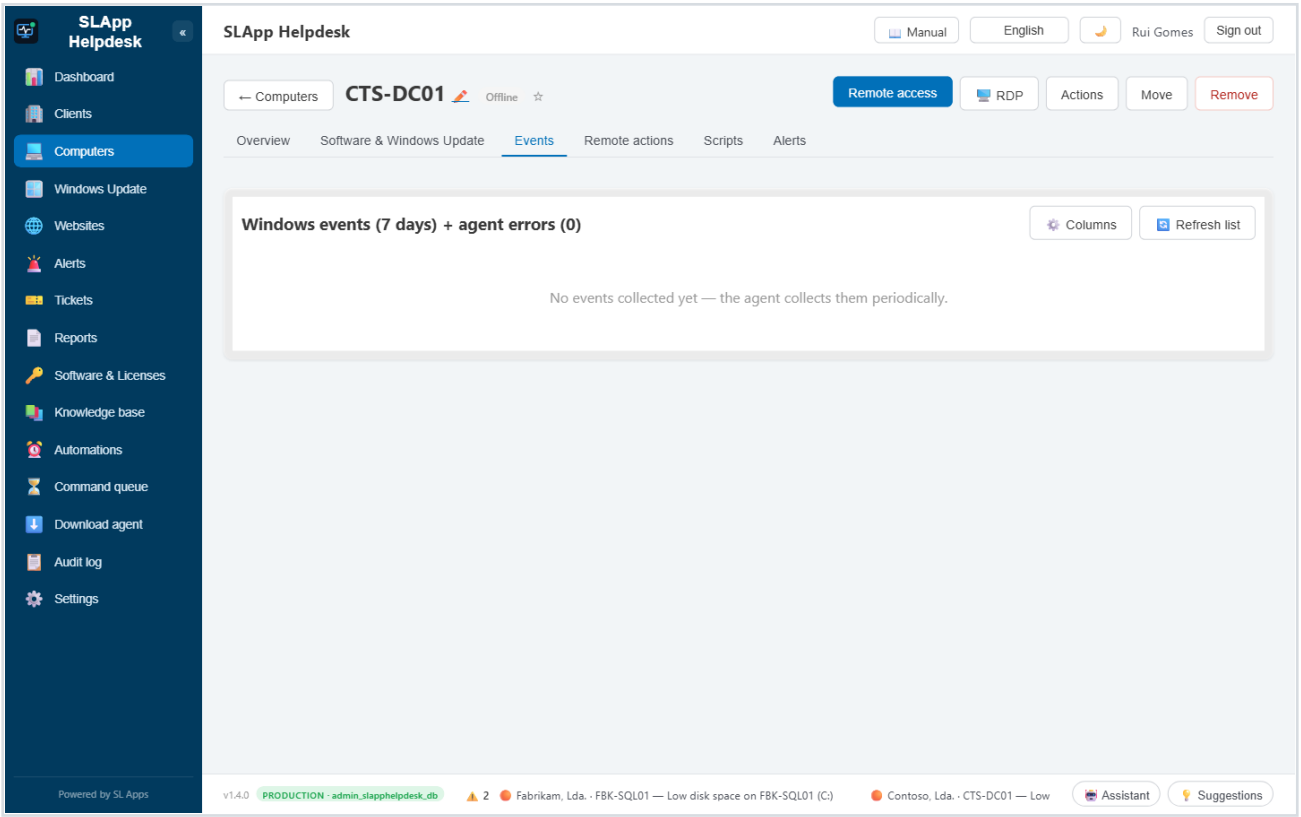
- The device supports SNMP v2c. It is the only version offered here, so a device set to v1 only or to v3 only stays silent.
- A firewall or a VLAN does not block UDP 161 between the computer that polls the device and the device. The requests come from the computer whose sheet you are on, not from your own PC.
- On a Dell iDRAC, use the community configured under the iDRAC's SNMP Agent service settings, not the one on the alert destination page.

While a device is offline its live readings (Temperature, Uptime, CPU, RAM, System and the expandable detail) are blank, but the last known state is kept. When the device answers again it is compared with that state, so a reboot, or an interface or VPN tunnel that went down during the outage, is still reported. Alerts that were already open stay open during the outage (while Raise alerts for this device is on) and are re-checked by the next poll that gets an answer.

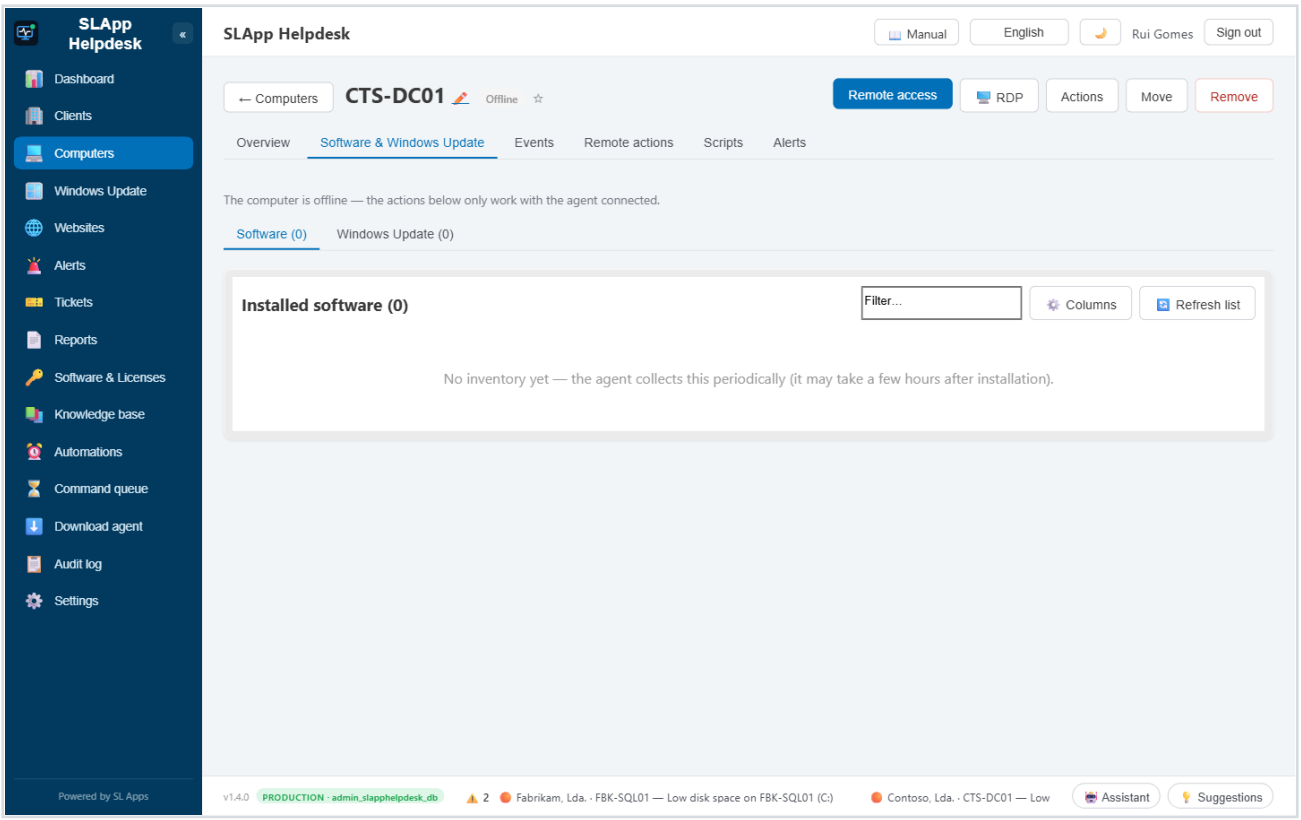
Besides the device as a whole, the agent reads tables from it, such as the interface, volume, disk and RAID tables. From agent 1.0.45, a table that times out is asked once more, about a second later, so a busy device that swallows one request in the middle no longer leaves that poll with a half-read table; the retry only happens while that poll still has time to spare. If the table stays silent, the agent gives up on it quietly: that is not treated as the device being offline and raises no No SNMP response alert, which is only for a device that does not answer at all. What could not be read is simply missing from that poll's readings until the next poll.

> ℹ After you fix the cause, click Poll now. At the first poll that gets an answer the device turns Online, Last error clears and the No SNMP response alert closes by itself.

Software & Windows Update and Events



The Events tab: Windows event log entries reported by the agent, each with an Ask AI button.



Installed software with available versions, and the Windows Update sub-tab with pending and installed updates.

Software & Windows Update has two sub-tabs, each showing its item count. Software lists Installed software with Name, Installed version and Available version; Update appears when a newer version exists, Remove uninstalls (after Confirm removal). Updates run one at a time on each computer, in the order they were asked for, so several Updates in a row (or an update of every application) wait for each other instead of colliding. Microsoft Edge and WebView2 are updated through Microsoft Edge Update, which keeps them, and a package that is already up to date counts as done. Windows framework packages (Windows App Runtime, VCLibs, UI.Xaml) are never offered: Windows and the apps that use them keep them up to date. Windows Update shows Pending (N) updates with Install and Installed (N) with Remove (also after Confirm removal); each KB number links to the Microsoft support article. Both have a filter box and Refresh list. The inventory is collected periodically by the agent (it can take a few hours after installation), and the action buttons only work while the computer is online.

In the Software list, an application that has a newer version and whose product is in CISA's catalog of Known Exploited Vulnerabilities carries a red KEV · N CVE badge after its name, N being the number of exploited vulnerabilities listed for that product. Hover it to read the date of the newest entry in the catalog, how many were added in the last 90 days, whether ransomware campaigns use them, how the application was matched and the first CVE numbers. These are the applications to update first; Update works for them as for any other application. While the computer has such an application it also has an Exploited vulnerability (KEV) alert, which is listed in the Alerts module and on the Dashboard rather than in the Open alerts card of the Overview. The Vulnerabilities section explains the catalog and the module that lists these applications for every computer.

> **i** Update, Remove and Install reply immediately with "Processing…" and keep running on the machine for minutes. Check the outcome in the Remote actions tab (command queue and history) and click Refresh list afterwards.

Events shows Windows events (7 days) + agent errors, split into All, Application, Security, Setup, System and Agent errors (the agent's own crash history, which is kept permanently), each with its count, and a Refresh list button. Columns: When, Level (Critical, Error, Warning, Information, Verbose),

Channel, Source, ID and Message; long messages expand when clicked. The Level dropdown and the Source, ID and Message boxes filter the table. Every row has an Ask the AI button that opens the assistant with the event already pasted, so it can explain the cause and, if needed, investigate on the machine with its tools.

The Processes tab

On a Windows computer the Processes tab is a live task manager. While it is open it shows every process running on the computer and refreshes every 2 seconds (the interval can be changed next to Pause); Live and “updated N s ago” say the list is current. Above the table are the computer’s total CPU, memory used out of the total, and the number of processes. The list stops updating when you leave the tab or the browser window is hidden, and starts again when you come back; Pause stops it by hand and Resume starts it again.

- For each process: Name, PID, User, CPU (share of the whole computer), Memory, Started and Description / Path. Click a column heading to sort (by CPU by default). Search processes… searches the name, PID, user, description and path, and the row of filters under the column headings narrows by Name, PID, User (several at once), CPU (>= 1, 5, 20 or 50 %), Memory (>= 100 MB, 500 MB or 1 GB) and Description / Path; they all combine and keep applying as the list updates. × Clear filters, in that row, clears the search and every filter.
- While the pointer or the keyboard focus is on the list, the order is held so rows do not jump while you read them (Order held while you are on the list).
- End task (the × of a row, or the right-click menu) asks for confirmation naming the process, then ends it. It is recorded in the Queue with the technician who asked, and a notice says whether it worked. The agent first checks that the PID still belongs to the same process, so it never ends another program that took over the number. Critical Windows processes and the agent itself cannot be ended: their button is disabled.
- The right-click menu also has Copy PID and Copy path.

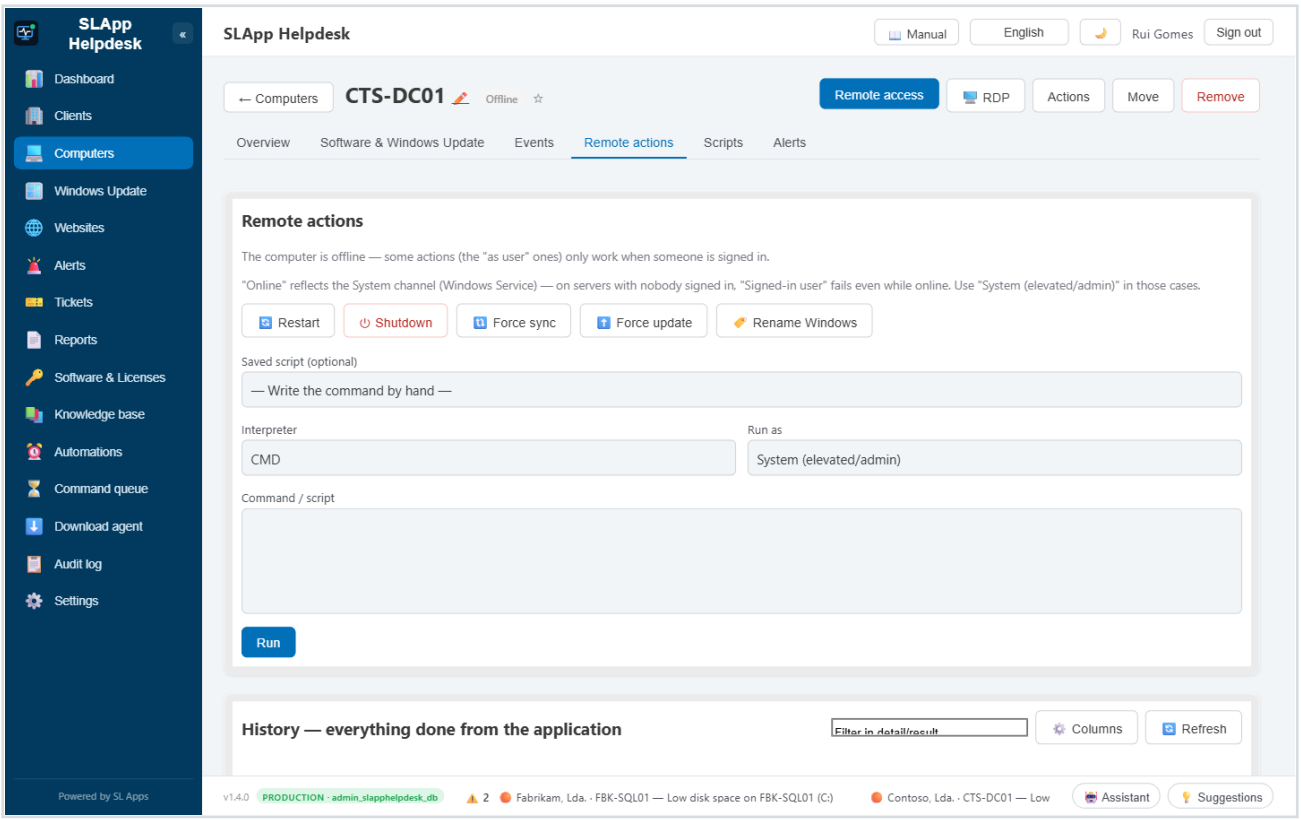
> **i** The live list needs agent 1.2.15 or newer, the computer’s service connection (the computer must be online) and Windows; other computers show a short notice instead. The list travels over the real-time connection and is not stored, so it does not fill the command history; several technicians looking at the same computer share one collection.

Windows services

The Services tab lists the Windows services of the machine — display name, service name, state, startup type and the account they run as. The list is collected only when you ask for it: click Scan services and the agent sends a fresh snapshot (nothing is polled on a timer; a real machine has hundreds of services). Each row has an actions menu with Start, Stop, Pause, Resume and Restart, plus a Startup type submenu to set Automatic, Manual or Disabled. Actions go through the command queue like any other remote command — an offline machine receives them when it comes back — and the row shows a spinner until the result arrives, with the real Windows error (for example a dependent service still running) if it fails.

> **i** To be alerted when a service stops, create a rule of type Windows service state in the Alerts module — the agent then watches that service continuously on that machine.

Remote actions, queue and history



The Remote actions tab: quick buttons, the command form, the command queue and the history.

The quick buttons are Restart (sent through the user's session so they see a warning with a countdown; if nobody is signed in it falls back to System and restarts without warning, and if the user accepts but their session is not allowed to shut the computer down — a standard user on Windows Server or Remote Desktop — it restarts through System straight away. The bulk Restart in Actions on N follows the same rule on every computer), Shutdown, Force sync (the agent re-sends its inventory now) and Force update (checks for a new agent version; if there is one the agent restarts itself). Restart, Shutdown and Force update ask you to confirm. For anything else use the command form.

1. Optionally pick a Saved script (optional) from the library — it fills in the command and the interpreter.
2. Choose the Interpreter: CMD or PowerShell on Windows; bash, zsh or sh on Linux/macOS.
3. Choose Run as: System (elevated/admin) runs through the service, which is always connected while the machine is online; Signed-in user runs in the user's session and fails when nobody is signed in.
4. Type or paste the Command / script and click Run. The output appears below the form when the machine replies.

Above the Command / script box, Generate with AI lets you describe what you want in plain language — for example "list the 20 largest files in the user's profile" — and the assistant writes a matching script into the box for the interpreter you chose. Read it, adjust it if needed, and run it like any other command; nothing is executed until you click Run.

If the computer is offline (or the chosen channel is not connected), the command is not lost: it goes into the Command queue for this computer, shown with Action, Channel (System or User), Status (Waiting, Sent, Running, Completed, Failed, Cancelled), Created and Result, and is delivered on its own as soon as the agent connects. You can Cancel a command while it is still Waiting, and Refresh reloads the queue. The History — everything done from the application card lists every action and remote access session with When, Type, Technician, Detail and Result (plus its own Refresh button); sort by clicking the When, Type or Technician headers, and filter by type, technician or free text. Clicking a detail or

result opens the full text with Ask the AI, Copy and Close.

> **i** On a Windows workstation that shows restart pending, the agent can also ask the signed-in user to restart every few hours, when Ask the user to restart is switched on (Settings › General › Computers › Restart prompt, agent 1.2.41 or later). The History card then lists Restart prompt shown, Restart postponed by the user and Restart accepted by the user, with agent in the Technician column and the Windows user in Detail; see Asking the user to restart in the Settings section.

> **i** On servers with nobody signed in, Signed-in user commands fail even while the machine shows Online. Use System (elevated/admin) — it is the default for that reason.

Remote access (disabled while offline) lets you choose SL App (native) — the built-in screen-sharing session —, Remote Desktop (through the agent) or a third-party tool detected on the machine: TeamViewer and AnyDesk open your local installation with the ID pre-filled, Remote Desktop Connection downloads a .rdp file for the computer's local address (so it only works inside the client's network, or over a VPN, as its tooltip says), Chrome Remote Desktop opens the web portal. Remote Desktop (through the agent) comes right after the SL App (native) block, with RDP tunnel through the agent: works outside the client's network. under it: it opens a Remote Desktop tunnel through the agent, so it works from outside the client's network. It is the same entry in every Remote access menu — the computer's page, the ticket's header and its Computer field, the client's list of computers — and in the Actions menu of the Computers list (also on right-click); it is not offered on a phone or tablet, which cannot run Remote Desktop. Up to version 1.14.73 it was a separate RDP button on the computer's page. On Windows the app creates the session and opens Remote Desktop right away, with nothing to download and no console window: Remote Desktop goes straight to the Windows sign-in popup, and the dialog says Remote Desktop is opening on this computer and connects to <target> through the agent. The first time, the browser asks whether the site may open the SL App agent: tick Always allow and confirm. Open again opens it once more. This needs SL App agent 1.2.40 or later on your own PC; for an older agent use Download launcher (.cmd), which downloads the small RDP-<computer>.cmd launcher of earlier versions: double-click it and Remote Desktop connects by itself, while its black "RDP Tunnel" window keeps the connection alive (close it when you are done; Download again fetches it again). On macOS or Linux the app downloads that launcher straight away. The agent opens only links whose tunnel goes through an SL App Helpdesk server it is connected to - any other link is refused with a message before Remote Desktop starts - and also refuses, opening nothing, a link created in an organization your PC does not belong to: if your PC is registered in another organization, use Download launcher (.cmd). When the tunnel is refused, a small SL App - Remote Desktop window says why: the link expired, you are no longer allowed to open tunnels, the agent service of the target is not connected, or someone else is connected. Either way the target's service channel must be connected, and the link or launcher stays valid for 8 hours. Opening an RDP or HTTP tunnel needs the tunnel permission (Settings › Users; every technician has it until an administrator blocks it): without it the entry is greyed out, with You are not allowed to open RDP/HTTP tunnels. Ask an administrator (Settings › Users). under it; it is also greyed out, with The agent service is not connected on this computer — the RDP tunnel cannot be opened., while the computer's agent service is not connected (or the computer is offline, in the Computers list). The Remote access menu checks both each time it opens; the Actions and right-click menus of the Computers list grey it out only while the computer is offline or when you may not open tunnels (as loaded with the page), and if the agent service is not connected the dialog shows that message. The server checks both again when it creates the tunnel. Each tunnel opened is recorded in the Audit log on the computer, and a failure shows the notice Could not open the RDP tunnel: …. Under SL App (native) the menu has two boxes for the session started now. Connect without asking for authorization appears only to a user whose profile has Can connect without asking for authorization (off for everyone, administrators included, until an administrator turns it on); without it the

computer's user is always asked, and if the request still asks to skip it the session starts asking for authorization with the notice You may not connect without authorization: the computer's user is asked to allow the session. Record session starts ticked when Settings › General › Remote control records every session automatically; only a user with Can turn off session recording may untick it. For everyone else it stays ticked and greyed out with Your company records every session; only a user with permission can turn recording off., and the Stop button of the session toolbar is greyed out while it records; if a start request still asks not to record (for example from a menu opened before the permission was removed), the session is recorded anyway, with the notice Your company records every session and you may not turn recording off: this session is recorded. On a browser that cannot record, such a user cannot start SL App (native): the entry is greyed out with Your company records every session and this browser cannot record: use Chrome, Edge or Firefox., and starting the session from elsewhere in the app shows the same message instead of the session. Ticking Record session for a session the company would not record is allowed to everyone.

A computer shows one remote session at a time, and the session already running is never interrupted. If another technician of your organization is already in an SL App (native) session on it, yours does not start and the window says who and since when, for example Ana Silva is already connected to this computer (since 14:05). Try again when that session ends. On a computer that is also managed by another organization, you see that organization's name and the time instead — never the name of its technician: Someone from Contoso is already connected to this computer (since 14:05). Try again later. (without the time when the agent does not know it). Opening the session again yourself (another window, a reload) still takes your own session over, as does starting one while the other is still waiting at its authorization popup or nobody is watching it any more. A session nobody is watching also ends on its own: closing its window ends it at once, and one whose page went away without ending it (a browser that closed unexpectedly, a computer gone to sleep) ends after two minutes, so the computer is never left busy for nobody. The name of the other organization and the time need agent 1.2.39; an older agent only says The computer is in a remote session of another organization. Try again later.

> **i** If the Windows session you are viewing is disconnected — someone signs in to the computer with Remote Desktop, or uses Switch user — the window keeps the last picture with the explanation The Windows session on this computer was disconnected (someone signed in with Remote Desktop or switched user). The picture comes back when that session is resumed., also shown as a notice, instead of going black. When that session is back you are told The Windows session on this computer is back - the picture resumed. and the session simply goes on. The system screen is not used meanwhile: it would show the session that took the console, which is not the one you were authorized to see (agent 1.2.39).

A Windows client computer (not a server) has only one interactive session at a time, so a Remote Desktop sign-in would disconnect the session someone is watching there. Remote Desktop (through the agent) to such a computer is therefore refused while another technician has a remote session on it: the dialog and the toast say who is connected and since when, followed by Remote Desktop would disconnect that session. A remote session of another organization is refused by the computer's agent itself (agent 1.2.39): the SL App - Remote Desktop window (with the .cmd launcher, its black RDP Tunnel window) then says Someone is already connected to this computer (a remote session of another organization). Remote Desktop would disconnect it - try again when that session ends. The other way round, while your RDP tunnel to the computer is open — and for one minute after its last connection closes, because Remote Desktop opens several connections while signing in — the computer counts as busy, so a remote session of another organization started meanwhile is told so instead of going black. Servers and other targets on the network are not affected.

The Scripts tab

The Scripts tab holds scripts written for this specific machine — unlike the shared script library

(Settings) and Automations, which target a client, a folder or the whole fleet, each script here carries its own body and belongs to this one computer. Typical uses: a cleanup routine tuned to one server, a vendor-specific health check, or a fix that only ever applies to this box.

Add script opens the editor: give it a name, pick the Interpreter (PowerShell or CMD), choose Run as (System (elevated/admin) or Signed-in user), and write the script body. The Schedule field decides how it runs: Manual only (the default — it only runs when you press Run), or Once, Daily, Weekly or Monthly together with a Next run date and time.

- Run now — runs the script immediately. While the machine is online you get the output right away, rendered in a terminal panel; if it is offline (or the command channel is not connected), the run is queued and delivered when the machine reconnects, and the message tells you which of the two happened.
- Schedule — scheduled runs are dispatched by the server every minute through the same store-and-forward command queue as Automations, so an offline machine still receives the run once it comes back. A Once schedule disarms itself after it fires; recurring ones move Next run forward automatically.
- Green dot / pause — click the dot next to a schedule to pause or resume it without touching anything else. Saving the edit form with a schedule always re-arms it.
- Next run, Last run and Last result columns show the schedule state at a glance; manual and scheduled runs also land in the computer's command history, like any other remote command.

> ℹ Scheduled times are entered in your local time zone. And the same caveat as Remote actions applies: on servers with nobody signed in, a Signed-in user script fails even while the machine is Online — use System (elevated/admin) unless the script genuinely needs the user's session.

The Backups tab

Administrators also see a Backups tab on each computer's detail page. It is the Backups module narrowed to this computer: Jobs lists the backup jobs that run on it, where you create, edit, run, pause and delete them exactly as in the module — a new job already starts on this computer — and Runs lists their runs, with the same detail window. View runs on a job switches to Runs on that job. Technicians who are not administrators do not see the tab.

Chatting with the person at the computer

The Chat button in the header of the detail page opens a conversation with whoever is signed in on that computer. Only a technician can start a conversation: your first message starts it, and the person at the computer answers it in a window the agent opens on their screen, but can never start one. Every technician can start a conversation; a computer has at most one open conversation at a time, shared by the technicians who open it.

- The line under the title says whether someone is signed in on the computer. If nobody is, or the agent is offline, the message waits and appears as soon as someone signs in. Each of your messages shows a check mark once the computer has received it and a clock while it has not.
- Your messages are on the right, the person's on the left. The window refreshes by itself every few seconds and marks the replies as read. Enter sends, Shift+Enter starts a new line.
- Attach — or Ctrl+V of a screenshot or of copied files — adds files to the next message: up to 5 per message, each within the size and the types allowed in Settings › General › Remote control › Files in conversations. A message can be just files. Images show as thumbnails (click to open them), other files as links to download.
- End conversation closes it: the window on the computer then says that the technician ended it and

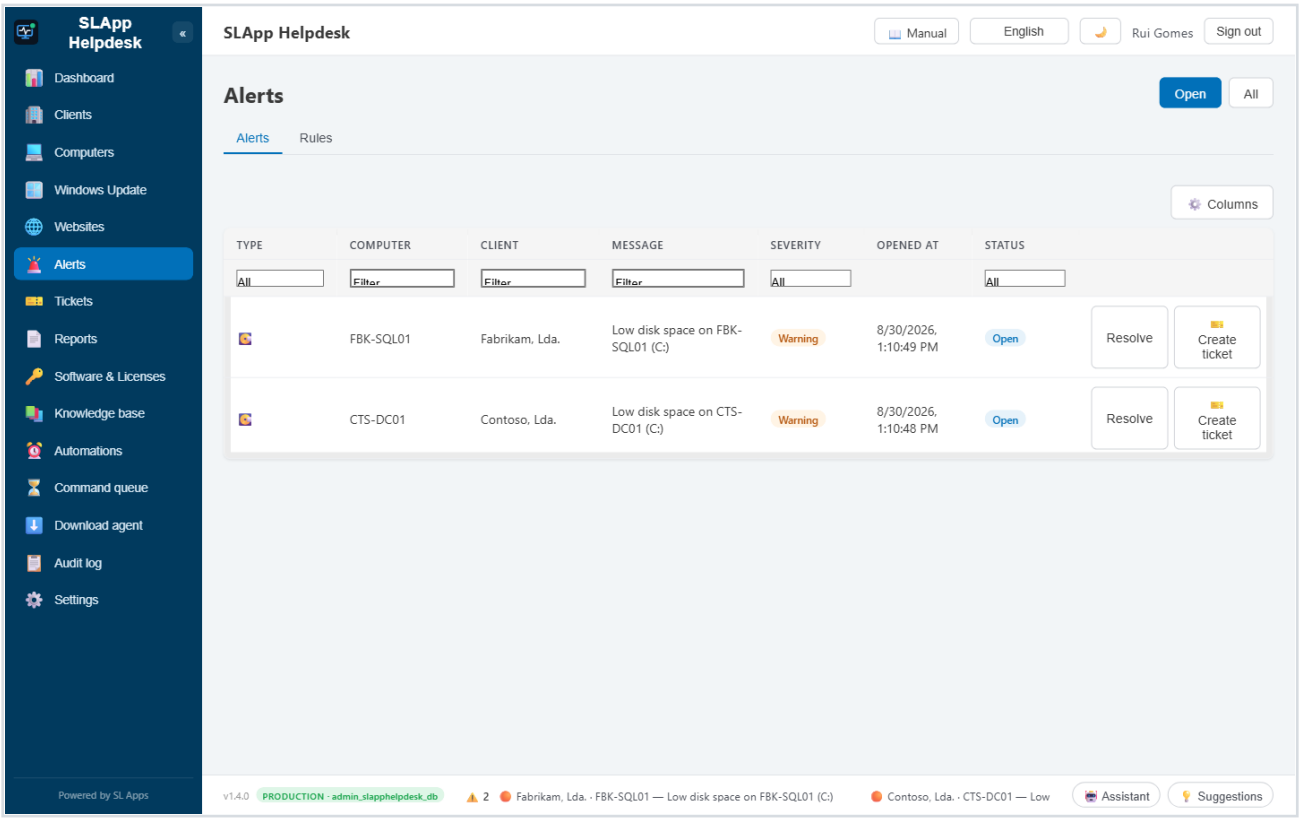
stops accepting replies.

- When a reply arrives while you are elsewhere in the app, a speech bubble counter appears at the top of the app and a notice pops up in the corner; clicking either opens that computer on its Conversations tab with the conversation open. The conversations you started are listed first.
- The Conversations tab of the detail page keeps every conversation with the computer: when it started, who started it, who answered on the computer, how many messages it has and when it ended. Click an ended one to read it; New conversation (or Open the current conversation) opens the chat.

> i On the computer, a file received in a conversation is never run: an image opens in the default viewer and anything else is only saved where the person chooses. Conversations need agent 1.2.2 or later on Windows; files need 1.2.4 or later.

6. Alerts

The Alerts module is where SLApp Helpdesk tells you that something on a monitored computer or network device needs attention: a disk filling up, a machine that stopped communicating, an antivirus that went silent, a NAS volume almost full. It also gathers the client alerts that other modules raise about your clients' domains, Microsoft 365, vulnerable software, hardware lifecycle, network device configurations, Active Directory and third-party backups. Open it from the sidebar entry Alerts. It is available to every user of the team and has two tabs: Alerts (the live feed) and Rules (what gets monitored, where, and how).



The Alerts feed with open computer and device alerts

The alert feed

The Alerts tab lists alerts in a table that refreshes itself every 30 seconds. Open alerts come first, newest at the top, followed by resolved ones when you choose to see them. Two buttons at the top of the page switch the view: Open (the default, only alerts still active) and All (open plus resolved history).

When there is nothing open you see the message No open alerts — all good; under All, if nothing was ever recorded, the message is No alerts recorded.

- The first column is an icon for the alert type: disk, memory, CPU, network, offline, antivirus, firewall, service state, event log, and for network devices also system, RAID, volume, temperature, interface, VPN, update and printer supply; website alerts (down, slow, SSL) use a globe, snail and lock icon, and each type of client alert has an icon of its own (see Client alerts below).
- Computer: the computer concerned. The name is a link that opens the computer's detail page directly, so you can check its metrics or act on it without leaving the feed. An alert that is not about a computer (a website alert, most client alerts) shows a dash here.
- Client: the client that owns the computer, or the client a client alert is about.
- Message: a short description with the measured value, for example Disk C: usage at 83.2%, or Computer stopped communicating. Device alerts start with a satellite-dish icon and the device name, for example NAS-01 — Volume volume1 at 91%. Client alerts start with the icon of their source (a link for a domain) and, unless they are about a computer, the item concerned, for example contoso.com — DMARC record missing.
- Severity: Critical (red badge) or Warning. While an alert stays open its severity and message are updated on every evaluation, so a Warning can become Critical without a new row appearing.
- Opened at: when the alert was first raised, in your local time.
- Status: Open or Resolved.
- Actions: Resolve (only on open alerts, and on a client alert only for administrators) and Create ticket.

Backup alerts have their own icons in the first column: a floppy disk for Backup failed, an alarm clock for Backup missed and a file cabinet for Backup destination space (see the Backups section).

Four more columns are hidden at first; turn them on with Columns: Value, Threshold, Resolved at and Source. Source says where each alert comes from (Computer, Device, Website, Domain, Microsoft 365, Vulnerability, Lifecycle, Network config, Active Directory or Third-party backup), and its filter, in the row of filters under the headers, shows the alerts of one source only.

Since version 1.14.81 alert messages are shown in the language of the app, here and wherever alerts are listed: the footer ticker, the Dashboard, a computer's sheet and its Alerts tab, a client's Alerts tab, Infrastructure › Overview and the open alerts in the details of each module. They are still stored in English: Copy message, in the row's right-click menu, copies the original English text, sorting by Message follows the English text, and the filter of the Message column finds an alert by the text you see or by the English one. Create ticket fills in the subject and the description with the message as you read it. A message the app does not know, such as a Syslog line or a text written by a device or by the Windows event log, is shown as it was written, and so are the names inside a message (computers, services, sites…). When you switch languages, Alerts, the ticker and the Dashboard follow at once; the other pages follow the next time they load. The alert emails are written in your company's email language, the message included.

Resolving an alert and creating a ticket from it

Most alerts resolve themselves: a usage alert closes as soon as the value drops back below the lowest threshold set in the rule, an offline alert closes when the computer reports in again, an antivirus or firewall alert closes when protection is detected again. Use Resolve when you have dealt with the situation and want the row out of the open list right away; the alert is marked Resolved and stays visible under All. Note that if the condition is still true at the next check, a fresh alert opens (and, if the rule sends emails, the email goes out again).

1. In the Alerts tab, find the alert and click Create ticket (the tooltip reads Opens Tickets with a new

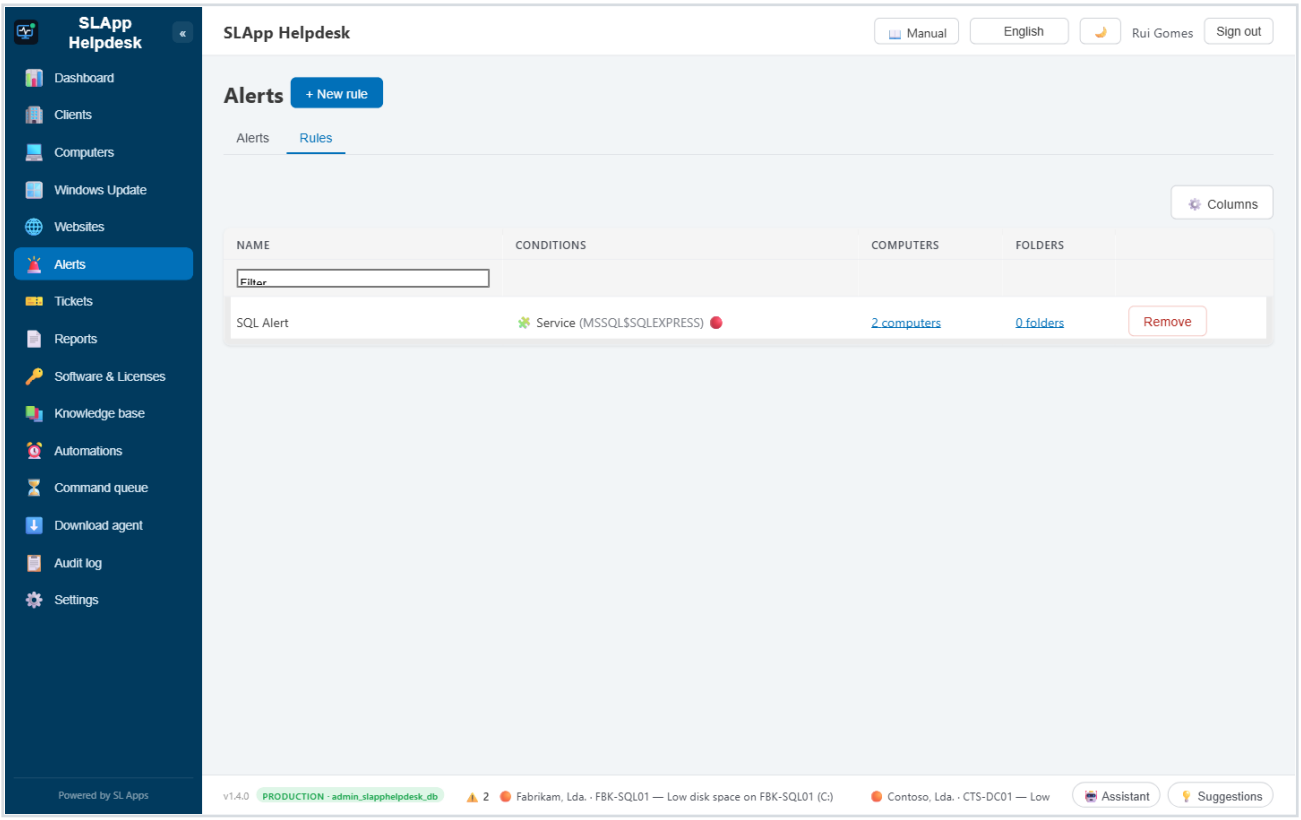
ticket already pre-filled from this alert).

2. You are taken to the Tickets module with a new ticket form already filled in: the client and computer are selected, the subject is the alert message followed by the computer name, and the description contains the computer, type, severity, message and the time the alert opened.
3. Adjust priority, assignee or description as needed and save the ticket as you would any other.
4. Back in Alerts, click Resolve on the alert if you no longer want it in the open list.

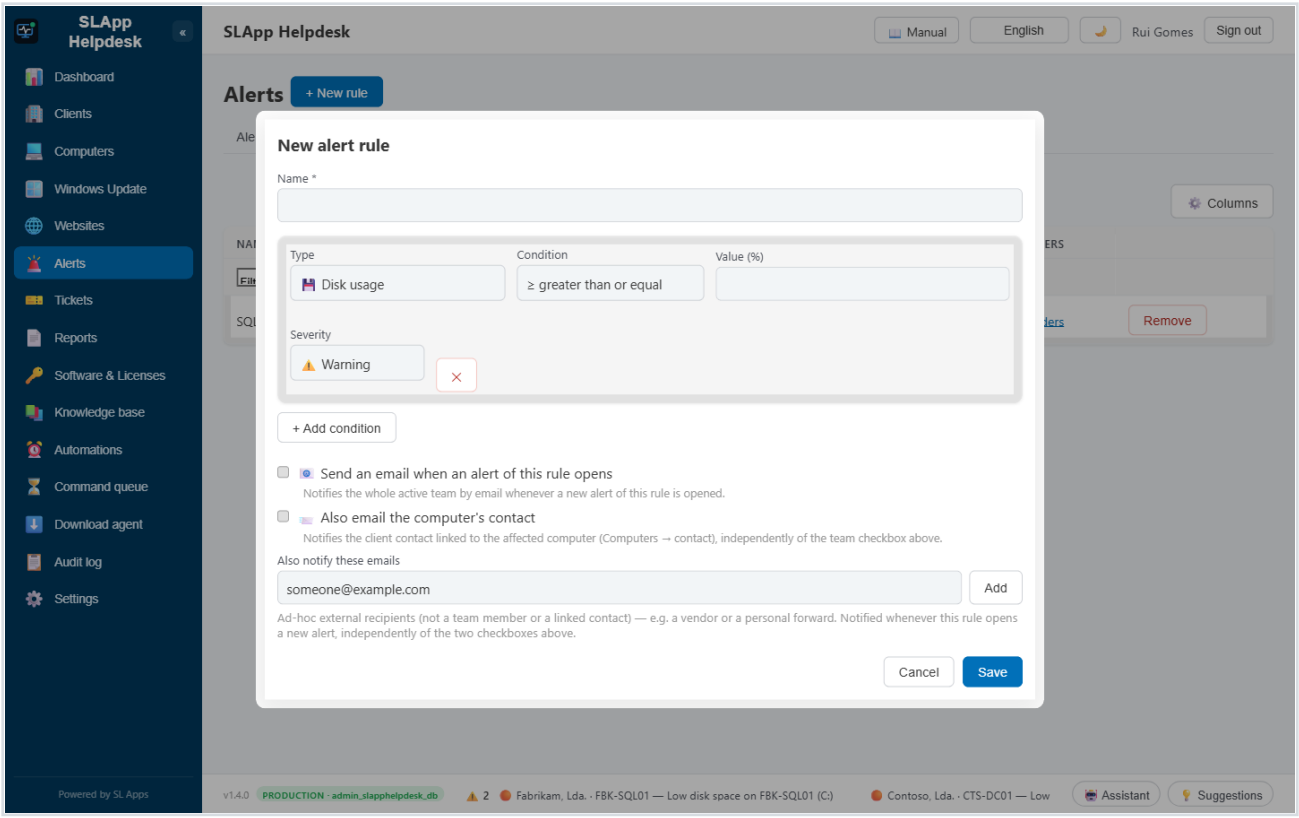
> **i** Creating a ticket does not resolve the alert automatically, and there is no option to open tickets automatically from a rule. Alerts and tickets are deliberately separate: the alert tells you what is happening now, the ticket tracks the work you decide to do about it.

Alert rules

Nothing is monitored until a rule exists and is linked to computers. The Rules tab lists the rules of your team, sorted by name, with the columns Name (with an envelope icon when the rule sends emails), Type, Warning, Critical, Computers and Folders, plus a Remove button. Until you create the first rule the tab only shows the message No rules — create one to start monitoring disk/memory/presence on some computers. Click the name, type or thresholds of a row to edit the rule; click the N computers or N folders link to change where it applies (hovering the folders link shows which folders are linked).



The Rules tab: each rule with its type, thresholds and the computers/folders it applies to.



Creating an alert rule

The rule form has these fields:

- Name *: a label of your choice, required (Save stays disabled until it is filled in).
- Type: one of Disk usage, CPU usage, Memory usage, Network traffic, Computer offline, Antivirus/EDR disabled, Firewall disabled, Windows service state or Windows event log.
- Warning threshold and Critical threshold: shown only for the usage types. Disk, CPU and memory are percentages (0 to 100). Network traffic is in Mbps and is the total traffic (sent + received) of all active adapters. You can leave one of the two empty; an alert opens when the value reaches the warning threshold and becomes Critical when it reaches the critical threshold. With only a critical threshold the alert opens directly as Critical; with both empty the rule never alerts.
- Send an email when an alert of this rule opens: when checked, every active user of your team with an email address who has access to Computers or Alerts (administrators always; see Who receives the alert emails below) receives a message the moment a new alert of this rule opens (not on every update, and not when it resolves).
- Also notify these emails: external addresses (added as chips, one at a time) that receive this rule's alert emails in addition to the team — for example a client's own IT contact. They receive them as typed, even when an address belongs to a technician without access to Computers.

Each type behaves slightly differently. Disk usage raises one alert per fixed drive (C:, D:, ...) and the same thresholds apply to all drives; an alert for a drive that was removed closes on its own. CPU usage is evaluated on the average of each batch of samples the agent sends, to smooth out momentary spikes, while Network traffic uses the peak of the batch. A Windows agent from 1.2.0 that could not reach the server sends the samples it kept once it can: they fill the charts, but only the current samples open, update or close a usage alert. Memory usage (and disk usage) is checked on every agent heartbeat. Computer offline has no thresholds: a server-side sweep runs every minute and opens a Critical alert for any linked computer that has reported at least once and has not reported for more than 5 minutes, closing it again as soon as the computer reports in. Antivirus/EDR disabled opens a Critical alert when the agent's security snapshot shows no enabled protection product, and Firewall disabled

does the same when Windows Firewall is off or only partially enabled across its profiles. Software updates available counts the applications for which winget reports a newer version (the Available column of the computer's Software tab) and is checked every time the agent sends its software inventory. The first inventory of a computer that says which applications have a newer version, such as that of a computer whose agent was just installed or of a Windows in another language once its agent reads winget in that language, is its starting point: the alerts it opens notify nobody (no email to the team, to the computer's contact or to the extra addresses), and later alerts notify as usual. An inventory that does not say it, such as the list of installed programs from the Windows registry sent because winget could not be read this time, is no reading: the alerts stay exactly as they are until an inventory that says it arrives. A computer whose inventories never say it, such as a Windows Server without winget, is checked on its list as before.

> **i** The firewall check only understands the Windows Firewall. If a computer runs a third-party firewall, its state cannot be read and the rule neither opens nor closes alerts for it.

Service state and event log rules

Windows service state watches a specific Windows service on a specific computer. In the condition you pick the computer and then the service — the picker lists the services of that machine, so use Scan services in its Services tab first if the list is empty. The agent only monitors services that a rule points at it (there is no blanket polling), reports their state with every heartbeat, and the alert opens when the service is not running — stopped, paused or not installed — and closes on its own when it runs again.

Windows event log fires on entries of the computer's event logs (Application, System, Security, Setup). Fill in any combination of Event ID, Source, Level (Critical, Error, Warning, Information) and a text the message must contain — all filled filters must match. This type must be the only condition of its rule. It is a point-in-time alert: the first matching event opens the alert and sends the emails, further matches increase a counter on the open alert, and it never resolves on its own — click Resolve when you have dealt with it. A Windows agent from 1.2.0 keeps the matching events it finds while the server cannot be reached and reports them when it is back, each one once, so an event that happened during an outage still opens its alert, up to 7 days back (see What the agent does on the computer in the Download Agent and installing on computers section). Two exceptions: right after the update from 1.1.x the agent starts only 5 minutes back, and a log that is flooded with matching events and falls more than an hour behind skips ahead, so those events are not reported.

Website alerts (down, slow response, SSL expiry) are not configured here: they are set per site in the Websites module, and simply appear in this feed alongside everything else. Backup alerts work the same way: the Backups module raises them by itself and there is nothing to set up here. There are three types, Backup failed, Backup missed and Backup destination space, raised when a backup fails or does not run, or when a destination is running out of room; the Backups section explains when they open, get worse and close, and who receives the email.

Choosing where a rule applies

A rule on its own does nothing: only the computers explicitly linked to it, directly or through folders, are evaluated. There are two ways to link, and you can combine them.

1. In the Rules tab, click + New rule, fill in the form and click Save.
2. In the new row, click the link in the Computers column (it reads 0 computers at first). The dialog Computers — rule name lists every enrolled computer with its client in parentheses. Tick the ones you want and click Save. Only the checked computers evaluate this rule.
3. Alternatively, or in addition, click the link in the Folders column. The dialog Folders — rule name lists every client folder you have created, each with its client in parentheses. Tick folders and click

Save: the rule applies to every computer inside the checked folders, including those that join them later.

4. To stop monitoring, untick the computer or folder and save, or click Remove on the rule and confirm. Open usage, antivirus and firewall alerts created by that rule are closed at the computer's next evaluation; an open Computer offline alert stays open until the computer reports in again or you click Resolve on it.

> **i** Folders are the easiest way to cover a whole client: create folders for the client in its detail page (Computers tab), place its computers in them, and link the rule to those folders. New computers added to a folder are picked up automatically, which direct links do not do.

> **i** If several rules of the same type reach the same computer (for example one by direct link and one by folder), the rule with the lowest warning threshold is the one applied. In the Rules table the thresholds of Network traffic rules are displayed with a % sign, but they are Mbps values.

Alerts in the client detail page

Each client also has an Alerts tab in its detail page, showing the alerts of that client's computers and devices, the alerts of the websites assigned to that client and the client alerts raised for that client, with the same columns as the main feed (minus the client). The heading carries a red badge with the number of open alerts, and the checkbox Include resolved adds the resolved history. You can click Resolve there too (on a client alert only if you are an administrator), and Create ticket sits next to it.

Create ticket is offered next to Resolve wherever an open alert is listed: the Alerts module, a computer's Alerts tab, a client's Alerts tab and the open alerts in the details of Active Directory, Exchange, SQL Server, Storage / iSCSI, Failover cluster, DHCP, IIS, Domains & email, Third-party backups, Microsoft 365 and Network configs (as a button and in the row's right-click menu). Unlike Resolve, every technician gets it. It never creates the ticket by itself: it switches to Tickets with a new ticket already filled in with the client, the computer when the alert is about one, a subject made of the message and the server, domain or device name, and a description with the type, severity, message and opening date; you review it and click Create.

Device (SNMP) alerts

Alerts for NAS units, routers and other SNMP devices are not configured in the Rules tab. They are set per device, in the Devices sub-tab of the computer that polls the device, when you add or edit the device. The alerts they produce appear in the same feed as computer alerts, linked to the polling computer and with the device name at the start of the message, and you can Resolve them or create tickets from them in exactly the same way. A new device starts with alerts on, Volume warning 85% and critical 95%, Temperature 60 °C, CPU off, the interface/VPN and firmware checks on, and email off. On top of these settings, an administrator can stop the alerts of every SNMP device at once with Generate alerts: on the Devices tab of a computer or a client for NAS units, routers and the other devices, and in the Printers module for printers (see Turning alert generation on or off per module below).

- Raise alerts for this device: the master switch. Turning it off closes every open alert of the device at the next poll.
- Critical health conditions always alert when the switch is on: no SNMP response, a system status other than Normal, a degraded RAID or storage pool, and a failed or crashed disk.
- No SNMP response: a device that does not reply in time is asked a second time (agent 1.0.43 or later) before it counts as offline. If it still stays silent, a Critical alert opens with a message that says what to check, for example "No SNMP response: no reply (timeout). Check SNMP is enabled on the device, the community, v2c and UDP 161 access." (the port is the device's own). When the device is

up but nothing listens on that UDP port, the message reads "No SNMP response: SNMP port unreachable: the device is up but nothing listens on UDP 161. Check SNMP is enabled on the device and the port." The alert closes by itself at the first poll that gets an answer. The Computers section has a checklist under When a device shows No SNMP response.

- Reboot detection: always on when the master switch is on, with no setting of its own. When a device reports a shorter uptime than at the previous poll, a Critical alert opens saying Device rebooted and how long it had been up. This also catches a reboot shorter than the polling interval, which never shows as offline. The alert closes by itself at the next poll. Being Critical, its alert e-mail is sent as CRITICAL.
- Volume — warning (%) and Volume — critical (%): usage thresholds for each volume. Leave a field empty to turn that threshold off.
- Temperature — warning (°C): applies to the chassis and to each disk.
- CPU — warning (%): off unless you enter a value.
- Interface/VPN that was up and goes down (routers): raises an alert only for an interface or VPN tunnel that was up at the previous poll and is now down, never for ports that are always empty. A WAN interface going down is Critical, other interfaces are Warning, a VPN tunnel is Critical. When they come back up the alert closes and, if emails are enabled, a recovery email is sent.
- Firmware update available: a Warning when the device reports an update.
- Email the team when an alert opens: sends an email, when new device alerts open, to the active users of your team who have access to the device's module — Printers for a printer, Computers for any other device — or to Alerts, and to the administrators (see Who receives the alert emails below).
- Printers (type Printer): the form shows only Supply — warning (%) and Supply — critical (%), 15 and 5 by default (leave a field empty to turn that level off), and Email the team when an alert opens. Each toner, ink cartridge or drum has its own alert: Warning below the first threshold, Critical below the second, and Critical as well when it is empty (Warning if only the warning threshold is set); it closes once the supply is back above the thresholds. A waste toner box never alerts, and a printer raises no reboot or interface alert; No SNMP response works as for any device. The supplies are read by the Windows agent from version 1.2.29 (see the Printers section).
- While a device is offline: with the master switch on, the alerts that are already open stay open, and nothing is reported as back up, until the next poll that gets an answer checks everything again. The last readings before the outage are kept as the point of comparison, so a reboot, or an interface or VPN tunnel that went down during the outage, is still reported when the device answers again.

> ℹ In the device detail you can edit the interfaces (Choose which ones alert + internal name) to give a port an internal name (shown in the alert message) or switch off the alert for a specific port, for example a backup WAN that is normally unplugged.

Client alerts

Several modules watch things that the alert rules do not cover (your clients' domains and Microsoft 365, the software and hardware lifecycle of their computers, their network devices, Active Directory and backup software) and raise alerts of their own, called client alerts in this manual. There is nothing to set up in the Rules tab: each module decides when its alerts open, get worse and close, and its section of this manual gives the details. They appear in this feed next to the other alerts, in the client's Alerts tab, in the Open alerts card of the Dashboard and in the footer ticker. The Source column tells them apart, and the Type filter uses the names below:

- Domain (Cloud & web › Domains): Domain expiring (Warning, Critical in the last 7 days) and Domain expired (Critical); SPF record and DMARC missing (Warning, when the domain has Alert on SPF

and DMARC problems on); MX changed and Nameservers changed (Critical); Registration lookup failing and DNS lookups failing (Warning, after 3 failed checks in a row).

- Microsoft 365 (Cloud & web › Microsoft 365): Administrator without MFA (Critical); Users without MFA, Licenses exhausted and Mailbox almost full (Warning); Microsoft 365 connection failing (Warning, after 3 failed syncs in a row); Microsoft 365 app secret expiring, when the secret could not renew itself (Warning, Critical in its last 7 days and once expired).
- Vulnerability (Automation & updates › Vulnerabilities): Exploited vulnerability (KEV), one Critical alert per computer that has outdated software with actively exploited vulnerabilities, listing the applications. It closes by itself once they are updated or removed.
- Lifecycle (Infrastructure › Lifecycle): OS support ending (Warning, within 90 days), OS support ended (Critical) and Warranty expiring (Warning, within the days set in Settings › General › Lifecycle).
- Network config (Infrastructure › Network configs): Network config backup failing (Warning, after two scheduled backups in a row), SSH host key changed (Critical) and Network configuration changed (Warning by default; the device's Alert when the configuration changes can make it Critical or Off, and a change announced with Expect a change… raises none).
- Active Directory (Infrastructure › Active Directory): AD: dcdiag failed, AD replication failing, AD service stopped, DC clock offset, SYSVOL not shared and New privileged account (all Critical); AD check failing (Warning, after 3 failed checks in a row).
- Third-party backup (Infrastructure › Third-party backups): Third-party backup result (Critical when a job failed, Warning when it finished with warnings), Third-party backup overdue (Critical, no successful run for longer than the job's expected frequency allows) and Backup check failing (Warning, after 3 failed checks in a row).
- Hyper-V (Infrastructure › Hyper-V): Hyper-V replication, one alert per virtual machine, Critical or Warning like its replication health (see Replication alerts in the Hyper-V section).

A vulnerability or lifecycle alert is about one computer: its Computer column links to it and Create ticket selects it, as for a computer alert. A Hyper-V replication alert is about a virtual machine: its message names the virtual machine, and its Computer column is the Hyper-V host that reported it, which Create ticket selects. The other client alerts are about a domain, a Microsoft 365 user or license, a network device, a domain controller or an AD domain, or a backup job: the Computer column shows a dash, the message names the item, and Create ticket selects only the client. On a network configuration alert the right-click menu also offers Open the network device and, for a configuration change, Show the changes, which opens the comparison with the previous version. Client alerts are not listed in a computer's own Alerts tab.

Client alerts are emailed to the administrators of your company (every active administrator with an email address, never a client contact; a technician only gets the Critical ones, after turning on Critical alerts in My profile › Options, and only for the modules they have access to — see Who receives the alert emails below): once when an alert opens, and once more if it goes from Warning to Critical. Refreshing an open alert or resolving it sends nothing. So that setting up a module does not flood your inbox, the first check of something new (a domain just added, a client's first Microsoft 365 sync, a network device's first backup, a domain controller or a backup job seen for the first time) opens its Warning alerts without an email (the Microsoft 365 connection and app secret alerts excepted), while Critical ones are emailed at once; vulnerability and lifecycle alerts found the first time a computer is evaluated, and Hyper-V replication alerts found the first time a host is evaluated, open without any email at all.

> ℹ When one pass opens several client alerts at once (typically the daily vulnerability pass after CISA lists a new exploited vulnerability, or the daily lifecycle evaluation), each administrator receives one summary email listing them, the first 50 and then how many more are waiting in the Alerts module, instead of one email per alert.

Most client alerts close by themselves once the condition is gone. Only administrators can resolve one by hand: technicians do not see Resolve on a client alert, and the server refuses it. Every manual resolve is written to the Audit log with the action alert_resolved, on the item the alert was about (on the computer, for a vulnerability or lifecycle alert). If the condition is still true at the next check, a new alert opens and is emailed again, with two exceptions: a vulnerability alert that comes back within 30 days with only applications it had already listed stays silent, and a Third-party backup overdue alert you resolved stays closed until the computer's next daily check. Alerts that report a change rather than a condition (MX changed, Nameservers changed, Network configuration changed and New privileged account) never come back by themselves: they close 7 days after the change, or when you resolve them, and every new change opens a new alert with its email.

Who receives the alert emails

Since version 1.14.80 an alert email only reaches the people of your team who have access to the module the alert belongs to. The rule applies to every list of your team's users that an alert emails: the administrators of the client alerts, the team of an alert rule, an SNMP device, a website or a backup job, and the users who turned on Critical alerts in My profile › Options. It only takes people out of those lists and never adds anyone: client alerts, for example, still go to the administrators, and to a technician only when they are Critical and the technician turned on Critical alerts. Who stays on the list:

- Administrators always receive the email: the permissions never apply to them.
- A technician whose modules include the module of the alert (listed below) receives it.
- A technician whose modules include Alerts receives every alert email: the Alerts module lists every alert of your company, so it counts as every module.
- A technician with no module rule (no group with modules chosen and no Allow or Deny of their own) has every regular module, Alerts included, and keeps receiving every alert email as before — the backup alerts too, through Alerts, although Backups itself is only given by explicit permission.

Each alert belongs to one module:

- The alert rules of the Rules tab, whatever their type (disk, CPU, memory, network traffic, computer offline, antivirus, firewall, Windows service, Windows event log…): Computers.
- SNMP devices: Printers for a device of the type Printer, Computers for a NAS, a router or any other device.
- Websites (down, slow response, SSL certificate, hosting expiry): Websites.
- Backup failed, Backup missed and Backup destination space: Backups.
- Client alerts: the module they come from — Vulnerabilities, Lifecycle, Network configs, Third-party backups, Active Directory, Exchange, SQL Server, Storage / iSCSI, Failover Cluster, DHCP, IIS, Domains, Microsoft 365 or Hyper-V.
- Syslog alerts belong to Syslog, but they send no email.
- Clients, the Dashboard and Infrastructure › Overview give no alert email, although a client's sheet shows its computers and the overview sums up every module.

ℹ For example, a technician whose group grants only Transactions, Websites, Tickets and Clients receives the website alert emails and none about Hyper-V, Active Directory, SQL Server, disks or printers. The addresses typed on an item are not filtered: those in Also notify these emails of an alert rule or of a website, and the computer's contact when a rule has Also email the computer's contact, receive the email as typed, even when the address belongs to a technician without the module. A change to the modules of a group or of a user applies to the next alert emails, at the latest 30 seconds later, and if someone's permissions cannot be read when an alert is sent, the email goes to them anyway rather than being lost.

Turning alert generation on or off per module

Since version 1.14.75 an administrator can choose, module by module, which modules generate alerts. Every alert that SLApp Helpdesk raises on its own belongs to one of these switches; the alert rules of the Rules tab do not, and keep their own Enabled / Paused switch. There is one switch for each of these modules, listed only when your plan includes the module:

- Vulnerabilities: Exploited vulnerability (KEV).
- Lifecycle: OS support ending, OS support ended and Warranty expiring.
- Network configs: Network config backup failing, SSH host key changed and Network configuration changed.
- Third-party backups: Third-party backup result, Third-party backup overdue and Backup check failing.
- Active Directory: the AD alerts of the domain controllers and the domains.
- Exchange: the Exchange alerts of the daily check.
- SQL Server: the SQL Server alerts of the daily check.
- Storage / iSCSI: the disk, iSCSI target and volume alerts of the daily check.
- Failover Cluster: the alerts of the cluster itself and those of each node (sources Failover Cluster and Failover Cluster (node)).
- DHCP: the alerts of the servers and of their failover relationships (sources DHCP and DHCP (failover)).
- IIS: the IIS alerts of the daily check.
- Hyper-V: Hyper-V replication, one alert per virtual machine whose replication health is Critical or Warning. While it is off, the module also stops asking the hosts for an inventory every half hour on its own (opening a host still refreshes it).
- Domains: the domain alerts (expiry, SPF and DMARC, MX and nameserver changes, lookups failing).
- Microsoft 365: the Microsoft 365 alerts (MFA, licenses, mailboxes, the connection and the app secret).
- Backups: Backup failed, Backup missed and Backup destination space.
- Syslog: the Syslog alerts of the collectors.
- Printers: the alerts of the SNMP devices of type Printer (supplies, No SNMP response).
- Network devices (SNMP): the alerts of every other SNMP device — NAS units, routers and firewalls, and Generic (SNMP) devices such as switches or UPS units. It has no page of its own: its switch is on the Devices tab of a computer and of a client, and in Settings.
- Websites: Website down, Slow response, SSL certificate and hosting expiry.

The switch is the Generate alerts checkbox at the top of each module's page, next to its buttons (on the Devices tab of a computer or a client for Network devices (SNMP), and next to the title of the list of

hosts for Hyper-V), and all of them are together in Settings › General › Alerts (see the Settings section). Every technician sees it, but only administrators can change it: for anyone else it is greyed out and its tooltip reads Only administrators can change this option. While a module is not generating alerts its checkbox stands out on an orange background, so that a quiet page is not taken for a healthy one.

Unticking the box asks Stop generating <module> alerts? and says how many open alerts the module has; Stop generating confirms (Cancel has the focus). At that moment every open alert of the module is resolved, in your name, and from then on the module opens no new alert, refreshes or escalates none, and sends no e-mail, notification or summary e-mail. The module itself keeps working: its checks still run and its page keeps showing the state of everything it watches, with its warnings and badges; only the alerts stop. A toast confirms the change and how many alerts were resolved, and the Alerts feed, the Dashboard, the Infrastructure overview and the footer ticker update at once. The change is recorded in the Audit log on your company, as alert_generation.<module> going from on to off (and from off to on when you turn it back on), with the number of alerts resolved.

Ticking the box again saves at once, with no question, and alerts open again the next time each module checks: within minutes for websites and SNMP devices, within half an hour for Hyper-V, at the next daily check, sync, pass or backup report for the others. An alert that was open when you turned the module off, and that this first check finds again, re-opens without any e-mail, notification or summary e-mail, since the team already heard about it; anything else that opens is new and is e-mailed as usual. This quiet re-open only applies to alerts that re-open within 26 hours of turning the module back on (1 hour for websites and SNMP devices, 2 hours for Hyper-V) and that the switch resolved less than 60 days earlier; past that, an alert that opens is treated as new. Alerts that report an event rather than a condition (MX changed, Nameservers changed, Network configuration changed, New privileged account, Backup missed and Device rebooted) never re-open quietly: each new occurrence is news.

> **i** The switches of each item keep working, and an alert only opens when both the module and the item allow it: a website's Alerts enabled, an SNMP device's Raise alerts for this device, a Syslog collector's Alert on critical/error, a network device's Scheduled daily backups, a domain's Alert on SPF and DMARC problems, a third-party backup job set to Ignore. To quiet a single item use its own switch, and to stop an alert rule pause it in the Rules tab: the module switches never touch the alert rules.

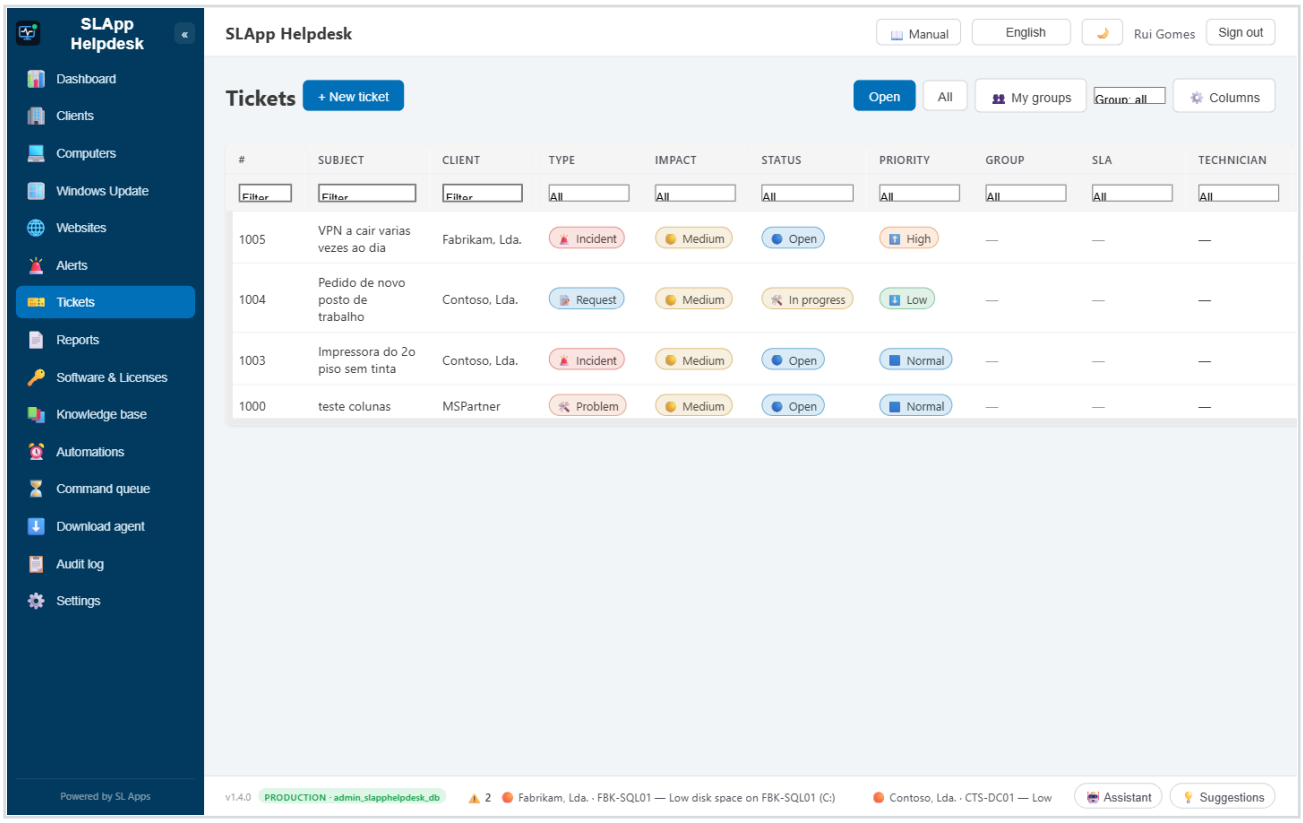
The footer ticker

Whatever module you are in, the footer shows a scrolling ticker of the open alerts of your team, preceded by a warning icon and the number of open alerts. Each item shows a red dot for Critical or an orange dot for Warning, then the client, the computer and the message (an alert that is not about a computer leaves the computer out). Hardware lifecycle alerts, which stay open for months and can be dozens at once, are not scrolled one by one: they are summed up in a single item at the end of the strip, for example 12 hardware lifecycle alerts, with a red dot when one of them is Critical; the number before the ticker still counts each of them. The ticker refreshes every minute and disappears entirely when nothing is open. It always moves at the same steady reading speed, however many alerts are open, so a long list simply takes longer to come round, and it stops while the mouse pointer is over it, so that you can read an item; when your operating system is set to reduce animations (in Windows 11, Animation effects turned off; in Windows 10, Show animations in Windows turned off), it stays still and shows what fits. Click anywhere on it to jump to the Alerts module.

7. Tickets

Tickets is where all support requests live: the ones your clients open in the Portal, the ones you open yourself, and the ones created from an alert. Each ticket carries a conversation with the client, internal notes, a task checklist, logged time, approval requests, links to related tickets and a full history of

changes. If the client has an SLA, the ticket also shows how much time you have left to respond and to resolve.



The ticket list with the Open/All, My groups and group filters

The ticket list

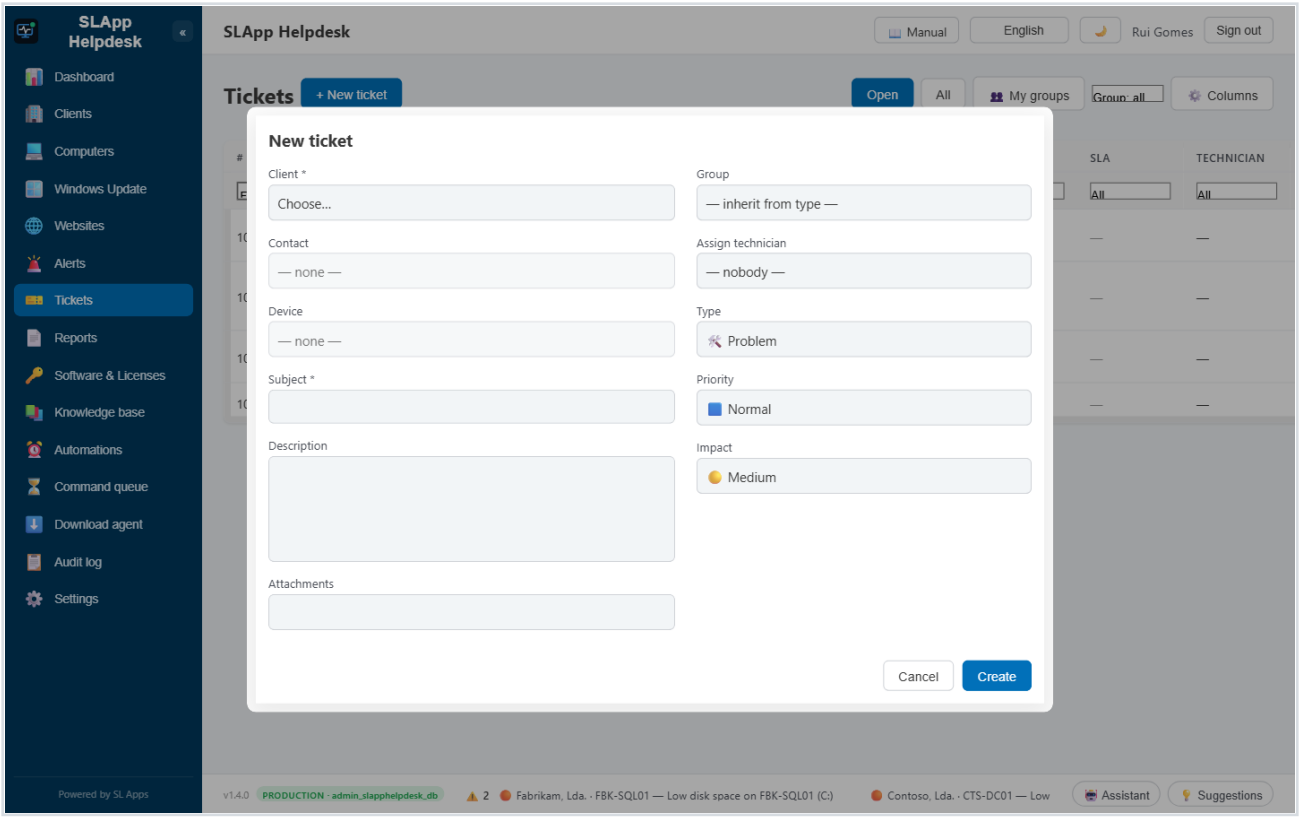
The list shows the most recently created tickets first. Click any row to open the ticket. The controls at the top of the page are:

- Open / All: Open (the default) hides tickets whose status is Resolved or Closed; All shows everything.
- My groups: a toggle that limits the list to tickets whose responsible group is one you are a member of. Hover it to see the hint "Tickets of the groups I am a member of". Click it again to turn it off.
- Group: all: a dropdown to show only the tickets of one specific responsible group (hover it to see "Filter by responsible group"). Choosing a group here replaces the My groups filter.
- + New ticket: opens the creation form (see below).

The columns are #, Subject, Client, Type, Impact, Status, Priority, Group, SLA, Technician and Updated. Type, Impact, Status, Priority and Group are shown as colored badges using the colors and icons configured in Settings. The SLA column shows two small chips (1st resp and Close) for tickets of clients with an SLA, or a dash when the client has none. Technician is the person the ticket is assigned to, or a dash when nobody is.

> There is no free-text search box on the list. To find a ticket, narrow the list with Open/All and the group filters, or open it directly from the cards on the Dashboard.

Creating a ticket



The New ticket form: request details on the left, assignment and classification on the right

Click + New ticket. The form has two columns. On the left you describe the request; on the right you classify and assign it.

- Client (required): choose the client. The Contact and Device lists depend on this choice and are reset if you change the client.
- Contact: the client contact the ticket is for. The primary contact is marked "(primary)". A public reply is only emailed to the client if the ticket has a contact with an email address.
- Device: a computer of that client, so the ticket is linked to the machine and you can open Remote access from the ticket.
- Subject (required) and Description. The Description accepts formatting (bold, italic, lists) and images you paste or drag into it, and you can make it bigger by dragging its bottom-right corner; a template written as plain text keeps its line breaks.
- Attachments: you can attach several files (up to 10) when opening the ticket.
- Group: leave it on "— inherit from type —" and the ticket is assigned to the group configured for its Type in Settings (a hint under the form tells you which group that is). Choosing a group here overrides that rule.
- Assign technician: a specific technician, or "— nobody —".
- Type (default Problem), Priority (default Normal) and Impact (default Medium). Impact is what selects the SLA targets; the lists of types, priorities and impacts are configurable in Settings.
- Product family and Product: only shown when at least one active product family has been set up in Settings (General, Tickets, under Product families and Products). Product appears once you choose a family that has active products, and it only offers those.

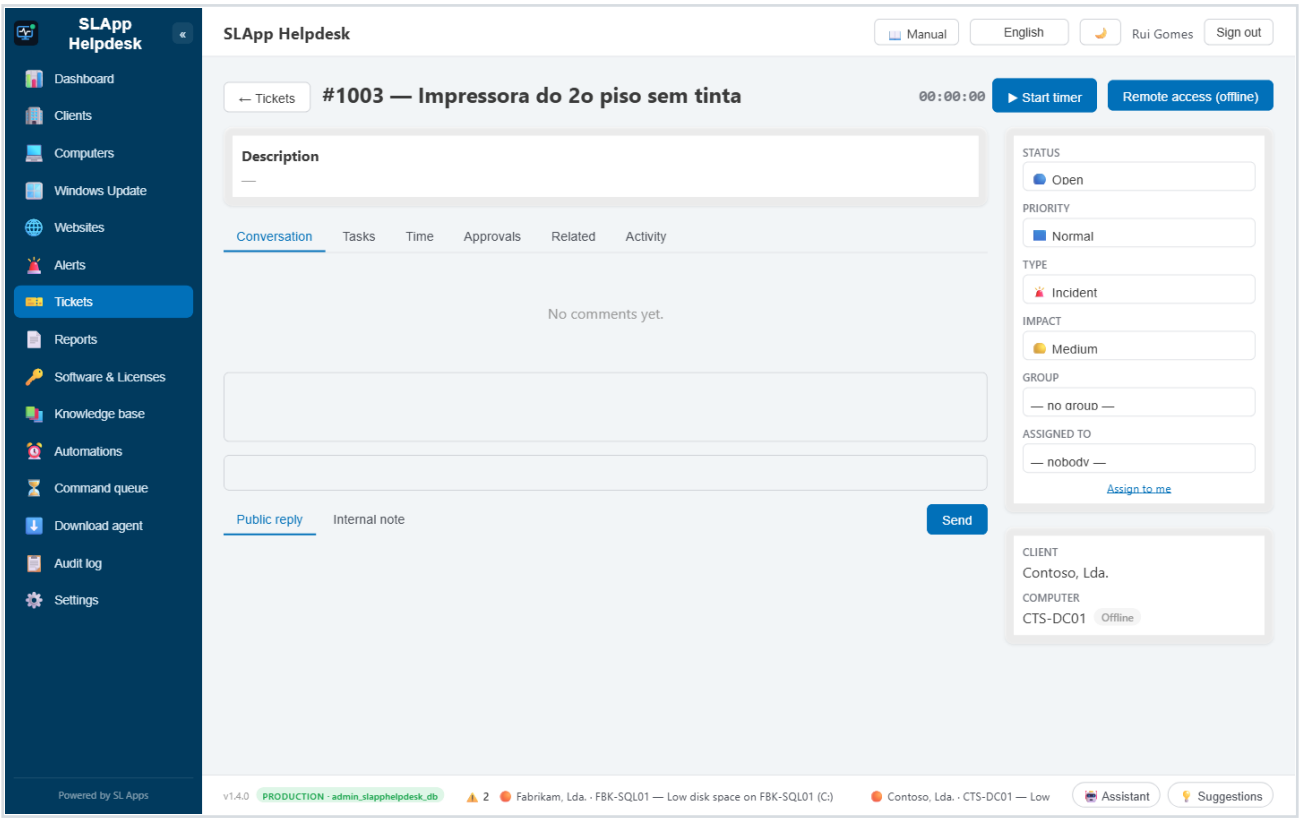
From template, at the top of the left column, only appears when your company has active general ticket templates (Settings › General › Tickets › Templates; see Ticket templates in the Settings chapter). Choosing one fills in the Subject and the Description and, when the template sets them, the Type, Priority and Impact (only values that are still active); a line under the field says how many tasks it adds

to the ticket, if any, and the right column lists those tasks in order (Template tasks), each with its group and note. Everything it filled in can still be changed. When you click Create, {{name}} becomes the name of the chosen contact, {{client}} the client's name and {{date}} today's date, and the template's tasks are copied into the ticket's Tasks tab with it. Choosing — None — again clears the subject and the description the template filled in, unless you changed them.

1. Click + New ticket.
2. Choose the Client, then optionally the Contact and Device.
3. Type the Subject and, if useful, a Description; add files under Attachments.
4. Set Type, Priority and Impact, and optionally a Group and a technician.
5. Click Create (it stays disabled until Client and Subject are filled in), or Cancel to discard the form.

When a new ticket ends up with a responsible group, every active member of that group who has an email address receives an email about it (except the person who created it). Tickets can also be created pre-filled from the Alerts module.

The ticket detail



Ticket detail: conversation and tabs on the left, SLA and properties on the right

The header shows the ticket number and subject, a "‹ Tickets" button to return to the list, the timer (a live clock with the total time logged on the ticket and a Start timer / Stop timer button), if a computer is linked, a Remote access button, Schedule visit and Ask the AI. Remote access is disabled and marked "(offline)" while the computer is not online; it opens the same menu as on the computer's page (SL App (native), with "Connect without asking for authorization" and "Record session" as your permissions allow, Remote Desktop (through the agent), Remote Desktop Connection, TeamViewer...), and the same Remote access button also sits next to the computer in the Computer field of the side panel. Schedule visit is disabled while the ticket is closed (see On-site visits below). The detail refreshes itself every 15 seconds, so replies from the client in the Portal and changes by colleagues appear without reloading.

Ask the AI opens the AI Assistant panel with a message already filled in from the ticket's own subject,

description, client, type, priority and status, asking it to help find the cause and a fix. When the ticket has a linked computer, its name and internal ID are included too, so the Assistant may use its own tools on that exact machine — for example running a command or restarting it — the same way it can from the Syslog module or a computer's Windows events. Review what it proposes to do before letting it act: nothing stops it from running a command the moment it decides to.

The main column starts with the ticket's Description, followed by seven tabs: Conversation, Tasks, Time, Visits, Approvals, Related and Activity. Each tab shows a counter in parentheses when it has content (number of messages, tasks, minutes logged, on-site visits that are not cancelled, approval requests, related tickets or recorded changes). Activity lists every change made to the ticket, newest first, starting with the "opened the ticket" event, including the on-site visits the client confirmed in the Client Portal.

A ticket opened by a staff change request, for a client's new employee or for an employee leaving, has one more tab, Staff change, right after Conversation and without a counter. It shows the request and runs its steps in Microsoft 365 and in SLApp, one click and one confirmation at a time; see the Staff changes chapter.

The right-hand column holds the ticket properties. Status, Priority, Type, Impact, Group and Assigned to are dropdowns that save immediately when changed. Under Assigned to there is an Assign to me shortcut whenever the ticket is not already yours. Below that you see the Client, the Contact (when the ticket has one) and the Computer (with an Online/Offline badge and a link to the computer's page, or "— none —").

- Status has five fixed values: Open, In progress, Pending client, Resolved and Closed. Resolved and Closed both record the resolution time used by the SLA; moving the ticket back to Open, In progress or Pending client clears it.
- Changing the Type without touching the Group moves the ticket to the new type's group automatically (or to no group, if that type has none configured). Changing the Group explicitly always wins.
- Every change to these fields is recorded in the Activity tab and also appears inline in the Conversation, so you can see when the status changed between two messages.

> **i** A Closed ticket is read-only everywhere: conversation, properties, tasks, time, approvals and related tickets; the exceptions are the Billable box of its time entries and, on a credit contract, their intervention type, which an administrator can still change (see Billable time and the hours bank below). Use Reopen ticket in the header (it takes the place of the timer button while the ticket is closed) or Reopen to reply under the conversation to unlock it; the ticket goes back to Open. Closing a ticket also stops any timer still running on it, from any technician, with the note "Stopped automatically when the ticket was closed".

Conversation: public replies and internal notes

The Conversation tab lists all messages in order, interleaved with the field changes. Your own messages are labelled Me, colleagues by their name, and the client by the contact's name (or Client). Internal notes have a shaded background and an Internal note badge. Attachments appear as a paperclip button; clicking one asks "Do you want to open the file or save it to your computer?" with Save and Open buttons.

1. Type your message in the box (placeholder "Write…"). You can also attach one file per message; a message needs text or a file.
2. Choose Public reply or Internal note. Internal notes show the reminder "Only visible to the team — the client never sees this in the Portal."
3. Click Send.

A Public reply is visible to the client in the Portal and, if the ticket has a contact with an email address, the contact receives an email saying there is a new reply. An Internal note never leaves the team and sends no email. The first public reply by a technician is what counts as the SLA first response; internal notes do not.

Tasks, Time, Approvals and Related

Tasks is an internal checklist for the ticket (the client does not see it). The summary reads "N to do / N completed". Click Add task to open the New task form, which has two tabs. Create a task asks for Title (required), Note, Group and Assign to; then click Create task. From a template adds every task of one of your company's active ticket templates to this ticket at once: choose the Template (templates of every kind are listed, General first, each with its number of tasks), check the list of the tasks it will add and click Add N tasks. They go after the tasks the ticket already has, in the template's order, each with its title, note, group (a group deactivated since is left out) and Staff change step, and a notice says how many were added. When your company has no active template, the tab says so and points to Settings › General › Tickets › Templates. Tick the checkbox to complete a task (it is struck through and shows "Completed" with the date; completed tasks sink to the bottom) or untick it to reopen. Each task has Edit and Delete buttons; Delete asks for confirmation in a Delete task dialog.

A ticket created from a ticket template, with From template or by a staff change request, starts with the template's tasks in its order, each with its title, note and group (a group deactivated since is left out). A task tied to a step of the Staff change tab shows a badge with that step, for example Microsoft 365: create the user, and is ticked by itself when the step is run there (hovering the badge says so); like any task, it can also be ticked by hand.

Time shows the Total time on ticket and the list of entries (duration, who, when, note). The timer itself lives in the header: Start timer begins counting for you; Stop timer opens a dialog showing the Time to log and asking "What was done during this time? (note)", then Stop and save writes the entry. You can only have one timer running per ticket, but several technicians can each run their own at the same time; a running timer appears in the list as "timer …" and cannot be edited or deleted until it is stopped. + New entry, next to the Time entries title, opens the New time entry window: enter Minutes and an optional note, untick Billable if that time must not be billed, and click Create. Finished entries can be edited with the pencil (or a double-click on the row), which opens the same window filled in, or deleted (bin), except the ones an on-site visit wrote (see On-site visits below); you can only change your own entries, unless you are an administrator. The Billable box and the Counts against bar at the top of the tab are explained in Billable time and the hours bank below.

Approvals lets you ask someone to sign off on something before you proceed. Click Request approval and fill in Name (defaults to "Approval request"), Mode (Everyone has to approve, or One is enough), Description, and tick the Approvers under Team (technicians) and Client (contacts of this ticket's client). Send request emails every approver that has an email address: team members get a link to the app, client contacts get a link to the Portal, where they decide. Each request card shows its mode, its status badge (Pending, Approved, Rejected or Cancelled), who created it, and every approver with their own state and comment (you are marked with a You badge). If you are one of the pending approvers you see Approve and Reject; either one opens a Comment (optional) box with Confirm approval and Confirm rejection buttons. A pending request can be withdrawn with Cancel request, which first asks "Cancel this request?" (Yes, cancel / No); the cancellation is only accepted from whoever created the request or an administrator.

Related links this ticket to others (duplicates, follow-ups, the same root cause). Click + Link ticket next to the Related tickets title, type the other ticket's number under Ticket number and click Link. Links work in both directions: the other ticket shows this one too. Each related ticket shows its number, subject, client and status; click it to open it, and Unlink removes the link. A ticket cannot be linked to itself or linked twice.

Billable time and the hours bank

Every time entry is either billable or not. Billable time counts against the client's hours bank when one of its contracts with an hours bank covers it, and goes to the billing in Reports › Hours & billing either way (the time no hours bank covers as time outside any hours bank); time that is not billable stays on the ticket and in its total, but never counts against an hours bank nor goes to billing. A stopped timer is billable, and so is time added by hand unless Billable was unticked before clicking Add. On a credit contract there is no Billable box: choose Not billable, the last option of the Intervention type list, instead of a type. As soon as a ticket has time that is not billable, a line under Total time on ticket splits the total, for example Billable 1h 30m · not billable 20m.

Each finished entry shows a Billable checkbox to whoever may change it: its author or an administrator while the ticket is open, and only an administrator once the ticket is closed, so time can be taken out of the billing, or put back, without reopening the ticket (nothing else of an entry can be changed on a closed ticket). Right-clicking an entry also offers Mark as not billable or Mark as billable. Everyone else sees a Not billable badge on the entries that are not billable.

When one of the client's contracts with an hours bank is in force today (see Hours bank of a contract in the Clients chapter), a bar at the top of the Time tab reads Counts against: followed by the contract the ticket's time counts against today and how much of that contract is used so far in its current period, by all of the client's tickets: for example 7.5 h of 10 h this month (this contract year for a yearly contract, in total for a one-time one), with the same coloured bar as the client's Contracts tab. The drop-down at the end of the bar chooses the contract:

- Automatic, the default, names the contract it picks (for example Automatic: Support 2026, or Automatic: no hours bank): each entry counts against the client's contract with an hours bank in force on the day of that entry, the day its timer was started or it was added by hand, and when several are in force that day, against the one that started most recently.
- Choosing a contract makes the ticket's time count against it on every day it is in force; the time of a day it does not cover still counts automatically. The list only offers the client's contracts in force today, active, with an hours bank and with dates that cover today: one that only starts later, or has already ended, cannot be chosen, and the app says why.
- When the contract chosen earlier is no longer in force, it stays selected, marked (not in force), and a note under the bar says that it is not in force today and which contract the time counts against automatically instead; when no other contract is in force either, the bar reads No hours bank and the time falls outside any hours bank. Choose Automatic or another contract to change it.

> **i** The contract can only be changed while the ticket is open. Like any change to the ticket's time, changing it updates the hours bank at once and can send the administrators an hours bank warning email (see Hours bank warning emails in the Clients chapter).

Time on a credit contract

When the ticket's time falls in a credit contract (see Credit contracts in the Clients chapter), the bar at the top of the Time tab shows that contract's balance, Balance: N credits or In debt: N credits, with its state, instead of the hours bar, and its drop-down also offers the client's credit contracts. The Stop timer dialog and the New time entry window then ask for the Intervention type, required and preselected with the contract's default type, each type with its rate (for example 1.5 cr/h), and show an estimate, ~ N credits, worked out from the minutes and the contract's rounding; Stop and save and Create stay greyed out until a type is chosen. Outside a credit contract the type is not asked.

Each entry of a credit contract shows its intervention type and the credits it debits, for example -0.75 cr; an entry without a type shows an amber no type badge and debits nothing until it gets one. Editing an

entry (pencil) also lets you change its type, and the entries of an on-site visit can change their type, note and Billable box. An entry keeps the rate its type had when the type was chosen, so changing a rate in Settings never changes time already logged, while choosing another type, or changing the minutes or the Billable box, changes the debit and the contract's balance at once. On a closed ticket an administrator can still change the type of an entry. Entries saved without a type, such as the timers stopped when the ticket was closed, get the contract's default type by themselves.

On-site visits

When the work needs someone at the client's premises, schedule an on-site visit on the ticket: Schedule visit in the header, or + Schedule visit on the Visits tab, opens the Agenda's New on-site visit window with the ticket's client, the ticket itself and the ticket's contact filled in, and you as the technician (see Scheduling a visit in the Agenda chapter). Both are disabled while the ticket is closed, with the hint Ticket closed — reopen it to schedule a visit. Once the visit is saved, the ticket shows its Visits tab.

The Visits tab lists the ticket's visits in order of their planned start, under a line such as To come: 1 · Done: 2. Each card shows when the visit is planned or, once it is done, the time actually spent at the client, marked at the client; then the technicians, the status (followed by an hourglass while the client's confirmation is awaited) and the address and, for a finished visit, the travel time, the kilometres and who signed, or Client not present (hover it for the reason). A cancelled visit is faded and struck through, with its reason. Click a card to open the visit's window, which offers every action described in the Agenda chapter, such as Start visit, Finish visit, Reschedule, Notify client and the visit's report; right-click it for Open visit and Copy address.

Finishing a visit writes its time into the Time tab: for each technician who did the visit, one entry with the time at the client, marked On-site visit, and, when there was travel, one entry with the travel time, marked Travel. Both are billable and dated with the arrival at the client, and they count against the hours bank like any other time, the travel as the contract's Travel and kilometres options say (see How a visit is billed in the Agenda chapter). These entries belong to the visit: they have no pencil or bin and cannot be deleted, because their time is changed on the visit itself, by an administrator with Correct the visit. On the ticket you can only switch their Billable box, under the usual rules, or right-click them for Mark as not billable or Mark as billable. The visit also leaves an internal note in the Conversation with its work summary and its PDF report attached.

> **i** A visit can only be finished while its ticket is open. Before closing a ticket, finish or cancel the visits still planned on it, or reopen the ticket to finish one later.

SLAs and at-risk tickets

SLAs are defined by an administrator in Settings (SLAs), with a first-response target and a resolution target for each Impact level, and assigned to a client from the SLA dropdown on the client's page. When a ticket belongs to a client with an active SLA, the deadlines are calculated from the ticket's creation time and the chosen Impact, and an SLA card appears at the top of the right-hand column with the SLA name and two chips, 1st response and Resolution (shortened to 1st resp and Close in the list).

- "Xh Ym left" (amber outline): the target is still open; hover to see the exact due date and time.
- "Met" (green): the first public reply was sent, or the ticket was resolved, within the target.
- "Breached" (red): the deadline passed before the response or resolution. The check uses the server clock, not your computer's.

If you set a ticket to Resolved or Closed while its resolution SLA is breached and no justification has been recorded yet, a dialog titled Resolution SLA breached asks for a Justification before it saves (Save and resolved / Save and closed, or Cancel). The justification is then shown on the SLA card as "Breach justification:".

> ℹ Changing a ticket's Impact changes which SLA target applies, so the chips are recalculated immediately. Tickets of clients without an SLA (or whose SLA has been deactivated) show a dash in the SLA column and no SLA card.

Email notifications

- New ticket with a responsible group: every active member of the group with an email address is emailed (the creator is not).
- Public reply: the ticket's contact is emailed, if they have an email address. Internal notes send nothing.
- Approval request: each chosen approver with an email address is emailed with a link to the app (team) or to the Portal (client contacts).
- Hours bank: when the ticket's time takes the client's contract to its warning percentage or to 100% of the hours it includes, the administrators are emailed (see Hours bank warning emails in the Clients chapter).
- On-site visits: the technicians of a visit are emailed the evening before it, and when the client confirms it in the Client Portal. The client is emailed a request to confirm the visit when it is scheduled and when its dates change (unless Ask the client to confirm by email is unticked), and the visit's details or its report when someone clicks Notify client or Email the report to the client in the visit's window (see the Agenda chapter).

8. Staff changes

When one of your clients takes on a new employee, or someone leaves, the same steps come back every time: the Microsoft 365 account and its licenses, the contact in SLApp and the equipment for a newcomer; for a departure, the blocked sign-in, the automatic reply, the licenses to free up and the equipment to recover. Staff changes turns each of these events into a request with its own ticket. The client asks for it in the Client Portal, or a technician records it from the client's sheet; the ticket is created from the active ticket template of that kind, with the template's checklist as its tasks and the request's details in its description, and its Staff change tab carries out the Microsoft 365 and SLApp steps one click at a time. Nothing is ever done automatically: each step runs only when a technician clicks it and confirms. Staff changes is not a module of the sidebar: it lives in the client's sheet, on the ticket and in the Portal.

Who can request

- In the Client Portal: the client's primary contact and the contacts allowed to request staff changes, who get a Staff changes tab there (see Staff changes in the Client Portal Guide). A technician gives a contact this permission with Can request staff changes in the contact's form (Clients › the client › Contacts), and the primary contact can give it from the Portal's Contacts tab; the primary contact always has it. Taking the permission away applies at once. A Portal session opened from the agent on a computer never has the tab, and Only the primary contact can open tickets in the Portal does not apply to these requests.
- In the app: any technician, from the client's Staff changes tab (Clients › the client › Staff changes), for example when the client asked by phone or by email.

The client's Staff changes tab

The Staff changes tab of a client's sheet lists the client's requests, newest first, with the columns Request (New employee or Employee leaving), Employee, Start / last day, Status (Open, Done or Cancelled), Requested by (followed by (portal) when a contact asked from the Portal), Ticket (its

number and its status) and Created; Ticket contact and Actions done (how many steps were run) can be added with Columns. The row under the headers filters the list and a click on a header sorts it. Click a request to open its ticket on the Staff change tab; right-click it for Open the ticket, Copy name and Copy ticket number. + New employee and + Employee leaving, at the top, record a new request.

Recording a request

+ New employee and + Employee leaving open the request form, whose first line reminds you that it creates a ticket with the tasks of the active template of that kind and that nothing changes in Microsoft 365 until a technician runs each step from the ticket. Requested by (contact) is optional: it is the client's contact the ticket is opened for, who then follows it in the Portal like any ticket of theirs. The rest of the form depends on the client's Microsoft 365: when the client is connected through its own app (see the Microsoft 365 chapter), the email address and the licenses are chosen from what the last sync read, and the person leaving can also be one of its Microsoft 365 users; otherwise the email address and the licenses are typed as free text, and the person leaving must be one of the client's contacts.

- New employee: First name *, Last name *, Job title, Department, Start date * and Manager (one of the client's contacts). Desired email is an alias, which follows the first and last name until you type in it, and one of the client's verified Microsoft 365 domains, the default one first; for a client that is not connected it is a free email address. Licences lists the client's paid Microsoft 365 licenses, each with how many are free, to tick, and Other licences or subscriptions takes anything else as text (for a client that is not connected, a Licences text field replaces both). Equipment needed, Access needed (groups, shared mailboxes, folders, applications) and Notes complete it.
- Employee leaving: Who is leaving *, a contact of the client and/or, for a connected client, one of its Microsoft 365 users (either is enough, and picking one selects the other when the email matches); Last working day * and an optional Time at which the access should end that day; Mailbox: Keep it as it is, Convert to a shared mailbox (with Give access to (optional), a contact) or Forward the email to a contact (Forward to * is then required); Automatic reply (optional: when it is empty, the Set automatic reply step proposes a short standard one); Remove the licences (ticked by default); Equipment to recover, the assets assigned to the contact leaving, all ticked; Data to keep; and Notes.

Create request stays disabled until the required fields are filled in. The server checks the request again, field by field: every contact, asset and Microsoft 365 user it names must be the client's, and the desired email of a connected client must be a valid alias at one of its verified domains; a field it refuses is outlined in red. The request then opens its ticket in one go, and a notice gives its number (Request created — ticket #128). The ticket gets:

- Subject: the template's subject, with {{name}} replaced by the person's name, {{client}} by the client's name and {{date}} by the start date or the last working day.
- Description: the template's description followed by a summary of the request, one line per field filled in, in your company's email language (the Email language under Settings › General › Connectors). The client sees it in the Portal.
- Type, Priority and Impact: the template's, as long as they are still active values (otherwise the usual defaults), and the responsible group of that type.
- Contact: the contact who asked in the Portal, or the one chosen in Requested by (contact).
- Tasks: the template's checklist, in its order, including the tasks tied to a step of the Staff change tab (see Tasks, Time, Approvals and Related in the Tickets chapter).

Creating a request sends exactly the emails a ticket created the same way sends: the members of the ticket's responsible group are emailed and, for a request from the Portal, so is the address set in Internal notifications go to (the Email card of Settings › General › Connectors). Nothing else is sent, and the client receives no email about it.

The Staff change tab of the ticket

A ticket opened by a request has a Staff change tab, right after Conversation. Its top bar shows the kind (New employee or Employee leaving), the request's status (Open, Done or Cancelled), the person, Starts on or Last day with the date, and who asked and when, as Requested by followed by the name and (client portal) or (technician), plus on behalf of and the contact when a technician chose one; while the request is open, Mark as done and Cancel request sit at its right. Below it, the Request card lists every field that was filled in, and the Actions card holds the steps.

The Actions card first says who the steps act on. The Microsoft 365 line shows the state of the client's connection and the user: for a new employee, the one created by the first step (— not created yet — until then); for a person leaving, the user chosen in the request or, failing that, the client's user whose address matches the person's email, marked found by the email address — check it is the right person. The user's licenses and an Active or Sign-in blocked badge follow. When no user is linked to a departure, choose one of the client's synced users in Choose the Microsoft 365 user… and click Link; Change picks another one. The SLApp contact line does the same for the contact: the one created for a new employee, or the contact leaving (or an active contact found by the person's email when the request named none). The steps follow, one button each:

- Create Microsoft 365 user (new employee) opens the New Microsoft 365 user window already filled in from the request: the names, the alias and domain of the desired email (or first.last at the default domain), the usage location of the client's organisation (PT when it has none) and the licenses asked for that still have free units. Check it, click Create user, and the Temporary password window shows the password once, as described in the Microsoft 365 chapter.
- Create contact in SLApp (new employee) asks for the Name *, the Email (the user just created, or the desired email) and the Job title, and creates the contact at the client, with access to the Portal. It is refused when an active contact of the client already has that email.
- Block sign-in and sign out everywhere (person leaving), after its confirmation, blocks the user's sign-in and closes every session, on every device. If signing out fails, the notice says so: open sessions then end within the hour.
- Set automatic reply (person leaving) opens the Out of office window prefilled with the automatic reply of the request or, when it gave none, a short one saying that the person no longer works at the client, sent to everyone outside the organisation too; adjust it and click Save and turn on. There is no Turn off here: switch the replies off in the client's Microsoft 365 tab.
- Remove licences (person leaving) asks to confirm the list of the user's licenses and removes them one by one. When some cannot be removed, for example a license that comes from a group, the notice names them with the reason and the tied task stays open: click again once that is sorted out.
- Reset password (person leaving) opens the Reset password window; the new temporary password is shown once.
- Deactivate contact in SLApp (person leaving), after its confirmation, removes the contact from the client's contacts, and with it its access to the Portal; its tickets and history are kept.

A step that cannot run now is disabled; hover it for the reason: the request is no longer open, the ticket is closed (reopen it first), the client is not connected to Microsoft 365 through its own app (said under the Microsoft 365 line too), no Microsoft 365 user is linked to the request, the user has no licenses, the user or the contact was already created, no SLApp contact was found for the person, or the contact is already inactive. A step already run shows Done, with the date and who did it, next to its button. Only one step runs at a time on a request: clicked meanwhile, another one answers Another action on this request is still running — try again in a moment.

Each step that works is recorded in four places: the ticket's tasks tied to that step are ticked (a notice says how many), an internal note starting with Staff change: is added to the conversation saying who

did what (never a password), the step is added to the History card at the bottom of the tab, and the Audit log gets a Staff change request entry, next to the step's own Microsoft 365 or Contact entries (see the Audit Log chapter). Linking a Microsoft 365 user by hand leaves an internal note too.

For a person leaving, an Equipment to recover card lists the assets the request named, as they are now: type, serial number, status and computer, or no longer in the inventory for an asset removed since. Recovering them is one of the ticket's tasks, not a button: Open the client's asset inventory takes you to the client's Asset inventory, where you take each asset back by clearing its Contact (owner).

Finishing or cancelling a request

When everything is done, click Mark as done: after its confirmation the request becomes Done and its steps are no longer available. Cancel request asks for an optional Reason and makes the request Cancelled. Either way an internal note records it (with the reason, for a cancellation), the Audit log gets an entry, and the ticket stays as it is: close it as usual. Both remain possible after the ticket has been closed, then without the note, so a request whose ticket was closed first can still be finished.

In the Portal, the client can cancel a request too, but only while nothing has been done on it: the request is still open, no step has been run and its ticket is still in the starting status for new tickets. Only the primary contact and the contact the ticket was opened for can. The request then becomes Cancelled, a reply from the client is added to the ticket (Request cancelled by followed by the name, and the reason when one was given) and the ticket is closed, with its running timers stopped and without any email. Once a technician has run a step or changed the ticket's status, the Portal no longer offers Cancel request, and the client writes in the ticket to change anything; a cancellation tried from a page opened before that is refused with a notice asking them to write in the ticket instead.

> ℹ A request about a person leaving is never shown in the Portal to that person, not even to a primary contact, unless they asked for it themselves or its ticket is theirs: a company sometimes decides a departure before telling the person.

The templates behind the requests

Every request opens its ticket from the active template of its kind in Settings › General › Tickets › Templates (see Ticket templates in the Settings chapter): the New employee template for a newcomer and the Employee leaving template for a departure. Each company starts with one of each, in its email language (type Request, priority Normal and impact Medium, when those values are active in your lists), and an administrator can adapt its subject, description, type, priority, impact and tasks. Their tasks are:

- New employee (onboarding): Create the Microsoft 365 account (tied to Create Microsoft 365 user), Assign the licence, Add to groups and shared mailboxes, Create the contact in SLApp (tied to Create contact in SLApp), Prepare the computer, Install the software, Set up email on the phone, Hand over the credentials and Follow up on the first day.
- Employee leaving (offboarding): Block sign-in and sign out everywhere (tied to the step of the same name), Set the automatic reply (tied to Set automatic reply), Forward or convert the mailbox, Remove the licences (tied to Remove licences), Recover the equipment, Remove from groups, Deactivate the contact in SLApp (tied to Deactivate contact in SLApp) and Final check.

> ℹ Converting a mailbox to a shared one, forwarding it and the groups are not steps of the tab: they are tasks, done in the Microsoft 365 or Exchange admin center as the request describes and ticked by hand.

9. Agenda

The Agenda plans the visits your technicians make to your clients' premises and records what happened there. Each on-site visit belongs to a ticket and has one or more technicians, a start and an end, an address and a contact, and it moves through five statuses: Planned, Confirmed, In progress, Done and Cancelled. When a visit is scheduled or moved to a new date, the client is emailed a request to confirm it, and confirming it with one click in the Client Portal turns it Confirmed (see The client's confirmation). On the day, the technician starts the visit on arrival and finishes it at the client, usually on a phone, with the time spent there, the travel time, the kilometres, a summary of the work and the client's signature. Finishing records that time on the ticket, where it counts against the client's hours bank and goes to Reports › Hours & billing like any other time, and files a PDF report of the visit on the ticket. Open the module from Agenda in the sidebar, right after Tickets. Every technician sees every visit of your company and can schedule, edit, reschedule, confirm and cancel visits and notify the client about them; starting and finishing a visit, and cancelling one that is in progress, is for the technicians of that visit and for administrators, and correcting a finished visit is for administrators only. Every change is recorded in the Audit log under On-site visit.

Views and filters

The header holds + New visit, My calendar (iCal) (see Your visits in your own calendar below) and five buttons that choose the view: My visits today, Day, Week, Month and List. On a computer the Agenda opens on Week; on a phone it opens on My visits today. Under them, the arrows either side of Today move to the previous or next day, week or month (Previous week, Next week and so on), Today comes back to the current one, and the period on screen is written next to them. The view you chose and the filters are remembered while the browser tab stays open, and the Agenda reloads itself every minute while the page is visible.

- Day and Week: an hour scale from 07:00 to 20:00, stretched when a visit starts earlier or ends later, with one column per day, Monday first. Each visit is a block with its times, its client and, when there is room, its technicians and its ticket; hover it to read all of it. A red line marks the current time on today's column. Click a day's heading to open that day on its own, and click a free half-hour to schedule a visit starting then. On a phone, Week shows one card per day instead, each with a + button that proposes a visit at 09:00.
- Month: six weeks on one page. Each day shows up to three visits and +N more when it has others; click the day number or +N more to open that day. The + that appears when you point at a day schedules a visit at 09:00. On a phone each day shows a coloured dot per visit and how many there are; tap the day to open it.
- List: the visits of one month as a table, the arrows moving a month at a time. When, Client, Ticket, Technicians, Address and Status are shown; Contact, Travel, Kilometres and Signature can be added with Columns. Type in the row under the headings to filter, click a heading to sort, and click Clear the column filters to start again. The line above the table reads, for example, 12 visits this month, or 3 of 12 visits while a filter is on.
- My visits today: your own visits of today, made for the phone (see On the day: My visits today below).

The colours are the statuses, shown in the legend next to the period: Planned in blue, Confirmed in green, In progress in orange, Done in grey and Cancelled in red. The legend ends with Awaiting the client's confirmation, drawn with a dashed edge because it is not a status: a Planned visit whose client was emailed a confirmation request and has not confirmed it yet shows an hourglass before its times in the Day, Week and Month views, and after its status in the List, in My visits today and on the ticket's Visits tab, where pointing at it says since when. The filters apply to every view except My visits today: All technicians lets you tick one or more technicians, All clients chooses one client, and Show cancelled

brings back the cancelled visits, which are otherwise hidden and are then shown faded and struck through. Clear filters removes all three. Times are shown in your browser's time zone.

> **i** Click a visit to open its window. Right-clicking it offers the same actions without opening it: Open visit, Open ticket, Mark as confirmed, Start visit, Finish visit, Notify client, Report (PDF), Copy address and Cancel visit, each one only when the visit's status and your permissions allow it.

Scheduling a visit

Click + New visit, click a free time in the Day or Week view, or use + on a day; on a ticket, use Schedule visit (see the Tickets chapter). The New on-site visit window opens with the start you clicked, or the next full hour, an end one hour later and you as the technician, or the technician the Agenda is filtered on when it is filtered on exactly one. A client the Agenda is filtered on is filled in too.

- Client *: the client visited. When the window is opened from a ticket, the client and the ticket are fixed.
- Ticket: New ticket (On-site visit), the default, or one of the client's tickets that are not closed. The hint under it says on which ticket the visit's time will be recorded.
- Start * and End *: the planned date and times. Moving Start moves End with it, keeping the duration, which is shown under End. A visit lasts at most 24 hours. Times in the past are accepted, so a visit that already took place can be recorded and then finished.
- Technicians *: tick from one to ten active technicians.
- Address: where the visit takes place, by default the client's address (street, postal code and city). In the visit window it becomes a link that opens a maps app.
- Contact: — none — or one of the client's active contacts, whose phone and email appear under the field. Choosing a ticket fills in the ticket's contact, unless you already picked one. The client emails described under Telling the client go to this contact.
- Ask the client to confirm by email: ticked by default, it emails the client a request to confirm the visit in the Client Portal as soon as you save (see Telling the client). The line under it says where the request goes: Sent to followed by the address when the visit's contact has one, otherwise to the ticket's contact or, failing that, to the client's primary contact, the first one with an email address. Untick it to tell the client nothing for now: you can use Notify client later. When the Start has already passed, the box is unticked and greyed out, since there is nothing left for the client to confirm.
- Notes: what to do, what to bring, who to ask for. They are shown in the visit window, in My visits today and in the technicians' calendars.

Click Schedule visit. The visit starts as Planned, and your browser's time zone is kept with it: its emails, its PDF report and the note it leaves on the ticket write its times in that zone. When the Ticket was left on New ticket (On-site visit), a ticket is opened for the client at the same time, with the subject On-site visit (in your company's email language, the Email language under Settings › General › Connectors), the visit's contact, the visit's first technician as its assigned technician, the type Problem with Normal priority and Medium impact, and the group set for that type; the visit's notes go into an internal note on that ticket rather than into its description, which the client can read in the Portal. Unlike a ticket opened by hand, nobody is emailed about it; the confirmation message gives its number. Nobody else in your team is emailed about the visit either: its technicians hear about it the evening before (see Emails to the technicians). The client, on the other hand, is sent the confirmation request at once, and the message shown after saving adds Confirmation request emailed to and the address, or The client was not asked to confirm — use Notify client when you want to if you unticked the box. When the request cannot go, the visit is saved all the same and a second message says why: there is nobody to email (neither the visit's contact, the ticket's nor the client's primary contact has an email address), or the email failed; use Notify client once that is sorted out.

The visit window

Clicking a visit opens its window, titled with the client's name and the status. It shows When (the planned times and their duration), Ticket (a link that opens the ticket on its Visits tab, followed by ticket closed when it is), Technicians, Contact (with links to call or email them), Address (a link that opens a maps app), Notes, Arrived while the visit is In progress, Scheduled by (who and when), Confirmation and Client notified (when the client was last sent the visit's details). Confirmation reads Confirmed by, the contact's name, (client) and the date when the client confirmed the visit in the Client Portal, or Confirmed by, the technician's name and the date when a technician used Mark as confirmed; while the visit is Planned after a confirmation request, it reads Awaiting the client's confirmation since and the date the request was emailed. A cancelled visit shows its reason. A finished visit adds a Visit report part with Time at the client, Travel time (round trip), Kilometres, Finished by, Report emailed to the client (when it was), Work summary and Client signature: the signature with Signed by and the date, or Client not present with the reason.

The buttons under it depend on the status and on who you are:

- Mark as confirmed (Planned only): the client agreed to the date, by phone for example. The visit turns Confirmed, with you as the one who confirmed it. Clients can also confirm it themselves in the Client Portal (see The client's confirmation).
- Start visit (Planned or Confirmed; the visit's technicians and administrators): you have arrived. The arrival time is recorded and the visit turns In progress.
- Finish visit (Planned, Confirmed or In progress; the visit's technicians and administrators): see Finishing a visit at the client. It is disabled while the ticket is closed, with the hint The ticket is closed: reopen it to finish the visit.
- Edit (Planned, Confirmed or In progress): opens Edit visit to change the technicians, the address, the contact and the notes, and the dates while the visit has not started. The client and the ticket never change.
- Reschedule (Planned or Confirmed): the same window, titled Reschedule visit, with the dates highlighted and ready to change.
- Notify client (Planned or Confirmed): emails the visit's details to the client and, while the visit is Planned and still to come, asks the client to confirm it in the Client Portal (see Telling the client).
- Report (PDF), Download report and Email the report to the client (Done only): see The visit report and Telling the client.
- Correct the visit (Done only; administrators): see Correcting a finished visit.
- Cancel visit (Planned or Confirmed for every technician; In progress for the visit's technicians and administrators): asks Cancel this visit?, reminds you that the client is not notified automatically and offers a Reason (optional), which is kept on the visit and in the Audit log. Click Confirm to cancel it.

Rescheduling, or changing the dates with Edit, moves the visit and sends its technicians the evening-before email again for the new day. The confirmation was for the old date: a Confirmed visit goes back to Planned (the window says so before you save) and who confirmed it is cleared. Reschedule visit shows the Ask the client to confirm by email box, ticked, from the start, and Edit visit shows it as soon as you change a date: on saving, the client is sent a new confirmation request for the new date, and the message shown after saving says to whom, as when scheduling. A visit that is In progress keeps its planned times. Visits are never deleted: a visit that will not happen is cancelled, and stays in the Agenda (with Show cancelled), on its ticket and in the Audit log.

On the day: My visits today

My visits today is made for the phone: it lists your own visits of today, the ones you are a technician of, in order of time, as large cards. Each card shows the planned times, the status, the client, the ticket, the

technicians, the address as a big link that opens the phone's maps app (Apple Maps on an iPhone or iPad, the maps app on Android, Google Maps in a new tab on a computer), the contact with their phone number, which you tap to call, and the notes. Start visit, Finish visit, Report (PDF) for a finished visit and Details, which opens the visit window, are big buttons. Cancelled visits are left out and finished ones are dimmed. Refresh reloads the list, and when you have no visit today, See the week switches to the Week view.

Finishing a visit at the client

Finish visit opens the Finish visit window, titled with the client's name, which takes the whole screen on a phone. Everything marked * is required:

- Time at the client: Arrived * and Left *. Arrived starts at the arrival recorded by Start visit, or at the planned start when the visit was not started, and Left at the current time. When that would make the visit longer than 24 hours (a visit recorded days later), Left starts at the planned end instead, or at Arrived plus the planned duration when the arrival came after the planned end; and when a visit that was not started is finished before its planned start, Arrived starts at the current time minus the planned duration. The time at the client is shown under them. Left cannot be in the future, and a visit lasts at most 24 hours.
- Travel time (minutes) *: the time to the client and back, in whole minutes from 0 to 1440. Enter 0 when there was none.
- Kilometres *: the round trip, from 0 to 10000, decimals allowed. Enter 0 when there were none.
- Technicians who did the visit *: every technician of the visit starts ticked; untick anyone who did not go. Each technician ticked gets the time at the client, and the travel, on the ticket.
- Work summary *: what was done, what is left, what the client should know, up to 4000 characters.
- Client signature: the client signs inside the box with a finger, a pen or the mouse (Clear starts again), and Signer's name * says who signed; the visit's contact, when it has one, is filled in. When nobody can sign, tick Client not present and give the Reason * instead.

Click Finish visit. The visit turns Done and, in the same step:

- The ticket's Time tab gets, for each technician ticked, one entry with the time at the client and, when there was travel, one more with the travel time. Both are billable and dated with the arrival time, and they count against the client's hours bank and go to Reports › Hours & billing like any other billable time, the travel and the kilometres as the contract says (see How a visit is billed). Their time can only be changed through the visit; on the ticket, only their Billable box can be switched (see On-site visits in the Tickets chapter).
- An internal note is added to the ticket's conversation with the work summary, the time at the client, the travel, the kilometres and who signed, or why the client did not, with the visit's PDF report attached. Being internal, it is never shown to the client in the Portal and sends no email.
- When the new time takes the client's contract to its warning level or past its hours, the administrators get the hours bank warning email (see Hours bank warning emails in the Clients chapter).

> ℹ A visit can only be finished while its ticket is open: reopen a closed ticket first. Start visit is not required before finishing: a visit recorded afterwards, or one finished from a computer, is finished straight from Planned or Confirmed.

On a credit contract (see Credit contracts in the Clients chapter), the Time at the client part of the window also asks for the Intervention type * of the work, preselected with the contract's default type, and, when there was travel and the contract's travel debits credits, the Intervention type for travel, preselected with the contract's travel type, otherwise the first type marked Use for travel, otherwise the

default one. Under each one, Debits N credits per hour gives its rate. The entries of every technician ticked get those types.

The visit report

Every finished visit has a PDF report: your company's logo on the left and the client's on the right, as in the client reports, then the client, the address, the contact, the technicians who did the visit, the time at the client and its duration, the travel time (round trip), the kilometres, the work summary and the client's signature with who signed and when, or Client not present and the reason. It is written in your company's email language, with the times in the visit's own time zone, and it is attached to the ticket's internal note as visit-report- followed by the ticket number and the date of the visit. In the visit window, Report (PDF) opens it in a new tab and Download report saves it.

Telling the client

SLApp Helpdesk writes to the client about a visit when it is scheduled and when its dates change, unless Ask the client to confirm by email was unticked, and when someone clicks Notify client or Email the report to the client. The visit's details go to the visit's contact or, when the visit has no contact with an email address, to the ticket's contact or, failing that, to the client's primary contact (the first one with an email address), who sees every visit of the company in the Client Portal; the report goes to the visit's contact or the ticket's. The visit's details are written in the client's Notification language (in the client's record), or in your company's email language when the client is on System default, and the report in your company's email language. Each email sent is recorded in the Audit log.

- The confirmation request, sent when the visit is scheduled, when its dates change and by Notify client while the visit is Planned and its start is still to come: an email titled Please confirm the on-site visit, with the date and times, the technicians, the address and the ticket, then Please confirm this date in the client portal, using the button below., a Confirm in the portal button that opens the visit in the portal's Agenda tab, and If this date does not suit you, please contact us to reschedule. The visit window then shows Client notified with the date and Awaiting the client's confirmation since that date; Notify client says Confirmation request emailed to and the address.
- Notify client on a Confirmed visit, or on a Planned visit whose start has passed: the same details in an email titled On-site visit scheduled, without the button, adding This visit is confirmed. for a Confirmed visit. Notify client then says Visit details emailed to and the address.
- Email the report to the client (Done): an email titled On-site visit report with the PDF report attached, the technicians and who signed, or why the client did not. The visit window then shows Report emailed to the client with the date.

When there is nobody to email (for the visit's details: neither the visit's contact, the ticket's nor the client's primary contact has an email address; for the report: neither the visit's contact nor the ticket's), a message says so and nothing is sent: before notifying, choose a contact with Edit; for the report, add an email address to the contact in Clients › Contacts, or download the report and send it yourself.

The client's confirmation

The Confirm in the portal button opens the Client Portal on its Agenda tab, on that visit, once the client has signed in. The Agenda tab lists the visits that concern the contact signed in (for the primary contact, every visit of the company; for any other contact, the visits whose contact they are or that belong to a ticket they opened): those still to come or under way, and those done or cancelled in the last 60 days. Each Planned visit whose start is still to come has a Confirm visit button, and one click confirms it. The visit turns Confirmed for your whole team at once, and:

- its window shows, under Confirmation, Confirmed by, the contact's name, (client) and the date;
- the ticket's Activity tab records that the contact (client) confirmed the on-site visit of that date in the

client portal;

- each technician of the visit is emailed Visit confirmed by the client, with the contact, the client, the date, the ticket and the address and a button that opens SLApp Helpdesk, unless they switched off A client replies on a ticket of mine in the Email notifications of their profile;
- the Audit log records client_confirmed, with the contact as its author.

The portal only confirms the date the client is looking at. When the visit was confirmed or cancelled in the meantime, has started, has passed or was rescheduled after the page was opened, the portal says so instead, and a rescheduled visit is shown with its new date, to be confirmed again. In the portal, a visit a technician confirmed is shown as confirmed by your company, under its name, never by the technician; and the client never sees the visit's notes.

> **i** Only the visits with the hourglass are waiting for the client: a Planned visit without it never had a confirmation request, for example one scheduled with Ask the client to confirm by email unticked. The client can still confirm it in the Agenda tab, and Notify client sends the request.

Correcting a finished visit (administrators)

A finished visit can no longer be edited. When its figures are wrong, an administrator clicks Correct the visit, which offers Arrived, Left, Travel time (minutes, round trip), Kilometres, the Client not present box with its Reason, the Signer's name and Replace the signature: tick it and have a new signature drawn, otherwise The signature collected at the visit is kept. Save stays disabled until something has changed. The visit's time entries on the ticket follow the new times and travel: a travel entry is added for each technician who did the visit when there was none, or removed when the travel becomes 0, and their Billable boxes stay as they were. The PDF report attached to the ticket is rewritten, and every change goes to the Audit log. On a closed ticket the times and the travel are locked, since that time is already part of the closed ticket, but the kilometres, the signature and Client not present can still be corrected.

> **i** The internal note written when the visit was finished keeps its original text: after a correction, the right figures are in the visit window and in the PDF report.

On a credit contract the correction also offers the intervention type of the work and of the travel. When the technicians' entries have different types the field reads Several, and they only change if you pick a type. A type chosen here applies to every entry of the visit, while the ticket's Time tab changes one entry at a time. The type can still be corrected on a closed ticket.

How a visit is billed

The time at the client is billable time like any other: it counts against the contract with an hours bank that the ticket's time counts against, the one chosen on the ticket or otherwise the client's contract in force on the day of the visit, and it falls outside any hours bank when there is none. The travel and the kilometres follow that same contract, according to the Travel and kilometres options of its form in the client's Contracts tab (see Travel and kilometres of on-site visits in the Clients chapter):

- Bill travel time, on by default. When it is off, the travel time is never counted nor billed; it only stays on the ticket.
- Travel counts against the hours included, on by default and only available while Bill travel time is on. When it is on, the travel counts in the hours used, like the work; when it is off, the travel is kept out of the hours used and billed apart, at the contract's Price of an extra hour (€).
- Price per km (€): the kilometres of the visits that count against the contract are billed at this price. Left empty, they are listed but not billed. The kilometres do not depend on Bill travel time.

A visit that no contract with an hours bank covers is billed as time outside any hours bank, travel included, and its kilometres are listed without a value. Reports › Hours & billing shows the travel and the kilometres of each month in columns of their own (see the Reports chapter), and a contract's Hours bank window marks with a car icon the travel it counts. The client report's Hours bank section only shows the travel a contract counts in its hours, never travel billed apart, kilometres or prices.

On a credit contract the time at the client debits credits instead of counting hours, at the rate of its intervention type, and the travel follows the contract's Travel option: Travel debits credits, Bill travel separately (at the Hourly price of travel billed separately) or Free. The kilometres are billed at the Price per km, as on a contract with an hours bank.

Emails to the technicians

From 18:00 on the evening before a visit, in the visit's time zone, every technician of a Planned or Confirmed visit receives one email, Your on-site visits tomorrow, that lists their visits of the next day in order, each with its times, the client, the address and the ticket, and a button that opens SLApp Helpdesk. The app checks every ten minutes, so a visit added later that evening for the next day arrives in an email of its own. Each technician is sent each visit once; a rescheduled visit is sent again for its new day, while cancelled visits and visits already started are not. Only active technicians with an email address receive it, in your company's email language, and nothing is sent while your company's account is suspended. They are also emailed when a client confirms one of their visits in the Client Portal (see The client's confirmation).

Your visits in your own calendar

My calendar (iCal), in the Agenda's header, gives you a private link to your own visits that Outlook, Google Calendar or a phone's calendar can subscribe to: your visits then show up there and keep up to date by themselves. The calendar is read-only; every change is made in the Agenda. It holds the visits you are a technician of, from 90 days ago to two years ahead, except the cancelled ones. Each one is named after the client and the ticket, with the address as its location and the ticket, the contact, the technicians and the notes in its description, and the visits still Planned show as tentative.

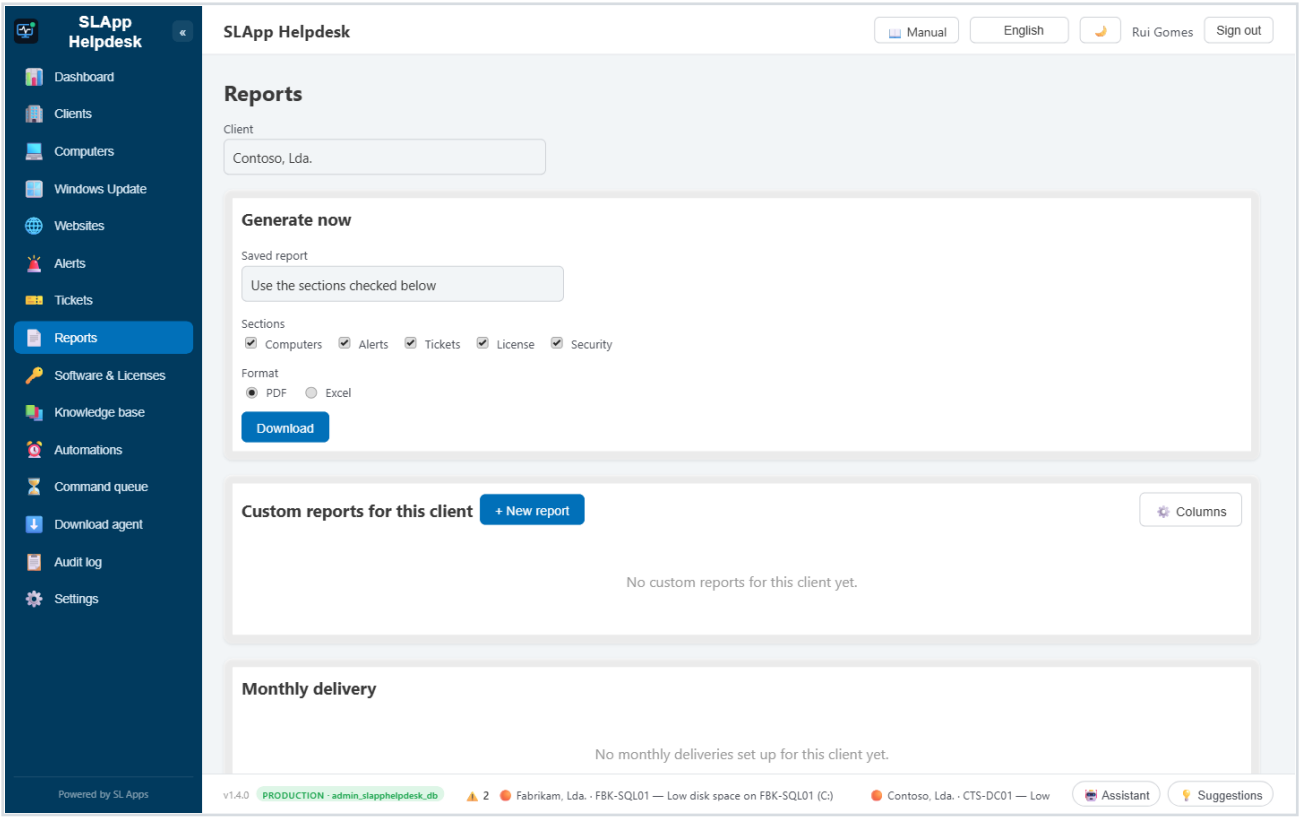
1. Click My calendar (iCal), then Create my link.
2. Copy the Calendar link with Copy, or click Add to calendar app to hand it to the calendar app of the device you are on.
3. Subscribe to it as How to subscribe explains: in Outlook, Add calendar › Subscribe from web; in Google Calendar on the web, Other calendars › + › From URL, after which it also shows in the Google Calendar app on Android; on an iPhone or iPad, Add to calendar app, or Settings › Calendar › Accounts › Add Account › Other › Add Subscribed Calendar.

Calendar apps read the link every few hours, and Google can take up to a day, so a change made in the Agenda takes a while to appear there. The link needs no sign-in: whoever has it sees your visits, with their clients, addresses and contacts, so do not share it. If it gets out, click New link, which stops the old link at once (subscribe again with the new one), or Switch off, which stops it until you create a new one; both ask for confirmation, and creating, replacing and switching off a link are recorded in the Audit log. The link also stops working when your user is deactivated.

10. Reports

The Reports module lets you produce a client-facing summary of what has been happening on a client's estate: its tickets with their satisfaction ratings, the time logged and the on-site visits, the hours and credits used from its contracts, its computers with their Windows updates, outdated software, security posture, vulnerabilities and hardware lifecycle, its assets and licenses, the alerts, websites,

domains, Microsoft 365 and network devices, its infrastructure servers (Active Directory, Exchange, SQL Server, storage, failover clusters, DHCP, IIS and Hyper-V) and its backups. You can download a report on the spot as a PDF, which opens with a cover and a one-page summary, or as an Excel workbook, save custom report layouts for a client, and set up an automatic monthly email so one of the client's own contacts receives the report without anyone having to remember. Open it from Reports in the sidebar; it is available to every technician, and only shows the clients of your own company. Administrators see two tabs at the top: Client reports, which holds everything described up to Monthly delivery, and Hours & billing, the month's hours banks of every client for invoicing (described at the end of this chapter). Everyone else sees the client reports alone, without tabs.



The Reports screen after choosing a client: Generate now, custom reports and monthly delivery.

Pick a client first

Everything in this module is tied to one client. At the top of the page there is a single Client drop-down (showing — choose — until you pick one). Until a client is selected the page only shows the message Pick a client above to generate or manage its reports. Once you choose a client, three cards appear below: Generate now, Custom reports for this client and Monthly delivery. Switching to another client reloads all three cards for that client and resets the Saved report choice in the Generate now card back to Use the sections checked below.

What goes into a report

A report is built from up to twenty-nine sections, grouped in five areas: Service, Computers and equipment, Monitoring, Infrastructure and Business continuity. In the PDF each section has a header, a one-line summary, the key figures, a chart where it helps and a table, and each area starts on a new page (see Cover, summary and layout); in the Excel file each section is a separate worksheet. Sections that belong to a module of the Complete plan (the infrastructure modules, Backups, Third-party backups, Network devices and the contract sections) only appear when your plan includes that module. The sections are:

- Computers — every active computer of the client with its Name, OS, whether it is Online right now (it reported in within the last couple of minutes) and its Last boot time. This section is a current

snapshot, not limited to the report period.

- Alerts — the alerts raised during the report period, with their messages in the report's language. In the PDF the alerts still open come first, one per row (Source, Computer / item, Severity, Message, Open since), the most severe and then the oldest first. The resolved ones follow, grouped: the repeats of an alert on the same computer or item share one row, with how many times it was raised (Times) and the first and last time (First seen, Last seen), the most repeated first. A value that changed between repeats shows as its range, for example a memory usage between 95.0 and 97.0%; alerts of one kind about several things of the same computer or device, such as its interfaces or its drives, are listed after the kind, the most frequent first. A note under the table says how many alerts the rows hold. The Excel sheet keeps every alert, one by one, with Computer, Type, Severity, Message, Opened and Resolved (unresolved alerts show Open). A message the app cannot translate, such as a device's own text or a log line, is shown as it was recorded. In the alerts of failed backups, what the agent reported for each copy that failed (for example a timeout connecting to an FTP server, or a SQL Server instance that cannot be reached) is also in the report's language, and whole even when the recorded text was cut short; words of Windows, SQL Server or an FTP server that the app does not know stay as they were recorded.
- Tickets — the tickets opened during the report period: # (the ticket number), Subject, Status, Priority and Opened date.
- License — the client's active licenses: Product, number of Computers and Users covered (Unlimited when no limit is set), Expiration and Periodicity. Current snapshot.
- Security — one row per active computer with its Firewall state, the Antivirus/EDR products detected (a product that is switched off is marked (off)) and the count of Pending updates. Current snapshot.
- Vulnerabilities — one row per active computer of the client, the most exposed first, with its Level as in the Vulnerabilities module (Critical, High, Medium, OK or No data), the exploited software (CISA KEV) with the installed and the available version of each application, marked [ransomware] when CISA knows its vulnerabilities to be used by ransomware, their CVEs, the number of Other outdated applications, the Pending updates with the age in days of the oldest one, and the protection: antivirus/EDR and firewall. The Excel sheet adds Why, the reason for the level, and has the full list of CVEs. Current snapshot.
- Hardware lifecycle — one row per active computer with its Type (Workstation or Server, with (VM) for a virtual machine), its Age in years and where it comes from (purchase date, warranty start, BIOS, OS installation or agent installation), the Operating system, the OS support (until when it is supported, or when it ended), the Warranty (its end date, with the manufacturer when it was looked up, or Expired) and the Recommendation (Replace, Upgrade OS, Renew warranty or OK). The Excel sheet adds the manufacturer, the model, the serial number and the date to replace the computer by. Never any price. Current snapshot.
- Websites — the client's monitored websites with their Name, URL, current Status (Up or Down), the Uptime and the Avg. response measured during the report period, and the date their SSL certificate expires (SSL expires).
- Domains — the client's monitored domains with the date their registration Expires and the Days left, the Registrar, the state of SPF, DKIM (with the selector found) and DMARC (with its policy), each OK, Warning or Error, and the Last check. Current snapshot.
- Active Directory — one row per domain controller of the client's Active Directory domains, with its Health (OK, Warning or Critical), the state of dcdiag, Replication and Services (OK, Warning or Error), its Time offset from the PDC emulator and the Last check. The Excel sheet adds the domain's account counts: Users enabled, Locked out, Password never expires, Inactive 90 days and Privileged accounts. Current snapshot.

- Third-party backups — one row per backup job found on the client's computers, such as a Veeam, Windows Server Backup or Acronis job, except the jobs whose expected frequency is set to Ignore, with the Computer, Product, Job, Last result (Success, Warning or Failed), Last success (Never when there is none, followed by (overdue) when it is older than the job's expected frequency allows), Last run and Expected (Daily, Weekdays or Weekly). The Excel sheet adds an Overdue column. Current snapshot.
- Hours bank — for each contract of the client with an hours bank in force during the report period, the hours used in the report period and in the contract period up to the end of the report period, with the hours remaining and the extra hours. Hours only, never prices or values; see The Hours bank section below.
- Credits — for each credit contract of the client in force during the report period, or that still has a balance, the credits bought, used and expired during the report period, the balance at its end (In debt when it is negative) and the next expiry. Credits only, never prices; see The Credits section below.
- Satisfaction — the ratings the client's contacts gave to tickets in the period: the average out of 5, how many ratings, the share rated 4 or 5, a chart of the 1 to 5 distribution and the list of ratings.
- Time logged — every time entry on the client's tickets in the period: the date, the ticket, the technician, the kind (work or travel), the duration and the note, with the total at the end. Hours only, never values.
- On-site visits — the visits in the period (cancelled ones left out): the date, the ticket, the technicians, the time on site, the kilometres, the status and the work summary.
- Windows updates — the computers with Windows updates pending, the most overdue first: how many, the oldest pending update, the month it was released and how many days it has been pending.
- Outdated software — every application with an update available on the client's computers, those with actively exploited vulnerabilities (KEV) first, with the installed and available versions.
- Assets — the client's asset register: type, name and serial number, status, who or what it is assigned to, warranty end and supplier. Never the purchase price.
- Microsoft 365 — the client's connected organisation as of its last sync: active users, licensed users, users without MFA and administrators, the licences bought with how many are assigned and free, and the users that need attention (administrators, no MFA, mailbox above 90%).
- Network devices — the SNMP devices of the client: model, firmware (flagged when an update is available), online status, uptime and CPU / memory / temperature, plus how many Syslog entries were received in the period by severity (never the messages themselves).
- Exchange, SQL Server, Storage and iSCSI, Failover clusters, DHCP and IIS — one section per infrastructure module, from its daily health check: each server with its health and the figures that matter for that module (databases, queues and certificates; databases, size, oldest full backup and failed jobs; iSCSI sessions, MPIO and low volumes; nodes, failed groups and resources and CSV free space; scopes and their use; stopped sites and pools, certificates and 5xx errors).
- Hyper-V — the client's Hyper-V hosts and every virtual machine with its state, vCPUs, memory and replication health.
- Backups — the backup jobs run by the agent with their runs in the period: how many ended OK, with warnings or failed, the last run and the volume copied. Only job names and numbers: never destinations, paths or credentials.

In the Generate now card you choose the Period (Last month, the default; This month; Last 3 months; or Custom, with a start and an end date) and the Report language (the language of your interface by default, or English, Portuguese or French). A monthly delivery always covers the previous full calendar

month and is written in the client's notification language, or in your company's email language when the client has none. The period is printed on the cover and on every page.

A section with nothing to show is never printed. The full report, with every section, leaves out every section without data for the client (Computers is always kept); a custom report lists the sections it asks for that had no data on one closing page, Sections without data, instead of printing empty pages.

Cover, summary and layout

The PDF is A4 portrait. The first page is a cover with your company's logo (Settings › Company), the report title (the custom report's name, or Monthly report - Service report when the period is not one calendar month), the client's name and logo, the period, when it was generated, who prepared it and your company's contacts, with "Powered by SLApp Helpdesk" at the foot. The second page is the summary: up to nine key figures, the points of attention found in the data (for example certificates about to expire, computers without antivirus, failed backups) and the table of contents with the page where each area and each section starts. The table of contents stays on the summary page whenever it fits there, in two columns when one does not fit; only a very long one moves to a page of its own. Every page after the cover has a header with your company and the client and period, and a footer with the generation date and "Page X of Y". Each area (Service, Computers and equipment, Monitoring, Infrastructure, Business continuity) starts on a new page, opened by a banner with the area's name, one sentence on what it covers, the report's sections in it (In this area: …) and an illustration of the area; above the name, Area 2 of 5 shows where the area sits in the report when it has more than one. The area's first section starts right under the banner and the next ones follow it, on the same page when there is room. The tables wrap long texts instead of cutting them; an infrastructure server that is not healthy says why in a Reason column. The colours are sober: neutral greys plus one accent colour taken from your company's logo - navy when that colour is a red, amber or green, so it is never mistaken for a status colour; the illustrations of the areas use shades of that accent colour only, never a status colour. The Excel file starts with an Overview sheet with the same logos and key figures.

The Hours bank section

The Hours bank section lists every contract of the client with an hours bank that was in force on at least one day of the report period, including one that ended during it. Each contract is shown in its contract period that holds the end of the report period (or holds its own last day, when it ended before that), with the contract's name, that period, the hours included, the hours used in the report period, the hours used in the contract period up to the end of the report period, the hours remaining at that point and the extra hours that arose in the report period; time logged after the end of the report period is left out. For a monthly contract the two figures of hours used are the same; for a yearly or one-time contract the second one also counts the earlier months of the contract period.

> **i** The section shows hours only, never the price of an extra hour nor any value, because the report goes to the client's contact; the amounts to invoice are in Hours & billing. The full report leaves the section out when the client has no contract with an hours bank in force during the report period; a custom report that includes it then lists it on the closing Sections without data page. The travel time of on-site visits is included only when the contract counts it against its hours; travel billed apart and kilometres never appear in the client report.

The Credits section

The Credits section has one row per credit contract of the client in force on at least one day of the report period, or that still has a balance (a credit contract that has ended keeps being used while it has valid credits). The PDF shows the Contract, the credits Bought, Used and Expired during the report period, the Balance at its end, written as N in debt when it is negative, and the Next expiry, with the credits that expire and the last day they can be used. The Excel sheet adds the credits Adjusted during

the period and the State (OK, Low or Negative), and a chart compares the credits bought and used.

> ℹ Like the Hours bank section, it never shows prices: the value of the packages is only in Hours & billing. The full report leaves it out when the client has no credit contract to show; a custom report that includes it then lists it on the closing Sections without data page.

PDF or Excel, and the logos

Both formats contain the same sections; for some of them the Excel sheet has a few more columns than fit on a PDF page (see What goes into a report). In the PDF a text longer than its column wraps onto up to three lines (it is only cut short with … beyond that), a row is never split across two pages, the column headings are repeated at the top of a new page, and a table stops after 300 rows with a note pointing to the Excel export; the Excel sheet always has everything in full. The Excel workbook starts with an Overview sheet with your company's and the client's logos, the report title, the client, the period, when it was generated, who prepared it, the key figures and points of attention of the summary and, for a custom report, the sections without data; then comes one sheet per section with data, named after it in the report language, which is handy when the client wants to filter or pivot the data themselves.

The cover shows your company's logo, large, at the top (taken from Settings › Company; the company's icon when there is no logo, or its name when there is neither) and the client's logo next to the client's name (taken from the client's record; its icon when it has no logo). The accent colour of the report is taken from your company's logo. A small image is printed smaller rather than blurred, so the company logo is printed at its full size on the cover when it is at least 500 px wide: Settings › Company keeps it up to 1200 px. Each image is optional and simply omitted when it is not set.

> ℹ Use PNG or JPEG logos. An SVG logo is never embedded, and GIF or BMP images may be dropped by the PDF or Excel renderer — in every case the image is silently skipped rather than failing the report. So if a logo is missing from the report, check the image format in Settings › Company or in the client's record. A company logo uploaded before the cover existed (version 1.14.66) was reduced to 320 px and prints small: upload it again in Settings › Company.

Generate a report now

The Generate now card produces a one-off download. It has a Saved report drop-down (default: Use the sections checked below), the Sections check-boxes grouped by area (all ticked by default), a Period drop-down (Last month, This month, Last 3 months, or Custom, which adds the From and To (inclusive) dates), a Report language drop-down (Interface language, English, Portuguese or French), a Format choice between PDF and Excel, and a Download button, which reads Generating… while the file is being built.

1. Choose the client in the Client drop-down at the top of the page.
2. In the Generate now card, either leave Saved report on Use the sections checked below and keep every section ticked for the full report, or pick one of the client's custom reports from Saved report — the check-boxes are then greyed out and the saved selection is used instead.
3. Pick the Period and the Report language. For a Custom period fill in both dates; the To date is included in the report.
4. Select PDF or Excel under Format.
5. Click Download. The file is saved by your browser as the client's name (spaces and punctuation replaced by hyphens) followed by -report- and today's date, with a .pdf or .xlsx extension, and a notification confirms it; if the report cannot be built, the error is shown in the card and in a notification. Each download is recorded in the audit trail.

> **i** On-demand generation only knows two shapes: the full report or a saved custom report. If you untick some sections without picking a saved report, Download stays disabled and the hint Pick a saved report above, or check every section, to enable the download. appears. To download a partial report, save that combination as a custom report first (see below) and select it in Saved report. Download also stays disabled while a Custom period is missing one of its dates.

Custom reports for this client

A custom report is a named selection of sections saved for one client — for example an executive summary with only Tickets and Security. Custom reports are per client; one you create for client A is not offered for client B. The card lists them in a table with the columns Name, Sections (the sections included, comma-separated) and Created, plus a Delete link on each row. Until you create one, the card shows No custom reports for this client yet.

1. With the client selected, click + New report in the Custom reports for this client card.
2. In the New custom report window, type a Name (the placeholder suggests e.g. Monthly executive summary).
3. Under Sections, tick the sections to include; all of them are pre-ticked, so untick the ones you do not want. Save stays disabled until there is a name and at least one section is ticked. Cancel closes the window without saving.
4. Click Save. The report now appears in the table and in the Saved report and Report drop-downs of the other two cards.

To change a custom report, click its Name, Sections or Created cell: the Edit custom report window opens with the current name and sections. Changes take effect immediately for any monthly delivery that uses this report — the next email will contain the updated selection, and the Report column of the Monthly delivery table shows the new name. To remove a custom report, click Delete and confirm in the Delete custom report dialog.

> **i** Deleting a custom report does not cancel monthly deliveries that used it. Those subscriptions are kept and switch to Full report (every section), so the contact keeps receiving a report — just a bigger one. Review the Monthly delivery table after deleting a definition.

Monthly delivery to a client contact

The Monthly delivery card automates sending. Each subscription emails one of the client's contacts a report for the previous month, on a day of your choosing, as a PDF or Excel attachment. This is opt-in: nothing is sent to anyone until you create a subscription here. The table shows, per subscription, the Recipient (contact name and email), the Report (a custom report name or Full report (every section)), the Format, the Day of month, the Status (Active or Paused) and Last sent (Never until the first delivery). Each row has Send now, Pause / Resume and Unsubscribe links. Until you create one, the card shows No monthly deliveries set up for this client yet.

1. With the client selected, scroll to Set up monthly delivery at the bottom of the Monthly delivery card.
2. Pick the recipient in the Contact drop-down. Contacts are listed with their email address; a contact without one is flagged no email and cannot be subscribed. Subscribe stays disabled until a contact is chosen.
3. In Report, keep Full report (every section) or choose one of the client's custom reports.
4. Choose PDF or Excel under Format.
5. Enter the Day of month on which the report should be sent (1 to 28; default 1).

6. Click Subscribe. The subscription appears in the table with Status Active and Last sent Never, and the form clears for another one.

Delivery happens automatically in the early morning (shortly after 06:00 UTC) of the chosen day. The contact receives an email whose subject is the client name in square brackets followed by Report — and the month covered (for example August 2026), with the file attached. The email and the report are written in the client's Notification language (in the client's record), or in your company's email language when the client is on System default. The attachment is named after the client followed by -report- and the year and month in which it was sent (so the report covering August, sent in September, carries 2026-09 in its name). Last sent shows the date of the latest email that actually went out (a send that failed, for example because email sending is not set up, leaves it as it was); each successful send is also recorded in the audit trail.

- Send now emails the subscription's report to its contact right away, after you confirm in the Send the report now dialog, which names the recipient, the month and the format. It is the same report the monthly delivery sends: the same sections, format and language, covering the previous month. The button shows Sending… while the report is built and sent, which takes a few seconds; a toast then confirms Report sent to …, or says why it could not be sent, and Last sent shows today's date. The monthly schedule does not change: the next automatic send still happens on its usual day. Send now also works on a paused subscription, and is greyed out when the contact has no email address.
- A contact deleted from the client no longer receives the monthly report. Its subscription stays in the table, with (removed) after the name, and Send now is greyed out; Unsubscribe removes the row.
- Pause stops future sends without deleting the subscription; the Status badge turns to Paused. Resume switches it back on and schedules the next send from today, so a subscription that was paused past its usual day does not fire immediately — it waits for the next occurrence of its day.
- Unsubscribe removes the subscription after you confirm Stop the monthly delivery to … in the Unsubscribe dialog.
- If the client has no contacts, the form is replaced by the message This client has no contacts yet — create one first in Clients › Contacts. The contact must also be active and have an email address; otherwise Subscribe is rejected with an explanatory error.

> **i** The day of month is capped at 28 so that every month has that day. The first report goes out at the next occurrence of that day at 06:00 UTC: if you subscribe after that moment (later that day, or later in the month) the first delivery is next month, not this one. Also note that a scheduled report is skipped (never sent) if the contact's email address has since been removed — keep contact details up to date in the client's record.

Hours & billing (administrators, or by explicit permission)

The Hours & billing tab, next to Client reports at the top of the module, gathers one month of the hours banks of every client of your company, ready for invoicing: per client and contract with an hours bank, the hours used against those included and the extra hours to bill, and per client the billable time that counts against no hours bank. Technicians without admin rights do not see the tab, and the server refuses them the data too. Going back to Client reports keeps the client and any half-filled form as you left them.

Pick the month at the top: the Month field starts on the previous month, the buttons either side of it (Previous month and Next month) move one month at a time, and no month after the current one can be chosen; the current month can be shown, and its figures grow as time is logged. Months are calendar months in UTC, as everywhere in the hours bank. The table has one row per contract with an hours bank in force during the month (two for a yearly contract whose anniversary falls in it, one for each contract year), plus an Outside any hours bank row for each client with billable time that no

contract with an hours bank covers. Its columns are:

- Client and Tax ID: the client and its tax ID.
- Contract: the contract's name, with its contract period underneath for a yearly or one-time contract, or Outside any hours bank.
- Periodicity and Included: the contract's periodicity and the hours it includes per period. Contract value, hidden at first, adds the contract's own Value when shown through Columns.
- Used: the hours counted in that month, each entry rounded up as its contract says. For a yearly or one-time contract a second line gives the hours used in the whole contract period, up to the end of the month.
- Extra hours: the extra hours that arose in that month, in red. For a yearly or one-time contract each month only shows the part of the period's overrun that happened during that month, so every extra hour is billed once, in the month it occurs, and adding up the months of a period gives its whole overrun. All the hours of an Outside any hours bank row are extra.
- Extra hour rate and Extra value: the contract's price of an extra hour, and the month's extra hours multiplied by it. For a yearly or one-time contract the cents are rounded over the whole period, so adding up the months of a period gives exactly the extra value of that period, and a month can be one cent off its extra hours times the price. A dash means there is no price, which is always the case on an Outside any hours bank row: agree those hours apart.
- Travel, Kilometres and Total value, at the end, for the on-site visits: Travel is the travel time of the row's visits, with in the hours used under it when it is already part of Used, or billed apart with its hours and, when the contract has a price of an extra hour, their value. Kilometres are those of the visits finished in the month, with the contract's price per km and their value under them when it has one. Total value is what the row bills for the month: the extra value plus the value of the travel billed apart and of the kilometres; a dash means the row has neither a price of an extra hour nor a price per km. Each contract's Travel and kilometres options decide how its travel and kilometres are billed (see How a visit is billed in the Agenda chapter). A client whose only item outside any hours bank is the kilometres of a visit still gets an Outside any hours bank row, with those kilometres and no value.

Type in the boxes under Client, Tax ID and Contract to filter the rows, click a heading to sort by it, and use Columns to choose the columns shown. The footer adds up Used, Extra hours, Extra value, Travel, Kilometres and Total value of the rows on screen, so filtering on one client gives that client's totals; when the month has visits with travel or kilometres, a note under the table recalls how they are billed. Below the table, Time entries (N) opens the list of every billable, finished time entry of the month, with Date, Client, Contract, Ticket, Technician, Time (rounded up as its contract says), Note, Type (Work, Work (visit), Travel or Kilometres) and Kilometres, its own filters and a Total of the entries shown, with their kilometres. Each on-site visit finished in the month with kilometres adds a Kilometres line, with the technicians who did the visit, its address as the note and no time; travel that a contract does not bill is left out. A click on a row of the table opens this list filtered to that row's client and contract; Show every entry clears the filters.

Export Excel downloads billing-YYYY-MM.xlsx with two sheets: Summary, the rows of the table (with the contract period, the month's Hours used and Extra hours, Used in period and Extra in period for the contract period up to the end of the month, the price and the Extra value, followed at the end by seven columns for the on-site visits, Travel hours in hours used, Travel hours billed apart, Travel value, Kilometres, Price per km, Kilometres value and Total value, and a Total row), and Entries, every time entry of the month with its date, client, contract, ticket, subject, technician, minutes and note, followed by its Type and Kilometres, with one Kilometres line per visit. Export CSV downloads billing-YYYY-MM.csv, the summary rows alone, with the same seven columns at the end, for an invoicing program: fields separated by a semicolon (;), UTF-8, decimals with a comma in Portuguese and French and with a point in English. Both files are written in the language of the app, always hold every row whatever the

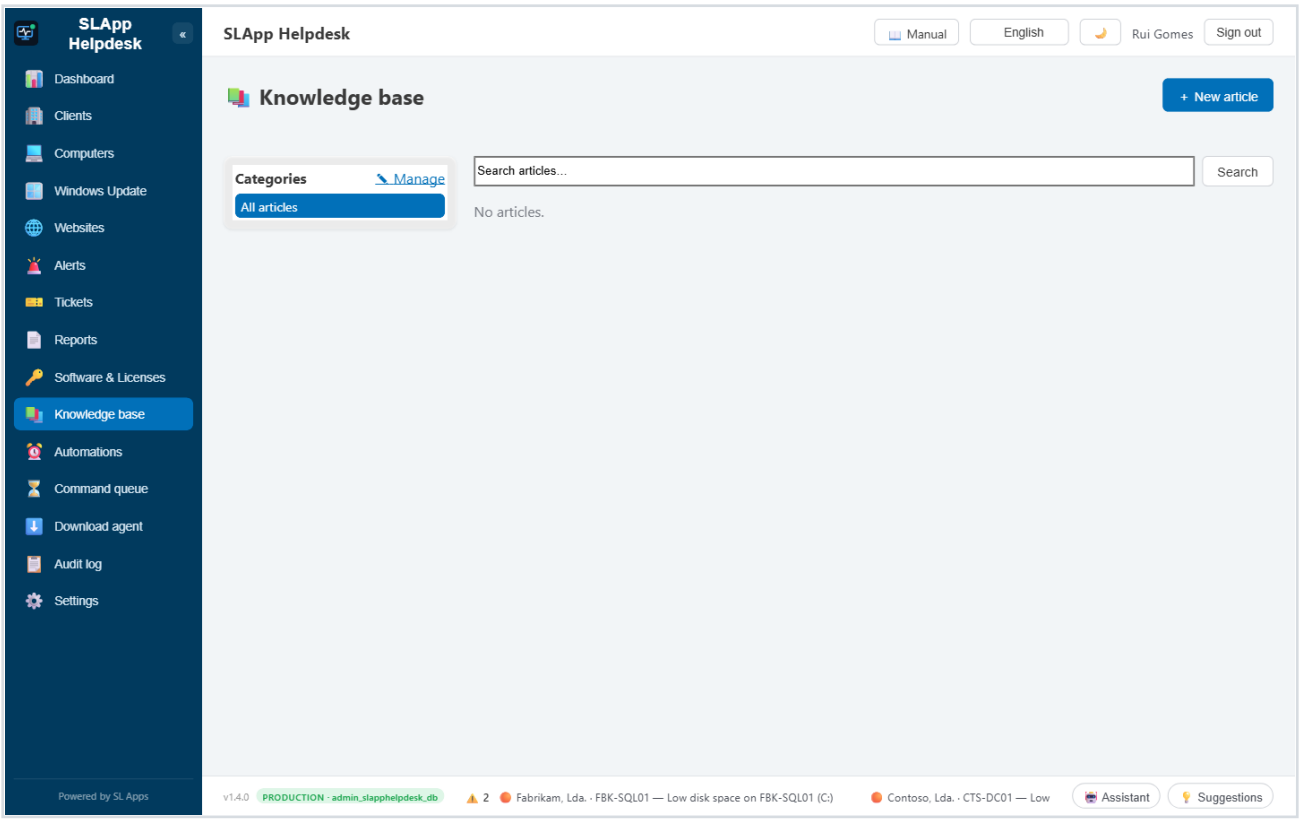
filters on screen, and each export is recorded in the Audit log. The columns for the on-site visits come after the ones that were there before them, so an invoicing program that reads the columns by their position keeps working.

A credit contract has one row per month, with no hours included nor extra hours (its Used is only for information). As soon as the month has one, the table ends with Model (Hours bank or Credits), Credits bought, Packages value, Credits used, Credits expired and Balance at month end, a dash on the other rows, with totals for Packages value, Credits used and Credits expired; the time entries gain Intervention type and Credits. The row bills the packages registered in that month (by the day they were recorded, not their purchase date) plus the travel billed separately and the kilometres, so its Total value is the Packages value plus those; a package voided in a later month counts as a negative value in that month and leaves the earlier one as it was. The Excel and CSV files add the same six columns at the very end of Summary and Intervention type and Credits at the end of Entries, so an invoicing program that reads the columns by their position keeps working.

> **i** The figures are worked out again from the current data each time a month is shown or exported. Editing a contract's hours, price, dates, rounding or travel and kilometres options, removing a contract, changing an entry's minutes or Billable box, or correcting a finished visit later also changes months already invoiced, so export the month when you invoice it and keep the file.

11. Knowledge Base

The Knowledge base module is where your team writes and organises how-to articles, procedures and reference notes, and decides which of them your clients can read in the client portal. Open it from the Knowledge base entry in the main menu. The screen has two parts: on the left, the tree of Categories and their sections; on the right, the list of articles with a search box. Every article is either Global (readable by every client) or belongs to one specific client, and only articles marked Published appear in the portal (the one exception: a contact who writes an article in the portal always sees their own drafts).



The Knowledge base module: category and section tree on the left, article list and search on the right.

How the knowledge base is organised

Content is arranged in three levels: Category, Section, Article. A category (for example "Email") contains sections (for example "Outlook" and "Webmail"), and each section holds articles. An article can also be left without a section; it then shows as No section in the list. Categories and sections can be managed here or in Settings, under Knowledge base; both places edit the same structure.

- All articles at the top of the tree removes the section filter and lists every article you are allowed to see.
- Clicking a section name filters the list to that section only. The filter is combined with whatever you type in the search box. When nothing matches, the list says No articles. or No articles in this section.
- Categories and sections created by a client in the portal show a lock badge with the client's name when you are in manage mode. Only that client sees them, and your team cannot rename, delete or add sections to them (you can still move them up or down in the tree).

Managing categories and sections

Click Manage at the top of the Categories card to switch the tree into management mode (the link changes to Done). Every category and section then shows Move up and Move down arrows; items created by your team also show three more small controls:

- Move up and Move down arrows reorder the item among its neighbours and are greyed out at the top and bottom of the list. Moving a category takes all its sections along with it.
- The people icon (Who sees it (audience)) opens the audience dialog described below.
- The pencil icon (Rename) opens a small window with the name filled in (a double-click on the name does the same): Save keeps the new name, Cancel closes the window (an empty name cannot be saved).
- The bin icon (Delete) asks Delete? with a tick to confirm and a cross to cancel, so there is no separate pop-up.

- + Section next to each category opens the New section window, and + New category at the top of the tree opens the New category window; both ask only for the name, and Create adds it.

> **i** Deleting a category removes all the sections inside it. The app refuses to delete a category that still contains sections a client created in the portal: the Delete? prompt simply stays and the category remains in the tree, with no message in this screen. Those clients must remove their own sections first. Articles are never deleted with their section, but a Global article left in a deleted section stops appearing in the portal until you open it and choose another section (or No section). Articles scoped to a client, including those written by the client's contacts, stay visible to that client.

Choosing who sees a category or section

By default a category or section created by your team is visible to all clients. To restrict it, enter manage mode and click the people icon next to it. The Who sees dialog states whether the item is visible to all clients or to a given number of clients, and offers an All clients checkbox, a Search client... box and one checkbox per client. Ticking individual clients makes the item visible only to them; ticking All clients clears the selection. All clients is ticked automatically whenever no client is selected, so to restrict the item simply tick the clients you want. Click Save to apply or Cancel to leave it unchanged. In the tree, manage mode shows a badge on each item: All, or the number of clients that can see it.

> **i** A section also inherits its category's audience. If a category is limited to two clients, a section inside it can never be seen by anyone else, even if the section itself is set to All clients. Only active clients can be picked as an audience.

Writing and publishing an article

Click New article at the top right to open the editor, or click any article card to edit it; the dialog is titled New article or Edit article accordingly. If a section was selected in the tree, a new article is pre-filled with that section. The editor contains these fields:

- Title (required): Save draft and Publish stay disabled until it is filled in.
- Section: a dropdown grouped by category, with a — No section — option.
- Scope: Global (all clients) or one specific client. A client-specific article is only ever shown to that client's contacts.
- Summary: one line shown under the title in the list and at the top of the article in the portal.
- Content: the body of the article. This is a plain multi-line text box; line breaks are preserved when the article is displayed, but there is no formatting toolbar.
- Tags: comma-separated keywords, used by the search and shown at the bottom of the article in the portal.

1. Click New article.
2. Type the Title, pick a Section and choose the Scope (Global or a client).
3. Write a one-line Summary, the Content and any Tags.
4. Click Save draft to keep it internal, or Publish to make it visible in the portal to the clients it applies to. Cancel (or clicking outside the dialog) closes the editor without saving.
5. To change its state later, open the article again and click the other button: Publish promotes a draft, Save draft takes a published article offline.

When editing an existing article a Delete button appears on the left; it asks Really delete? and you confirm with Yes, delete or back out with No. The list itself shows, for each article, its Published or Draft badge, the summary, the category and section (or No section), a globe with Global or a building icon

with the client name, and the author. Articles are sorted with the most recently updated first.

Searching

Type in the Search articles... box and press Enter or click Search. The search looks in the title, summary, tags and the full content of every article, and respects the section currently selected in the tree. Select All articles to search across everything.

What clients see in the portal

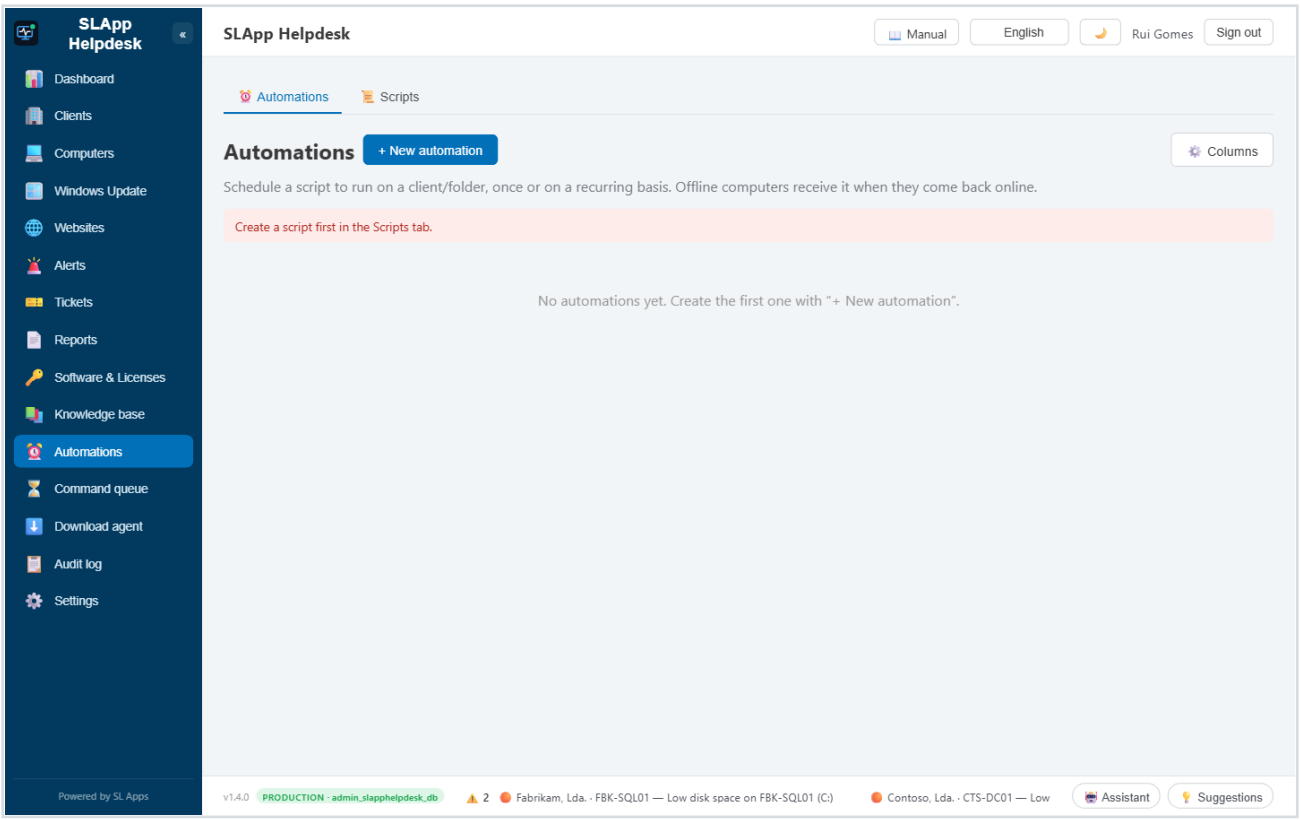
In the client portal, a contact sees the same tree restricted to the categories and sections whose audience includes their client, plus any categories and sections their own company created. The article list shows published Global articles (only if the section and category are visible to that client), published articles scoped to their client, and the contact's own drafts. Global articles are marked General in the portal, and an article with no section is also listed under General in place of the category and section. Articles written by your team are read-only for clients.

Contacts can also write their own articles from the portal. Those articles are always bound to the contact's client, so no other client ever sees them, and they appear in your module with the client's name and the contact as author; your team can open, edit, re-publish or delete them. In the portal a contact can edit or delete an article only if it was written by a contact of their client and either they wrote it themselves or they are the client's primary contact. Likewise, contacts can add their own private categories and sections (shown with a lock icon on both sides); they may add such a section under one of your team's categories that is visible to them, which is what prevents you from deleting that category until they remove it. Contacts may place their articles in any section visible to them.

> **i** For a Global article two settings control visibility and both must allow it: the article's Scope and the audience of the section and category it sits in. A Global article placed in a section restricted to one client is seen only by that client, while a Global article with No section is visible to every client as soon as it is published. An article scoped to a client is always shown to that client, whatever the audience of its section. Keep sensitive client procedures scoped to that client to be safe.

12. Automations

Automations let you schedule a script from your Scripts library to run on a set of computers, either once or on a recurring basis (daily, weekly or monthly). Typical uses are a weekly temp-file cleanup, a monthly inventory refresh, or a one-off fix you want to push to every machine of a client at a quiet hour. Open the module from Automation & updates › Automations in the sidebar; it is available to every logged-in technician. It has two tabs: Automations, the schedule described here, and Scripts, the library the automations run (see Scripts in the Settings chapter).



The Automations list with the schedule, status and last result of each automation.

Every automation is built on a script that already exists in the Scripts library (the Scripts tab of this module). If no script exists yet, the + New automation button is disabled and the screen shows the notice Create a script first in the Scripts tab. Until you create your first automation, the list area shows the message No automations yet. Create the first one with “+ New automation”.

The automations list

The table lists all automations of your company, active ones first and then ordered by their next run. Each row shows:

- Name — the name you gave the automation.
- Script — the script it runs, followed by (PS) for PowerShell scripts or (CMD) for command-prompt scripts.
- Target — the client it runs on, or All clients, followed by the folder name when one was chosen (for example Acme · Head office).
- Recurrence — Once, Daily, Weekly or Monthly.
- Next run — the date and time of the next execution in your local time zone, or a dash when nothing is scheduled any more (for example a one-off automation that has already run).
- Status — Active or Inactive. Only active automations are triggered by the scheduler.
- Last result — a short summary of the last dispatch: how many computers the script was sent to and when (the time is shown as a UTC timestamp). Hover over it to read the full text.
- The action buttons at the end of the row: Run now (play icon), Deactivate or Activate (pause / check icon), Edit (pencil) and Delete (bin).

Creating an automation

1. Click + New automation. The New automation dialog opens.
2. Enter a Name (for example Weekly cleanup).

3. In Script, replace the — choose — entry with the script to run. The list shows every script of your Scripts library with its type, PowerShell or CMD.
4. In Client, pick the client whose computers should run the script, or leave All clients to target every computer of every client.
5. In Folder, optionally narrow the target to one folder of that client, or leave All folders. The Folder list is disabled until a client is chosen, and it is cleared whenever you change the client.
6. In When, set the date and time of the first run (in your local time).
7. In Recurrence, choose Once, Daily, Weekly or Monthly.
8. In Run as, choose System (elevated/admin) to run the script with administrator rights under the local System account, or Logged-in user to run it inside the session of the user currently logged in.
9. Leave Status as Active so the scheduler picks it up, or set it to Inactive to prepare it now and enable it later.
10. Click Save. The button stays disabled until Name, Script and When are all filled in. Click Cancel (or anywhere outside the dialog) to discard the form.

If you left Client at All clients, a confirmation titled Run across the whole fleet? appears before saving. It reminds you that the script will run on all computers of all clients, states whether it runs as System or in the user's session and, for recurring automations, the recurrence. Confirm only if that is really what you want.

> **i** Scripts are sent to the computers listed in the Computers module that belong to the chosen client (and folder, if any); computers you have deleted are excluded. Computers that are switched off or offline at that moment are not skipped: the command waits in the Command queue and is delivered as soon as the computer reconnects. If a previous run of the same script is still waiting in the queue for a computer (for example because it has been offline since), a new run does not add a second copy: the waiting command is reused, so the computer runs the script once when it reconnects.

> **i** The automation always runs the current version of the script: if you edit the script in the Scripts tab, the next run uses the new content. A script set to Logged-in user can only run while someone is logged in; if the computer is online but has no user session, the command is handed to the agent service after about two minutes and runs as System instead, and the result is prefixed with a warning saying so.

How the schedule works

The server checks for due automations every minute. When the time in Next run is reached, the script is queued on every target computer and the Last result column is updated with a summary like Sent to 12 computer(s), including the dispatch time. What happens next depends on the recurrence:

- Once — the automation runs a single time and is then automatically set to Inactive, with no next run.
- Daily, Weekly, Monthly — the next run is moved forward by one day, one week or one month from the previous scheduled time, always to a moment in the future. The automation stays Active.

> **i** If the server was unavailable for some time, a recurring automation does not fire several times to catch up: it runs once and its next run is pushed forward until it lands in the future. Likewise, if you set When to a moment in the past, the automation is triggered within about a minute of saving (as long as it is Active).

Running, pausing, editing and deleting

- Run now — queues the script on the target computers immediately, without changing the schedule. The recurrence continues exactly as before, and Last result shows Manual run followed by the

dispatch summary. Run now also works on an Inactive automation, so you can use it for ad-hoc executions.

- Deactivate / Activate — quick toggle of the Status without opening the form. An Inactive automation keeps its settings and next run date, but the scheduler ignores it until you activate it again. If its next run date is already in the past when you reactivate it, it fires at the next scheduler pass.
- Edit — opens the Edit automation dialog with all fields prefilled; change anything and click Save. Removing the client here triggers the same whole-fleet confirmation as when creating.
- Delete — opens the Delete automation confirmation (Delete the automation "name"?) and removes it permanently. Commands that were already queued on computers are not cancelled.

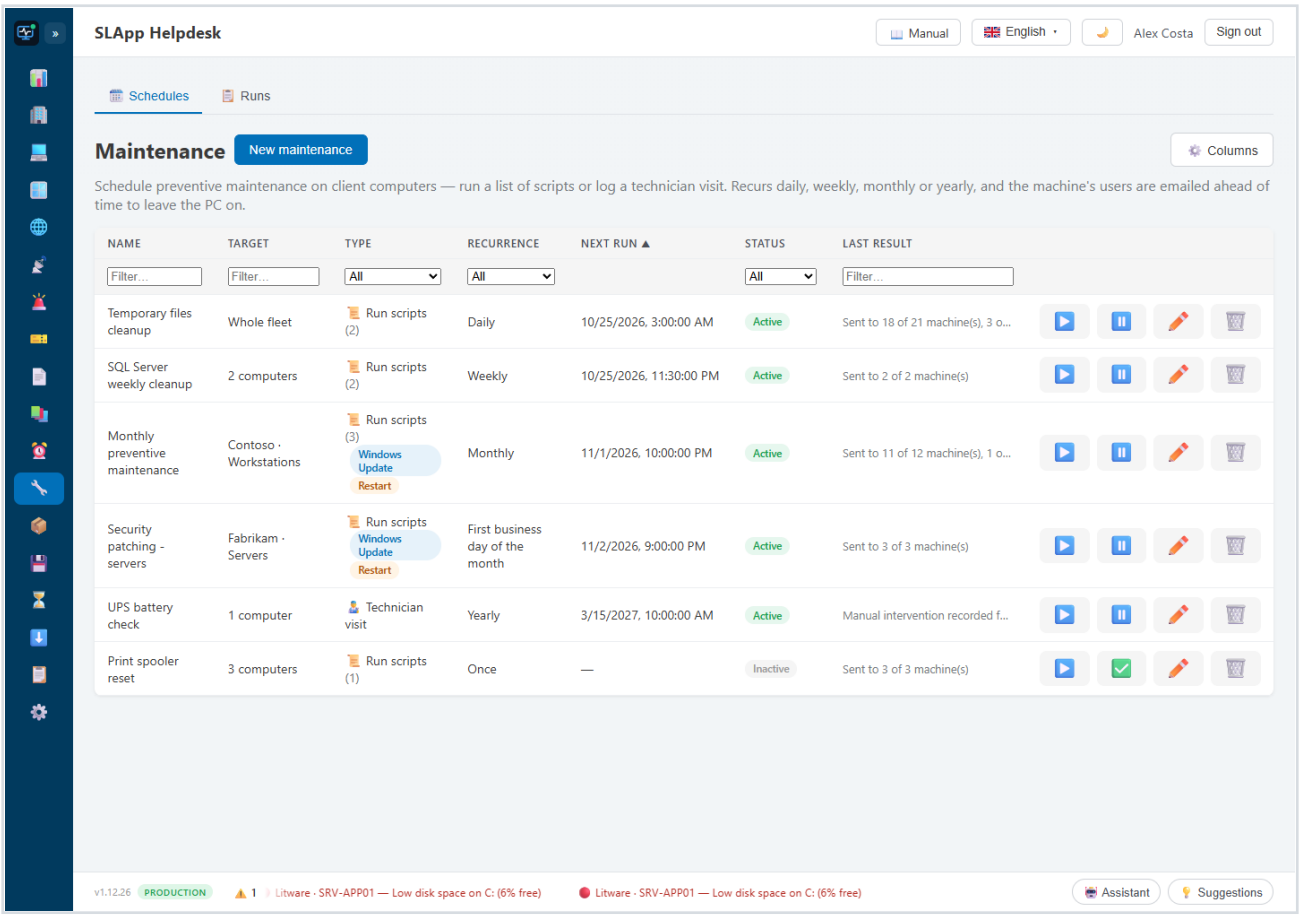
Where to see the script output

The Last result column only tells you how many computers the script was sent to and when, or whether some commands could not be queued. It does not contain the output of the script itself. To see what each computer actually did, open the Command queue module or the computer's detail page, where every queued command shows its state and result once the agent has executed it.

> **i** Every creation, change, activation, deactivation and deletion of an automation is recorded in the Audit log (an administrators-only module) with the user who did it, which is useful when a script unexpectedly runs across many machines. Run now is not written to the Audit log; its trace is the Manual run entry in Last result and the commands in the Command queue.

13. Maintenances

The Maintenances module schedules preventive maintenance windows on your clients' computers. A maintenance can either run an ordered list of scripts automatically or record a human intervention that a technician carries out and closes by hand. Before each window the people who use the machines are emailed to leave their PC on, the client's main contact is always kept informed, pending Windows updates can be installed and the computer restarted at the end, and every occurrence is written to a run history so you can prove exactly what happened and when. Open the module from Automation & updates › Maintenance in the sidebar; it is available to every logged-in technician, and the scripts a maintenance runs come from the shared Scripts library (the Scripts tab of Automation & updates › Automations).



The Maintenances module: the Schedules tab with each maintenance's target, type, recurrence, next run and last result.

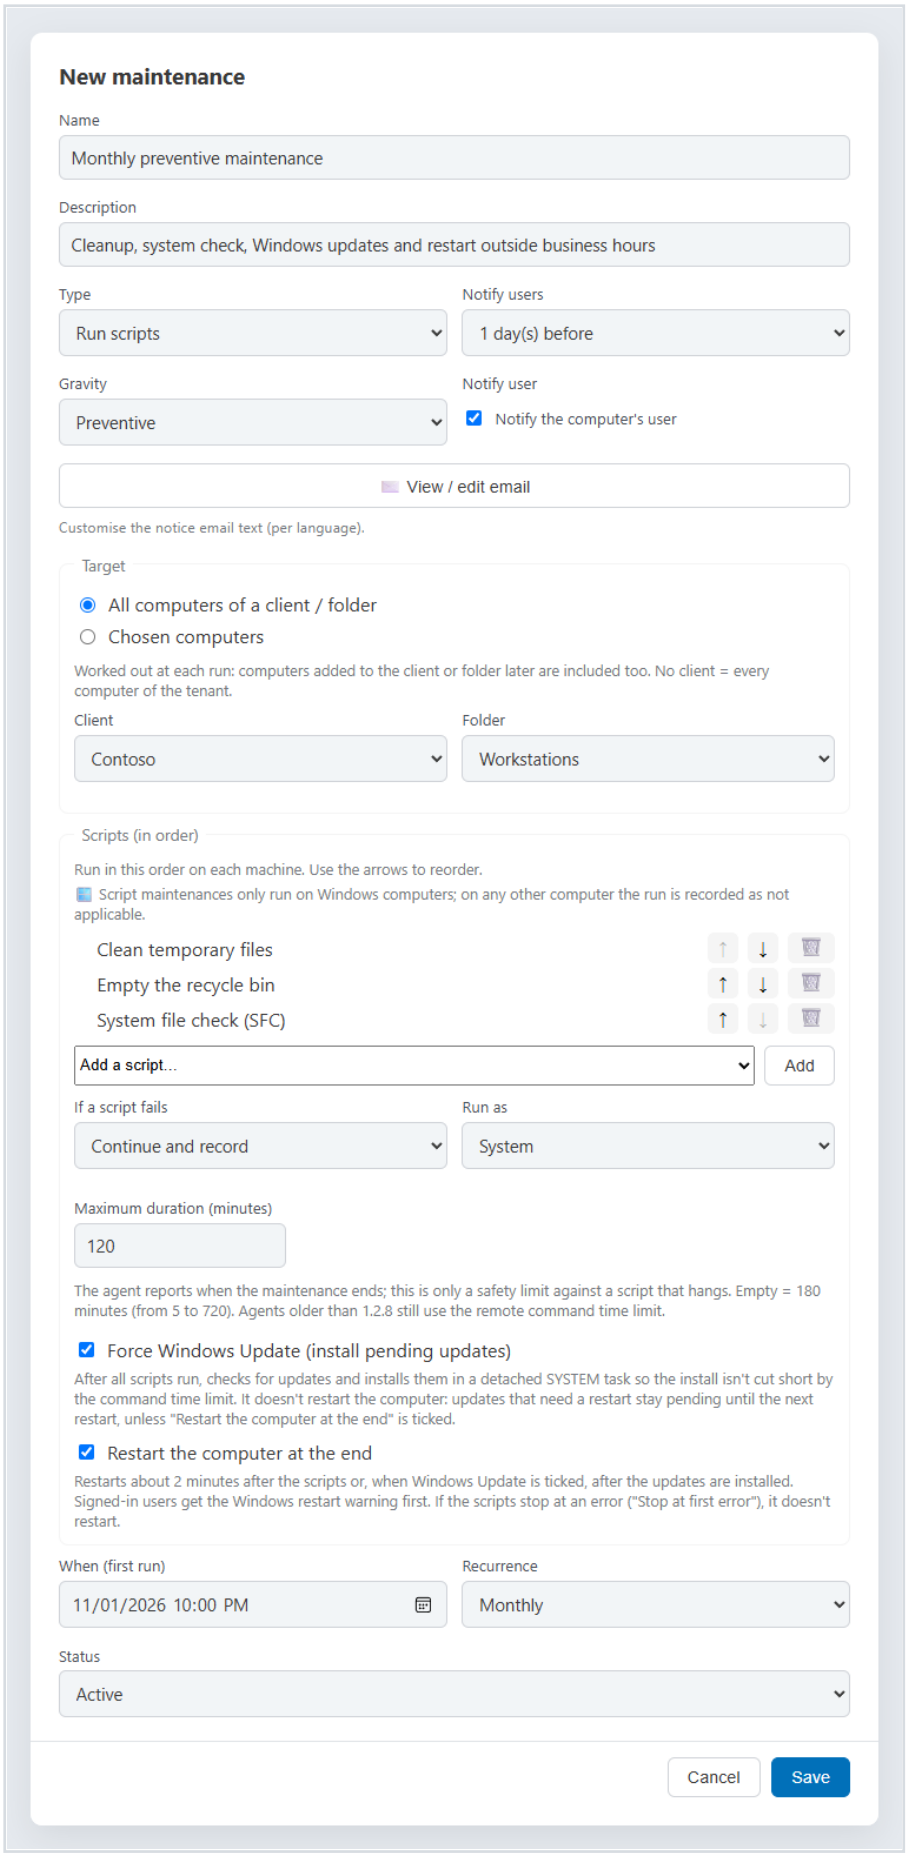
The module has two tabs. Schedules is where you create and manage the maintenance plans; Runs is the history of every occurrence that has already fired, machine by machine. Until you create your first plan the Schedules tab shows No maintenances yet. Create the first one with “+ New maintenance”.

The schedules list

The table lists every maintenance of your company, active ones first and then ordered by their next run. Each row shows:

- Name — the name you gave the maintenance.
- Type — Scripts (it runs a list of scripts) or Human (a technician carries it out by hand). A Scripts maintenance also shows a Windows Update tag and/or a Restart tag when those options are ticked.
- Target — for a maintenance on a client / folder, the client followed by the folder when one was chosen, or Whole fleet when no client was chosen; for a maintenance on chosen computers, how many computers it has (for example 3 computers).
- Recurrence — Once, Daily, Weekly, Monthly or Yearly.
- Severity — Critical, Preventive or Communication; this sets the coloured banner and tone of the notification emails.
- Next run — the date and time of the next occurrence in your local time zone, or a dash when nothing is scheduled any more.
- Status — Active or Inactive. Only active maintenances are triggered by the scheduler.
- The action buttons at the end of the row: Run now (play icon), the email editor (envelope icon), Edit (pencil), Duplicate (two overlapping squares) and Delete (bin).

Creating a maintenance



The New maintenance dialog: target, recurrence, the ordered script list, forced Windows Update, severity and notification options.

1. Click + New maintenance. The New maintenance dialog opens.
2. Enter a Name (for example Monthly preventive maintenance).
3. Choose the Type: Run scripts to execute an ordered list of scripts automatically, or Human intervention to record a maintenance that a technician performs and closes by hand from the Runs tab.
4. Under Target, choose one of the two modes — All computers of a client / folder or Chosen computers — and set it up (see Choosing the target below). Scripts and human-intervention maintenances choose their target the same way.
5. In When, set the date and time of the first run (in your local time).
6. In Recurrence, choose Once, Daily, Weekly, Monthly or Yearly.
7. For a scripts maintenance, add the scripts to run from the Scripts library and put them in the order you want — they execute in that order on each machine. Choose the On error policy (see below) and Run as System (elevated/admin) if the scripts need administrator rights.
8. Optionally tick Force Windows Update (install pending updates) to install the pending Windows updates after the scripts have run, and/or Restart the computer at the end to restart the machine when the maintenance has finished. The two options are independent (see Windows Update and restart). When either is ticked, Include Hyper-V hosts (off by default) decides whether Hyper-V hosts get them too (see Hyper-V hosts).
9. Set the Severity — Critical, Preventive or Communication — which controls the banner and tone of the emails.
10. Set the notification options: Notify user, how many hours before the run the notices go out, and, with the envelope button, the email text itself (see Notifications).
11. Leave Status as Active so the scheduler picks it up, or set it to Inactive to prepare it now and enable it later, then click Save.

The On error policy is chosen per maintenance and decides what happens when one script in the list fails on a machine: Continue and record runs the remaining scripts anyway and records the failure, while Stop at first error leaves the rest of the list unrun on that machine. It is set independently for each maintenance, so a cleanup routine can carry on past a non-critical error while a sensitive sequence stops the moment something goes wrong. Stopping at an error also cancels what was due after the scripts on that machine: neither Windows Update nor the restart happens there.

> **i** The scripts of a maintenance run in the order you set them on each machine, and machines are processed in parallel. A maintenance always uses the current version of each script: if you edit a script in the Scripts tab of Automation & updates › Automations, the next run uses the new content. Scripts are sent through the same store-and-forward command queue as Automations, so a machine that is online receives them at the scheduled time.

Script maintenances run only on Windows computers. When the target also covers a computer with another operating system (Linux or macOS, for example under All clients, in a folder that mixes systems or ticked among the Chosen computers), that computer is skipped: nothing is sent to it, it is left out of the advance notices and of the could-not-be-contacted email, and it is not counted as offline. Its entry in the Runs tab is recorded as Cancelled, with a Result saying the maintenance is not applicable because the computer is not Windows. Human-intervention maintenances are not affected: they apply to every computer in the target.

Maximum duration (minutes) is a safety limit, not a time slot. The agent (1.2.8 or later) tells the server when the scripts have finished, however long they take, and sends the result then; if at that moment the server is being updated or the computer has no network, the result is kept and sent as soon as

possible. Only a script that hangs is stopped when the limit is reached (empty = 180 minutes, from 5 to 720), and the run is then recorded as Failed with a note saying so. If the agent itself stops in the middle (an update, a restart of the computer), the run is recorded as Failed with that explanation. Computers with an older agent still run the maintenance as a remote command, within the time limit of Settings -> General -> Remote control.

Choosing the target

The Target box of the dialog has two modes, and a maintenance uses one or the other, never both. Pick the mode with its radio button; the box then shows only the fields of that mode:

- All computers of a client / folder — choose a Client and, optionally, one of its folders in Folder (Folder stays disabled until a client is chosen; All folders keeps every folder of the client). The target is worked out again at every run, so computers added to that client or folder later are included automatically. If you leave Client on All clients, the maintenance targets every computer of your company, including the ones added later; because that is the widest choice there is, Save first opens the Maintenance for the whole fleet confirmation, and the list then shows Whole fleet in the Target column.
- Chosen computers — exactly the computers you tick in the list, and only those. The selection is fixed: a computer added later to the same client is not included until you edit the maintenance and tick it.

In Chosen computers mode the list has its own filters to help you find the machines: Client (All clients or one client), Folder (the folders of that client, with All folders; disabled while Client is on All clients) and a Filter… text box that matches the computer's name, client or folder. Each computer in the list shows its name followed by its client and folder. These filters only narrow the list — they are not saved with the maintenance, and changing them never unticks anything. Between the filters and the list:

- Select all (N) ticks every computer the filters currently show, N being how many are visible. It adds them to the computers already chosen, and it is greyed out when every visible computer is already ticked.
- Clear unticks every chosen computer, including the ones the current filters are hiding.
- The counter (for example 12 chosen) counts the whole selection. A chosen computer hidden by the current filters stays chosen and is still counted; the counter then also says how many of them are hidden by the filters.
- While no computer is ticked, the box shows Choose at least one computer. and Save stays unavailable.

> **i** To reach most of a client's computers but not all of them, set the Client filter to that client, click Select all (N) and untick the few machines to leave out. To build a list across several clients, change the Client filter and click Select all (N) again — each click adds to the computers already chosen.

> **i** A chosen computer that is later deleted in the Computers module drops out of the maintenance. If every chosen computer has been deleted, the maintenance runs on no computer at all — it never falls back to the whole fleet; edit it and choose new computers, or delete it.

Windows Update and restart

When Force Windows Update (install pending updates) is ticked, the machine searches for pending updates after the maintenance scripts have finished and installs them. On its own it does not restart the computer: updates that need a restart stay pending until the machine is next restarted, unless Restart the computer at the end is ticked as well. It can be combined with a script list, or used on its own as a

plain monthly Windows-update window with no scripts at all.

> ℹ The update runs inside a detached scheduled task under the SYSTEM account, on purpose: a long cumulative update can take far longer than the few minutes a normal remote command is allowed before it is stopped, so it is handed off to run to completion on its own.

Restart the computer at the end is a separate option. You can tick it with or without Force Windows Update, and also on a scripts maintenance with no scripts at all — for example a weekly restart window. When it is ticked:

Shut the computer down at the end is the alternative to the restart: the computer shuts down about 2 minutes after the scripts or, with Force Windows Update ticked, once the updates are installed, instead of restarting. The two cannot be ticked together, so ticking one clears the other. It is skipped when the scripts stop on an error, and the Hyper-V hosts skip it like the restart unless Include Hyper-V hosts is ticked. The notice email lists it as Computer shutdown, and the row shows a Shutdown badge. Turning the computer back on needs someone on site or Wake-on-LAN.

- Without Windows Update, the computer restarts about 2 minutes after the scripts have finished, so the maintenance result reaches the server first.
- With Windows Update, the computer restarts once the updates have been installed — always, even when there was nothing to install.
- Users who are signed in to the computer get the Windows restart warning before it restarts.
- If the scripts stopped at an error (On error set to Stop at first error), neither Windows Update nor the restart happens on that machine, and the Result in the Runs tab says so.

> ℹ Maintenances that already had Force Windows Update ticked before the restart became a separate option were switched to Restart the computer at the end as well, so they keep restarting after the updates as before. Untick it if you only want the updates installed.

Hyper-V hosts

Restarting a computer that hosts Hyper-V virtual machines is not like restarting a workstation: when the host goes down its virtual machines are saved and, after the restart, they may fail to restore from that saved state, leaving them off or forcing them to start from scratch. Installing Windows updates usually ends in such a restart. That is why the form has Include Hyper-V hosts, just under Restart the computer at the end. It is off by default, also on every maintenance created before the option existed, and it only matters when Force Windows Update or Restart the computer at the end is ticked. While it is off, on a Hyper-V host of the target:

- The maintenance scripts run as written, in the same order and with the same On error policy. Only the two built-in options below are left out: a script that itself restarts the computer or installs updates still does so on the host. So that a script can opt out, the scripts of such a host run with the environment variable SLAPP_HYPERV_HOST set to 1 (in PowerShell: if ($env:SLAPP_HYPERV_HOST -eq '1') { ... }).
- The forced Windows Update and the restart at the end are skipped. The host's Result says which of the two was skipped and why.
- If the maintenance has no scripts (only Windows Update and/or the restart), nothing is sent to the host. Its entry in the Runs tab is recorded as Cancelled, with a Result saying so, and it is not counted as offline.

SLApp Helpdesk knows a computer is a Hyper-V host when it has been added to Infrastructure › Hyper-V, or when its Services tab (after Scan services) lists the Hyper-V Virtual Machine Management service

(vmms). On top of that, before its scripts run, the maintenance looks for that service on the computer itself, so a host SLApp Helpdesk did not know about yet is spared as well; its Result then says the host was found on the computer itself. Virtual machines are not affected: a VM is an ordinary computer and gets the update and the restart as usual.

> **i** Tick Include Hyper-V hosts only for a window you have prepared the hosts for, for example with their virtual machines shut down cleanly or moved to another host beforehand. To keep updating hosts in a controlled way, a separate maintenance with only the hosts as Chosen computers and the option ticked makes the choice explicit.

Notifications

Ahead of every run — by the number of hours you set in the maintenance — SLApp Helpdesk sends the notices. Two audiences are handled separately:

- The client's main contact is always emailed. Their message lists all the computers that will have the maintenance and the names of the scripts that will run, so the client has the full picture.
- The person who uses each machine is emailed only when Notify user is ticked. Their message shows only their own computer, together with the scripts, and asks them to leave the PC on. The recipient is the contact linked to that computer on its detail page; a machine with no linked contact simply has nobody extra to notify.

Every maintenance carries a Severity — Critical, Preventive or Communication — shown as a coloured banner at the top of the email so the reader immediately grasps how important it is. The emails are sent in the client's notification language, which is set per client in the Clients module (Editing a client › Notification language); each client can therefore receive its notices in Portuguese, English or French regardless of your own interface language.

The body of the notice is a template you can edit. Click the envelope button on the maintenance row (or the same control in the form) to open the email editor. It has a language dropdown so you can write the text for each of the three languages, a live Preview, and a Restore default button. The template understands these variables, which are filled in for each recipient when the email is sent:

- {computers} — the list of computers (the whole list for the main contact, only the recipient's own machine for the user notice).
- {scripts} — the names of the scripts that will run, followed by Windows Update and the computer restart when those options are ticked. When Include Hyper-V hosts is off, a recipient whose computers are all Hyper-V hosts is not told about Windows Update or the restart, a recipient with hosts among other computers reads which hosts skip them, and with no scripts to run the hosts are left out of the notice altogether.
- {date} — the scheduled date and time.
- {client} — the client's name.

> **i** If you leave the template empty for a language, SLApp Helpdesk uses its built-in default text for that language, so a client whose language you have not customised still receives a complete, sensible message.

If a computer cannot be contacted at the scheduled time — it is switched off or offline — the maintenance is not silently skipped. The occurrence is recorded as not performed and a Critical email is sent stating that the maintenance was not carried out because the computer could not be reached, so nobody assumes it happened when it did not.

How the schedule works

The server checks for due maintenances every minute. Two things are timed off the same schedule: the advance notices go out once, the configured number of hours before the run, and the maintenance itself is dispatched when its Next run time is reached. Each occurrence is claimed atomically, so even with several server processes a maintenance is never notified twice or run twice. What happens to Next run afterwards depends on the recurrence:

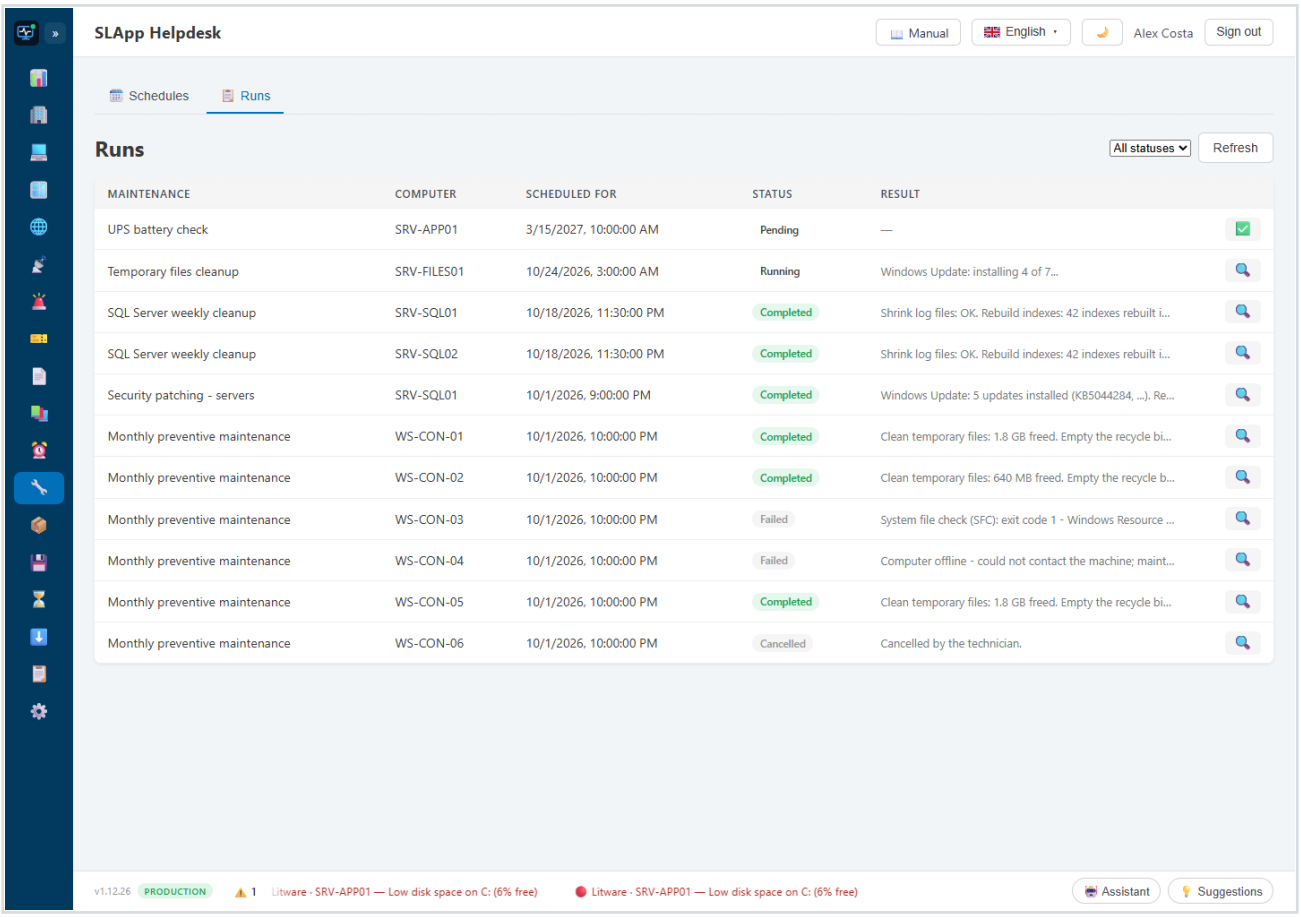
- Once — the maintenance runs a single time and is then automatically set to Inactive, with no next run.
- Daily, Weekly, Monthly, Yearly — the next run is moved forward by one day, week, month or year from the previous scheduled time, always to a moment in the future. The maintenance stays Active.

> **i** If the server was unavailable for a while, a recurring maintenance does not fire several times to catch up: it advances its next run forward until it lands in the future and runs once from there.

Human-intervention maintenances

A Human intervention maintenance sends the same advance notices and creates the same run records, but it does not execute anything on the machines — it is a reminder and a checklist for a technician who does the work in person or remotely. When the time comes, open the Runs tab, find the occurrence and mark it done once you have finished. The notices still go out so the client and the users know the technician is coming and can leave the machines available.

The Runs tab



The Runs tab: every occurrence, machine by machine, with its status and result.

The Runs tab is the history of every occurrence that has fired, one entry per computer, with the maintenance it belongs to, When it ran, its Status (Waiting, Running, Completed, Failed and so on —

the same vocabulary as the Command queue) and the Result. For a scripts maintenance the result is the combined output of the scripts (and the Windows Update and restart steps when they were ticked); for a human maintenance you set the status yourself when you close it. The entry of a computer that is not Windows shows Cancelled, with a Result saying the scripts maintenance does not apply to it. A Hyper-V host whose forced Windows Update or restart was skipped shows a Hyper-V host tag next to its Status; hovering it, or opening the result, explains what was skipped and that Include Hyper-V hosts includes it. This is where you confirm that a maintenance really ran, see which machine was offline, and read exactly what each script did.

The per-computer output of a scripts maintenance also lands in that computer's own command history (Remote actions tab), alongside every other remote command, so you can review it from either place.

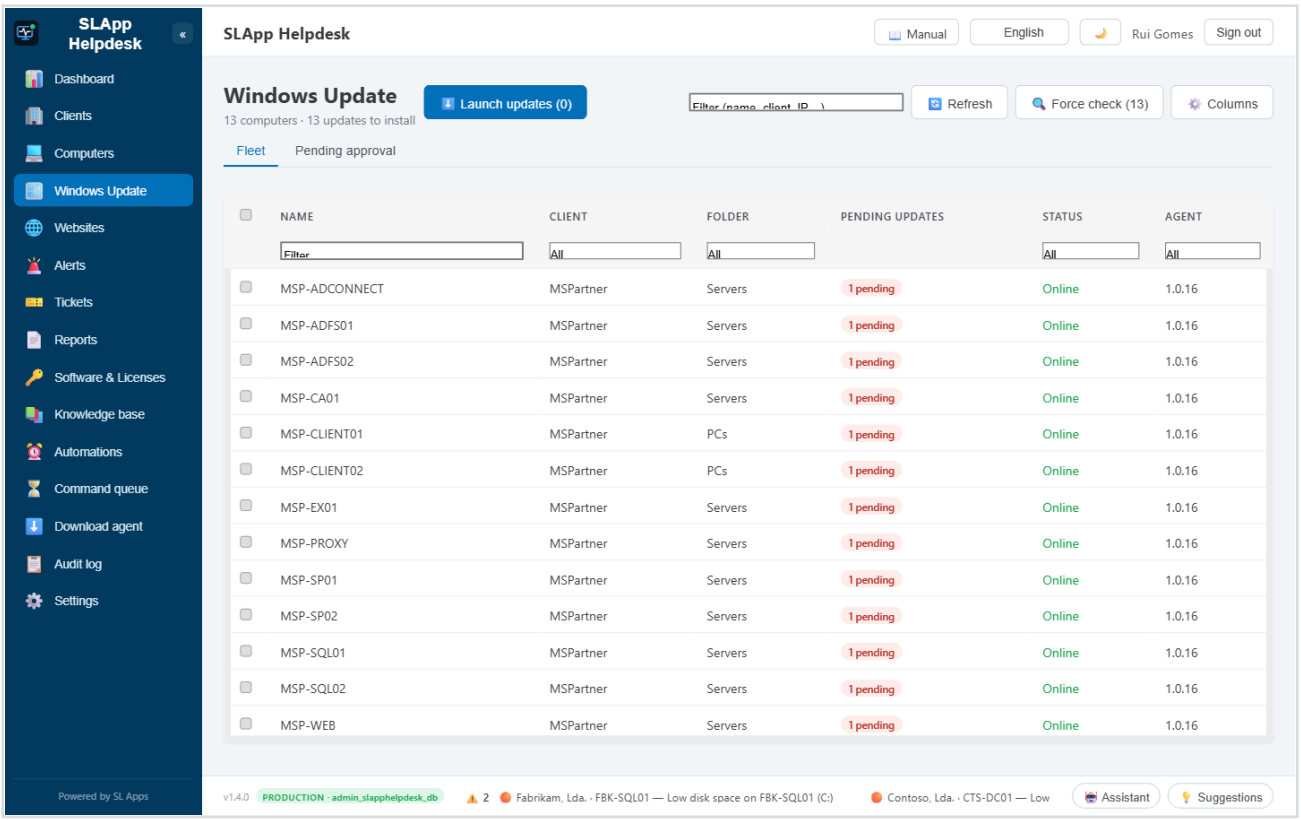
Running, editing and deleting

- Run now — dispatches the maintenance on its target computers immediately, without changing the schedule. Use it to test a new plan or to run an ad-hoc maintenance out of cycle.
- Edit — opens the maintenance with all fields prefilled; change anything and Save. Editing the schedule re-arms both the run and its advance notice.
- Duplicate (two overlapping squares) — opens the form filled with a copy of the maintenance, named "(copy)", with the same target, scripts, options, email and schedule. Nothing is created until you Save, so change the name, the target or the date first. Use it to repeat a maintenance for another client or at another time.
- Delete — removes the maintenance permanently, after a confirmation. Its run history is removed with it; occurrences that were already dispatched to computers keep running as queued commands.

> **i** Creating, changing and deleting a maintenance is recorded in the Audit log (an administrators-only module) with the user who did it, which is useful when a maintenance unexpectedly touches many machines.

14. Windows Update and patch policies

The Windows Update module gives you a fleet-wide view of every computer that still has Windows Updates to install, and lets you decide which updates are allowed and push the installation to many machines at once. It complements the per-computer view in the computer sheet: here you work across clients, not one machine at a time. Patch policies, which live in Settings, then take over the routine work by installing the updates you approved during recurring maintenance windows.



The Windows Update module, Fleet tab: every computer with pending updates, ready to be selected.

Opening the module and reading the header

Click Automation & updates › Windows Update in the left sidebar. The header shows a one-line summary such as "12 computers · 37 updates to install" (or "Nothing pending" when the fleet is clean); the computer count covers the whole list, while the updates total follows whatever filters you have applied. Next to it sits a free-text filter box labelled "Filter (name, client, IP…)", which matches the computer name, client, IP address and folder, and three buttons: Refresh reloads the list; Force check asks the online machines to re-scan right now, with the number of machines it will reach in parentheses (it is greyed out when none of them is online); and Launch updates starts the installation on the computers you selected. Below the header are two tabs: Fleet and Pending approval.

The Fleet tab

The Fleet tab lists only computers that report at least one update left to install. By default they are sorted with the most pending updates first; click any column header to sort by that column instead (a second click reverses the order). A second header row holds per-column filters; when nothing matches, the table shows "No results with these filters."

- Checkbox column: tick the computers you want to act on. The checkbox in the header is "Select all online" and only selects the machines in the current (filtered) list that are online. Offline computers cannot be ticked and show "Offline — cannot launch" when you hover them.
- Name: the computer's name (or hostname), as a link that opens its computer sheet. The filter under it is a text box.
- Client and Folder: where the computer belongs. Both filters are drop-downs listing the values present in the list; computers with no client show "— no client —" and computers with no folder show "—".
- Pending updates: a badge with the count, for example "5 pending".
- Status: Online or Offline. The filter lets you show only one of the two.
- Agent: the version of the agent installed on the machine, with a drop-down filter per version.

> The counts come from the agent's periodic inventory, which runs roughly every six hours. If Windows has installed updates on its own since then, the list is stale. Click Force check: it re-runs the scan on the online machines in the filtered list (or only on the ones you selected), shows a notice such as "Update check launched on 8 online machine(s)" above the table (close it with ×), and the page reloads itself after a few seconds.

Launching updates on several computers

1. Use the filters to narrow the list, for example to one client or one folder.
2. Tick the computers you want to update, or use the header checkbox to select every online machine in the current list.
3. Click Launch updates (the number in parentheses is how many computers are selected).
4. Read the confirmation, titled "Launch Windows Updates": it warns that ALL pending Windows Updates will be installed on those computers, that the installation runs in the background and that the machines may restart. Click Confirm to go ahead or Cancel to back out.
5. A Launch result card appears and fills in one line per computer: a tick with "Launched — installing in the background", or a cross with the error if the command could not be delivered. While the commands are going out it reads "Sending to the computers…".
6. When the last command has been sent, the selection is cleared and the list reloads. Leave the page as it is; the pending counts refresh when each computer reports its inventory again.

> The manual launch installs every pending update on the machine, regardless of any approval decision. Approvals only govern what patch policies install automatically. Because the command must reach the agent's service, which runs with elevated rights, it can only be sent to computers that are online right now, which is why offline rows cannot be selected.

The Pending approval tab

The same update (KB) is usually pending on many machines at once, so this tab groups updates by KB rather than by computer: you decide once per KB and the decision applies to every computer that reports it. Each row is one KB for one client, so the same KB appears on separate rows when several clients have it pending. A drop-down above the table, which reads "Client: All" by default and only appears when at least one row belongs to a client, filters the rows to a single client. It is only a filter: which scope a decision is recorded for is given by the Client column of the row you act on, not by the drop-down. When nothing is waiting, the tab shows "Nothing pending approval right now."

- KB: the Microsoft knowledge-base number of the update.
- Name: the update's title as reported by Windows ("—" when Windows gave none).
- Client: the client the row applies to. A row showing "— no client —" groups the computers that are not assigned to any client.
- Computers affected: how many computers of that client still have the update pending.
- Status: Pending, Approved or Rejected.
- Approve / Reject: record the decision for that KB and that client. The button matching the current status is disabled, so you can always switch a decision the other way; while the row is Pending both buttons are available.

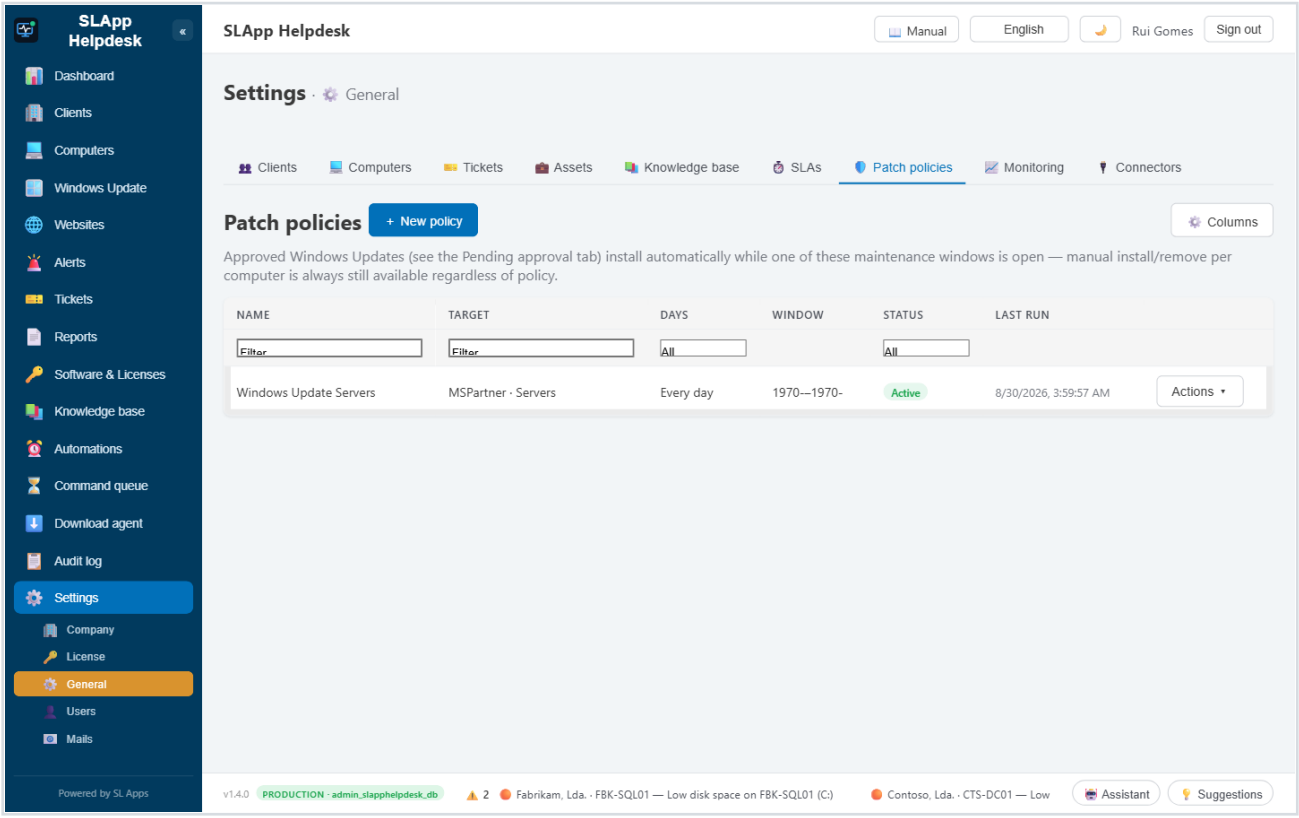
Approving or rejecting the "— no client —" row creates a decision for every client. It also acts as the fallback for any client that has no decision of its own for that KB: a client-specific decision always wins over the all-clients one, and a computer with no client only ever follows the all-clients decision. Only

administrators can record decisions; other users see the buttons but the save is refused and an error message appears above the table.

1. Open the Pending approval tab and, if useful, pick a client in the drop-down.
2. Review each KB and its title, and check Computers affected to gauge the impact.
3. Click Approve to allow patch policies to install it, or Reject to keep it from being installed automatically. An update left on Pending is not installed by policies either: only an explicit Approve lets it through.
4. The status badge updates immediately; rows stay in the list while the update is still pending somewhere, so you can revise a decision later.

Patch policies: where they live and what they do

A patch policy is a recurring maintenance window. While a window is open, the system looks at every computer the policy targets, finds the updates that are still not installed and have been approved for that computer's client (or for all clients), and queues their installation one KB at a time. Updates that are Pending or Rejected are left alone. The check runs every minute, and a KB that is already queued or being installed on a computer is never queued twice. Policies are managed in Settings, which is only visible to administrators: open Settings in the sidebar, choose General, then the Patch policies tab. Without at least one active policy, approved updates only get installed when a technician launches them manually.



Settings › General › Patch policies: the list of maintenance windows and the policy editor.

The list shows active policies first, then the inactive ones, each group in alphabetical order.

- Name: the label you gave the policy.
- Target: the client it covers (or All clients), followed by the folder when one was chosen, for example "Acme · Servers".
- Days: the days of the week it runs on, or "Every day" when all seven are selected.
- Window: the start and end time, for example 02:00–04:00.

- Status: Active or Inactive. Inactive policies are kept but never run.
- Last run: the last time the policy was evaluated while its window was open, or "—" if it has never run.
- Action buttons on each row: the pause/tick icon to Deactivate or Activate without opening the form, the pencil to Edit, and the bin to Delete (a confirmation titled "Delete patch policy" reminds you that the computers it targets will stop having approved updates auto-installed under it).

Creating a maintenance window

1. In Settings › General › Patch policies, click New policy. The form opens as "New patch policy" with all seven days selected and the window pre-filled from 02:00 to 04:00.
2. Type a Name, for example "Weekday nights".
3. Choose the Client, or leave All clients to cover every computer of every client. When a client is chosen you can further narrow it to one Folder (the default is All folders); otherwise the folder selector stays disabled.
4. Click the day buttons (Mon to Sun) to toggle which days the window is open; at least one day must stay selected.
5. Set Window start and Window end. The window may span midnight, for example 22:00 to 02:00.
6. Leave Status on Active (or pick Inactive to prepare the policy without running it yet) and click Save. Save stays disabled until a Name is typed and at least one day is selected.
7. If you left the client on All clients, a confirmation titled "Run across the whole fleet?" explains that approved updates will auto-install on any computer of any client during this window; click Confirm to accept or Cancel to go back to the form.

> **i** Schedule windows outside working hours: installing updates may restart the machine. Windows are evaluated on the server's clock in UTC, not in the local time of each computer, so allow for the time-zone offset when you set the start and end times.

> **i** A policy does not require the computer to be online at that exact moment: the installation command waits in the queue and is delivered as soon as the agent connects. You can follow those commands, like any other, in the Command queue module. Manual install or removal of individual updates from a computer's sheet remains available whatever the policies say.

15. Vulnerabilities

The Vulnerabilities module tells you what to patch first across all your clients. It puts at the top the outdated applications whose product is in CISA's catalog of Known Exploited Vulnerabilities (KEV), the vulnerabilities that attackers are exploiting right now, and then lists the other outdated applications, the pending Windows updates and the computers without a working antivirus/EDR or firewall. From that it gives every computer and every client a level, from Critical to OK, with the reasons behind it, opens an alert for each computer that has exploited software to update, and lets you start the updates through winget. Nothing new runs on the computers: the module works from the software inventory, the Windows Update list and the security snapshot the agents already send, and works the levels out again each time you open or refresh it. Open it from Automation & updates › Vulnerabilities in the sidebar. Every technician can use it and can update an application computer by computer; only administrators can update an application on many computers at once.

The CISA catalog and when an application counts

CISA, the US cybersecurity agency, publishes a list of the vulnerabilities (CVEs) that are known to be

exploited in real attacks, with the date each one was added, the action CISA requires and the date by which US federal agencies must have applied it, and whether ransomware campaigns are known to use it. SLApp Helpdesk keeps one copy of that catalog for all its users and refreshes it from CISA once a day, or every hour while it has never been downloaded or the last attempt failed. A download that fails, or that looks incomplete, never replaces the copy already in use. The line under the introduction of the page tells you which copy is in use, for example CISA KEV catalog 2026.10.29 · 1,812 exploited vulnerabilities · checked 30/10/2026, 07:12. If the last download failed it adds that the previous copy is still in use, and before the first download it says The CISA KEV catalog has not been downloaded yet — exploited vulnerabilities are matched as soon as it is (the server tries every hour). Until then every outdated application counts as an ordinary outdated one.

- Only an outdated application can count. An application is outdated when winget reports a newer version for it, the Available version of the computer's Software tab. Having a product that appears in the catalog installed is never enough: an application that is up to date, or whose installed version is unknown, is not flagged.
- Its product must have entries in the catalog. SLApp Helpdesk knows the winget ids of the common applications that CISA lists under names that cannot be guessed, for example Google Chrome and the other Chromium browsers (Edge, Brave, Opera, Vivaldi), Microsoft Edge WebView2, Mozilla Firefox and Thunderbird, Adobe Acrobat and Reader, Oracle Java, 7-Zip, WinRAR, VMware Tools and Workstation, Citrix Workspace, Cisco Secure Client, FortiClient, Zoom, Microsoft Teams, OneDrive and Office, and Apple iTunes and iCloud. Any other application is matched by its name or winget id, and only when every word of the vendor and of the product in the catalog appears there. The tooltip of such a match says Matched by the application's name — check that the CVEs are about it.
- The catalog does not say which versions of a product are affected. An outdated application of a product in the catalog is therefore the first thing to update, not a proof that the version installed is vulnerable.
- Only applications of the software inventory are compared with the catalog. The fixes of Windows itself count through the pending Windows updates.

Levels and their reasons

Every computer gets one level, the most serious of its reasons. Hover the level badge to read what the level means and, below it, each reason of that computer, for example 2 applications with actively exploited vulnerabilities: Google Chrome, 7-Zip or 3 Windows updates pending, the oldest for 45 days.

- Critical: an application with actively exploited vulnerabilities (CISA KEV) needs updating, or no antivirus/EDR is enabled. The antivirus reason only counts once the agent has reported the computer's security; a computer that never reported it shows Not reported in the Antivirus/EDR column and is not raised by it.
- High: Windows updates have been pending for more than 30 days, or the Windows Firewall is off in at least one profile. The age of an update is read from the year and month at the start of its title, as in 2026-10 Cumulative Update for Windows 11, and counted from the second Tuesday of that month, the day Microsoft releases them. Updates without a date in their title, such as Defender definitions and drivers, count as pending but never as old.
- Medium: other outdated applications, or pending Windows updates that are 30 days old or less, or have no date in their title.
- OK: nothing outdated, no pending Windows updates and the protection on.
- No data: the agent has not reported software, updates or protection yet.

A client's level is the worst level of its computers, leaving out the ones with No data. At the top of the

page one button per level counts the computers of each level, for the client chosen in the client list or for all of them. Click a button to show those computers in the By computer view, and click it again to remove that filter.

The top of the page

- Update (administrators, By application view only): updates the applications you ticked on all the computers that still need them (see Updating on many computers). With applications ticked it reads Update on all N computers.
- Excel: downloads the view on screen as an Excel file (see Exporting to Excel).
- Refresh: works the lists out again at once. The page also refreshes itself every 5 minutes.
- Columns: opens the List columns window, where you choose the columns of the current view and drag them into the order you prefer. Each view has its own choice, saved for your user.
- The three tabs By application, By computer and By client switch the view. The view you chose is remembered while the browser tab stays open.
- The client list (All clients, No client or one client that has computers) narrows the By application and By computer views; the By client view always shows every client. The level list (All levels or one level) narrows the By computer and By client views.
- Exploited (KEV) only keeps the applications with exploited vulnerabilities, the computers that have such an application and the clients that have such a computer.
- The search box looks in the application names and winget ids (By application), in the computer names, their clients and every one of their outdated applications (By computer) and in the client names (By client). When the column filters hide rows, Showing N of M says how many are left.

In every view a click on a column header sorts by that column and a second click reverses the order; until you click one, the most urgent rows come first. The row under the headers holds a filter for some columns, and a right-click on a row opens its menu. When a view has nothing to show it says so, for example No outdated application in this view., or Nothing matches these filters. when the filters hide everything.

By application

By application has one row for every outdated application, grouped by its winget id, with the applications that have exploited vulnerabilities first: those used in ransomware campaigns, then those with the most entries added in the last 90 days, then those with the most CVEs, and after them the others, by the number of computers to update. The columns are Application (the name, with the winget id in grey), Exploited (KEV), Newest KEV entry, Ransomware, Computers, Installed versions (up to five, newest first) and Latest version; winget id, New in 90 days, Online, Clients and CVEs are available through Columns. The Exploited (KEV) column shows a red badge such as KEV · 12 CVE; hover it to read the date of the newest entry in the catalog, how many were added in the last 90 days, whether ransomware campaigns use them, how the application was matched and the first CVE numbers. Exploited (KEV) can be filtered on With KEV or Without KEV, and Ransomware on Yes or No. Click a row to open the application window. Its right-click menu offers Open details, Copy application name and Copy winget id and, for administrators, Select or Deselect and Update on all its computers.

By computer

By computer has one row for every active computer, the most exposed first. The columns are Computer, Client, Level, Exploited apps (how many, with the first two names; hover for each one with its versions and its number of CVEs), Other outdated (hover for the names), Pending updates (how many, with the age of the oldest, in orange beyond 30 days), Antivirus/EDR (Protected, None enabled or Not reported), Firewall (On, Off or Unknown) and Status (Online or Offline); Why, Oldest pending, Folder

and Operating system are available through Columns. Level, Antivirus/EDR, Firewall and Status have a list filter, and Computer, Client and Exploited apps a text filter. Click a row to open the computer window. The right-click menu offers Open details, Open the computer, which opens the computer's page, Copy computer name and Copy client name.

By client

By client has one row for every client that has computers, plus one for the computers without a client, the most exposed first. The columns are Client, Level (hover for how many computers of each level), Computers, one column per level (Critical, High, Medium and OK, with No data through Columns), With exploited apps (the number of computers that have one), Pending updates and No antivirus; Exploited apps, Other outdated apps and Firewall off are available through Columns. Click a row to see that client's computers in the By computer view. The right-click menu offers Show its computers, Show its applications, which opens By application for that client, and Copy client name.

The computer window

The computer window shows everything behind the level of one computer. The title is the computer's name, a link that opens its page, with the level badge; under it come the client, Online or Offline, the operating system and Last software inventory with its date (or No software inventory yet), then the list of reasons. Applications with actively exploited vulnerabilities (CISA KEV) lists each application with its installed and available versions, its KEV and Ransomware badges, Show the CVEs and an Update button. Show the CVEs opens a table with each CVE, a link to its page in the NVD (nvd.nist.gov), the name of the vulnerability (hover it for CISA's description and required action), the date it was Added, the CISA due date and whether ransomware uses it; Hide the CVEs closes it. Other outdated applications lists the rest, each with its own Update button. Pending Windows updates lists every pending update with its KB number, which links to Microsoft's article, its title and its month, after a line such as The oldest is from August 2026 (75 days). Protection shows the Antivirus/EDR state with each product the agent found, ticked when it is enabled, and the Firewall state with the text the agent reported. Close closes the window.

The application window

The application window shows one application on every computer that still needs its update, within the client chosen in the client list. The top shows the winget id, Latest version, the KEV and Ransomware badges and, for an application with exploited vulnerabilities, Show the CVEs with the same table as the computer window. The table N computers to update gives each computer, a link to its page, with its Client, its Version (installed › available), its Status, its Level and an Update button. For administrators the top also has Update on all N computers, which updates the application on every computer of the table (see Updating on many computers). When every computer has been updated meanwhile, the window says No computer in this view still needs this update.

Updating one application

The Update button of the computer window and of the application window, and Update in the right-click menu of an application there, ask the agent of that computer to update that application through winget. It is the same command as the Update button of the computer's Software tab, open to every technician, and it shows in the Command queue as Update software. The button then reads Processing… and a notice says Update of Google Chrome started — it runs in the background. If the computer is offline, the notice says that the update is queued and runs when the computer connects. winget works in the background, so the application leaves the lists only after the agent has sent its next software inventory; the computer's Software tab has Force inventory now to send one at once.

Updating on many computers

1. As an administrator, open the By application view. To work on one client, choose it in the client list first: the update reaches only the computers of the client chosen there, or of every client when All clients is chosen.
2. Tick the applications to update in the first column. The tick box in the header ticks or unticks every application listed; you can also use Select in the right-click menu of a row.
3. Click Update on all N computers at the top of the page, where N is the number of different computers that need at least one of the ticked applications. For a single application you can also use Update on all its computers in its right-click menu, or the button of the same name in its application window.
4. The Update applications window names the applications and asks Queue the update of … on N computers? Offline computers run it when they come back online. Click Confirm; the button reads Queuing… while the updates are queued. Cancel, which has the focus when the window opens, closes it without doing anything.
5. Notices tell you how it went: Update queued on N computers, then, when it applies, N of them are offline — they run the update when they come back online and N skipped — an update of the same application is already queued or running there. If nothing needed to be queued, the notice is Nothing to queue — no computer in this view still needs these updates. The list is read again a few seconds later.

- Only computers whose inventory still shows a newer version receive the command, each with the winget id exactly as that computer reported it.
- At most 50 applications and 2,000 updates (one per application and computer) can be queued at once. Over either limit the button is disabled and a red line says why: At most 50 applications at once — 62 are selected., or Too many updates at once: at most 2,000 (one per application and computer). Filter by client or select fewer applications.
- A computer that already has an update of the same application waiting, sent or running, queued in the last 12 hours, is skipped, so a second click never runs winget twice.
- An offline computer keeps its update in the Command queue and runs it when its agent reconnects.
- Each application queued is written to the Audit log under the entity type Vulnerability, with the number of computers, how many were offline, the client when one was chosen and how many already had it queued.
- Technicians who are not administrators see neither the tick boxes nor these buttons; the By application view tells them Only administrators can update on many computers at once — open an application to update it computer by computer.

Exporting to Excel

Excel downloads the current view as the grid shows it: the same client, level, KEV and search filters, the column filters, the rows in the order on screen and only the visible columns, in their order. The file is written in the language of the app and named after the view and the day, for example vulnerabilities-computers-2026-10-30.xlsx, and a notice says Excel file created. Every technician can export; the button is disabled while the view has no rows. When more than 10,000 rows are on screen, the file holds the whole view in its default order, without the column filters.

Alerts and emails

The module raises one alert per computer while that computer has at least one outdated application with exploited vulnerabilities. It does not come from a rule, so there is nothing to set up in the Rules tab of Alerts. It is Critical, its type is Exploited vulnerability (KEV), shown with a bug icon, its source is Vulnerability, and its message lists the applications with their versions, for example Outdated software with actively exploited vulnerabilities (CISA KEV): Google Chrome 129.0 › 130.0, 7-Zip 23.01 › 24.08,

with +N more when they do not all fit. It is linked to the computer in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, and you can create a ticket from it like from any other alert. An administrator can stop these alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section); the levels and lists of the module are not affected.

- When it is checked: every time the computer's agent sends its software inventory, and once a day for every computer, right after the server's daily check of the CISA catalog, so that a CVE added today flags software that nobody has touched. Nothing is checked while no catalog has been downloaded.
- When it closes: on its own, at the next check after the last exploited application was updated or removed. While it stays open, its message follows the applications that are still outdated. An inventory that does not say which applications have a newer version changes nothing: when the agent could not read winget this time and sent the list of installed programs from the Windows registry, or could not read winget's list of updates, the alert stays exactly as it is, at that inventory and at the daily check, until an inventory that says it arrives.
- Who is emailed: the active administrators of your company who have an email address, when the alert opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the names of the applications and the vulnerabilities stay as they were written).
- No email is sent for the alerts opened the first time the daily check runs for your company, nor for a computer that the daily check has not yet seen once, such as one whose agent was installed today. Those alerts open without an email. The same goes for a computer whose software inventory has only just started to say which applications have a newer version, for example a Windows in Portuguese or French once its agent reads winget in that language: that first inventory is its starting point, and the computer counts as not yet seen until the next complete daily check.
- An alert that opens again for applications that were all in alerts of that computer closed during the last 30 days opens without an email, so a browser that falls one version behind after each release does not email you every time. A different exploited application emails as usual.
- When the daily check opens several alerts at once, for example after CISA adds a vulnerability of an application installed on many computers, each administrator receives one email, Vulnerability alerts: N new, that lists the computers, up to 50, instead of one email per computer. The alerts that open without an email, above, are not listed in it.
- Only an administrator can close the alert by hand, with Resolve; it is written to the Audit log of that computer. If the applications are still outdated, the next check opens it again, without an email when they are the same applications.

> **i** The same information reaches two other places. In the Software tab of a computer, an outdated application with exploited vulnerabilities carries the same red KEV · N CVE badge after its name. And the client report has a Vulnerabilities section (PDF and Excel) with the level of each computer, its reasons and its exploited applications with their CVEs; the full report leaves it out for a client that has no computers.

16. Software packs (administrators only)

The Software packs module keeps software installed on your clients' Windows computers. You build a pack — a named list of packages taken from the winget catalog (winget is the Windows package manager) — and assign it to a client's folder or to a single computer. SLApp Helpdesk then installs whatever is missing on every Windows computer the pack covers, and keeps doing so for computers that join the folder later. Open the module from Automation & updates › Software packs in the sidebar. Only administrators see it. It has two tabs: Packs, where you create and assign packs and follow their

progress, and Catalog, where you browse the winget packages you can add.

The Catalog tab

The Catalog tab is a searchable copy, kept on the server, of the official winget package index. Type in Search to look by name, ID, publisher or tag; every word you type has to match, so a second word narrows the result. The table shows ID, Name, Version (the latest one), Publisher and Tags, fifty results at a time — click Load more at the bottom for the next ones. Click a row to open its details: Latest version, Publisher, Short name, Homepage, Tags and Description. Publisher, homepage and description are read from the public winget-pkgs manifest the first time you open a package, so for some packages the homepage and description stay empty (the publisher then falls back to the part of the winget ID before the first dot). From the Actions menu at the end of a row, or from the right-click menu, you can:

- Open Details, which shows the window described above, with the winget ID and a small Copy button next to it.
- Choose Add to pack… to open the Add to pack dialog. Pick an existing pack in Pack (inactive packs are marked), or pick + New pack… and type a New pack name, then click Add. A message confirms it and offers an Open pack link; a package that is already in the chosen pack is refused with a message.
- Copy the winget ID or the name, in the right-click menu only.

Above the table, a summary line shows how many packages the catalog holds, its index version, when it was last refreshed and when it was last checked. SLApp Helpdesk checks Microsoft's index by itself, about a minute after the server starts and then roughly every eight hours (and again after about half an hour when a check fails), and only downloads it again when it has changed. Refresh catalog does it on demand: the button shows Refreshing… and the list reloads by itself when the refresh ends, usually within a minute. A full re-download is done at most once every five minutes; pressing the button more often only checks whether a newer index has been published. If the last refresh failed, a line starting with Last refresh error gives the reason and the previous catalog stays in use. On a brand-new installation the tab shows The catalog is empty. Click “Refresh catalog” to download the winget index.

The Packs tab

The Packs tab lists every pack of your company with the columns Name, Description, Packages, Folders, Computers, Install state, Active and Updated; click a header to sort by it. Folders and Computers count the assignments. Install state shows how many package installs are Installed (green) and how many Failed (red) across the covered computers. Search and the Status filter (All, Active, Inactive) narrow the list, and the Active checkbox switches a pack on or off. Click a row to open the pack. The Actions menu at the end of a row, or a right-click, offers Edit, Assign…, View status, Install now / retry, Activate pack or Deactivate pack, and Delete. To create a pack and assign it:

1. Click + New pack. The New software pack dialog opens.
2. Type a Name (it must not be used by another pack of your company) and, if you like, a Description, then click Create. The pack opens in its editor.
3. In the Packages card, type at least two characters of a name, ID, publisher or tag in the search box and click Add next to each package you want; Details shows more about a package first. Each change is saved as soon as you make it. You can also add a package from the Catalog tab with Add to pack….
4. In the Assignments card, click Assign to folders / computers…. The Assign dialog opens.
5. Choose a Client, then tick the folders and the computers you want. Choose another client and tick more if you need to: the selections add up and appear under Selected. The dialog also tells you

how many Windows computers the selection covers and how many of them are new.

6. Click Save. If the change adds a folder or a computer, the Assign pack? question tells you how many more Windows computers will start receiving the software (or that none do yet); click Confirm to go ahead.
7. Open the Status tab of the pack to follow the installations, which start shortly.

The pack editor has two views, Configuration and Status. In Configuration, the General card holds the Name, the Description and the Active checkbox; Save and Discard changes appear once you change something. A new pack is Active. An inactive pack is not applied at all: nothing is installed or retried for it, and software it already installed is left alone. In the Packages card the table lists each package with its winget ID, and Remove takes it out of the pack. A pack holds up to 200 packages. There are no version settings: a pack installs the version winget offers, and a computer that already has the package is simply marked Installed — the pack never updates software that is already there.

Assigning means keeping installed

Assigning a pack is not a one-off installation. It tells SLApp Helpdesk to keep the pack's software installed: every Windows computer the pack covers gets whatever is missing, now and in the future. The server looks again every minute, so there is nothing to run by hand.

- Assigning a folder covers every Windows computer in it, now and whenever another one joins it later. Folders are created in the client's Computers tab.
- Assigning a computer covers just that computer, whichever folder it is in.
- A package you add later to an assigned pack is installed on the covered computers shortly.
- Each computer receives one package at a time, because Windows Installer runs one installation at once. The next package is sent within about a minute of the previous one finishing, even when the packages belong to different packs.
- Other systems are skipped silently: in the Assign dialog they are dimmed and marked not Windows.
- To change what a pack covers, click the cross (Unassign) on a folder or computer chip in the Assignments card, or reopen the Assign dialog, untick it and click Save.

The status view

The Status tab of a pack, or View status in the Actions menu, has one row per computer and package the pack currently covers. The columns are Computer, Client, Folder, Package (with its winget ID underneath), Status, Attempts and Updated. Above the table, one chip per status shows how many rows have it; click a chip to filter on it and click it again to clear. Search and the Status and Client filters narrow the rows, and a click on a header sorts them. When a row has output, click it (or use Show output in the right-click menu) to read the Last output of the command — the place to look when something failed. The view reloads itself every 15 seconds while a row is still Pending, Installing or Uninstalling; Refresh does it at once, and the line next to it shows when the view was last updated. The statuses mean:

- Pending — the package is wanted on the computer but has not been sent yet, usually because the computer is still busy with another package or the next one-minute check has not come yet.
- Installing — the install command has been sent to the Command queue and is waiting or running. A computer that is offline keeps this status until it reconnects and runs the command.
- Installed — the install succeeded, or winget already had the package. SLApp does not check an Installed package again later, so software removed by hand is not put back.
- Failed — the install did not succeed; read the Last output. SLApp tries again on its own, up to three attempts in total, waiting about 30 minutes after the first failure and about an hour after the second.

After that the row stays Failed until you click Retry.

- Unsupported — winget is not available on the computer, which is typical of Windows Server. It is not an error: SLApp checks again once a day, and Retry checks right away.
- Uninstalling — an uninstall you asked for is queued or running. The row disappears when it succeeds and turns Failed if it does not; SLApp then retries the uninstall on its own on the same schedule as an install, and Retry tries it again at once.

Retry, at the end of a Failed or Unsupported row (and in the right-click menu), starts a new attempt on that computer straight away and resets its Attempts. Install now / retry — at the top of the Status view, in the editor header and in the Actions menu of the Packs tab — does the same for the whole pack: it sends the next package to every covered computer that is free and gives the Failed and Unsupported rows a fresh start. A message tells you how many installs were queued, or that there is nothing to install because everything is installed or no computer is covered yet.

Removing a package, unassigning and deleting

Nothing is ever uninstalled silently. Three actions can leave a pack's software on computers where nobody manages it any more: removing a package from a pack (Remove), unassigning a folder or a computer (the cross on its chip, or unticking it in the Assign dialog and saving), and deleting the pack (Delete, confirmed first). When one of them would affect software that was installed through the pack, SLApp changes nothing and asks first, in a dialog titled Uninstall the removed software?, Uninstall from computers no longer covered? or Uninstall the software before deleting the pack?. The dialog lists what is affected and on how many computers, and you answer:

- Uninstall too — queues an uninstall of the software on those computers; their rows show Uninstalling until it is done.
- Keep installed — leaves the software as it is; it simply stops being managed by the pack.
- Cancel — nothing changes at all: the package stays in the pack, the assignment stays and the pack is not deleted.

> ℹ The question is always asked and the answer is never assumed. It only concerns software that was installed through the pack and that the pack would no longer keep (a removed package, or the pack's software on computers that would no longer be covered), so you are not asked when nothing is installed or when a computer is still covered by another assignment of the same pack. A package that another pack still wants on the same computer is never uninstalled. Moving a computer out of an assigned folder, or deactivating a pack, never uninstalls anything either. If you delete a pack with Uninstall too, or while uninstalls chosen earlier are still running, the pack leaves the list at once and the uninstalls carry on in the background.

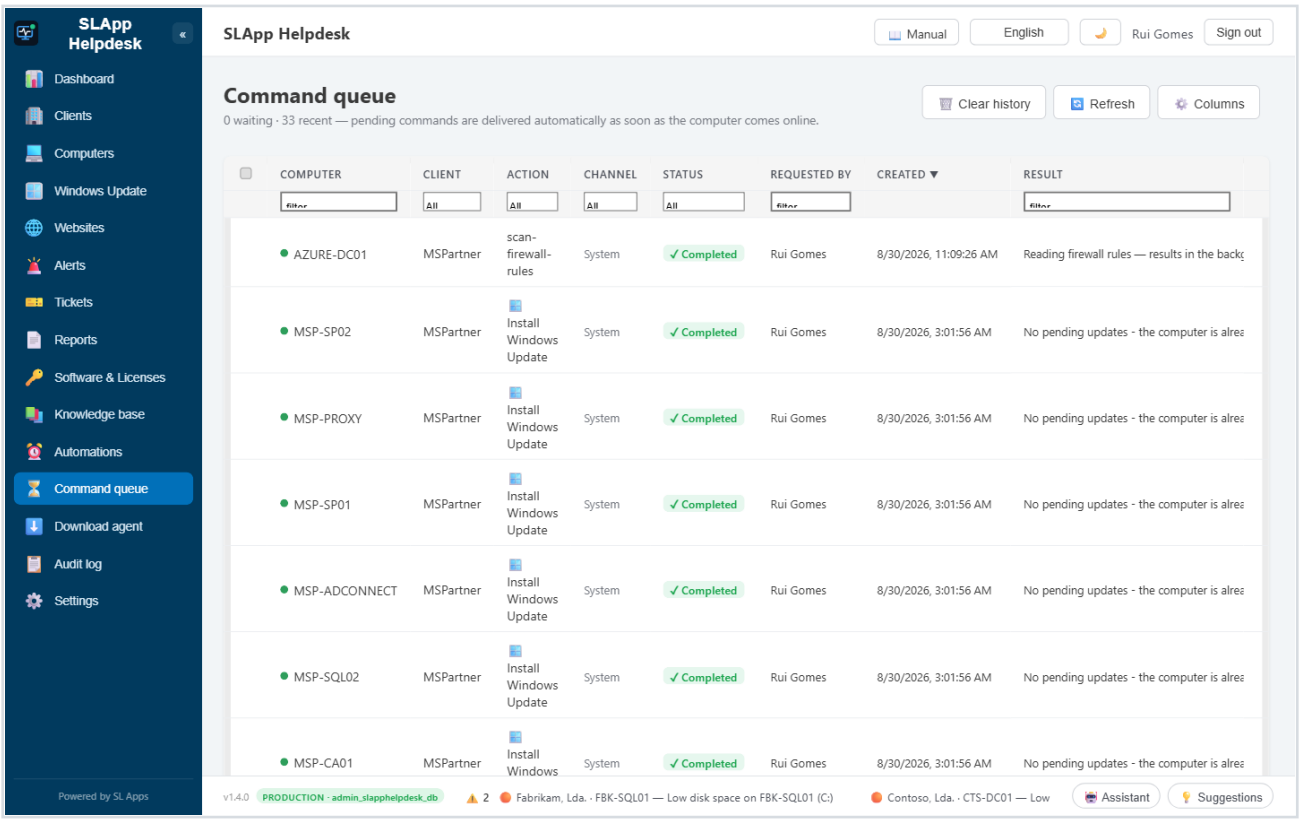
Limits

- Only Windows computers are covered. A Windows computer where winget is not available, typically Windows Server, shows as Unsupported.
- Installs and uninstalls are sent as ordinary elevated commands through the Command queue (shown there as Command/script on the System channel), so you can follow them like any other command. A computer that is offline receives them when it reconnects, and its rows stay Installing until then.
- An install obeys the command timeout. By default a command is cancelled after 300 seconds, so a very long installer can end up Failed. If that happens, raise Agent limit (seconds) in the Command timeout card of Settings › General › Remote control, and keep Server wait for reply (seconds) at or above it. The setting applies to every remote command of your company, not only to packs.

- Creating, changing, assigning, running, retrying and deleting a pack, and refreshing the catalog, are recorded in the Audit log.

17. Command Queue

Every remote action you trigger on a computer (a restart, a script, a software update, a network scan, and so on) goes through the Command queue. The queue is a store-and-forward mechanism: the command is recorded on the server and handed to the computer's agent the next time it checks in. If the computer is offline, the command simply waits and is delivered automatically as soon as the computer comes back online. The Command queue module, available to every user from Automation & updates › Command queue in the sidebar, shows you all of these commands for your whole company in one place, so you can see what is still waiting, what is running right now, and what the outcome of finished commands was.



The Command queue with the summary line, the filter row and per-command status badges.

What you see on the screen

Under the title, a summary line tells you how many commands are waiting, how many are running (only shown when there is at least one) and how many recent commands are listed, followed by the reminder that pending commands are delivered automatically as soon as the computer comes online. The list refreshes itself every 8 seconds, and you can force an update at any moment with the Refresh button. The list holds at most 300 commands: the ones that are Waiting or Sent are always included first, and the remaining places go to the most recent finished commands, so older history is not shown here.

Each row in the table has the following columns:

- A checkbox at the start of the row, shown only for commands that are still Waiting; it lets you cancel several at once.
- Computer: the computer's name, preceded by a small dot that is green when the computer is online (its agent reported in within the last 2 minutes) and red when it is offline; hover over the dot to read Online or Offline. The name is a link that opens the computer's page.

- Client: the client the computer belongs to (a dash if it has none).
- Action: what the command does, for example Restart, Shut down, Command/script, Sync, Update agent, Update software, Remove software, Install Windows Update, Remove Windows Update, Network scan, Immediate sample or Printer scan. For Command/script the first 30 characters of the command are shown after the label, followed by an ellipsis when the command is longer.
- Channel: System when the command runs with elevated rights through the agent service, or User when it runs in the logged-in user's session.
- Status: the current stage of the command (see the next section).
- Requested by: the technician who triggered the command. Automatic agent updates that nobody requested show auto; any other command without a requester shows a dash.
- Created: the date and time the command was queued, in your local time.
- Result: the output or result message returned by the agent once the command finishes (a dash while there is none yet). Long results are cut off in the table; hover over the cell to read the full text in a tooltip.
- A Cancel button at the end of the row, available only while the command is Waiting.

Commands sent by the Backups module appear in this list too, under their own names and without an icon: backup-sync (sent by Run now) and backup-cancel (sent by Cancel run), both only when the computer is connected at that moment (an offline computer finds the request by itself when its agent reconnects), and backup-test-destination (sent by Test on a destination). They use the System channel and carry only identifiers, never passwords.

The configuration backups of the Network configs module appear as Network configuration backup, on the System channel of the device's collector computer. Their text is only the identifier of the backup run, never a password: the device and the outcome are shown in that module. Likewise, a scan started with Find printers in Printers appears as Printer scan, on the System channel of the computer that scans: its text is only the identifier of the scan, never the SNMP community, and it is Running while the agent scans; the printers found are shown in the Find printers window.

The daily check of the Active Directory and Third-party backups modules goes through this queue too: once a day every Windows computer gets a Command/script command on the System channel whose text starts with # SLApp health check (see The daily check in the Active Directory section). So that one command per computer and per day does not crowd out your own commands, the list shows these checks only while they are Waiting, Sent or Running; once they have finished they are hidden, and they do not take any of the 300 places, unless you tick Show daily checks in the header, next to Clear history. Hover the box for a reminder. Wherever they are listed, a check shows only the first line of its script and the first 500 characters of its result, followed by […]. Requested by is a dash for the scheduled checks and the technician's name for a Check now. The Command queue for this computer card on a computer's Remote actions tab always lists that computer's checks, finished ones included. Since version 1.14.77 the commands of the Hyper-V module are shown the same way (the first line of its script, the line naming the operation and the first 500 characters of the result), and the inventories it asks for on its own (every half hour on each online Hyper-V host, and after each change) count as daily checks: hidden here once finished unless Show daily checks is ticked, and deleted 3 days later.

Click any column header to sort by that column; click again to reverse the order. By default the list is sorted by Created with the most recent commands first. Sorting by Status orders the rows by stage (Waiting, Sent, Running, Completed, Failed, Cancelled) rather than alphabetically. The second header row holds a filter for every column except Created: Client, Action, Channel and Status offer a dropdown of the values present in the list (with All to remove the filter), while Computer, Requested by and Result accept free text and match any row containing what you type. If nothing matches, the table shows No command matches the filters. When there are no commands at all, the table is replaced by the

message No commands in the queue.

> **i** Finished daily checks are deleted automatically 3 days after they finish, and Clear history removes them as well, shown or not. A check that a computer has not picked up within 6 hours, because it is offline, is cancelled with the result Health check not delivered within 6 hours (computer offline) - it is queued again when the computer is back. Cancelling a waiting check yourself is harmless: the computer simply runs its next check at the usual time.

Command statuses

- Waiting: the command is stored on the server and has not yet been picked up. This is the normal state for a command aimed at an offline computer; it will be delivered when the computer reconnects.
- Sent: the agent has received the command but has not reported a result yet.
- Running: the agent has started executing the command and it is still in progress (typical for longer tasks such as installing updates).
- Completed: the command finished successfully; check the Result column for its output.
- Failed: the command finished with an error; the Result column usually explains why.
- Cancelled: a technician cancelled the command before it was delivered.

> **i** Only commands in the Waiting state can be cancelled. Once a command has been Sent to the agent it is already on the computer and will run; cancelling it from the queue is no longer possible.

Cancelling commands that are still waiting

1. Open Automation & updates › Command queue from the sidebar.
2. Optionally narrow the list with the filters, for example by choosing Waiting in the Status dropdown or typing part of a computer name in the Computer filter.
3. To cancel a single command, click Cancel at the end of its row.
4. To cancel several, tick the checkboxes of the commands you want to drop, or tick the checkbox in the header to select all waiting commands currently visible (hidden rows are not selected).
5. Click the red Cancel selected (n) button that appears in the header as soon as at least one waiting command is ticked; the number shows how many commands will be cancelled.
6. The rows switch to Cancelled and the selection is cleared.

> **i** If a computer comes online and picks up a command between the moment you click Cancel and the moment the server processes it, the cancellation is ignored and the row simply shows its new status after the next refresh.

Clearing the history

Finished commands stay in the list as a history so you can review their results. When the list becomes cluttered you can remove them with the Clear history button. The button is disabled while there is nothing finished to delete. Clearing only removes commands that are Completed, Failed or Cancelled; commands that are Waiting, Sent or Running are never touched, because they are still in flight.

1. Click Clear history in the header.
2. Read the Clear queue history dialog. It confirms that only the finished commands (completed, failed and cancelled) will be deleted and tells you how many waiting commands will be kept. This action cannot be undone.

3. Click Clear history in the dialog to confirm, or Cancel (or click outside the dialog) to keep everything. The button reads Clearing… while the deletion runs, and the dialog cannot be dismissed until it finishes.
4. The list reloads without the finished commands.

> **i** Clear history removes the finished commands of every computer in your company, not just the ones currently visible after filtering, and this module is the only place where history can be cleared: the Command queue for this computer card on a computer's Remote actions tab has no Clear history button. If you only want to find one computer's commands, use the Computer filter instead of clearing.

Seeing the queue of one computer

The same information is also available per computer. Open a computer from the Computers module and go to its Remote actions tab: when that computer has commands in the queue, a Command queue for this computer card lists them with the same Action, Channel, Status, Created and Result columns, the same status badges and the same Cancel button for waiting commands. The card has its own Refresh button and also refreshes itself every 8 seconds. It includes a reminder that, if the computer is offline, the commands stay waiting and are delivered on their own as soon as it comes back online. Two differences from the global module: in the card, a result is a clickable link that opens a View full message popup instead of a hover tooltip, and there are no checkboxes, filters, bulk cancel or Clear history. Use the global Command queue module when you need the overview across all clients and computers, and the computer's Remote actions tab when you are working on one machine.

> **i** Because commands wait patiently for offline machines, you can queue work ahead of time, for example a script for a laptop that is currently out of the office, and simply check the Result column later. Keep an eye on the waiting count in the summary line: a command that stays Waiting for a long time usually means the computer has not come online, which you can confirm from the red dot next to its name.

18. Infrastructure overview

The Infrastructure part of the sidebar gathers the modules that watch your clients' servers and network: Hyper-V, Active Directory, Exchange, SQL Server, Storage / iSCSI, Failover Cluster, DHCP, IIS, Backups, Third-party backups, Syslog, Network configs, Printers and Lifecycle. Most of the server modules are fed by the same read-only daily check described in the Active Directory section, and each of them has its own monitoring setting in the health block of the computer's page (automatic detection, always monitor or never monitor). The sections that follow describe each module; this one describes the page that sums them all up.

Infrastructure › Overview is the first entry of the Infrastructure menu, and clicking the Infrastructure group header in the sidebar opens it. It shows the state of every Infrastructure module at a glance and only reads what SLApp Helpdesk already has: opening or refreshing it never asks the agents anything and never starts a check.

Totals and cards

At the top, four totals: Critical alerts and Warnings (the open alerts of all the modules together), Modules OK (shown as OK out of the modules on the page) and Modules with problems (a module counts as having problems when it has an open alert, a server in warning or error, or could not be read). Below them there is one card per module, in the order of the menu.

- The module's name and how many servers, hosts, controllers, clusters, jobs, sources, devices or computers it watches.

- A health bar split into OK, Warning, Error and Unknown, with the count of each underneath.
- The open alerts of the module (for example 2 critical, 1 warning) or No open alerts.
- Two figures that matter for that module, such as Domains and DCs failing for Active Directory, Stopped pools for IIS, Fullest scope (%) for DHCP or Overdue for Third-party backups.
- Last check with the date and time of the most recent result, or Never checked.
- Up to three top problems, the open alerts first, each with its client; +N more says how many others there are. Clicking a problem, or anywhere on the card, opens the module.
- A module where nothing has been detected yet shows a small card with Nothing detected yet. If a module's data cannot be read, its card says Could not read this module, while the rest of the page keeps working.

Most urgent open alerts

Below the cards, Most urgent open alerts lists the 10 most urgent open infrastructure alerts, critical ones first and then the newest, with Severity, Alert, Module, Client, Computer and Since. Clicking a row or its module opens that module, the computer name opens the computer, and See all alerts opens the Alerts page. When nothing is open the list says No open infrastructure alerts.

Plan, permissions and refresh

- Modules your plan does not include appear dimmed with Available in the Complete plan and Change plan in Settings › License; clicking one only shows that note, and their data is never read.
- The Backups card is for administrators: other technicians do not see it, and its alerts are left out of their totals and of the urgent list.
- The page refreshes itself silently every 2 minutes while it is visible, and again when you come back to it after a longer while. The refresh button (Refresh now) refreshes it at once and confirms with Overview updated.

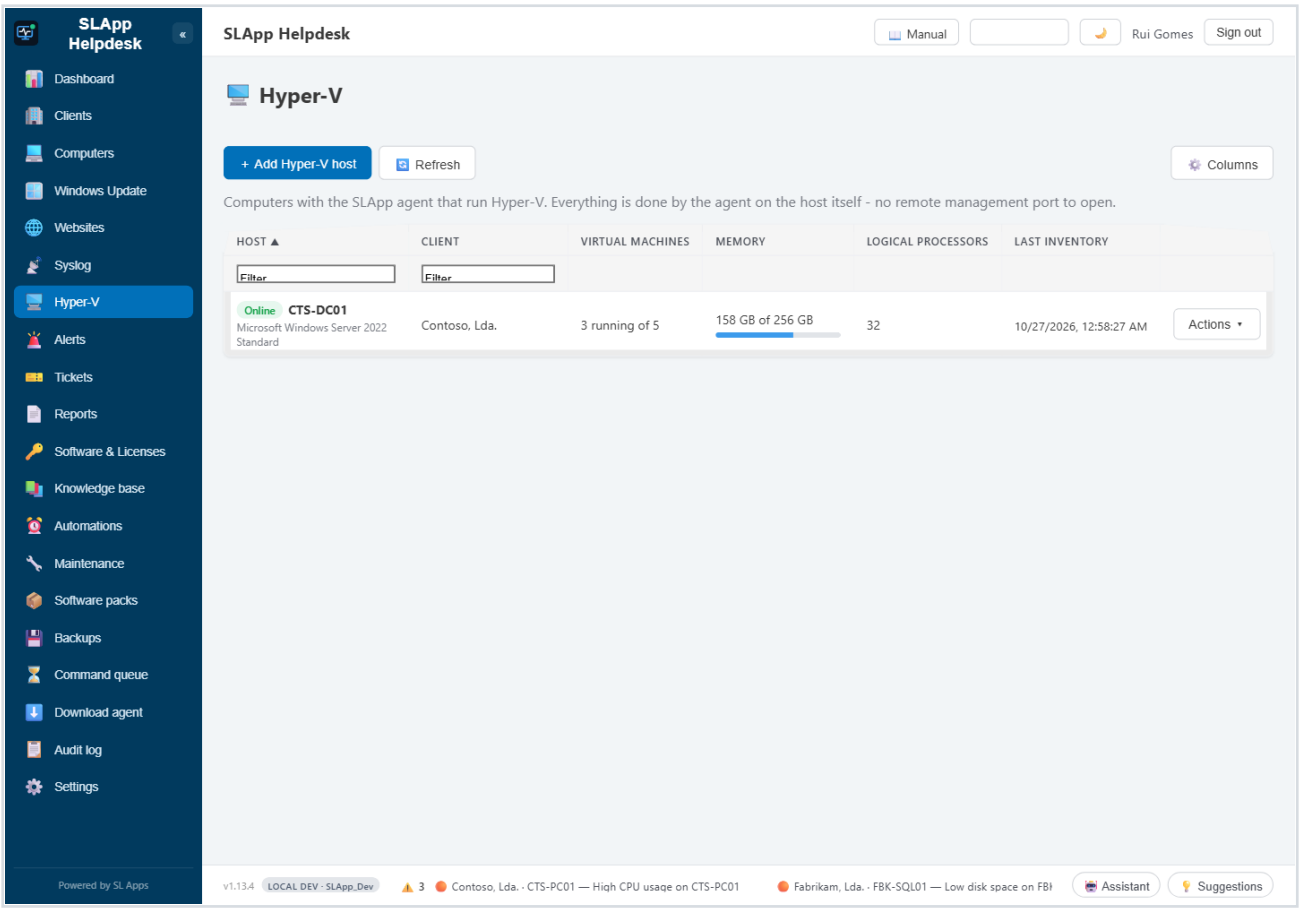
> **i** Start the day here: the totals tell you whether anything needs attention, and the top problems and the urgent list take you straight to the module and server concerned.

19. Hyper-V

The Hyper-V module manages Hyper-V hosts from the browser: their virtual machines, virtual switches and Hyper-V settings, including replication to another host. A host is any Windows computer that already runs the SLApp agent; everything is done by that agent on the host itself, as the system account, through the same command queue as remote actions — no remote management port has to be opened and no host password is stored. Open it from Infrastructure › Hyper-V in the sidebar. Every technician can see the hosts, the virtual machines and the history; only administrators can change anything, and every change is written to the Audit log.

Adding a host

1. Click + Add Hyper-V host (administrators only).
2. Search for the computer by name, client or system and click it. Only Windows computers with the agent that are not hosts yet are offered; servers are listed first.
3. The host is added and asks for its first inventory. While it answers, the list shows First inventory…; a computer that does not run Hyper-V shows No Hyper-V instead (while Generate alerts is on, the module asks it again once a day).



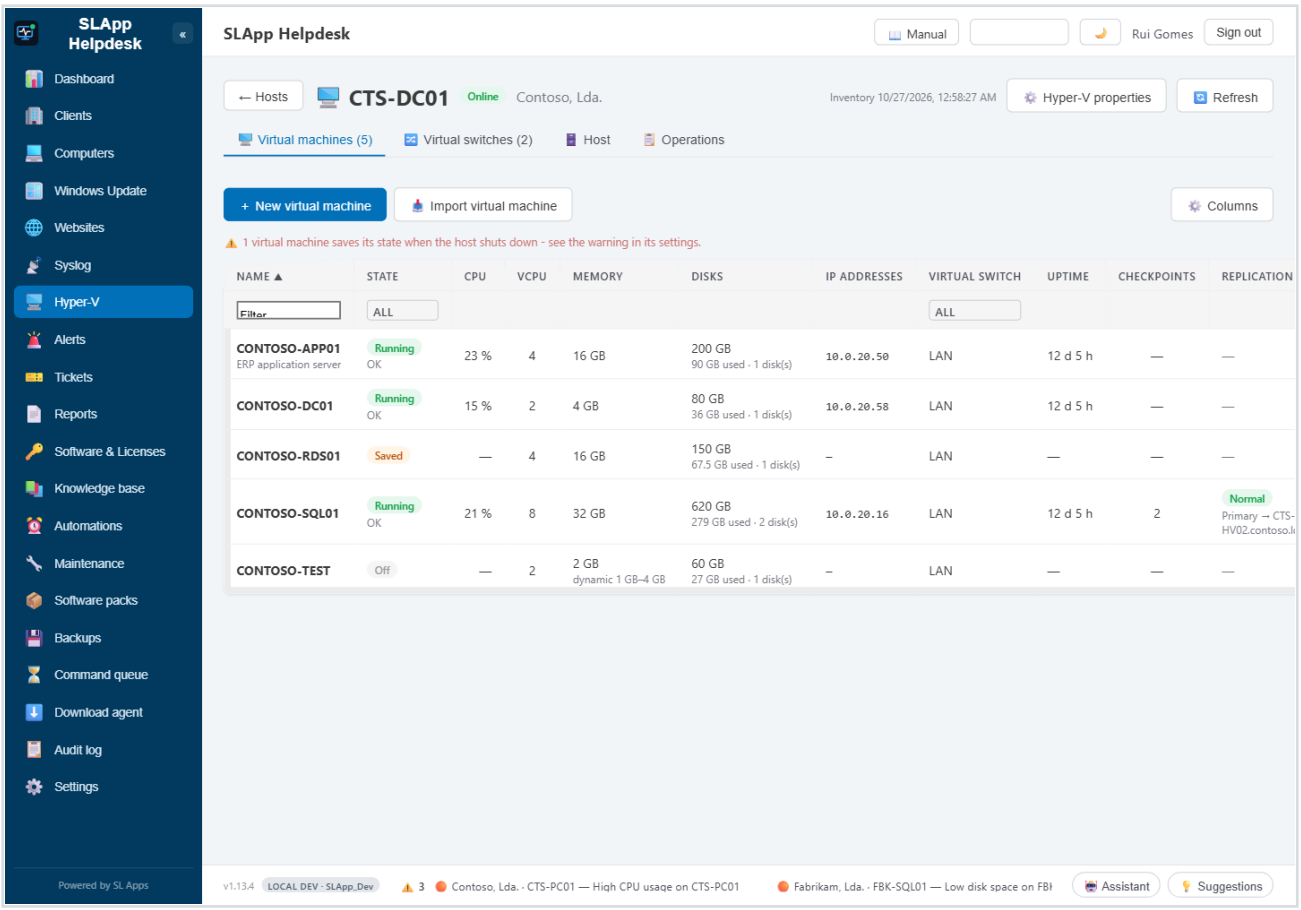
The list of Hyper-V hosts: online state, client, virtual machines running, memory used and the time of the last inventory.

The list shows each host with its online state, operating system, client, the virtual machines running out of the total, the memory in use, the logical processors and the time of the last inventory. The Actions menu has Open, Hyper-V properties (opens the host straight on its Host tab), Refresh (asks for a new inventory) and Remove from the module, which only takes the host and its operation history out of SLApp — nothing changes on the host and its virtual machines keep running. Next to the title, the Generate alerts box turns the replication alerts on or off (see Replication alerts below).

Inventory and operations

What the module shows is the latest inventory the host sent, with its time in the header of the host. Opening a host whose inventory is older than two minutes asks for a new one on its own; Refresh asks for one at any time. On top of that, every online host is asked for a new inventory every half hour without anyone opening it, unless the Generate alerts box is off; once finished, these automatic inventories are not listed in the Operations tab, and they are deleted after 3 days. Every change is an operation: the toast at the bottom right says it was sent, that it is running on the host and how it ended, and a new inventory is asked for as soon as it finishes so the screen shows the new state. Operations still on their way are listed under the tabs, and a virtual machine with one running cannot receive a second one until it ends. Moving, exporting, importing and the initial copy of a replication can take hours: they run in the background on the host, outside the agent's time limit, and the toast only says they started — follow them in the Operations tab.

Virtual machines



A host's Virtual machines tab: state, CPU, virtual processors, memory, disks, IP addresses, switch, uptime, checkpoints and replication health of each VM.

The Virtual machines tab lists every VM of the host with its state, CPU use, virtual processors, memory (assigned while running, startup memory otherwise, and the dynamic range when dynamic memory is on), disks (total size and space used), IP addresses, virtual switch, uptime, checkpoints and replication health. Filter by name, IP or notes, by state and by switch; Columns chooses and orders the columns. A warning above the list counts the VMs that save their state when the host shuts down.

The Actions menu of each VM, and the right-click menu, offer the power actions that fit its state — Start, Shut down (asks the guest to shut down), Restart, Save state, Pause, Resume and Turn off (like pulling the plug) — plus Properties, Create checkpoint, Move to another host, Export and Delete virtual machine. Shut down, Restart and Turn off ask for confirmation first.

> **i** When the host shuts down, a VM set to Save state keeps its memory on disk and restores it on the next start. After a processor or firmware update of the host that saved state may not restore and the VM has to start from scratch — for servers, prefer Shut down the guest (General tab, When the host shuts down).

Creating a virtual machine

1. Click + New virtual machine.
2. Under Basics give it a Name, choose the Generation (2, UEFI, for current systems), the virtual processors and the startup memory, with dynamic memory and its minimum and maximum if you want it. The form shows how many logical processors and how much free memory the host has.
3. Under Disk choose New disk (size and type: dynamically expanding or fixed size), Use an existing file, or No disk. Left empty, a new disk is created as <name>.vhdx in the host's default folder for virtual hard disks.
4. Under Network and installation choose the virtual switch, a VLAN if needed and, optionally, an ISO

image on the host — a generation 2 VM boots from it first. Choose the Secure Boot template (Windows, Linux or off).

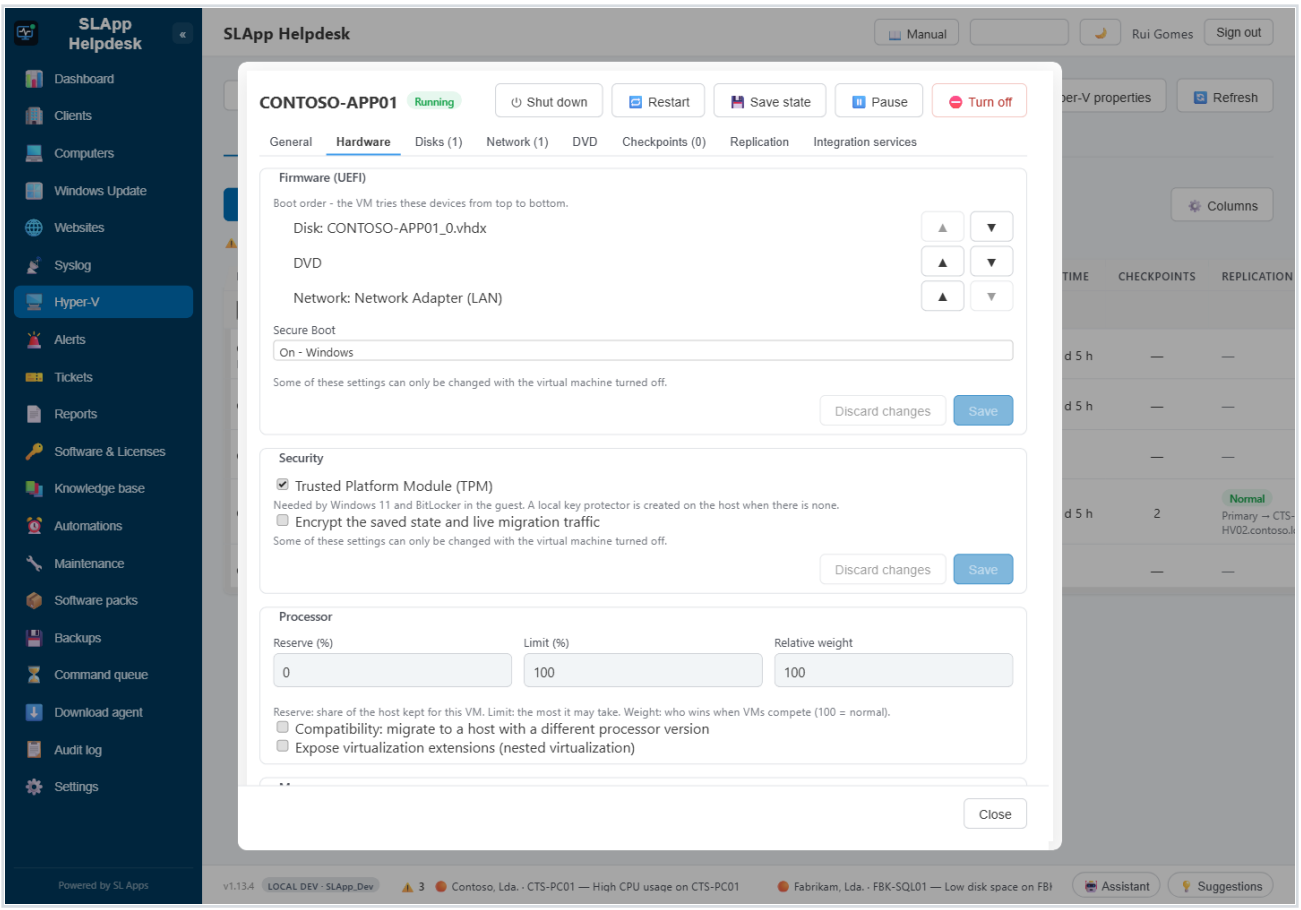
5. Under More options choose the folder for the VM configuration, what to do when the host starts and when it shuts down (Shut down the guest is the default), add notes, and tick Start it when created if you want.
6. Click Create.

Every path field has a Browse… button that lists the folders of the host itself (not of the computer you are using), so a disk, an ISO or a folder can be picked without typing. New folder, next to Go, creates a folder inside the one that is open and opens it, for example a folder for a new virtual machine; it is recorded in the Audit log.

Properties of a virtual machine

Click a VM, or choose Properties in its Actions menu, to open its properties. The power actions that fit its state sit at the top; the tabs below follow the Settings of Hyper-V Manager:

- General — uptime, memory in use, generation and configuration version, heartbeat, location, creation date and ID, and the settings: name, virtual processors, startup memory, dynamic memory with its minimum and maximum, what to do when the host starts (with a start delay) and shuts down, the checkpoint type and notes. The number of processors and turning dynamic memory on or off can only be changed with the VM turned off.
- Hardware — the boot order (move devices with the up/down arrows; for a generation 1 VM, the BIOS order), Secure Boot and its template, the TPM (a local key protector is created on the host when the VM has none) and the encryption of the saved state and migration traffic, the processor reserve, limit and relative weight, compatibility for migration to a host with a different processor and nested virtualization, the memory buffer and weight, the folders of the checkpoint files and of the smart paging file, and automatic checkpoints.
- Disks — the virtual disks with controller, file, size, space used and type. Add a new or existing disk, Grow a disk (only larger; extend the partition inside the guest afterwards; not possible while the VM has checkpoints) or remove one, keeping or deleting its file. A physical or iSCSI disk of the host given whole to the VM (pass-through) is shown as such; removing it only detaches it. For a running VM, the section Inside the guest shows the disks the VM's own Windows sees, read by the SLApp agent installed in it: type (iSCSI disks are flagged), size, free space, drives and iSCSI sessions — including iSCSI disks the guest connects itself, which Hyper-V never shows.
- Network — each adapter with its switch, VLAN, MAC address and IP addresses. Change the switch or the VLAN and Save, add or remove an adapter. The settings button opens the advanced settings: static or dynamic MAC, MAC address spoofing, DHCP guard, router guard, device naming, the VLAN mode (none, access or trunk with its allowed VLANs and native VLAN) and the maximum bandwidth and minimum weight.
- DVD — mount an ISO of the host (a DVD drive is added when the VM has none) or eject it.
- Checkpoints — the checkpoints as a tree, the current one marked. Create one (empty name = name and date), Apply one (the VM returns to that moment), rename or delete it, with or without the ones that came after it.
- Replication — see Replication of a virtual machine below.
- Integration services — what the host offers the guest (operating system shutdown, time synchronization, data exchange, heartbeat, backup through volume shadow copy, guest services), with their state; tick or untick Offered to turn each one on or off.



The Hardware tab of a VM's properties: boot order, Secure Boot, TPM, processor, memory and file locations.

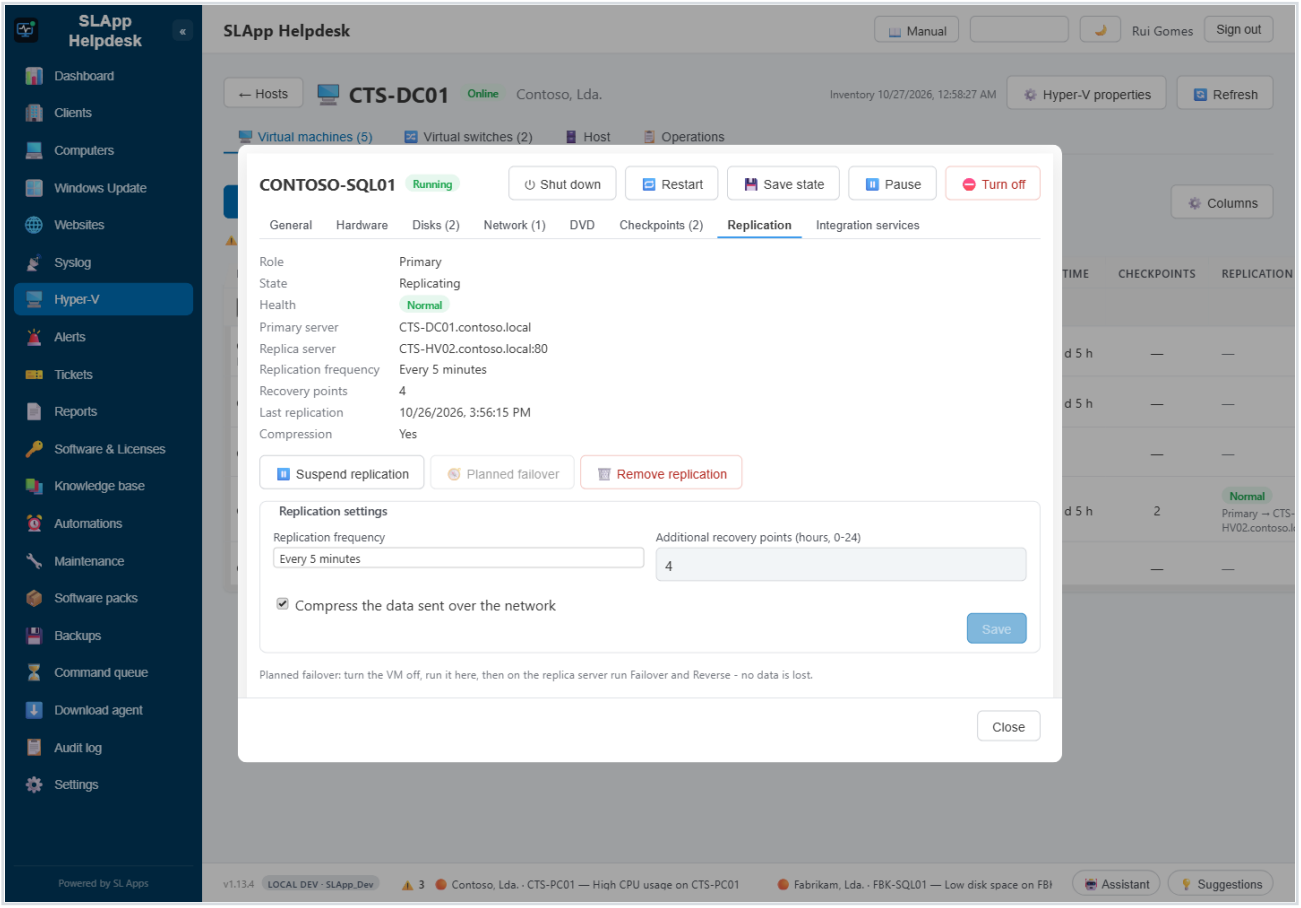
Deleting, moving, exporting and importing

- Delete virtual machine asks you to type the VM's name. A running VM must be turned off first (tick the option to do it). Its checkpoints are merged into the disks first. The disk files stay on the host unless you tick Also delete its disk files — and a file another VM still uses, or a physical disk, is never deleted.
- Move to another host offers two modes: move the VM's files (configuration, disks, checkpoints) to another folder of the same host while it keeps running, or live-migrate the VM to another Hyper-V host, with or without its disks. Moving to another host runs as the host itself, so both hosts must be in the same domain with live migration turned on and Kerberos with constrained delegation set up between them.
- Export writes the VM (configuration, disks, checkpoints) to a folder of the host, inside a folder with the VM's name.
- Import virtual machine (button above the list) takes the folder of an export or a .vmcx file. Copy it gives the VM a new ID and copies its files to the folders you choose, leaving the export untouched; Register it where it is uses the files in place with the same ID. A new name is optional.

Replication of a virtual machine

Hyper-V Replica keeps a copy of a VM on another Hyper-V host, updated every 30 seconds, 5 or 15 minutes, ready to start if the first host is lost. The other host must accept replication first: open it, go to its Host tab and, under Hyper-V Replica, enable it as a replica server with its default folder for the replica files (its firewall rule is opened at the same time). Only Kerberos (HTTP) authentication is set up from here, so both hosts must be in the same domain; certificate authentication is configured on the host itself.

1. Open the VM's properties and go to the Replication tab.
2. Under Replica server choose the other host — the field suggests the full names of the other hosts in the module.
3. Choose the replication frequency, the additional recovery points (hours) and, with recovery points, how often an application-consistent copy is taken; leave Compress the data sent over the network on.
4. Leave Start the initial copy over the network right away ticked: once replication is enabled, the whole disks are sent in the background on the host — it can take hours.
5. If the VM's own Windows connects iSCSI disks (see Inside the guest on the Disks tab), tick Turn off the Backup (volume shadow copy) integration service — see the tip below.
6. Click Enable replication.



The Replication tab of a replicated VM: role, state, health, servers, frequency, recovery points and the actions of the primary.

Once replicated, the tab shows the role (primary or replica), state, health, both servers, frequency, recovery points, last replication and compression, and the list shows the replication health of each VM. On the primary: suspend and resume, resynchronize, change the frequency, recovery points and compression, prepare a planned failover (with the VM turned off — the last changes are sent so nothing is lost) and remove the replication. On the replica server: test failover (a separate copy named <name> - Test starts without network, without touching replication) and stop it, failover (only when the primary is lost or after a planned failover), complete or cancel it, and reverse the replication so the replica becomes the primary. Removing, failing over, preparing a planned failover, reversing, completing and resynchronizing ask for confirmation.

> **i** Hyper-V Replica only replicates the VM's virtual disks. A physical (pass-through) disk attached to the VM blocks it completely. A VM whose Windows connects iSCSI disks itself is refused with "one or more iSCSI disks are attached" while the Backup (volume shadow copy) integration service is offered, because the host then asks the guest for its disks; with that service turned off the VM is accepted and only its virtual disks are replicated, the iSCSI disks being left alone. The VM's backups and checkpoints then become crash-consistent instead of application-consistent.

Replication alerts

Since version 1.14.77 the module raises an alert when the replication health of a virtual machine is Critical or Warning: Hyper-V replication, with the same severity as the health, in the Alerts module (source Hyper-V), the client's Alerts tab, the Dashboard, the Infrastructure overview and the footer ticker, under the client of the host that reported it. Its message names the virtual machine, the host and the replication state, and its Computer column is that host, which Create ticket selects. There is one alert per virtual machine, even when its primary, its replica and the nodes of its cluster are all in the module: what the primary reports counts first, then the replica, and between two nodes the most recent inventory. Only hosts that are online, whose last inventory is less than 75 minutes old and did not fail afterwards are taken into account: what an offline host reported last opens and closes nothing. While the primary's host cannot be read, what the replica reports can open the alert or make it Critical, but for up to a day it never closes or lowers it. The Infrastructure overview reads each virtual machine the same way, so a replica's Critical is not listed there while the primary reports Normal.

The alert is e-mailed to the administrators when it opens and again when a Warning becomes Critical; when several open in the same pass (a replica server going down turns every virtual machine it receives Critical at once), each administrator receives one summary e-mail. The first time a host is evaluated (a host just added, and every host already in the module when version 1.14.77 arrives), what it finds opens without any e-mail. The alert closes by itself when the health is Normal again or the replication is removed, and when the virtual machine no longer appears on any host of the module once they have all been read again; removing a host from the module closes at once the alerts of the virtual machines only it listed, and removing its computer closes them within minutes. A health of Not applicable, a test failover virtual machine or a virtual machine the host could not read leave the alert as it is. Unticking Generate alerts resolves them and also stops the automatic inventories (see Turning alert generation on or off per module in the Alerts section).

> **i** An open alert whose host is turned off stays open until the host is read again: nothing is known about it in the meantime. A host that is normally turned off never opens an alert from its old inventory; if it is retired for good, remove it from the module and its alerts close.

Virtual switches

The Virtual switches tab lists the host's switches with their type, physical network adapter, whether the host shares it and the VMs using it. + New virtual switch creates an External switch (bound to a physical adapter, so the VMs reach the real network — an adapter already bound to another external switch cannot be chosen), an Internal one (between the VMs and the host) or a Private one (between the VMs only); click a switch to rename it, change its type or adapter, or edit its notes. Deleting a switch used by VMs disconnects them.

> **i** Binding a switch to a network adapter, changing it or deleting a switch the host shares drops the host's network on that adapter for a few seconds. If the agent reaches the host through that adapter, the result may take a while to come back.

The host (Hyper-V properties)

The Host tab — also opened by Hyper-V properties in the header — shows the memory in use and assigned to running VMs, the logical processors and the virtual processors of running VMs, the VMs running out of the total, the system, domain and uptime, the volumes with their free space and the physical network adapters. Below, the Hyper-V settings: the default folders for virtual hard disks and VM configuration, live migrations (allowed or not, how many at a time, authentication, any network or only the listed live migration networks, which you add and remove by subnet), NUMA spanning, enhanced session mode, the pool of dynamic MAC addresses and Hyper-V Replica (this host as a replica server). The configuration versions the host supports are listed at the bottom.

Operations

The Operations tab lists the latest 100 operations of the host: when, what, on which VM, its state (waiting for the host, running, done or failed, with its duration), the outcome — the error of a failed one, the files deleted with a VM, the folder of an export — and who asked for it. Inventories and folder listings are hidden unless they failed; tick Show inventories and folder listings to see them. The inventories the module asks for on its own (every half hour, and after each change) are left out once they have finished, even with that box ticked, so they never push the real operations out of the 100. An operation the host never picks up (offline) is given up after 24 hours; the command itself stays in the Command queue.

20. Active Directory

The Active Directory module watches the health of your clients' Active Directory domains. Once a day the agent of every domain controller runs a read-only check that covers dcdiag, replication, the AD services, the clock against the PDC emulator, the SYSVOL and NETLOGON shares, the FSMO roles, the user accounts and the members of the privileged groups, and the module turns the results into one row per domain, a health status for every part and alerts when something goes wrong, including a new member of Domain Admins. Nothing has to be installed or configured: domain controllers are recognised by the check itself. Open it from Infrastructure › Active Directory in the sidebar. Every technician can use it, mark or unmark a domain controller and run a check; only administrators can close an alert by hand.

The daily check

The same check also feeds the Third-party backups module. SLApp Helpdesk queues it once a day on every Windows computer of your clients, as a PowerShell script that the agent already installed runs with system rights through the Command queue, so no agent update is needed (Windows PowerShell 3.0 or later). On every computer it finds out whether the computer is a domain controller and reads the event logs of the supported backup software; on a domain controller it also checks Active Directory. It only reads: it installs nothing, loads no module and writes nothing to the directory, to the disk or to an event log.

- When: at the same time of day, every day. The first checks are spread out so that a whole fleet never checks at once: within two hours for servers, within a day for workstations and within an hour for a computer added in the last day. The computer's page shows the time of the next one.
- Only while the computer is online. A computer that is offline at its time is tried again every few minutes until it is back, and that never counts as a failure. A check queued for an offline computer, for example with Check now, is cancelled after 6 hours and queued again once the computer is back. A check that gives no result within 14 hours counts as a failed check.
- The results are collected as soon as the check has finished, within five minutes, or at once when somebody opens one of these pages.

- In the Command queue the check is a Command/script command on the System channel whose text starts with # SLApp health check. The finished ones are hidden from the queue module unless you tick Show daily checks, and they are deleted 3 days after they finish (see the Command queue section).
- Check now runs it at once (see Check now). Linux and macOS computers have no such check.

Domain controllers: detection and marking

The check recognises a domain controller by the role Windows reports for the computer (a backup or primary domain controller), or by the presence of the NTDS service when the role cannot be read. The primary domain controller is the PDC emulator of its domain, the time reference that the other domain controllers are compared with. A domain controller appears in the module once its first check has been read, with its domain. How each Windows computer is monitored is a setting with three values, in the Monitoring list of each domain controller in the domain window and in the Monitoring as a domain controller list of the computer's Overview:

- Automatic detection, the starting value: the computer is checked as a domain controller whenever its daily check says it is one.
- Always monitor as a domain controller: the Active Directory part runs even when the detection does not see a domain controller. A check is queued at once.
- Never monitor as a domain controller: no Active Directory checks or alerts for this computer. Choosing it for a computer that is monitored as a domain controller asks Stop monitoring DC01 as a domain controller? Its AD alerts are resolved; its daily check keeps running for the third-party backups. Confirm closes its AD alerts, and a domain that has no other monitored domain controller leaves the list with its own alerts.
- Going back to Automatic detection stops the monitoring at once when the last check did not detect a domain controller. Every change is written to the Audit log under the entity type Active Directory.

1. To monitor a server that the detection misses, click + Mark a domain controller at the top of the module.
2. In the Mark a domain controller window, choose the computer in Server. The list holds the Windows servers that are not in the module yet, grouped by client, with (Offline) after the ones that are offline. When there is none left it says No Windows server left to mark — every server with the agent is already in the list.
3. Click Mark as domain controller. Its check is queued at once and the server appears in the list as soon as the result is in, a few minutes when it is online.

What the check reads on a domain controller

- dcdiag: Windows' own domain controller diagnostic, run in quiet mode, which only prints the tests that fail. It is given 150 seconds. Its output is kept (the last 8,000 characters) and shown.
- Replication: every replication partner of the domain controller, for each directory partition, with its last success, last attempt, consecutive failures and last result, read from Windows' replication API.
- AD services: NTDS (Active Directory Domain Services), DNS, Netlogon, Kdc (Kerberos Key Distribution Center), W32Time, DFSR and NtFrs, with their state and start mode.
- Time: the clock offset from the PDC emulator, measured with three samples of w32tm. The PDC emulator itself is the reference and is never measured.
- SYSVOL and NETLOGON: whether both shares exist, and the SysvolReady flag that Windows sets once SYSVOL has replicated.
- The domain: DNS and NetBIOS names, forest, forest root domain, domain SID, the domain and forest functional levels and the holders of the five FSMO roles.

- Accounts: the numbers of enabled users, disabled users, locked-out users, enabled users whose password never expires, and enabled users inactive for 90 days (no logon in the last 90 days, or never logged on and created more than 90 days ago). The last logon of an account replicates with a delay of up to 14 days, so this figure is approximate.
- Privileged groups: the members of Domain Admins of the domain and of Enterprise Admins and Schema Admins of the forest root domain, including the members of nested groups and the accounts whose primary group it is, with their account name, display name, type and whether they are enabled.
- Each part is read on its own, so one that fails never hides the others; a part that could not be read shows its error.

Health and statuses

Every part of a domain controller gets a status, shown as a small chip: OK, Warning, Problem or Unknown (the part could not be read), or Not checked yet. The health of a domain controller is Critical when a part has a Problem, Warning when a part has a Warning or is Unknown, or when three checks in a row gave no readable Active Directory result, and OK otherwise. The health of a domain, and the status of each of its parts in the list, is the worst one of its domain controllers.

- dcdiag: Problem when a test fails, Warning when the only tests that fail are SystemLog, DFSREvent, KccEvent or FrsEvent. Those four fail on any recent error or warning in an event log, a reboot is enough, so they are shown but never alerted.
- Replication: Problem when a partner is failing and has had no successful replication for more than 24 hours, Warning when a partner is failing or has not replicated for more than 24 hours. A domain controller without partners, the only one of its domain, is OK.
- Services: Problem when one of the services above is set to start automatically and is not running.
- Time: Problem beyond 5 minutes from the PDC emulator, the limit beyond which Kerberos sign-ins fail, Warning beyond 60 seconds. The PDC emulator is always OK.
- SYSVOL/NETLOGON: Problem when either share is missing, Warning while SysvolReady is not set, which is normal during the first or a non-authoritative SYSVOL synchronisation.

The list of domains

The module lists one row per Active Directory domain and client, so two clients that both use corp.local are two rows. The top of the page has + Mark a domain controller and Columns. Under the introduction come a client list (All clients, — no client — or a client that has a domain), a search box that looks in the domain and NetBIOS names and in the domain controllers' names, and the Show list: Show all, With problems or open alerts, New privileged accounts and Not checked yet. Until the first domain controller is found the page explains that every Windows computer runs the check once a day. The list refreshes itself every minute.

- Client and Domain: the DNS name in bold with the NetBIOS name; hover it for the forest.
- Domain controllers: one chip per domain controller with its health, and pending while a check is waiting for its result. Hover a chip for its full name, PDC emulator when it is one, its health and whether it is offline.
- Health, dcdiag, Replication, Services and Time: the status chips described above. Hover the dcdiag, Replication, Services or Time chip for what is wrong on which domain controller, for example the failed tests or the failing partners. Time also shows the largest offset of the domain.
- FSMO (PDC): the short name of the PDC emulator; hover it for the holders of the five roles.
- Admins: how many different accounts the three privileged groups hold together, with a red crown while a New privileged account alert is open; hover it for when the groups last changed.

- Locked out, Inactive 90 days, Last check and Alerts (the open alerts of the domain and of its domain controllers).
- SYSVOL/NETLOGON, Functional level, Enabled users and Password never expires are available through Columns.

Click a header to sort by that column and again to reverse it; the list starts sorted by client. The row under the headers filters Client, Domain, Domain controllers, FSMO (PDC) and Functional level by text, the status columns by status, Admins on New accounts and Alerts on With alerts or Without alerts. Click a row to open the domain window. The refresh button at the end of the row checks all the domain controllers of the domain now, and the right-click menu offers Open details, Check all domain controllers now, Copy domain name and Copy PDC emulator.

The domain window

The title is the DNS name of the domain with its NetBIOS name, followed by the client, the time of the last check, the health of the domain and the button Check all domain controllers now. Three blocks come first. Domain shows the DNS name, NetBIOS name, forest, forest root domain, domain SID and the domain and forest functional levels, for example Windows Server 2016. FSMO roles shows the PDC emulator, RID master, Infrastructure master, Schema master and Domain naming master, and Accounts the numbers of enabled, disabled, locked-out, password-never-expires and inactive users; both end with the domain controller and the time they were read from.

Privileged groups shows Domain Admins, Enterprise Admins and Schema Admins with the number of members of each and a table with the Account, Name, Type (User, Group, Computer, managed service account and so on), Enabled and Seen since, the date the account was first seen in the group. An account added after the monitoring started carries a red New badge for 7 days, and the line under the groups says since when they are monitored. When a group could not be read on a domain controller, a red line says so and the members shown are the ones known before. Right-click an account to copy its name or its SID. The members known to the domain are those reported by the domain controllers that answered in the last 7 days, so a domain controller whose replication is stuck neither drops nor adds anybody.

Domain controllers has one block per domain controller: the online dot, its name (a link to the computer's page), its full name, PDC emulator when it is one, its health, Check in progress while a check waits for its result, the Monitoring list described above and a Check now button. Under them come the time of the last check and of the last readable result, the number of checks in a row without a readable result and the last error, then the parts of the last result: dcdiag with the failed tests in red and the event log tests in grey, and dcdiag output to read the full text; Replication with a table of the partners (Partner, Partition, Last success, Last attempt, Failures, Result), the failing ones in red; AD services with each service's state and start mode; Time with the offset and the reference, and w32tm output; and SYSVOL / NETLOGON with both shares and SysvolReady. History (N) opens the last 30 checks of that domain controller, each with its health, the status of every part, the time offset and a note naming what failed. The right-click menu of a block offers Check now, Open computer, Copy computer name and Copy FQDN.

Open alerts lists the open alerts of the domain and of its domain controllers with their Severity, Type, Message and Opened at; administrators also have a Resolve button and a right-click Resolve. When the last domain controller of the domain stops being monitored while the window is open, the window closes and says that the domain has no monitored domain controller left.

Check now

Check now queues the daily check of one computer at once: in the block of a domain controller, in the computer's Overview, or for every domain controller of a domain with Check all domain controllers now,

which skips the ones already waiting for a result. A notice says Check of DC01 queued — the results show up as soon as the computer answers, or that the computer is offline and the check runs as soon as it is back online. The check runs on the computer, where dcdiag alone can take a couple of minutes, so the result is not immediate: the domain window looks again every 10 seconds while a check is waiting, and the list a little after Check all domain controllers now and then every minute. A computer can be checked at most once every 5 minutes, and never while a check of it is still waiting; otherwise the notice says This computer was checked less than 5 minutes ago — try again in N s, or that a check is already waiting for its result. Check now does not move the time of the daily check.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is Active Directory; they appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, with the name of the domain controller, or of the domain, in front of the message, and you can create a ticket from them like from any other alert. They are evaluated each time a check of a domain controller is read, and a part that could not be read leaves its alert as it is. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- AD: dcdiag failed (Critical): a dcdiag test failed, except the four event log tests. The message names the tests. It closes when dcdiag passes again.
- AD replication failing (Critical): a partner has been failing with no successful replication for more than 24 hours. It closes when no partner is in that state any more.
- AD service stopped (Critical): one of the AD services set to start automatically is not running. It closes when they all run.
- DC clock offset (Critical): the clock of the domain controller is more than 5 minutes off the PDC emulator. It closes when the offset is back under 5 minutes.
- SYSVOL not shared (Critical): SYSVOL or NETLOGON is not shared. It closes when both are.
- AD check failing (Warning): three checks in a row gave no readable Active Directory result. It closes as soon as a check reads Active Directory again.
- New privileged account (Critical): an account appeared in Domain Admins, Enterprise Admins or Schema Admins that the domain did not have before. It is about the domain, names who was added to which group and closes by itself 7 days after it opened, or when an administrator resolves it. Each new change closes the previous alert and opens a new one, with its own email. The first reading of a domain, or of a group, is the starting point and raises nothing.
- Emails go to the active administrators of your company who have an email address, when an alert opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the names in it, such as the dcdiag tests, stay as they were written). On the first check of a domain controller, a Warning opens without an email; Critical alerts always send one.
- The alerts of a domain controller close when it stops being monitored or is removed, and follow it when it is renamed or moved to another client. Only an administrator can close an alert by hand, with Resolve, and that is written to the Audit log under the entity type Active Directory; if the condition is still true, the next check opens it again.

> **i** The computer's Overview shows the role the check detected and, for a domain controller, its health and the status of each part, under Active Directory & third-party backups (see the Computers section). The client report has an Active Directory section (PDF and Excel) with every domain controller, its health and parts, and the account figures and number of privileged accounts of its domain; the full report leaves it out for a client without a domain controller.

21. Exchange

The Exchange module watches the health of your clients' on-premises Microsoft Exchange servers (Exchange 2013, 2016, 2019 and Subscription Edition). Once a day, as part of the same read-only daily check that serves Active Directory and the third-party backups (see The daily check in the Active Directory section), the agent of every Exchange server reads its version, services, mailbox databases and their copies, transport queues, certificates, server components, disk space and health sets. The module turns the results into one row per server, a status for every part and alerts when something goes wrong. Nothing has to be installed or configured: Exchange servers are recognised by the check itself. Open it from Infrastructure › Exchange in the sidebar. Every technician can use it, mark or unmark an Exchange server and run a check; only administrators can close an alert by hand. Exchange is part of the Complete plan: on Essential the menu, the page and the Exchange part of the daily check are not available.

Detection and the monitoring mode

The check recognises an Exchange server from the Windows registry: a computer where Exchange setup recorded a server role. A computer with only the Exchange Management Tools, such as an admin or Remote Desktop server, does not count. The server appears in the module once its first check has been read. How each Windows computer is treated is a setting with three values, in the Exchange list of the computer's Overview (in the Active Directory & third-party backups block) and in the Monitoring list of the server's detail window:

- Automatic detection, the starting value: the computer is checked as an Exchange server whenever its daily check finds Exchange installed. Under the list, Exchange detected by the last check. confirms it.
- Always monitor as an Exchange server: the Exchange part runs even when the detection does not see Exchange. A check is queued at once.
- Never monitor as an Exchange server: no Exchange checks or alerts on this computer. For a monitored server it first asks Stop monitoring EX01 as an Exchange server? Its Exchange alerts are resolved; its daily check keeps running for the rest. (in the detail window, the question ends with: Its Exchange alerts are resolved and it leaves this list.). Confirm closes its Exchange alerts and removes the server from the list.
- Going back to Automatic detection on a server that the last check did not detect asks a similar question first, because it stops the monitoring at once. Every change is written to the Audit log under the entity type Exchange.

1. To monitor a server that the detection misses, click + Mark an Exchange server at the top of the module.
2. In the Mark an Exchange server window, choose the computer in Server. The list holds the Windows servers that are not in the module yet, grouped by client, with (Offline) after the ones that are offline. When there is none left it says There are no more Windows servers to mark.
3. Click Mark as an Exchange server. Its check is queued at once and the server appears in the list as soon as the result is in.

What the check reads

Some parts are read without Exchange's own tools, so they are there even when Exchange is in trouble: the install (the version of ExSetup.exe, the roles and the install folder), the Exchange services together with IIS (W3SVC and WAS), the free space of the volumes, and the TLS certificates bound to ports 443 and 444. The rest is read through the Exchange Management Shell, with Get and Test cmdlets only, in a separate session limited to 240 seconds: the server, the mailbox databases

(mounted, size, internal free space, mailboxes, last full backup), the DAG copies and their content index, the transport queues, the Exchange certificates and the auth certificate, the server components (maintenance mode), Test-ServiceHealth and the unhealthy health sets with their failing monitors. A part that finishes in time is kept even if a later one does not, and slow parts are skipped near the limit. The shell needs 64-bit PowerShell, and only one Exchange check runs on a computer at a time. On an Edge Transport server the database parts are shown as not applicable. The check only reads: it changes nothing on the server.

Health and statuses

Every part gets a status chip: OK, Warning, Problem or Unknown (the part could not be read), or Not checked yet. The health of a server is Critical when Services, Databases, Queues, Certificates, Components or Disks has a Problem, Warning when one of them has a Warning, and OK otherwise. When the Exchange Management Shell could not be read, the health is at most Warning (Unknown when nothing at all was read), and three such checks in a row make it Warning. Version and Health sets are shown only: they never change the health and never raise an alert.

- Version: Warning when the version is out of support (of the versions above, Microsoft only supports Exchange Server Subscription Edition) or older than the latest known build of its line, shown with the Out of support or Outdated badge.
- Services: Problem when an Exchange service (MSExchange…), W3SVC or WAS set to start automatically is not running, or when Test-ServiceHealth reports required services not running. Optional services switched off on purpose (EdgeSync, POP3, IMAP4 and their back ends, set to Disabled) do not count.
- Databases: Problem when a mailbox database that should be mounted on this server is dismounted, or a copy is Failed or FailedAndSuspended; Warning when a content index is failed or unknown, or a copy or replay queue is above 10. Recovery databases are ignored. A last full backup older than 7 days is highlighted but never alerted.
- Queues: the messages waiting for delivery, without the shadow, submission and poison queues. Problem from 500 messages, Warning from 100 or when the poison queue holds messages.
- Certificates: the certificates in use (with services assigned) and the current auth certificate. Problem when one expires within 7 days or has expired, Warning within 30 days; a self-signed certificate other than the auth certificate is at most a Warning. Old auth certificates and Windows' own MS-Organization-Access and MS-Organization-P2P-Access certificates are ignored. A port 443 or 444 binding to a certificate that no longer exists or has no private key, or with no certificate at all, is a Problem; a Mailbox server with TLS on 443 and nothing bound to 444 is a Warning.
- Components: Warning when a server component is not Active (maintenance mode), Problem when ServerWideOffline is inactive, which takes the server out of service. ForwardSyncDaemon and ProvisioningRps are ignored.
- Disks: the volumes that hold the databases, the logs or the install folder (all volumes when none can be matched). Problem below 5% free, Warning below 15%.
- Health sets: Warning when Exchange's managed availability reports an unhealthy health set.

The list of servers

The page lists one row per Exchange server, sorted by client and then by name. Above it come a client list (All clients, No client or a client), a search box (Search server, FQDN or version…) and the Show list: Show all or Only with problems (a server with a Warning or Critical health, or with open alerts). The columns are Client, Server (with the online dot, and pending while a check is waiting), Version, Health, Services, Databases (with how many are mounted), Queues, Certificates (hover it for the next certificate in use to expire), Components, Disks, Mailboxes, Last check and Alerts; Health sets, Edition

and Roles are available through Columns. Click a header to sort by that column, and use the row under the headers to filter. Until the first server is found the page explains that the daily check detects Exchange servers by itself. The list refreshes itself every minute.

Click a row to open the server's detail window. At the end of the row, refresh checks the server now and, when the AI assistant is available and the server has a problem, assistant is Ask the AI. The right-click menu offers Open detail, Check now, Ask the AI, Open computer and copying the server name or FQDN.

The detail window

The title shows the server, whether it is online, the client, the last check and the health, followed by the Monitoring list, Ask the AI (when the server has a problem) and Check now. Notes give the last successful check, the number of failed checks in a row, the last error, and whether reading through the Exchange shell timed out or failed. Then come Server (name, FQDN, edition, roles, version, build, latest build, site and DAG, install folder), Services, Databases with their copies, Queues, Certificates with the table Certificates bound to ports 443 / 444 (IIS) and how to fix a broken binding, Server components, Disks (the volumes watched with alerts are marked Exchange), Unhealthy health sets with their failing monitors, Open alerts and History (N), the last 30 checks with the health and the status of every part. Administrators have a Resolve button on each open alert.

Check now and Ask the AI

Check now queues the daily check of the server at once, from the list, the detail window or the computer's Overview. It follows the same rules as in the Active Directory module: at most once every 5 minutes per computer and never while a check is still waiting, and an offline computer is checked as soon as it is back online. The Exchange part can take a few minutes, so the detail window looks again every 10 seconds while a check is waiting. Ask the AI sends the AI assistant what the last check found on the server, with its open alerts, and asks it to find the cause and propose the fix, using read-only commands and waiting for your confirmation before anything that changes the server.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is Exchange, the server's name comes in front of the message, and they appear in the Alerts feed, on the Dashboard and in the Alerts tab of the client like every other alert. They are evaluated each time a check of the server is read, and a part that could not be read leaves its alert as it is. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- Exchange service stopped (Critical): a watched service is not running. It closes when they all run.
- Exchange database dismounted (Critical): a mailbox database is dismounted. It closes when none is.
- Exchange database copy failed (Critical): a database copy is Failed or FailedAndSuspended. It closes when none is.
- Exchange queues building up: Warning from 100 messages waiting, Critical from 500. It closes below 100.
- Exchange certificate expiring: a certificate in use expires within 30 days (Warning), or within 7 days or has expired (Critical; a self-signed one other than the auth certificate stays a Warning). It closes when no certificate in use is in that state.
- IIS certificate missing (Exchange) (Critical): a port 443 or 444 binding points at a certificate that is gone or has no private key, or has no certificate; it is a Warning when the only problem is the missing 444 binding. It closes when every binding could be read and is correct.
- Exchange component inactive: Warning for an inactive component, Critical when ServerWideOffline

is inactive. It closes when they are all active.

- Low disk space on Exchange: Warning below 15% free, Critical below 5%. It closes when every watched volume has enough space again.
- Exchange check failing (Warning): three checks in a row could not read Exchange. It closes after a successful check.
- Emails go to the active administrators of your company who have an email address, when an alert opens. On the first check of a server, Warnings open without an email; Critical alerts always send one. The alerts close when the server stops being monitored, when the computer is removed or when the plan no longer includes Exchange. Only an administrator can close one by hand, with Resolve; if the problem is still there, the next check opens it again.

> **i** An Exchange server also runs IIS, so it appears in Infrastructure › IIS as well, where the problems with the certificates bound to its sites on ports 443 and 444 show too.

22. SQL Server

The SQL Server module watches the Microsoft SQL Server instances of your clients. Once a day the agent reads every instance of a computer that has SQL Server, with a check that only reads, and the module turns the results into one row per server, a status for every part and alerts when something goes wrong: a stopped service, a database that is not online, a full transaction log, missing backups, a failed SQL Server Agent job, low disk space or an Express database near its 10 GB limit. Nothing has to be installed or configured: servers with SQL Server are recognised by the check itself. Open it from Infrastructure › SQL Server in the sidebar. SQL Server is a module of the Complete plan. Every technician can use it, mark or unmark a server and run a check; only administrators can close an alert by hand.

What the check reads

The SQL Server part belongs to the same daily check as Active Directory, run by the agent already installed on every Windows computer (see The daily check in the Active Directory section). It finds the instances in the Windows registry, up to 16 per computer, and connects to each one with Windows authentication as the computer's SYSTEM account (NT AUTHORITY\SYSTEM). It only runs SELECT queries: it changes nothing in SQL Server, on the disk or in the permissions. Each connection gets 10 seconds, each query 30 seconds (60 for the backup history) and the whole SQL Server part 150 seconds; a part cut by that limit is simply not read this time. Per instance it reads:

- Services: the database engine and the SQL Server Agent, with their state, start mode and account, read from Windows, so they are known even when the instance cannot be connected to.
- Version: the product (SQL Server 2008 to 2025), the build, the update level and the edition, with the end of Microsoft's extended support.
- Databases: every database with its state, recovery model, data size and log size, and the percentage of each transaction log that is in use.
- Backups: the last full, differential and log backup of every database, from the backup history kept in msdb.
- Jobs: the SQL Server Agent jobs, with whether they are enabled and the outcome, time, duration and message of their last run. SQL Server Express has no SQL Server Agent, so there are no jobs to read there.
- Disks: the volumes that hold the data and log files of the databases, with their size and free space.
- Each part is read on its own, so one that fails or is refused never hides the others.

Detection and marking

The check recognises a computer with SQL Server from the instances registered in the Windows registry. A server appears in the module once its first check has been read. How each Windows computer is monitored is a setting with three values, in the Monitoring list of the server window and in the SQL Server list of the computer's Overview, in the Active Directory & third-party backups block:

- Automatic detection, the starting value: the computer is checked as a SQL Server whenever its daily check finds an instance. When the last check detected one, the Overview adds SQL Server detected by the last check.
- Always monitor as a SQL Server: the SQL Server part runs even when the detection does not find SQL Server. A check is queued at once.
- Never monitor as a SQL Server: no SQL Server checks or alerts on this computer. Choosing it for a server that is monitored asks for confirmation; on the computer's Overview the question is Stop monitoring SRV01 as a SQL Server? Its SQL Server alerts are resolved; its daily check keeps running for the rest. In the server window it ends with Its SQL Server alerts are resolved and it leaves this list.
- Going back to Automatic detection on a server that the last check did not detect also asks first: Detection does not see SQL Server on SRV01, and continuing stops its monitoring and resolves its SQL Server alerts. Every change of this setting is written to the Audit log under the entity type SQL Server.

1. To monitor a server that the detection misses, click + Mark a SQL Server at the top of the module.
2. In the Mark a SQL Server window, pick the computer in Server. The list holds the Windows servers that are not in the module yet, grouped by client, with (Offline) after the ones that are offline; when there is none it says There are no more Windows servers to mark.
3. Click Mark as a SQL Server. Its check is queued at once and the server appears in the list as soon as its result is in.

Health and statuses

Every part of a server gets a status, shown as a small chip: OK, Warning, Problem, Unknown or Not checked yet; the status of a server with several instances is the worst one of them. Unknown means that the part could not be read on some instance, because the instance could not be connected to or a read was refused: it is never shown as OK and never closes an alert. A part that does not apply, such as Jobs on an Express instance, shows (n/a). An instance whose engine is stopped on purpose, with the start mode Manual or Disabled (an instance nobody uses, or the passive node of a clustered instance), counts as idle: nothing is expected from it, so it raises no alert and leaves the other parts as they are. The health of the server is Critical when a part has a Problem and Warning when a part has a Warning; when the last check could not read every instance, or three checks in a row failed, it is at least Warning, or Unknown when nothing could be read; otherwise it is OK. Version and Access are shown only: they never set the health and never raise an alert.

- Services: Problem when an engine service set to start automatically is not running, Warning when the SQL Server Agent is set to start automatically and is not running (never on Express).
- Databases: Problem when a database is SUSPECT, RECOVERY_PENDING or EMERGENCY, Warning when it is OFFLINE or RECOVERING. Database snapshots are left out.
- Log space: Problem when the transaction log of a database is more than 90% used. tempdb is left out, because its log empties and refills on its own.
- Backups: Warning when a database has no full backup in the last 7 days, or never had one, and when a database in FULL or BULK_LOGGED recovery that has a full backup has no log backup in

the last 24 hours. Only online databases are judged; tempdb, model, snapshots, databases in an Always On availability group (their backups may run on another replica) and databases created less than 24 hours ago are left out. master and msdb without a recent full backup are only a note in the server window, shown in grey: they never make the server a Warning nor open an alert.

- Jobs: Warning when an enabled SQL Server Agent job failed on its last run within the last 48 hours.
- Disks: Warning when a volume that holds data or log files has less than 15% free, Problem below 5%.
- Express size: only on SQL Server Express, whose databases are limited to 10 GB of data each. Warning when the data of a user database reaches 90% of the limit, Problem from 98%.
- Version: Out of support once Microsoft's extended support for the product has ended, or for a version older than SQL Server 2008, and Support ends soon within a year of that date. A product whose end of support has not been published shows Not announced yet.
- Access: Warning when the check was refused some reads; the server window then shows the Read access missing panel.

Read access missing

The SYSTEM account is often not a sysadmin of the instance, and a part it may not read is left Unknown, never OK. The server window then shows Read access missing, with each instance concerned, what the check could not read (for example the backup history (msdb) or the volumes) and the T-SQL script that grants the read access to NT AUTHORITY\SYSTEM: it creates the login if it does not exist, grants VIEW SERVER STATE and VIEW ANY DEFINITION, and in msdb adds the account to db_datareader and SQLAgentReaderRole. The script can be run more than once without harm. Copy puts it on the clipboard; run it in SQL Server Management Studio as an administrator of that instance and press Check now. SLApp never runs it and never changes permissions on your servers: it only shows the script, and running it is up to you. Permissions the check needs (T-SQL), at the bottom of the window, shows the same script for any instance, for when you want to grant the access ahead of time.

The list of SQL Servers

The module lists one row per computer with SQL Server, with all its instances. The top of the page has + Mark a SQL Server and Columns. Under the introduction come a client list (All clients, No client or one client), a search box that looks in the server and instance names, the version and the edition, and the Show list: Show all, Only with problems and Only with read access missing. When filters hide servers, the count of the ones shown is displayed. Until the first server is found the page explains that the daily check detects them by itself. The list refreshes itself every minute.

- Client and Server: the online dot, the computer name and pending while a check is waiting for its result.
- Instances: the names of the instances, MSSQLSERVER for the default one.
- Version: the product and the build, with an Out of support or Support ends soon badge.
- Health, Services, Log, Disks and Express size: the status chips described above. Hover Services for the services that are stopped.
- Databases: how many databases are online out of the total, with the status chip.
- Backups: the date of the oldest last full backup, with the status chip.
- Jobs: how many jobs failed, with the status chip.
- Access: Read access missing when some reads were refused, hover it for what could not be read; otherwise the status chip.
- Last check, in red with the number of failed checks in a row when the last ones failed, and Alerts, the number of open alerts, in red when one of them is Critical.

- Edition and Data size (data and logs of all the databases) are available through Columns.

Click a header to sort by that column and again to reverse it; the list starts sorted by client. The row under the headers filters Client, Server and Instances by text, Version on Supported, Support ending or ended or Out of support, the status columns by status, Access on Read access missing or Full read access, and Alerts on With alerts or Without alerts. Click a row to open the server window. At the end of the row, refresh checks the server now and assistant asks the AI assistant about it. The right-click menu offers Open detail, Check now, Ask the AI, Open computer and Copy server name.

The server window

The title is the computer name, a link to its page, with its instances. Under it come the online state, the client, the time of the last check, the health, Check in progress while a check waits for its result, the Monitoring list, Ask the AI and Check now, then one chip per part: Version, Services, Databases, Log space, Backups, Jobs, Disks, Express size and Access. Notes follow when they apply: the last successful check, the failed checks in a row and the last error, another SLApp check reading SQL Server on the same computer at that moment (the reading is skipped and asked again shortly), and the time limit reached. Then come the Read access missing panel and Problems found, one line per problem, for example a database with no full backup in the last 7 days or a volume with little free space.

Instances (N) shows one instance at a time, with one button per instance when there are several (warning when its engine is down or it could not be connected to). Each instance shows its Version, Edition, End of support, Running since, the Check login, the Permissions of that login (sysadmin, or whether it has VIEW SERVER STATE and VIEW ANY DEFINITION), Cluster, Collation and Machine. Below come Services (Service, Role, State, Start mode, Account); Databases (Database, State, Recovery, Data, Log, Log used, Last full, Last differential, Last log), with tags for system databases, snapshots, availability groups (AG) and read-only databases, the share of the 10 GB limit on Express, and old or missing backups highlighted; SQL Server Agent jobs (Job, Enabled, Last outcome, Last run, Duration, Message), the recent failures highlighted; and Volumes of the database files (Volume, Label, Files, Size, Free, Free %). Right-click an instance, a database or a job to copy its name, its connection name or the job message.

Open alerts lists the open SQL Server alerts of the server with their Severity, Type, Message and Opened at; administrators also have a Resolve button and a right-click Resolve. History (N) opens the last 30 checks, each with its health, the status of every part, the databases online and the failed jobs, or the error of a check that could not read anything.

Check now and Ask the AI

Check now queues the daily check of the server at once: with refresh at the end of its row, from the right-click menu, in the server window or in the computer's Overview. It works as described under Check now in the Active Directory section: the result arrives once the computer has run it, the server window looks again every 10 seconds while a check is waiting, the list a little after and then every minute, and a computer can be checked at most once every 5 minutes and never while a check of it is still waiting.

Ask the AI appears when the AI assistant is available to you and the server has a Warning, is Critical or has open alerts: as assistant in the row, in the right-click menu and in the server window. It sends the assistant what the last check found on that server, the instances, the problems and the open alerts, and asks it to find the likely cause of each failure and propose the steps to fix it, using only read-only commands on the server and explaining any command that would change something and waiting for your confirmation first. Review what it proposes before you let it act.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is

SQL Server; they appear like the Active Directory alerts, with the name of the server in front of the message, and you can create a ticket from them like from any other alert. There is one alert of each type per server, and the message names the instances and databases concerned. They are evaluated each time a check of the server is read. An alert closes by itself only when its part was read on every instance and nothing was found, so a part that is Unknown leaves its alert as it is. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- SQL Server service stopped: Critical when an engine service set to start automatically is not running, Warning when only the SQL Server Agent is. It closes when they all run.
- SQL Server database not online: Critical for a database that is SUSPECT, RECOVERY_PENDING or EMERGENCY, Warning for one that is OFFLINE or RECOVERING. It closes when every database is online.
- SQL Server transaction log almost full (Critical): a transaction log more than 90% used. It closes when every log is back under 90%.
- SQL Server full backup missing or old (Warning): no full backup in the last 7 days, or never. It closes when every database judged has a recent full backup.
- SQL Server log backup missing or old (Warning): a database in FULL or BULK_LOGGED recovery with no log backup in the last 24 hours. It closes when they all have one.
- SQL Server Agent job failed (Warning): an enabled job whose last run failed within the last 48 hours. It closes when no job is in that state any more.
- Low disk space on SQL Server: Warning under 15% free on a volume with data or log files, Critical under 5%. It closes when every volume is back above 15%.
- SQL Server Express near its size limit: Warning from 90% of the 10 GB limit, Critical from 98%. It closes when every database is back under 90%.
- SQL Server check failing (Warning): three checks in a row could not read the server. An instance that refuses the SYSTEM login does not count: it is shown with the grant script instead. It closes as soon as a check reads the server again.
- Emails go to the administrators as for the Active Directory alerts, when an alert opens or goes from Warning to Critical. The first evaluation of a server is silent: its Warnings open without an email, while Critical alerts always send one.
- The alerts of a server close when it stops being monitored and follow it when it is renamed or moved to another client. Only an administrator can close an alert by hand, with Resolve, which asks Resolve the alert by hand? If the problem is still there, the next check opens it again.

23. Storage / iSCSI

The Storage / iSCSI module watches the SAN and server storage of your clients' Windows computers: the iSCSI sessions and targets, the disks, the Multipath I/O (MPIO) paths, the Storage Spaces pools with their virtual and physical disks, and the free space of the storage volumes. It is a part of the daily check (see The daily check in the Active Directory section), which reads the storage through WMI queries only: SLApp never brings a disk online, reconnects a session or changes MPIO. The module turns the results into one row per server, a status for each part and alerts when something goes wrong. Nothing has to be installed or configured: storage servers are recognised by the check itself. Open it from Infrastructure › Storage / iSCSI in the sidebar; it is a module of the Complete plan. Every technician can use it, mark or unmark a storage server and run a check; only administrators can close an alert by hand.

What the check reads

- iSCSI: the state and start mode of the Microsoft iSCSI Initiator service, the initiator name, the sessions (target, target portals, whether connected, connections and the disks they carry), the persistent logins (favourite targets) and the discovered targets.
- Disks: number, name, serial number, bus (iSCSI, Fibre Channel, SAS, Storage Spaces and so on), health, status, whether the disk is offline and why, read-only, system, boot or clustered, and its size.
- MPIO: for each disk managed by Multipath I/O, the number of paths and the DSM.
- Storage Spaces: the pools, the virtual disks (with their resiliency and detached reason) and the physical disks (with their media, usage and health). They are read on servers only.
- Volumes: drive letter or path, label, file system, health, size and free space.
- Each read has a time limit of 30 seconds and the whole storage part 120 seconds. A part that is not in use on the server, for example MPIO when Multipath I/O is not installed, is skipped and shown as not applicable; a part that could not be read is Unknown, never OK, and shows its error. On Windows Server 2008 R2 and older the disks and volumes are read from older classes, with fewer details.

Storage servers: detection and marking

The check recognises a storage server when at least one of these is in use: the Microsoft iSCSI Initiator service is running and has sessions, persistent logins or discovered targets; the Multipath I/O driver (the Microsoft DSM) is installed; or, on a server, there is a Storage Spaces pool other than the primordial one. A server appears in the module once its first check has been read, and the In use column shows which of iSCSI, MPIO and Storage Spaces it uses. A detection that fails never removes a server. How each Windows computer is monitored is a setting with three values, in the Monitoring list of the detail window and in the Storage / iSCSI list of the computer's Overview, under Active Directory & third-party backups:

- Automatic detection, the starting value: the computer is monitored whenever its daily check detects iSCSI, MPIO or Storage Spaces. The Overview then adds iSCSI / MPIO / Storage Spaces detected by the last check.
- Always monitor as a storage server: monitored even when the detection does not see any of them. A check is queued at once.
- Never monitor as a storage server: no storage checks or alerts on this computer. Choosing it for a monitored computer asks Stop monitoring SRV01 as a storage server? Its storage alerts are resolved; its daily check keeps running for the rest (in the detail window: Its storage alerts are resolved and it leaves this list). Confirm closes its storage alerts and the server leaves the list; its history is kept and comes back if it is monitored again.
- Going back to Automatic detection on a server that the last check did not detect asks first: Detection does not see iSCSI, MPIO or Storage Spaces on SRV01: automatic detection stops monitoring it as a storage server and resolves its storage alerts. Continue? Every change is written to the Audit log under the entity type Storage / iSCSI.

1. To monitor a server that the detection misses, click + Mark a storage server at the top of the module.
2. In the Mark a storage server window, choose the computer in Server. The list holds the Windows servers that are not in the module yet, grouped by client, with (Offline) after the ones that are offline. When there is none left it says There are no more Windows servers to mark.
3. Click Mark as storage server. The server is set to Always monitor as a storage server, its check is queued at once and it appears in the list as soon as the result is in.

Health and statuses

Every part (iSCSI, Disks, MPIO, Pools, Virtual disks, Physical disks and Volumes) gets a status, shown as a small chip: OK, Warning, Problem or Unknown (the part could not be read), or Not checked yet; in the detail window a part that does not apply to the server shows (n/a). The health of a server is Critical when a part has a Problem, Warning when a part has a Warning, and OK otherwise. When the storage itself could not be read (the check failed or the disks could not be read) the health is never OK: it is Unknown, or Warning when some parts were still read, and it becomes at least Warning after three checks in a row without a readable result. A part that could not be read is never counted as fixed: its alerts stay as they are.

The list of servers

The module lists one row per storage server, sorted by client and then by name, servers without a client last. The top of the page has + Mark a storage server and Columns. Under the introduction come a client list (All clients, No client or one client), a search box that looks in the server, client and operating system, and the Show list: Show all, Only with problems, Only with iSCSI, Only with MPIO and Only with Storage Spaces. Until the first server is found the page explains that the daily check detects them by itself. The list refreshes itself every minute.

- Client and Server: the online dot and the computer name, with pending while a check is waiting for its result and, under it, the name the server reports when it is different.
- In use: one chip each for iSCSI, MPIO and Storage Spaces.
- Health: the health of the server described above.
- iSCSI, Disks, MPIO, Pools and Volumes: the status chip of each part with a short figure: the sessions connected out of the targets known (for example 2/2 connected), the number of SAN disks (3 SAN), the fewest MPIO paths of any disk (min. 2 paths, hover it for the number of MPIO disks), the number of pools and the number of storage volumes low on space (2 low).
- Last check, in red when the last checks failed (hover it for how many), and Alerts, the open alerts of the server.
- Operating system, Virtual disks, Physical disks and Disk count are available through Columns.

Click a header to sort by that column and again to reverse it. The row under the headers filters Client and Server by text, In use by feature, Health and each part by status, and Alerts on With alerts or Without alerts. Click a row to open the detail window. At the end of the row, assistant asks the AI (see below) and refresh runs Check now; the right-click menu offers Open detail, Check now, Ask the AI, Open computer and Copy server name.

The detail window

The title is the server name, a link to the computer's page, with the iSCSI, MPIO and Storage Spaces chips. Under it come the online state, the client, the operating system, the time of the last check, the health, Check in progress while a check waits for its result, the Monitoring list described above, Ask the AI and Check now. Then a chip for each part, the note Read-only: SLApp never brings a disk online, reconnects a session or changes MPIO, the time of the last successful check, the number of failed checks in a row and the last error. A note says so when the check reached its time limit, or when another SLApp check was reading the storage of the computer at the same time; that reading is skipped and asked for again shortly. What the last check found lists every problem in a plain sentence, for example Disk 2 (iSCSI) is offline, and Detection shows what the detection saw: In use, the iSCSI sessions, persistent logins and targets, whether Multipath I/O is installed and the number of Storage Spaces pools.

The parts follow, with the rows that have a problem in red. iSCSI shows the service, the initiator name

and the tables Sessions, Persistent logins and Discovered targets, each target with its Last seen connected (Never (last 30 checks) for one never used). Disks lists every disk with its bus, health, status, offline reason, read-only and size, marked SAN, System, Boot or Cluster. MPIO shows the Paths of each disk next to Before, the most paths it had in the last 30 checks, and Disk no longer under MPIO for one that left it. Storage Spaces shows the Pools, Virtual disks and Physical disks, and Volumes the free space of each volume, with Not evaluated on the local ones. Open alerts lists the open alerts of the server with their Severity, Type, Message and Opened at; administrators also have a Resolve button and a right-click Resolve, which asks, for example, Resolve the alert “Disk offline” by hand? If the problem is still there, the next check opens it again. History (N) opens the last 30 checks, each with its health, the status of every part, the iSCSI sessions, the fewest MPIO paths and the low volumes, or the error of a check that gave no readable result. The right-click menu of the title offers Check now, Open computer and Copy server name.

Check now and Ask the AI

Check now queues the daily check of the server at once, from the row, the detail window or the computer's Overview, and a notice says Check of SRV01 queued — the results show up as soon as the computer answers, or that the computer is offline and the check runs as soon as it is back online. The result is not immediate: the detail window looks again every 10 seconds while a check is waiting, and the list 20 seconds after Check now and then every minute. A computer can be checked at most once every 5 minutes, and never while a check of it is still waiting; otherwise the notice says This computer was checked less than 5 minutes ago — try again in N s, or that a check is already waiting for its result. Check now does not move the time of the daily check.

Ask the AI appears when the AI Assistant is available to you and a server has a Warning or Critical health or open alerts: as assistant at the end of its row, in its right-click menu and as Ask the AI in the detail window, which then closes. It opens the AI Assistant with a ready-made request holding the server, its client, operating system, the storage in use, its health, what the last check found and its open alerts, and asks for the likely cause of each failure and the steps to fix it, in order. The assistant may use its tools on the server for read-only diagnostic commands such as Get-IscsiSession, Get-Disk, Get-StoragePool or mpclaim -s -d; the request asks it to explain any command that would change something, such as bringing a disk online or connecting a target, and to wait for your confirmation first.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is Storage / iSCSI; they appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, with the name of the server in front of the message, and you can create a ticket from them like from any other alert. There is one alert of each type per server, whose message names the disks, targets or volumes concerned and whose severity is that of the worst of them. They are evaluated each time a check of the server is read, and each one closes on the first check that reads its part and no longer finds the problem. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- iSCSI target not connected (Critical): a persistent login has no connected session, or the Microsoft iSCSI Initiator service is not running while it is set to start automatically, has persistent logins or had connected targets. It closes when every persistent target is connected again and the service runs.
- iSCSI target without a session (Warning): a target that was connected in one of the last 30 checks now has no session. It closes when the target is connected again, or once it has not been connected in any of the last 30 checks.
- Disk offline (Critical, or Warning when the disk is offline because of the SAN policy): a disk is offline. The message gives the reason. It closes when the disk is back online.

- Disk read-only (Critical): a disk is online but read-only. It closes when it is writable again.
- Disk unhealthy (Critical when a disk's health is Unhealthy or its status is Failed, Lost communication or No contact; Warning when its health is Warning). It closes when the disk is healthy again.
- MPIO path lost (Critical when a disk has no path left, including a disk that had 2 or more paths and is no longer under MPIO while the disk itself is still there; Warning when a disk is down to 1 path after having had 2 or more in the last 30 checks). It closes when the paths are back.
- Storage pool degraded (Critical when a pool is Unhealthy or read-only; Warning when its health is Warning or its status Degraded).
- Virtual disk degraded (Critical when a virtual disk is Unhealthy or Detached; Warning when its health is Warning or its status Degraded or Incomplete).
- Physical disk failing (Critical when a physical disk is Unhealthy or its status is Lost communication, No contact, Error, Non-recoverable error or Failed; Warning when its health is Warning, its status Predictive failure or it has been retired).
- Low space on storage volume (Warning under 15 % free, Critical under 5 % free; also Warning or Critical when the volume's own health is Warning or Unhealthy). Only storage volumes count: those on iSCSI, Fibre Channel, Storage Spaces or MPIO disks, on the disks of an iSCSI session, and Cluster Shared Volumes. It closes when every storage volume is back above 15 % free.
- Storage check failing (Warning): three checks in a row gave no readable storage result. The message gives the error. It closes as soon as a check reads the storage again.
- Emails go to the active administrators of your company who have an email address, when an alert opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the names in it, such as the disks and targets, and an error text quoted from the server, stay as they were written). On the first evaluation of a server, when it starts being monitored, a Warning opens without an email; Critical alerts always send one.
- The alerts of a server close when it stops being monitored or is removed, and follow it when it is renamed or moved to another client. Only an administrator can close an alert by hand, with Resolve; if the condition is still true, the next check opens it again.
- Shown but never alerted: a clustered disk that is offline or read-only, since it belongs to the cluster and is owned by another node; its health is still checked.
- Shown but never alerted: a disk offline because it is a redundant path of a disk that is already online, or because it is a snapshot. A disk offline by SAN policy is only a Warning.
- Shown but never alerted: a virtual disk detached by policy, such as a cluster disk owned by another node, or set to be attached by hand. A read-only pool that belongs to a cluster is not a Problem either.
- Shown but never alerted: a discovered target that this server never connected to, since a portal often advertises targets the server does not use, and an MPIO disk that has always had a single path.
- Shown but never alerted: the volumes of local disks, marked Not evaluated; their free space is watched by the disk alert rules.

> **i** Storage / iSCSI belongs to the Complete plan. When a company moves to Essential its open storage alerts are resolved and the module and the Storage / iSCSI list of the Overview disappear; moving back to Complete resumes the monitoring with the next daily check, with the history kept.

24. Failover Cluster

The Failover Cluster module watches the Windows failover clusters of your clients: the quorum and its witness, the nodes, the roles, the resources, the Cluster Shared Volumes (CSVs) and the cluster networks with their interfaces. It is fed by the same daily read-only check as Active Directory (see The daily check in the Active Directory section): on a cluster node that check also reads the cluster, through WMI queries only. SLApp never pauses or drains a node, moves a role, brings a resource online or offline, puts a CSV in maintenance or changes the quorum. Nothing has to be installed or configured: cluster nodes are recognised by the check itself. Open it from Infrastructure › Failover Cluster in the sidebar. It is part of the Complete plan. Every technician can use it, mark or unmark a cluster node and run a check; only administrators can close an alert by hand.

What the check reads

- Cluster Service: the state and start mode of the node's own Cluster Service. While it is not running, nothing else is read on that node.
- Quorum: the cluster name, FQDN, id and functional level, the quorum type and path, dynamic quorum and the witness resource with its state. A cluster without a witness is shown as No witness (node majority).
- Nodes: every node of the cluster with its state (Up, Down, Paused, Joining) and its drain status.
- Roles: every role (resource group) with its type, state, owner node and whether an administrator took it offline on purpose.
- Resources: all of them are counted, but only the ones that are not online are listed, so a large Hyper-V cluster stays readable.
- CSVs: each Cluster Shared Volume with its state, whether it is redirected, its fault state (no direct I/O, no access, in maintenance, dismounted), its owner node, size and free space.
- Networks: every cluster network with its state and role, and every network interface of every node with its state.
- Each part is read on its own, each read is given 20 seconds and the cluster part 90 seconds in all; a part that could not be read shows its error and never hides the others.

Cluster nodes: detection and marking

The check sees a computer as a cluster node when the Cluster Service is installed and the computer is a member of a cluster: the cluster registry hive is loaded, the cluster answers through WMI, or the Cluster Service is set to start automatically. A computer without the Failover Clustering feature reads nothing more. How each Windows computer is monitored is a setting with three values, in the Failover cluster monitoring list of the computer's Overview (in the Active Directory & third-party backups card) and in the Monitoring list of each node in the cluster window:

- Automatic detection, the starting value: the computer is monitored as a cluster node whenever its daily check detects one. The hint adds Cluster node detected by the last check and Node of CLUSTER01 once it has been read.
- Always monitor as a cluster node: the cluster part runs even when the detection does not see a node. A check is queued at once.
- Never monitor as a cluster node: no cluster checks or alerts from this computer. For a computer that is monitored as a node it asks Stop monitoring NODE01 as a cluster node? Its node alerts are resolved (and its cluster's, when it is the last monitored node); its daily check keeps running for the rest.
- Going back to Automatic detection on a node the last check did not detect asks Detection does not

see NODE01 as a cluster node: automatic detection stops monitoring it and resolves its node alerts (and its cluster's, when it is the last monitored node); its daily check keeps running for the rest. Continue? Every change is written to the Audit log under the entity type Failover Cluster.

1. To monitor a server that the detection misses, click + Mark a cluster node at the top of the module.
2. In the Mark a cluster node window, choose the computer in Server. Only Windows servers not already monitored as cluster nodes are listed; when there is none left it says Every Windows server is already monitored here.
3. Click Mark as cluster node. Its check runs at once and the node appears under its cluster as soon as the result is in. Until then it is listed under Nodes without a cluster read yet, at the bottom of the page, with its last error if the check failed.

Health and statuses

Every part gets a status chip: OK, Warning, Problem or Unknown (not read). A part that does not apply, such as CSVs on a cluster that has none, is shown as Not used on this cluster. The health of a cluster is Critical when a part has a Problem or when the Cluster Service of one of its monitored nodes is not running, Warning when a part has a Warning, and OK otherwise. When no node has read the cluster in the last 48 hours, the cluster shows No recent report: what you see is the last view read and its health is Unknown.

- Nodes: Problem when a node is Down; Warning when a node is Paused, Joining or in an unknown state, or when draining it failed.
- Roles: Problem when a role has Failed; Warning when it is Offline while the cluster expects it online, or Partially online. A role that an administrator took offline on purpose carries the badge offline on purpose and is never alerted, nor is the Available Storage group, which is normally offline. Pending roles are only shown.
- Resources: Problem when a resource has failed. A failed resource of a failed role is counted under the role, the witness under the quorum and the CSVs under the CSVs.
- CSVs: Problem when a CSV has failed, has no access or is dismounted; Warning when it is not online, redirected or in maintenance (set on purpose); free space below 15% is a Warning and below 5% a Problem.
- Networks: Problem when a cluster network is down, partitioned or unavailable, or an interface has failed or is unreachable; Warning when an interface is unavailable. Networks not used by the cluster, and the interfaces of a node that is down, are shown but not evaluated.
- Quorum: Problem when the witness has failed or is offline, Warning while it is pending, coming online or going offline.

The list of clusters

Every node runs the same check and sees the same cluster, so the module groups the nodes and lists one row per cluster, whichever node reported it. Under the introduction come a client list (All clients, No client or a client), a search box that looks in the cluster, node and client names, and the Show list: Show all, Only with problems and Only without a recent report. The columns are Cluster (name and FQDN), Client, Nodes (for example 2/2 up, with the nodes that are down or paused), Monitored nodes (how many nodes have the agent and are monitored, with one chip per node coloured by its own Cluster Service), Roles (online of all, and how many failed), Resources, CSVs (how many, with the lowest free space), Networks, Quorum (type and witness state), Health, Last report (when, and from which node) and Alerts. Columns chooses the columns, a click on a header sorts, the row under the headers filters each column, and the list refreshes itself every minute.

Click a row to open the cluster window. The refresh button at the end of the row is Check cluster now,

and Ask the AI appears when the cluster has a problem. The right-click menu offers Open the cluster, Check cluster now, Ask the AI, Copy cluster name and Copy FQDN.

The cluster window

The title is the cluster name, followed by the client, Reported by NODE01 at … with the time of the last view read, the health, Check in progress while a check waits, and the buttons Ask the AI and Check cluster now. Under the part chips come the problems found, then Summary (cluster name, FQDN, cluster id, functional level, Last view read by, quorum type, quorum path, dynamic quorum and witness), Nodes with their state, drain, vote and build and whether each one has the SLApp agent (monitored, agent, not monitored or no agent), Roles with type, state and owner node, Resources (only those not online), Cluster Shared Volumes with path, state, redirected, fault state, size and free space, and Networks and Network interfaces.

Monitored nodes lists each node with the agent: its Cluster Service, health, last check (and how many checks in a row failed), the Monitoring list described above, magnifying glass to open the node and Check now. Open alerts lists the alerts of the whole cluster and of its nodes, with About (Whole cluster or the node); administrators also have Resolve. History merges the last checks of all the nodes. A node window shows the same for one node: its own Detection (Cluster Service, Cluster registry hive, Cluster name (registry), Cluster read through WMI), its parts, errors, alerts and history, and Open the cluster.

Check now and Ask the AI

Check cluster now queues the daily check on every online monitored node of the cluster and skips the ones already waiting; the notice says Check requested on 2 nodes, or that no monitored node is online or a check is already running. Check now on a node checks only that node. The same limits as every daily check apply: a computer can be checked at most once every 5 minutes and never while a check of it is still waiting. Ask the AI opens the AI assistant with what the last check found on the cluster and its open alerts, and names a monitored node it can work through for read-only diagnostic commands such as Get-ClusterNode, Get-ClusterGroup or Get-ClusterLog; the question asks the assistant to explain what a command would change and wait for your confirmation before running it.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Problems of the cluster itself have the source Failover Cluster and are raised ONCE per cluster, with the cluster name as target: a second node reporting the same problem only refreshes the open alert, with no second email. Problems of one node have the source Failover Cluster (node) and are raised per node. Cluster alerts are evaluated each time a node reports a real read of the cluster; a part that could not be read leaves its alerts as they are. They appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- Cluster node down (Critical): a node is Down. It closes when no node is down.
- Cluster node paused (Warning): a node is Paused, Joining or in an unknown state, or draining it failed.
- Cluster role failed (Critical): a role has failed; the message names its owner node and its failed resources.
- Cluster role offline (Warning): a role is offline while the cluster expects it online, or partially online.
- Cluster resource failed (Critical): a resource has failed outside a failed role, the witness and the CSVs.
- Cluster Shared Volume failed (Critical): a CSV has failed, has no access or is dismounted. Cluster

Shared Volume redirected (Warning): a CSV is not online, redirected or in maintenance.

- Low space on Cluster Shared Volume: Warning below 15% free, Critical below 5%. It closes only when the free space of every CSV was read and all are above 15%.
- Cluster network down: Critical for a network down, partitioned or unavailable, or an interface failed or unreachable; Warning for an interface unavailable.
- Quorum witness problem: Critical when the witness has failed or is offline, Warning while it is pending.
- Cluster Service stopped (Critical, per node): the Cluster Service is set to start automatically and is not running. Cluster check failing (Warning, per node): three checks in a row could not read the cluster; it closes as soon as a check succeeds.
- Each alert closes when the next read of its part no longer finds the problem. Emails go to the active administrators of your company who have an email address, when an alert opens. On the first evaluation of a cluster or a node, a Warning opens without an email; Critical alerts always send one. After that, an email goes only when an alert opens again.
- A node's alerts close when it stops being monitored or is removed; a cluster's alerts close when its last monitored node stops being monitored or joins another cluster. Moving to the Essential plan resolves them all. Only an administrator can close an alert by hand, with Resolve; if the problem is still there, the next check opens it again.

25. DHCP

The DHCP module watches your clients' Windows DHCP servers. Once a day the agent of every DHCP server reads, without changing anything, the DHCP Server service, the authorization of the server in Active Directory, the use of every IPv4 scope, the IPv6 scopes, the failover relationships and the server statistics, and the module turns the results into one row per server, a status for every part and alerts when a scope is almost full, the service is stopped, the server is not authorized or a failover relationship is not in its normal state. Nothing has to be installed or configured: DHCP servers are recognised by the check itself. Open it from Infrastructure › DHCP in the sidebar. The module is part of the Complete plan. Every technician can use it, mark a server, change how a computer is monitored and run a check; only administrators can close an alert by hand.

What the check reads

The DHCP part belongs to the daily check that every Windows computer of your clients runs (see The daily check in the Active Directory section). It recognises a DHCP server by the DHCP Server service together with the DHCP Server role. It only reads, through WMI queries: it never activates or deactivates a scope, deletes a lease or reservation, forces a failover state, touches the service or changes the authorization in Active Directory. Each read is given 20 seconds and the whole DHCP part 60 seconds, and only one DHCP check runs at a time on a computer; a result skipped because another check was already reading the server leaves everything as it was, and the check is asked for again later the same day. While the DHCP Server service is not running, nothing else is read.

- Service: the state and start mode of the DHCP Server service.
- Server settings: whether the server is in a domain, whether it is authorized in Active Directory, the list of authorized servers in Active Directory, the conflict detection attempts, dynamic BOOTP and whether policies are active.
- IPv4 scopes: every scope with its name, mask, range, state, lease duration and superscope, and its statistics: free, in use, reserved and pending addresses and the percentage in use.
- IPv6 scopes: only how many there are.

- Failover relationships: the name, partner, mode (Load balance or Hot standby), state, role of this server, scopes and settings of each one. The failover shared secret is never read.
- Server statistics: the totals of scopes and addresses, the addresses in use and available, the counts of discovers, offers, requests, acks, NAKs, declines and releases, the pending offers and the time the server started.

Detection and monitoring mode

How each Windows computer is monitored as a DHCP server is a setting with three values, in the Monitoring list at the top of the server window and in the DHCP Server monitoring list of the computer's Overview, under Active Directory & third-party backups (see the Computers section). A notice says DHCP server monitoring of DHCP01 saved, and every change is written to the Audit log under the entity type DHCP.

- Automatic detection, the starting value: the computer is monitored as a DHCP server whenever its daily check finds the DHCP Server role. When it has, the computer's Overview says DHCP server detected by the last check.
- Always monitor as a DHCP server: the DHCP part runs even when the detection does not see a DHCP server. A check is queued at once.
- Never monitor as a DHCP server: no DHCP checks or alerts from this computer. Choosing it for a monitored server asks Stop monitoring DHCP01 as a DHCP server? Its DHCP alerts are resolved (and a failover relationship's, when it is the last monitored server listing it); its daily check keeps running for the other parts. In the server window the end of the question says instead that it leaves this list.
- Going back to Automatic detection on a server where the detection does not see the DHCP Server role asks first too: Detection does not see the DHCP Server role on DHCP01: automatic detection stops monitoring it and resolves its DHCP alerts (and a failover relationship's, when it is the last monitored server listing it). Continue?

1. To monitor a server that the detection misses, click + Mark a DHCP server at the top of the module.
2. In the Mark a DHCP server window, choose the computer in Server. The list holds the Windows servers that are not in the module yet, with (Offline) after the ones that are offline; when there is none left it says There are no more Windows servers to mark.
3. Click Mark as DHCP server. This sets the computer to Always monitor as a DHCP server, its check is queued at once and the server appears in the list as soon as the result is in.

Health and statuses

Every part of a server (Service, Authorization, Scopes and Failover) gets a status, shown as a small chip: OK, Warning, Problem or Unknown (the part could not be read), or n/a when the part is not used on this server. The health of a server is Critical when a part has a Problem, Warning when a part has a Warning, when a part could not be read or when three checks in a row gave no readable result, and OK otherwise.

- Service: Problem when the DHCP Server service is set to start automatically and is not running. A service stopped on purpose (Manual or Disabled) is shown but is not a problem.
- Authorization: Problem when the server is in a domain but is not authorized in Active Directory, because such a server hands out no addresses. A server in a workgroup needs no authorization and shows Workgroup, with n/a.
- Scopes: only the active scopes are evaluated. Warning from 85% of a scope's addresses in use, Problem from 95% or when a scope in use has no free address left. Inactive scopes hand out no address: they are shown, never alerted. When the scope statistics cannot be read, the use of the

scopes is not evaluated.

- Failover: OK in the Normal state, Problem when the partner is down (Partner down), Warning in any other state, such as Communication interrupted or Recover wait. A server without failover relationships shows n/a.
- DHCP tools missing: when the DHCP Server tools (RSAT-DHCP) are not installed on the server, SLApp cannot read its scopes or failover. The server is shown with the DHCP tools missing badge and raises no alert; install the tools on the server to see them. SLApp installs nothing.

The DHCP page

The top of the page has + Mark a DHCP server and Columns. Under the introduction come a client list (All clients, No client or one client), a search box that looks in the server, scope, failover and client names, and the Show list: Show all, Only with problems, Only with failover and Only without the DHCP tools. Until the first server is found the page explains that the daily check recognises a server with the DHCP Server role by itself. The list refreshes itself every minute. Its columns are Client; Server, with pending while a check is waiting for its result; Service; Authorization (Authorized, NOT authorized, Workgroup or Unknown); Scopes, the active IPv4 scopes over all of them, for example 3 active / 4; Scope use, a usage bar for each of the three fullest active scopes, amber or red as described above, and +N more for the rest (hover a bar for the addresses in use and free); Addresses, in use and free; Failover, a chip per relationship with its state (hover it for the partner and mode); Health; Last check; and Alerts. Operating system is available through Columns.

Click a header to sort by that column and again to reverse it; the row under the headers filters Client and Server by text, the status columns by status and Alerts on With alerts or Without alerts. Click a row to open the server window. The refresh button at the end of the row checks the server now, and the right-click menu offers Open detail, Check now, Ask the AI, Open computer and Copy server name. Below the servers, Failover relationships (N) lists each relationship once, even though both partners report it: Relationship, Client, Mode, State, Servers (each partner with its role; a monitored partner is a link that opens its server window, the other one carries not monitored), Scopes, Reported by (which partner gave the newest report, and when) and Alerts. Right-click a relationship to copy its name.

The server window

The title is the name of the server (a link to the computer's page) with its authorization badge, followed by the online dot, the client, the operating system, the time of the last check, the health, Check in progress while a check waits for its result, the Monitoring list, Ask the AI and Check now. A line recalls that it is read-only: SLApp never changes scopes, leases, failover or the AD authorization. Then come the last successful read, the failed checks in a row, the last error and what the last check found, in plain sentences. Summary shows the service, the authorization in AD, In a domain, Listed in AD, the authorized servers in AD, the conflict detection attempts, when the server started, the IPv6 scopes and other settings. IPv4 scopes (N) is a table with Scope, Name, Mask, Range, State, Lease, Use (a usage bar), Free, In use, Reserved and Pending; inactive scopes are marked inactive - not monitored. Failover (N) has one block per relationship with its state, Open alert when it has one, Partner, Mode, Role of this server, Scopes, Maximum client lead time, State switch interval, Load balance or Reserve addresses, Message authentication, the servers and who reported it. Server statistics shows the counters listed above.

Open alerts lists the open alerts of the server and of its failover relationships with their Severity, Type, Message and Opened at; administrators also have a Resolve button and a right-click Resolve, which asks Resolve the alert by hand? If the problem is still there, the next check opens it again. History (N) opens the last 30 checks, each with its health, the status of every part, the scopes, the fullest scope, the addresses and the failover relationships that were not normal. Right-click the title for Check now, Open computer and Copy server name. When the server stops being monitored while the window is

open, the window closes and says so.

Check now and Ask the AI

Check now queues the daily check of the server at once, from the refresh button of its row, its right-click menu or the server window. It behaves as in the Active Directory section: a notice says that the check is queued, or that the computer is offline and the check runs as soon as it is back online; a computer can be checked at most once every 5 minutes and never while a check of it is still waiting for its result, otherwise the notice says This computer was checked less than 5 minutes ago — try again in N s or that a check is already waiting. The server window looks again every 10 seconds while a check is waiting, and the list 20 seconds after Check now and then every minute.

When the AI assistant is available and a server has a warning, a problem or an open alert, a Ask the AI button appears on its row, in its right-click menu and in the server window. It sends the AI assistant what the last check found on the server and its open alerts, and asks it to find the likely cause and propose the fix step by step. The assistant may run read-only diagnostic commands on the server, and explains and waits for your confirmation before any command that changes anything, such as starting the service, authorizing the server or changing a scope or a failover relationship.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is DHCP for the alerts of a server and DHCP (failover) for the alerts of a failover relationship; they appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, and you can create a ticket from them like from any other alert. They are evaluated each time a check of a server is read, and a part that could not be read leaves its alert as it is. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- DHCP Server service stopped (Critical): the service is set to start automatically and is not running. It closes when it runs again.
- DHCP server not authorized in AD (Critical): the server is in a domain but is not authorized in Active Directory. It closes when it is authorized.
- DHCP scope almost full: Warning when an active scope has 85% or more of its addresses in use, Critical when one has 95% or more or no free address left. The message names up to five scopes, the fullest first. It closes when every active scope is back under 85%.
- DHCP check failing (Warning): three checks in a row gave no readable DHCP result. It closes as soon as a check reads the server again.
- DHCP failover not normal: Warning while a relationship is in any state other than Normal, Critical when the partner is down. It is raised once per relationship, not once per server: both partners run the check and refresh the same alert, which carries the name of the relationship, follows the newest report and never sends a second email. It closes when a check shows the relationship Normal again, or when no monitored server lists it any more.
- Emails go to the active administrators of your company who have an email address, when an alert opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the names in it, such as the scopes, stay as they were written). The first evaluation of a server, or of a failover relationship, is silent for warnings: a Warning opens without an email, while Critical alerts always send one.
- The alerts of a server close when it stops being monitored or is removed; a failover alert closes with the last monitored server that lists the relationship. When your company moves to a plan without the DHCP module, all its DHCP alerts are closed. Only an administrator can close an alert by hand, with Resolve, and that is written to the Audit log under the entity type DHCP; if the condition is still true,

the next check opens it again.

> **i** A server whose DHCP Server service is stopped on purpose, or that you are retiring, does not need to stay in the list: set it to Never monitor as a DHCP server. Its daily check keeps running for the other parts, such as Active Directory and third-party backups.

26. IIS

The IIS module watches your clients' IIS web servers. Once a day the daily check of every Windows Server that runs IIS reads the World Wide Web Publishing (W3SVC) and Windows Process Activation (WAS) services, every site with its bindings and applications, every application pool, the state each site and pool is running in, the HTTPS certificates and the HTTP 5xx errors of the last 24 hours, and the module turns that into one row per server, a status for each part and alerts when a site or pool is stopped, a certificate is missing, expiring or does not cover the site's name, or a site returns too many server errors. Open it from Infrastructure › IIS in the sidebar. It is part of the Complete plan. Every technician can use it, mark a server, change how a computer is monitored and run a check; only administrators can close an alert by hand.

What the check reads

The IIS part runs inside the daily check (see The daily check in the Active Directory section) and only reads: it uses no IIS management module or tool, opens the IIS configuration file and the log files for reading only and never starts, stops or recycles a site or pool, nor touches a binding, a certificate or a service. Passwords stored in the configuration, such as application pool identities, are never read. It is given 45 seconds.

- Services: the state and start mode of W3SVC and WAS.
- Configuration: the sites with their ID, bindings, physical path, application pool, automatic start and log settings, their applications and the application pools (start mode, .NET CLR, pipeline, 32-bit, identity type), read from the IIS configuration file. A server that uses shared configuration is read from the shared location and shows a Shared configuration badge.
- Runtime state: whether each site and application pool is started or stopped, read from the IIS performance counters, and the running worker processes. When the counters cannot be read the state is Unknown and nothing about it is alerted; when W3SVC is not running the state is not read at all.
- Certificates: the certificate HTTP.sys hands out for every HTTPS binding, on any port (by IP address, by SNI host name or from the central certificate store), matched to the machine's certificate store with its subject, DNS names, issuer, expiry and whether its private key is there. Certificates from the IIS central certificate store are listed but not checked.
- HTTP errors: the requests and HTTP 5xx responses of the last 24 hours, counted from the end of each site's W3C log: the two newest files of each site, at most 4 MB of each, 50 sites and 24 MB in all. When only part of the day could be read the count is marked partial (its percentage stays valid). A site that logs in another format, has logging off or writes to the server's central log is not counted.

IIS servers: detection and marking

A computer is recognised as an IIS web server when its check finds the World Wide Web Publishing service (W3SVC) on Windows Server, and appears in the module once its first result has been read. A workstation (Windows 10 or 11) where IIS is installed, often by another program, is not detected: it appears only when it is set to Always monitor as an IIS web server. Exchange servers run IIS, so they appear here too, and a certificate problem on their ports 443 and 444 shows both in the Exchange

module and here. How each Windows computer is monitored is a setting with three values, in the IIS web server monitoring list of the computer's Overview and in the Monitoring list of the server's window:

- Automatic detection, the starting value: the computer is checked as an IIS web server whenever its daily check finds W3SVC on Windows Server (never on a workstation).
- Always monitor as an IIS web server: the IIS part runs even when the detection does not see W3SVC. A check is queued at once.
- Never monitor as an IIS web server: no IIS checks or alerts for this computer. Choosing it for a monitored server asks, in the computer's Overview, Stop monitoring WEB01 as an IIS web server? Its IIS alerts are resolved; its daily check keeps running for the other parts. — and in the module, Stop monitoring WEB01 as an IIS web server? Its IIS alerts are resolved and it leaves this list. Confirm closes its IIS alerts and takes it off the list.
- Going back to Automatic detection on a server whose last check did not see W3SVC asks first too, because it stops the monitoring at once and resolves its IIS alerts. Every change is written to the Audit log under the entity type IIS.

1. To monitor a computer that the detection misses, click + Mark an IIS server at the top of the module.
2. In the Mark an IIS server window, choose it in Computer. Only Windows computers not yet in the list are shown, grouped by client, with (Offline) after the ones that are offline.
3. Click Mark as IIS server. The computer is set to Always monitor as an IIS web server, its check runs right away and it shows up in the list as soon as its result is in.

Health and statuses

Every server has five parts, each with a status chip: Services, Sites, Application pools, Certificates and HTTP errors, each OK, Warning, Problem, Unknown (not read) or Not checked yet; Certificates shows n/a on a server without HTTPS bindings. The health of the server is Critical when a part has a Problem, Warning when a part has a Warning, and OK otherwise; it is also at least Warning after three checks in a row that could not read IIS. A site or application pool whose automatic start is off has been stopped on purpose (IIS Manager turns it off when someone stops it): it shows as Stopped (manual) and is never alerted.

- Services: Problem when W3SVC or WAS is set to start automatically and is not running. A service set to Manual or Disabled is shown but not judged.
- Sites: Problem when a site with an HTTP or HTTPS binding and automatic start on is stopped. Sites with only FTP or net.tcp, net.pipe or msmq bindings are never judged, and the state must be known: Unknown never counts as stopped.
- Application pools: Warning when a pool with automatic start on is stopped.
- Certificates: Problem when an HTTPS binding has no usable certificate (no HTTP.sys binding, no certificate on it, not in the certificate store, or no private key) or a certificate expires in under 7 days or has expired; Warning when one expires in under 30 days or does not cover the binding's host name. Only the bindings of sites with automatic start on are judged; a wildcard certificate covers one level of name.
- HTTP errors: Warning when a site had more than 100 HTTP 5xx responses in the last 24 hours that are also more than 5% of its requests.

The list of servers

The top of the page has + Mark an IIS server and Columns. Under the introduction come a client list, a search box (server, site, binding, certificate or client) and the Show list: Show all, Only with problems, Only with certificates expiring within 30 days, Only with HTTP 5xx errors and Only with shared

configuration. The columns are Client, Server (with pending while a check waits for its result), Services, Sites (started out of the total and how many are stopped), Application pools, HTTPS bindings, Next certificate expiry (days left: green from 30 days, amber under 30, red under 7 or expired), HTTP 5xx (24 h), Health, Last check and Alerts; Operating system is available through Columns. Click a header to sort, use the row under the headers to filter, and click a row to open the server's window. The refresh button at the end of the row runs Check now, and the right-click menu offers Open detail, Check now, Open computer and Copy server name. The list refreshes itself every minute.

The server window

The title is the server's name, a link to the computer's page, followed by its online state, client, operating system, last check, health, the Monitoring list, Ask the AI and Check now, then the status chips of the five parts and a reminder that SLApp only reads. Notes say when the last check failed, how many checks failed in a row, when the state of sites and pools could not be read, or when the check reached its time limit. Findings lists every problem found in plain words, and the sites and pools stopped on purpose are named under it. Then come Summary (the two services, the configuration file and shared configuration, when the configuration last changed, logging, the counts, the next certificate expiry, the requests and HTTP 5xx of the last 24 hours and the worker processes); Sites, with each site's ID, state, automatic start, application pool, physical path, bindings and applications (up to 100 sites); Application pools, with their state, start mode, .NET CLR, pipeline, 32-bit, identity, worker processes, worker failures and recycles; Certificates, one row per HTTPS binding with its HTTP.sys endpoint, where the certificate comes from, thumbprint, subject, DNS names, issuer, self-signed, expiry, store and private key; HTTP errors (24 h), per site with the requests, HTTP 5xx, the count by status, since when it was counted and the files read; Open alerts, with Resolve for administrators; and History (N), the last 30 checks. Right-click a site, pool or certificate to copy its name, physical path, thumbprint or subject.

Check now and Ask the AI

Check now queues the daily check of the server at once, from its window, the refresh button of its row or the right-click menu, with the same limits as in the Active Directory section (at most once every 5 minutes per computer, never while a check is still waiting); the window looks again every 10 seconds while the check is in progress. It does not change the time of the daily check. When the AI assistant is available and a server is in Warning or Critical or has open alerts, a Ask the AI button appears in its row, in its window and in the right-click menu: it opens the assistant with what the last check found on that server and asks it to find the likely cause of each problem and propose the steps to fix it. The assistant is asked to use only read-only commands on the server and to explain any command that would change something, such as starting a site or recycling a pool, and wait for your confirmation.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is IIS; they appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, with the server's name in front of the message, and you can create a ticket from them like from any other alert. They are evaluated each time a check of the server is read; a part that could not be read leaves its alert as it is. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- IIS service stopped (Critical): W3SVC or WAS is set to start automatically and is not running. It closes when both run.
- IIS site stopped (Critical): a site with automatic start on is stopped. It closes when no such site is stopped.

- IIS application pool stopped (Warning): an application pool with automatic start on is stopped; the message gives its worker process failures. It closes when it runs again.
- IIS HTTPS binding without certificate (Critical): an HTTPS binding has no usable certificate, with the reason. It closes when every binding has one.
- IIS certificate expiring (Warning under 30 days, Critical under 7 days or expired): it names the certificate and its expiry date, and closes once the certificate is replaced.
- IIS certificate name mismatch (Warning): the certificate of an HTTPS binding does not cover its host name. It closes when it does.
- IIS HTTP 5xx errors (Warning): a site had more than 100 HTTP 5xx responses in 24 hours that are also more than 5% of its requests. It closes when no site is over the threshold.
- IIS check failing (Warning): three checks in a row could not read IIS. It closes as soon as a check reads it again.
- The first check read after a server starts being monitored never sends an email, not even for a Critical alert, so the first run over a server with many sites cannot send a burst; its alerts still open and show in Alerts. From then on, an alert emails the active administrators of your company who have an email address when it opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the names in it, such as the sites and certificates, stay as they were written).
- The alerts of a server close when it stops being monitored or is removed, and when the company moves to a plan without IIS. Only an administrator can close an alert by hand, with Resolve; if the problem is still there, the next check opens it again.

27. Backups

The Backups module copies folders and SQL Server databases from your clients' Windows computers to a safe place, on a schedule, and keeps a history of every copy. You describe a backup once, as a job: which computer runs it, what to back up (the sources), where to put the copies (the destinations), when to run it and how many copies to keep. From then on the agent installed on that computer does the work by itself: it builds a ZIP file for each folder and a .bak file for each database (a .bak can go into a ZIP too, and a ZIP can be protected with a password), copies the result to every destination, checks it, deletes the oldest copies according to your retention rules and reports back. SLApp Helpdesk records every run, opens an alert when a backup fails or does not run, or when a destination is running out of room, and shows how much free space each destination has. Open the module from Infrastructure › Backups in the sidebar. It has four tabs: Jobs, Calendar, Runs and Destinations. The Calendar shows every backup, past and planned, by day, week or month. The module makes, checks and tidies the copies but does not restore them: to bring data back you take the file from the destination and extract it, or restore the .bak with SQL Server (see Restoring from a backup).

> **i** To watch the backups that other software makes on your clients' computers, such as Veeam Backup & Replication, Veeam Agent for Windows, Windows Server Backup, Acronis or any program that writes to a Windows event log, use Infrastructure › Third-party backups instead, which every technician can open (see the Third-party backups section).

Who can use it and what it needs

- An administrator account. The Backups entry is hidden from other technicians and the server refuses their requests, so they can neither see the jobs nor read the destinations.
- A Windows computer whose agent is version 1.1.0 or later. Linux and macOS computers cannot run backups. In the job editor, the Computer (agent) list shows a dot in front of every computer, green

when it is online and dark when it is offline, and marks a computer that is not Windows, or whose agent is older, as not supported; you cannot choose it. Agents update themselves every 6 hours, so an older agent normally becomes usable on its own. You can save a job for an offline computer and it starts working as soon as the agent reconnects, but testing a destination needs a computer that is online (green dot). The archive options of a source (a ZIP of the SQL backup, a compression level other than Normal, a password) need agent 1.2.0 or later (see The job editor: archive options).

- A computer that is switched on at the scheduled time. The backup runs on the computer itself, not on the server: a computer that is off or asleep at that moment makes no copy (see When the computer is off). A computer that is on but cannot reach SLApp Helpdesk still runs its jobs from its saved definitions and delivers the reports when the connection returns; only a destination it cannot reach makes the backup fail.
- Sources and destinations that this computer can reach: folders on its own disks or on network shares, SQL Server instances installed on the same computer as the agent, and destinations such as a share or an FTP server. The data goes straight from the computer to the destination and never passes through SLApp Helpdesk; the server only receives the reports.
- The agent service runs as the Windows SYSTEM account, and three things follow from that. SYSTEM does not see mapped drive letters, so network locations must be written as \\server\share. SYSTEM has no user name and password of its own to present to a NAS, so a share that asks for a login needs a user and password in the destination. And SYSTEM often has no rights in SQL Server, so a SQL Server login is usually the easiest way in.

Key concepts

- Job: one scheduled backup. It belongs to one computer, has one or more sources and one or more destinations, a schedule, a retention rule and a few options. A job is either enabled or paused.
- Source: something to back up. It is either a Folder (everything under a path, optionally filtered with patterns) or a SQL Server instance with its databases. Sources belong to the job, because their paths exist on that job's computer.
- Destination: a place the copies are written to: a folder on a disk or a network share, or an FTP or FTPS server. Destinations live in a library shared by all the jobs of your company. One destination can serve several jobs and one job can write to several destinations; every source is copied to each of them.
- Run: one execution of a job. A run is started by the schedule, by Run now or as a late catch-up, and is recorded in the Runs tab with one entry, called an item, for every file the run produced and every destination (one ZIP per folder source and one .bak, or .bak.zip, per database of a SQL Server source).
- Series: the copies of one source, or of one database of a SQL Server source, at one destination. Retention counts and deletes copies series by series.

Your first backup, step by step

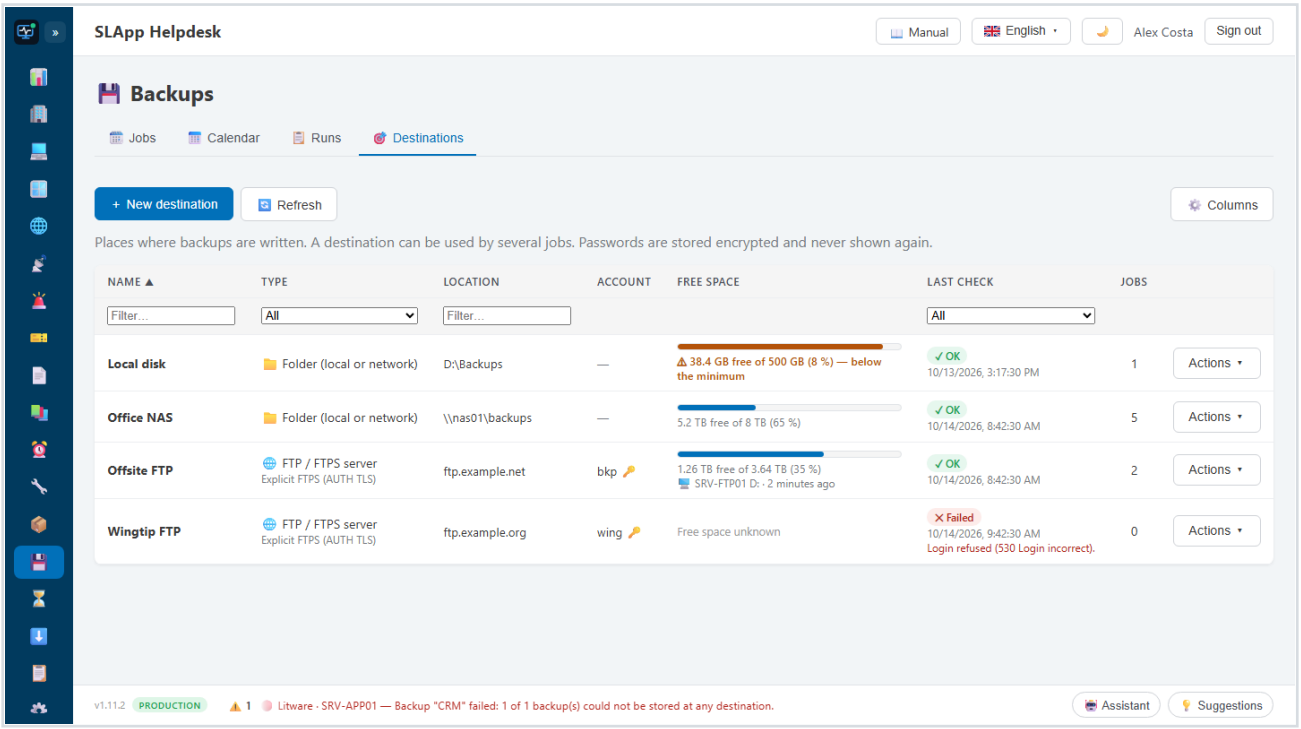
1. Open Infrastructure › Backups in the sidebar and go to the Destinations tab. Click + New destination.
2. Type a Name, for example NAS Porto. Leave Type on Folder (local or network) and write the Folder as a full network path such as \\nas\backups. If the share asks for a login, fill in User and Password under Network share credentials. Click Save. The window closes and a notice such as Destination “NAS Porto” saved. appears above the list.
3. Click Test in the destination's Actions menu (or click the row to reopen the editor and use Test there). Choose the computer that will run the backup in Run the test from, click Run test and wait. You want to see Destination reachable and the free-space bar. If you see The test failed, read the

message and fix it before you go on (see Troubleshooting).

4. Go to the Jobs tab and click + New job. Type a Name, choose the Computer (agent) and add a source with + Add folder: write the Folder to back up, for example D:\Data, and a Source name.
5. Under Destinations tick the destination you just created. Under Schedule choose Daily and a Time such as 02:00; leave the Time zone as it is. Under Retention and options leave Keep the newest (copies) at 7.
6. Click Save. The window closes, a notice such as Job “Server files” saved. appears above the list, and the job is in the list with its Next run.
7. To make the first copy now, open the job's Actions menu and click Run now. Nothing asks you to confirm: the tab answers Run requested. The agent starts it in a moment; follow it in the Runs tab. The run starts within about a minute. (A daily, weekly or monthly job that you save after its time of day has passed does not run for that day, so Run now is the way to get the first copy at once.)
8. Open the Runs tab. The run appears as Running and then becomes Completed (or Completed with warnings, or Failed). Click it to see one card for every file and destination, with the size, the duration, whether the copy was verified and the file name.
9. Open the destination in Windows Explorer. Inside the folder you gave there is a folder named after the job and, in it, one ZIP file per source. Open the ZIP once to see that it holds what you expect, and do a real test restore soon (see Restoring from a backup).

The Destinations tab

The Destinations tab is the library of places where backups are written. You create a destination once, here or from inside the job editor, and then tick it in as many jobs as you like. The top of the tab has + New destination, Refresh and Columns, and a line under them reminds you that passwords are stored encrypted and never shown again. Until you create the first one the tab says No destinations yet.



The Destinations tab: each destination with its type, location, account, free-space bar (read from SRV-FTP01 D: for Offsite FTP), last check and the number of jobs that use it.

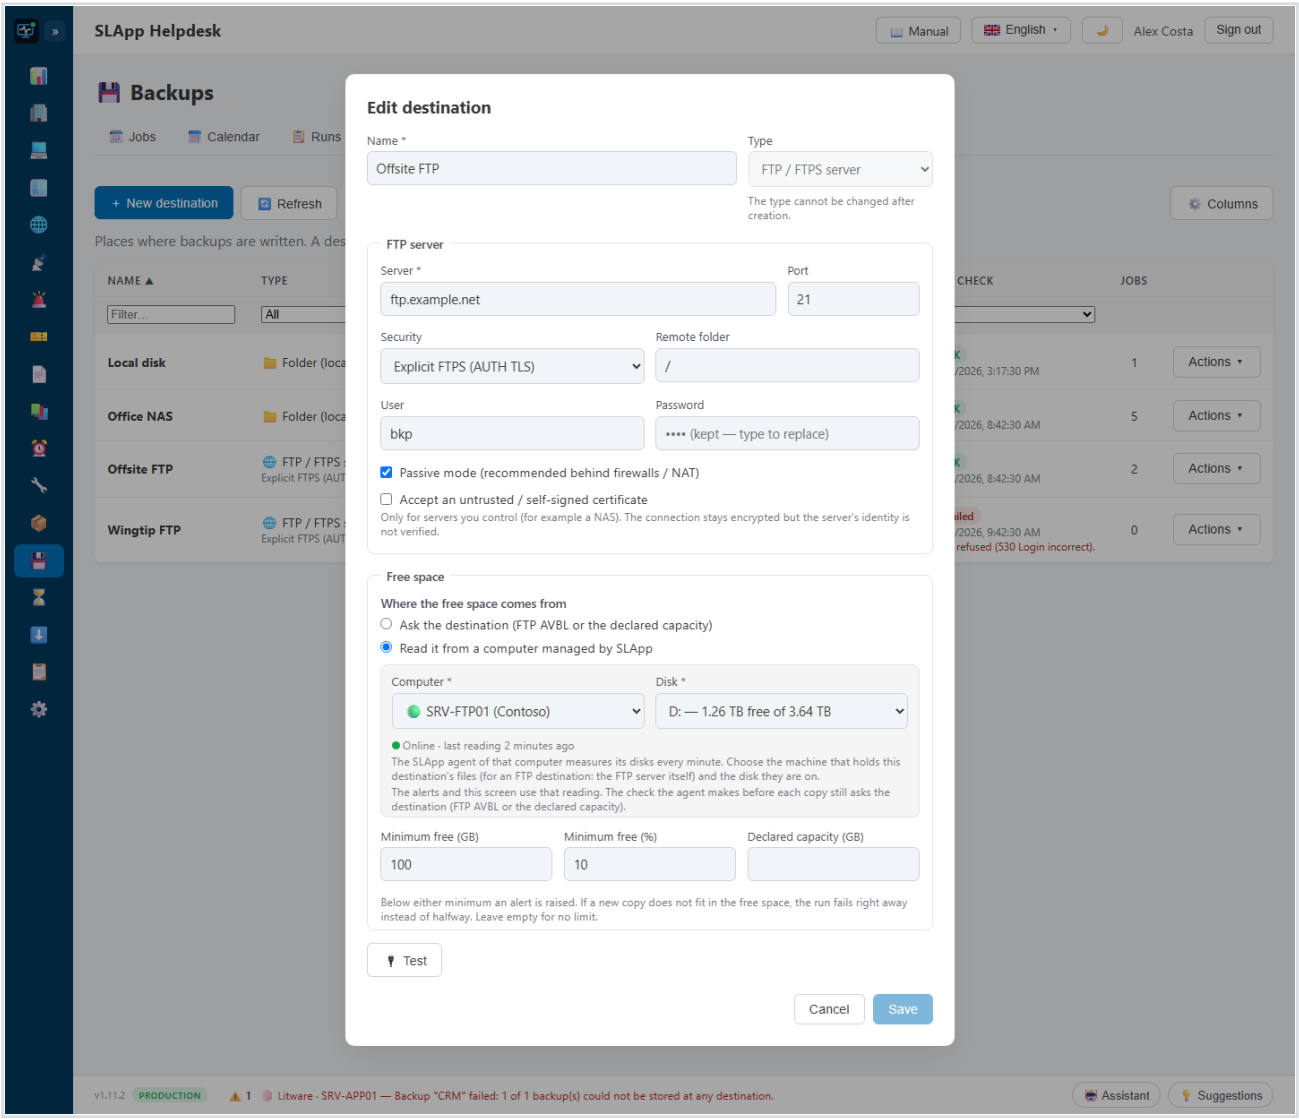
- Name.
- Type: Folder (local or network) or FTP / FTPS server. For FTP the security mode is written underneath: None (plain FTP) or Explicit FTPS (AUTH TLS).

- Location: the folder path, or for FTP the server (with the port when it is not 21) and the remote folder.
- Account: the user name, with a small key beside it when a password is stored.
- Free space: the free-space bar described under Testing a destination and, for a destination that reads its free space from a computer, a line under it such as SRV-FTP01 D: · 2 minutes ago (see Reading the free space from a managed computer).
- Last check: OK or Failed, with the date and time of the last test or of the last backup that checked the destination, and the error underneath when it failed. It says Never checked when nothing has used the destination yet.
- Jobs: how many jobs use the destination.
- Minimum free and Declared capacity (GB) are available through Columns and hidden at first.

The Columns button opens the List columns window, where you tick the columns to show and drag the rows into the order you prefer. Your choice is saved for your user, separately in the Jobs, Runs and Destinations tabs (the Calendar has no Columns button); the filters under the headers and the sort order are not saved. A row of filters under the headers narrows the list by Name, Type, Location and Last check, and a click on a header sorts by that column. Click a row to edit the destination. The Actions menu at the end of the row, or a right-click, offers Test, Edit and Delete; the right-click menu can also copy the name, the location and the last error. Delete asks for a confirmation and then removes the destination from the library; the list shows Destination “NAS Porto” deleted. If jobs still use the destination, the confirmation says it is used by N job(s) and that the server refuses to delete a destination that an active job uses, and asks Try anyway? The server then refuses and names the jobs (up to five, then several). Every job that has not been deleted counts, paused jobs included, so open those jobs first and untick the destination. Deleting a destination never deletes the files already written there.

Creating or editing a destination

+ New destination opens the destination editor, and clicking a row opens the same editor for an existing destination. Every destination has a Name of up to 120 characters and a Type, which is either Folder (local or network) or FTP / FTPS server. The Type cannot be changed after the destination is saved; create a new destination instead. Save stays disabled until the form is valid. Click Save and the window closes, with a notice such as Destination “NAS Porto” saved. above the list. Cancel closes the editor at once, without asking, even if you have typed changes.



The destination editor for an FTP / FTPS server: connection, security, the free space read from a disk of a managed computer, the free-space limits and the Test button.

- Folder destination, Folder: a full path, either a drive letter such as E:\Backups or a network share such as \\nas\backups. A drive letter is a disk of the computer that runs the job: two jobs on two computers that both use E:\Backups write to two different disks. Mapped drive letters such as Z: do not work, because the agent runs as SYSTEM and does not see them; always use the \ \server\share form. The folder is created if it does not exist, and so is the job's own sub-folder inside it.
- Folder destination, Network share credentials (User and Password): only needed for a share that asks for a login, such as a NAS or a workgroup computer. Write the user as DOMAIN\user or as the local user name of the NAS, and make sure the account may create and write files. Leave both empty to use the computer's own access. Two destinations on the same share must use the same user name.
- FTP destination, Server and Port: the host name or IP address only, without ftp:// and without a folder. The port is 21 unless you change it.
- FTP destination, Security: a new destination starts on None (plain FTP), which sends the password and the data unencrypted; the editor warns you with Plain FTP sends the password and the data unencrypted. Prefer FTPS. Change it to Explicit FTPS (AUTH TLS), which opens a normal connection and upgrades it to TLS and is the recommended choice. SFTP, which is a different protocol that runs over SSH, and implicit FTPS (usually port 990) are not supported.

- FTP destination, Remote folder: where on the server the job folders are created. It must start with /; it is / by default. User and Password complete the login; with User empty the agent logs in as anonymous.
- FTP destination, Passive mode (recommended behind firewalls / NAT): ticked by default. In passive mode the agent opens every connection to the server, which works through most firewalls and routers. Untick it only for a server that requires active mode.
- FTP destination, Accept an untrusted / self-signed certificate: unticked by default and shown only with Explicit FTPS. The connection stays encrypted but the identity of the server is not verified, so use it only for servers you control, such as a NAS. Without it the certificate must have been issued by a trusted authority for the name you typed in Server.

Passwords are write-only. Once a destination is saved nobody can read its password back, you included: the field shows a placeholder saying the password is kept, leaving it empty keeps the stored password and typing a new one replaces it. A password is only kept together with a user, so if you empty the User field, saving removes the stored password and the editor says so. The server stores passwords encrypted and hands them over HTTPS only to the agents of your own company that need them: the computer that runs a job that uses them, or the computer you choose to run a test. The commands the app sends to the agent carry only identifiers, and the agent removes passwords from the error texts it reports. On the computer, the agent keeps its copy of the job definitions in a file encrypted with Windows data protection, in a folder that only SYSTEM and administrators can open, so scheduled jobs keep running when SLApp Helpdesk cannot be reached.

The last part of the editor, Free space, starts with Where the free space comes from. Ask the destination (FTP AVBL or the declared capacity) is the choice of a new destination and works as this section describes; Read it from a computer managed by SLApp takes the figures from a disk of a computer that has an SLApp agent (see Reading the free space from a managed computer). Minimum free (GB) and Minimum free (%) raise a Backup destination space warning when the free space falls below either value; leave a field empty for no limit. The percentage needs the total size of the disk or share, which the agent measures for folders; for an FTP server it works only when the free space is read from a computer or when you also fill in Declared capacity (GB). Declared capacity (GB) is meant for FTP servers that do not report their free space: the agent then calculates the free space as the capacity you declared minus the size of everything stored in the remote folder, so keep that folder for backups only. It must be greater than zero. Whatever you set, before copying each file the agent also checks that the destination has room for that file, and if the free space is smaller than the file the copy fails at once, with a clear message, instead of failing halfway. That check always asks the destination itself, also when the free space shown is read from a computer.

Testing a destination and reading its free space

The Test button checks a saved destination the way a backup would use it. It is in the destination editor and in the Actions menu of the row; in the editor it is disabled while there are unsaved changes, because the test uses the saved settings. Click it, choose a computer in Run the test from and click Run test. The wait can be up to 90 seconds, and the window says so while it works. A destination that is too slow to answer is reported by the agent itself after 75 seconds as The destination did not answer within 75 seconds. The test is carried out by the agent of the computer you choose, not by the server: a share, a NAS or an FTP server in a client's office can only be reached from that office, and what matters is what the computer that will run the backup can see, with its network, name resolution, firewall and account. So choose the computer that runs the job. The window opens with a computer already selected: the one you used last in this browser, otherwise the first one that is online and supports backups. Computers that are offline or not supported are listed but cannot be chosen. If none qualifies the window says No online computer with backup support is available to run the test.

For a folder or share the agent connects with the credentials, creates the folder if needed, writes a

small temporary file of 4 KB, reads it back, deletes it, and reads the free space. For an FTP server it connects, with TLS when Explicit FTPS is chosen, creates the remote folder if needed, uploads a small temporary file the way a real backup does (under a name ending in .part, then renamed), reads it back, deletes it, and asks the server for its free space. A good result reads Destination reachable, with the free-space bar and a line saying what was tested. A failure reads The test failed with the reason in a box that has a Copy button. If the agent of that computer is not connected at the moment, or does not answer within 90 seconds, you get a message saying so instead. Every answer also updates the Last check and Free space columns (for a destination that reads its free space from a computer, the Free space column goes on showing that computer's figures while they are fresh).

The free-space bar, in the Free space column and in the test result, shows how much of the disk is used, with a line such as 120 GB free of 500 GB (24 %). Next to each destination in the job editor it is a short bar with a shorter line, 120 GB free (24 %), and the full line in its tooltip. When the free space is known but the total size is not, only the line 120 GB free is shown, without a bar. It changes to a warning look, with the words below the minimum, when the free space is under Minimum free (GB) or Minimum free (%), and to a critical look, with no free space, when nothing is left. The figures are not live: they are what the agent saw at the last test or in the check made before the last copy to the destination, unless the destination reads its free space from a computer (see Reading the free space from a managed computer). In the cases below the Free space column, the test result and the job editor say Free space unknown and no bar is drawn; the module then never raises a space alert for that destination and never refuses a copy because of it.

- Nothing has checked the destination yet.
- The destination is an FTP server that does not offer the AVBL command, which is most of them, IIS FTP included, Declared capacity (GB) is empty and the free space is not read from a computer. The Destinations tab then adds Free space unknown: in the destination, read it from a computer managed by SLApp or declare the capacity.
- The destination is an FTP server with a declared capacity, but the remote folder is too big to measure quickly: the agent stops counting after 20,000 entries, ten levels of sub-folders or 45 seconds.
- The destination is a folder or share where Windows cannot report the space for a reason other than the path being unreachable. Some NAS devices answer the question with access denied. The destination still works and the test still says Destination reachable, only without numbers.

Reading the free space from a managed computer

Most FTP servers, IIS FTP among them, never tell the agent how much room they have, and a declared capacity minus what the remote folder holds is only an estimate when other data shares the disk. When the machine that holds the destination's files is itself a computer with an SLApp agent, that agent's own reading of its disks can be used instead, for a folder destination as well as for an FTP one. In the destination editor, under Where the free space comes from, choose Read it from a computer managed by SLApp, then a Computer and a Disk; both are required. Computer lists the active computers of your company, each with its online dot, its name and its client. Disk lists the disks of that computer's latest reading with their free and total space, for example D: — 1.26 TB free of 3.64 TB: drive letters for Windows, mount paths such as /srv for Linux and macOS. Under the two lists the editor says whether the computer is online and how old its last reading is, for example Online · last reading 2 minutes ago, or no disk reading received yet. Choose the machine that holds the files (for an FTP destination, the FTP server itself) and the disk they are on, and click Save. To go back, choose Ask the destination (FTP AVBL or the declared capacity).

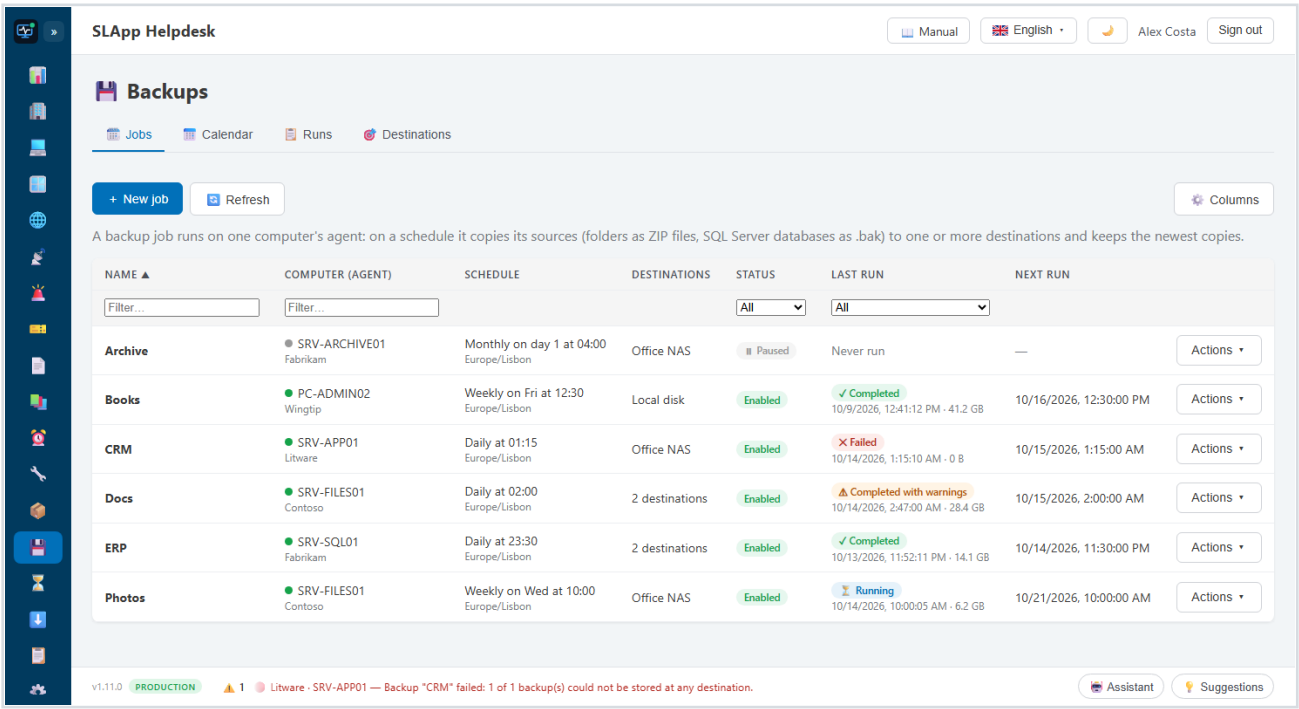
The agent of that computer sends a reading of its disks every 30 to 60 seconds with the default settings (with somebody signed in it follows Sending to the server (seconds), otherwise Heartbeat interval

(seconds), in Settings › General › Monitoring), and a reading is used while it is at most 30 minutes old. The Free space column of the Destinations tab and the destination list of the job editor then show that computer's figures, with a line under them such as SRV-FTP01 D: · 2 minutes ago. When the reading cannot be used the line starts with a warning sign and says why: SRV-FTP01 D: · last reading 3 hours ago when it is older (the computer is off or offline, for instance), D: is not in the latest reading of SRV-FTP01 when the disk is gone or was renamed, or No free-space reading of D: from SRV-FTP01; a computer that was removed since reads SRV-FTP01 (removed). In those cases the figures above the line come from the destination's own check again and, when there are any, a second line says so, for example Value shown: the destination's own check, 10 minutes ago.

The computer's reading is what the Destinations tab, the job editor and the Backup destination space alert use (see Alerts and e-mails). It does not change the check the agent makes before each copy, which still asks the destination itself (FTP AVBL or the declared capacity) and so cannot see the disk of an FTP server. The Test window of such a destination says so under its result: the free space is read from SRV-FTP01 D:, and the value the window shows is what the destination itself answered. Setting, changing or removing the link closes the open Backup destination space alerts of the destination, which are then judged again from the new source. A mount path longer than 16 characters is listed as name too long to be linked, and a Windows volume mounted in a folder rather than on a drive letter is not listed.

The Jobs tab

The Jobs tab lists every backup job of your company. The top of the tab has + New job, Refresh and Columns, and a line under them explains what a job is. Until you create the first job the tab says No backup jobs yet. Click a row to open the job in the editor.



The Jobs tab: each job with its computer, schedule, destinations, status, last run and next run.

- Name.
- Computer (agent): the computer, whose name is a link that opens the computer's page, with a dot that is green when it is online and dark when it is offline, the client underneath, and a not supported badge when the computer can no longer run backups (an agent older than 1.1.0, for example).
- Schedule: a summary such as Daily at 02:00, with the job's time zone underneath.
- Destinations: the destination's name, or N destinations when there are several (hover to read their

names).

- Status: Enabled or Paused.
- Last run: the status of the most recent run, which can also be Missed, with its date and size, or Never run.
- Next run: when the job will run next, in your own time zone, with the job's own clock time underneath when its time zone is different. A dash appears while the job is paused or when its computer has been removed. For an Every N minutes job it is an estimate (see the schedule).
- Through Columns you can also show Sources (how many), Retention, Time zone, Destination folder (the job's fixed folder name) and Revision.

The filters under the headers narrow the list by Name, Computer, Status and Last run. While a job is running, and for two minutes after you click Run now, the list refreshes itself every 10 seconds. The Actions menu at the end of each row, or a right-click, offers Run now, Pause or Resume, Edit, View runs and Delete.

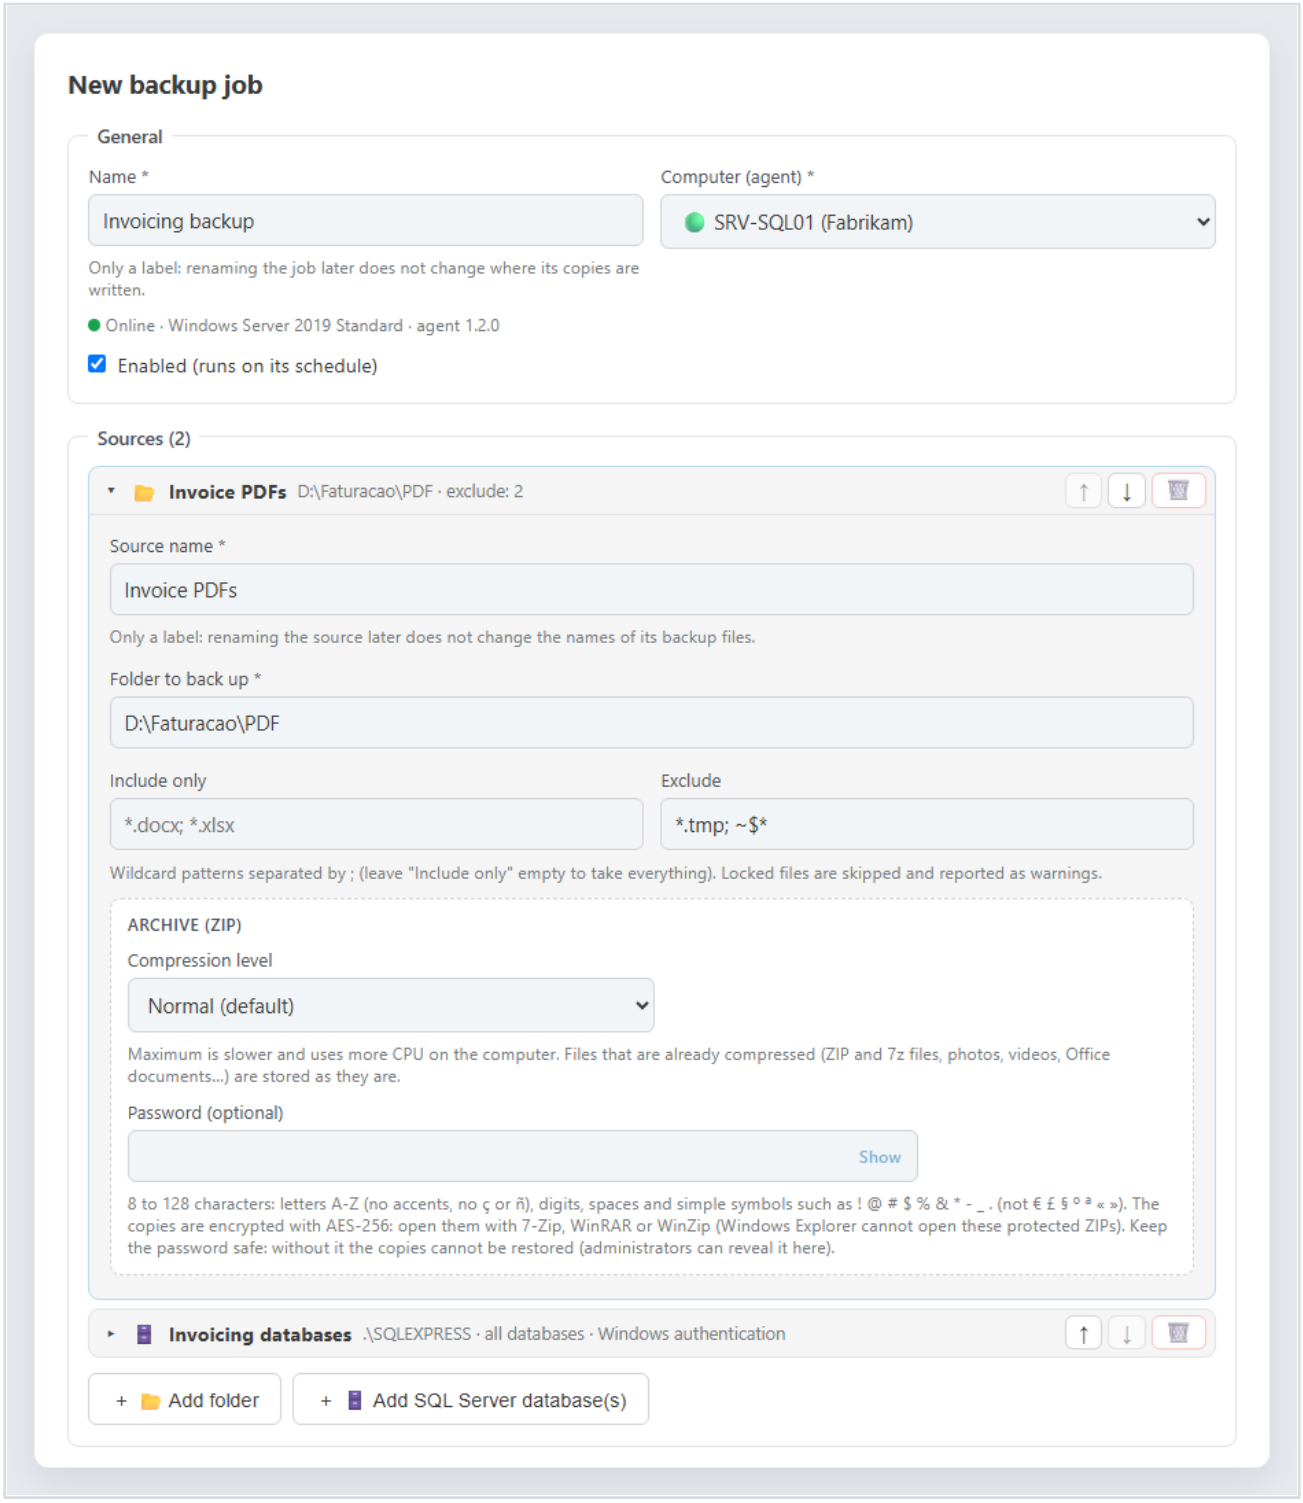
- Run now asks the agent to run the job as soon as possible with its saved settings, and asks for no confirmation. The Jobs tab (or the editor) answers “Run requested. The agent starts it in a moment; follow it in the Runs tab.” or, when the computer is offline, “Run requested, but the computer is offline: it starts when the agent reconnects.” The agent looks for requests every minute, and the app nudges it at once when it is connected, so the run starts within about a minute. A paused job answers “This job is paused. Resume it before running it.”, and a job that is already running answers “A backup of this job is already running.” If the computer is offline the request waits and the run starts when the agent reconnects, as long as that happens within 24 hours. A run started this way shows as Manual in the Runs tab.
- Pause stops the schedule without losing anything and Resume turns it back on. They are one item of the job's Actions menu (and of its right-click menu), whose label changes between Pause and Resume, because the editor never changes this: for an existing job it only tells you whether the job is enabled or paused. The list answers with a notice such as Job “Server files” paused. or Job “Server files” resumed. A paused job does not run, does not try to catch up what it missed while it was paused, and raises no missed-backup alerts. Pause does not interrupt a run that is already in progress: it carries on to the end. Use Cancel run in the Runs tab to stop it.
- View runs opens the Runs tab already filtered on that job.
- Delete asks for a confirmation (Delete backup job) and removes the job from the list, so that it never runs again; a run already in progress is not stopped, so cancel it first in the Runs tab. Its run history is kept: the runs stay in the Runs tab, but the Job filter there no longer lists the deleted job, so clear that filter to see them. Its Backup failed and Backup missed alerts are closed, and the files already written to the destinations are not deleted. Nothing in the module ever deletes copies except the retention rule of a running job.

> ℹ Each computer's detail page also has a Backups tab with that computer's jobs and runs, where a new job already starts on that computer (see The Backups tab in the Computers section).

The job editor: General

+ New job opens the job editor, and clicking a row opens the same editor for an existing job. It has five sections: General, Sources, Destinations, Schedule, and Retention and options. For an existing job Save stays disabled until you have changed something. While something is still missing or wrong, Save can be pressed but does not save: when the problem is in a source, it opens that source with the cursor in the field to fix. Just above the buttons the editor lists what is still missing after Cannot save yet: it names the first three reasons and adds (+N) for the rest. Save closes the window, and the Jobs tab

shows a notice such as Job “Server files” saved. above the list. If you click Cancel with unsaved changes, the editor asks Discard changes?, telling you that your edits to this job have not been saved. If somebody else saved the job after you opened it, Save is refused with “Somebody else changed this job after you opened it, so your changes were not saved. Reload the job (your unsaved changes are lost) and make them again.” The Reload the job button then loads the current version into the editor, and you have to make your changes again. For an existing job the editor also has Run now and View runs buttons; Run now answers as described under the Jobs tab, and is disabled while there are unsaved changes, with the hint Save your changes first: a run uses the saved settings. The picture below shows the first half of the editor, General and Sources; the second half, Destinations, Schedule, and Retention and options, is shown further down, under Destinations. General holds these fields.



The first half of the job editor: General, and Sources with the folder source Invoice PDFs open on its fields and its Archive (ZIP)

options, and the SQL Server source Invoicing databases closed to one row.

- Name: up to 120 characters. It is only a label: renaming a job later does not change where its copies are written.
- Computer (agent): the computer that runs the job. Each option shows the dot, the name and the client. Below the field the editor shows whether the computer is online, its operating system and its agent version. You can change the computer of an existing job: the job keeps its settings and keeps writing into the same folder at the destinations, but the paths of its sources must exist on the new computer. When a source uses archive options that the computer's agent is too old for, the message about it appears under this field (see The job editor: archive options).
- Destination folder: read-only, shown for existing jobs. It is the folder, inside every destination, that holds this job's copies (see Where the files land).
- Enabled (runs on its schedule): shown only when you create a job. Afterwards the editor just says whether the job is enabled or paused, and you use Pause and Resume in the job's Actions menu in the list.

The job editor: Sources

A job needs at least one source and can have up to 50; the title of the section counts them, as in Sources (2). Add them with + Add folder and + Add SQL Server database(s). The sources form a list with one row each, processed from top to bottom. A row shows the icon of the type, the Source name (Source 1, Source 2 and so on until you type one) and a one-line summary: the folder, with include: N and exclude: N when patterns are set, or the instance, all databases or N databases, and Windows authentication or SQL login. After it come the archive options that differ from the default (ZIP, the compression level, and a lock with the word password), then three small buttons whose tooltips read Move up, Move down and Remove source. Click the row (its tooltip says Show settings) to open the fields of that source under it, and click it again to close them. Only one source is open at a time, so opening one closes the other; a new source opens at once with the cursor in its first empty field, and a saved job with a single source opens on it. A closed source that keeps the job from being saved shows a badge such as 2 problems. With more than 8 sources a filter box, Filter by name, folder or instance…, appears above the list: the open source and every source with a problem stay listed whatever you type, Move up and Move down move a source among the ones listed, and No source matches the filter. says when none does. Remove source takes the source out at once, except a source with a saved archive password: Remove this source? first reminds you that its copies stay encrypted with that password and that retention never deletes the copies of a removed source, so reveal the password and note it down before you confirm. Every source has a Source name of up to 120 characters, which is only a label, and, once saved, a read-only File name prefix that is used in the names of its files. A Folder source has these fields.

- Folder to back up: a full path, with a drive letter (D:\Data) or as a network path (\\server\share\folder). Everything under it is backed up. The agent reads it as SYSTEM, so a network folder must be readable by the computer's own account. If you leave Source name empty, the editor suggests the last part of the path.
- Include only: wildcard patterns separated by semicolons, for example *.docx; *.xlsx. Leave it empty to take everything. It applies to files; folders are always entered.
- Exclude: patterns separated by semicolons, for example *.tmp; ~$*. It applies to files and to folders.

The patterns use * and ? as wildcards and ignore capitals. A pattern without a slash is compared with the name of every file and folder at any depth, so node_modules leaves out every folder of that name. A pattern with a slash is compared with the path relative to the source, as in Docs/Old/*.bak, and ** crosses folders. A full path such as D:\Data\Temp excludes exactly that file or folder. Each pattern can have up to 200 characters and a field up to 1,000 in all.

The agent writes one ZIP per folder source per run. The ZIP holds the files with their names, folder structure and modified dates; it does not hold NTFS permissions. While it works, the agent follows these rules.

- It always leaves out $Recycle.Bin, System Volume Information, pagefile.sys, hiberfil.sys and swapfile.sys.
- It does not follow symbolic links, junctions and mount points, which avoids endless loops. Each link that leads outside the source is a warning that names its target, so you know what is not in the backup.
- It skips OneDrive online-only files, because reading them would download them, and counts them in a warning.
- A file that another program has open is still read when Windows allows sharing. A file that cannot be read (locked exclusively, access denied, or deleted during the backup) is skipped with a warning that names it; the first 25 problems are listed and the rest summarised. A file that is being written to while it is copied is taken as it is at that moment and can come out inconsistent; from agent 1.2.14, when its size or modification time changes during the copy, its item carries a warning that names it (“This file changed while it was being copied”). A file that is only open, with nobody writing to it, is not flagged. Only when part of a file cannot be read does its item carry a warning, “Only part of this file could be read”.
- There is no snapshot of the disk. Files that are open for writing, such as the files of a database, can be copied in an inconsistent state. Back databases up with a SQL Server source, not by copying their files.
- A source that ends up with no file at all, because the folder is empty, the patterns match nothing or the folder cannot be read, is a failure with the message “nothing to back up”, not a copy with warnings. Nothing is uploaded and the earlier copies are kept.

A SQL Server source has these fields.

- SQL Server instance: an instance installed on the same computer as the agent, for example .\SQLEXPRESS, (local) or localhost. An instance on another computer is refused.
- Authentication: Windows (the agent's account) or SQL Server login. The agent runs as SYSTEM, which often has no rights in SQL Server, so if the backup fails with a permission error, use a SQL Server login. With a SQL Server login the simple answer is the sysadmin role; the minimum is the right to back up the databases (db_backupoperator or db_owner) together with the right to see the instance's default backup folder (sysadmin, or VIEW ANY DEFINITION on versions before SQL Server 2019).
- SQL login and SQL password: shown with SQL Server login. The password is write-only, like the passwords of the destinations: once it is saved, nobody can read it back.
- Databases: All user databases, or Only these databases with one name per line, spelled exactly as SQL Server has it (a name may contain spaces, commas or semicolons; at most 200 databases per source and 128 characters per name, beyond which the editor blocks Save and tells you to choose All user databases instead). A name is matched exactly first and then ignoring capitals. A listed database that does not exist makes the whole source fail before anything is backed up. All user databases leaves out master, model, msdb and tempdb; to back up master, model or msdb, choose Only these databases and list them. tempdb and database snapshots cannot be backed up. With All user databases, a database that is not online is handled by its state: SUSPECT, RECOVERY_PENDING and EMERGENCY databases are damaged, so they count as not backed up and fail the run; OFFLINE, RECOVERING and COPYING databases are skipped with a warning; RESTORING databases are skipped silently. A database that you list by name and that is not online always fails.

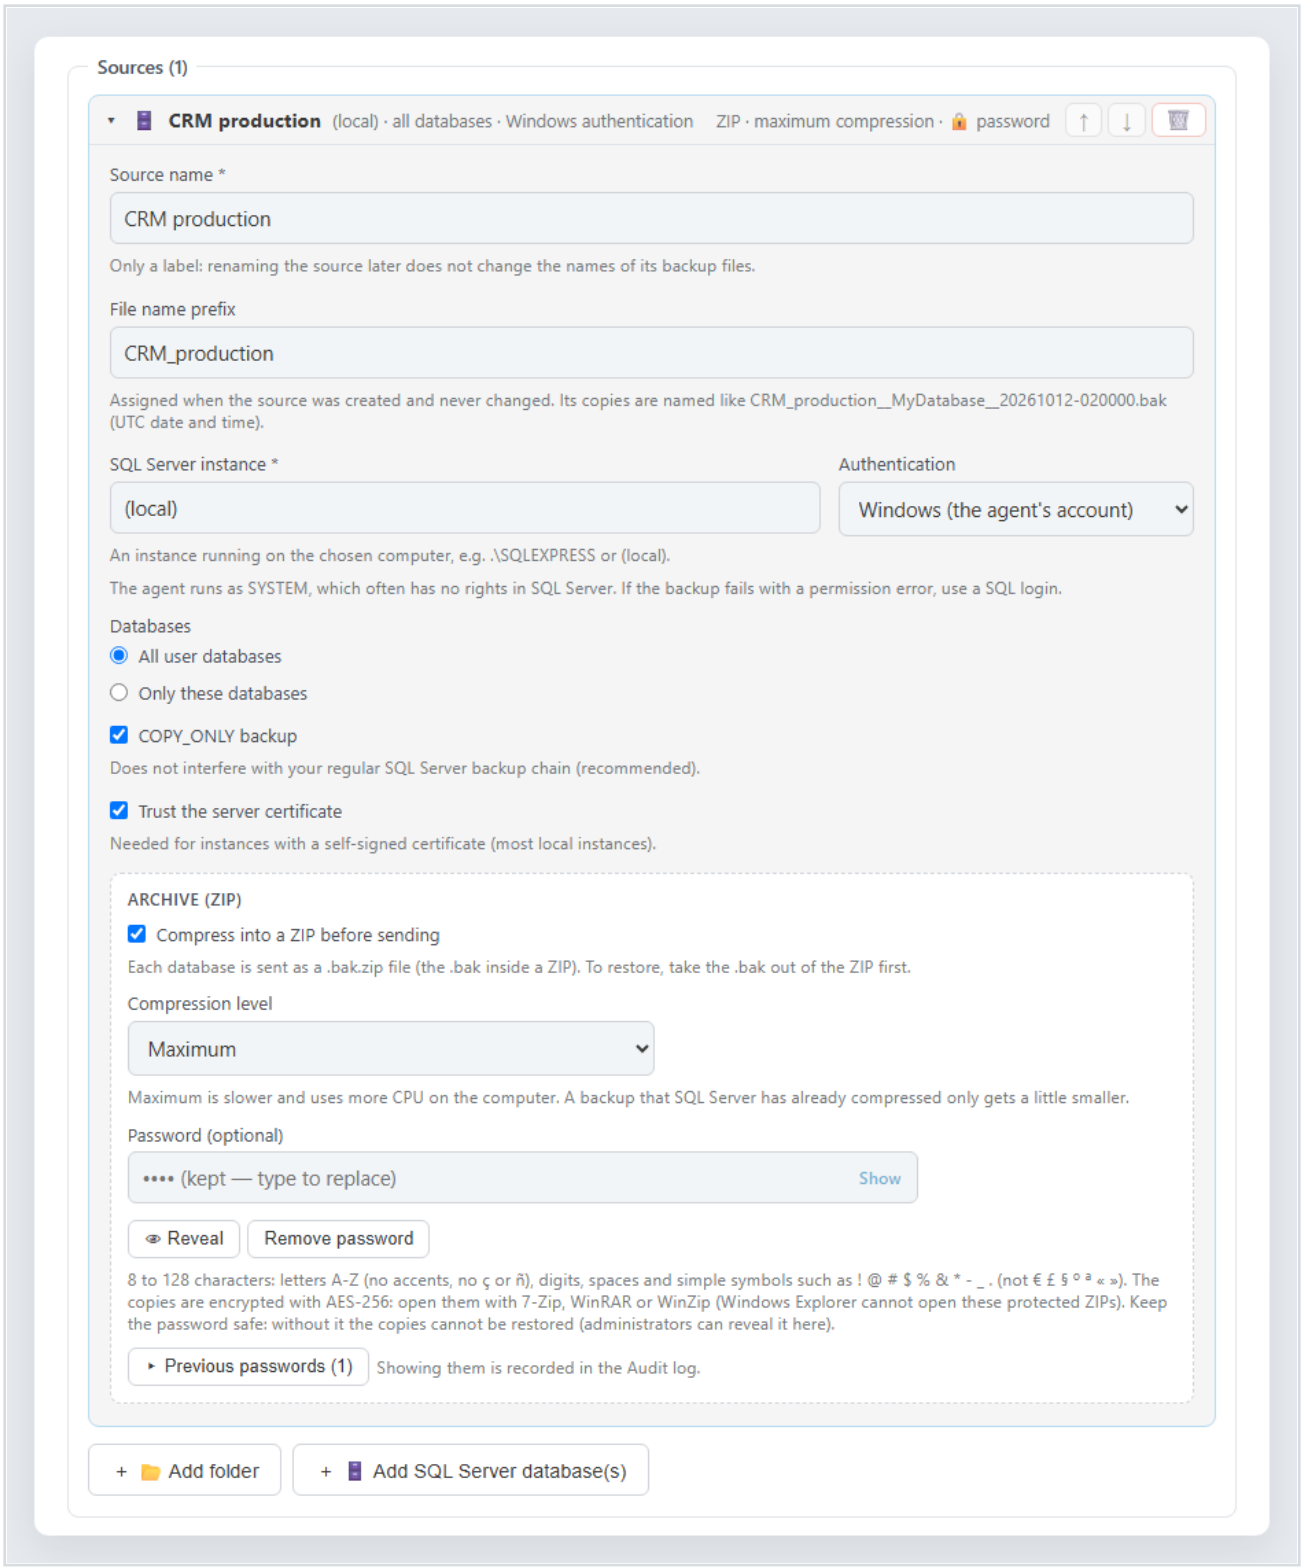
- COPY_ONLY backup: ticked by default. A copy-only backup is a complete copy of the database that SQL Server does not count in its backup sequence, so it does not become the base of the differential backups of the customer's own backup plan. Leave it ticked unless this job is meant to be that plan's full backup.
- Trust the server certificate: ticked by default. It is needed for instances with a self-signed certificate, which is the case for most local instances.

For each database the agent asks SQL Server to run BACKUP DATABASE with CHECKSUM, and with compression when the edition has it; Express and Web editions back up without compression, which is normal and not a warning. SQL Server itself writes the .bak into a work folder called SLApp inside the instance's default backup folder, so the account of the SQL Server service needs permission to write there. The agent then reads the file, checks it when Verify each copy after writing it is on, puts it into a ZIP when the source asks for it (see The job editor: archive options), uploads it and deletes it. Each database becomes one .bak file, or one .bak.zip file. The module takes full backups only, with no transaction log backups, so a restore returns the database to the moment of the backup and not to a later point in time. Cancelling a SQL Server backup that is running needs the ALTER ANY CONNECTION permission, which sysadmin has; without it the statement finishes first and its file is then deleted.

The agent also has to be able to read and delete the file that SQL Server writes. If SYSTEM and Administrators have no read and delete permission on the SLApp folder inside the instance backup folder, the database fails with a message asking you to grant it; the same happens when the instance backup folder is not a local path of the same computer. The .bak files of all the databases of a source are written first, then uploaded, and deleted only when the source is finished, so the volume of the instance backup folder needs room for all of them at once. Before each database the agent checks that room and refuses it with Not enough free disk space on the SQL Server backup volume when it does not fit: it expects at least 30 % of the used size of the data with compression, the whole used size without it, plus 64 MB. Leftover .bak files of an interrupted run in that SLApp folder are removed by a later SQL Server backup after 72 hours.

The job editor: archive options

The fields of every source end with an Archive (ZIP) group, which decides how the copies of that source are packed: how hard they are compressed and, if you want, a password. Nothing changes for a source until you change these options, and they need agent 1.2.0 or later on the computer that runs the job (see below). A Folder source is always a ZIP, so the group shows the level and the password at once. A SQL Server source starts with Compress into a ZIP before sending unticked, and its .bak files are sent as they are; tick it and each database is sent as a .bak.zip file, the .bak inside a ZIP, and the level and the password appear.



A SQL Server source open in the job editor, with its Archive (ZIP) options: Compress into a ZIP before sending ticked, the Maximum compression level, a saved password with Reveal and Remove password, and Previous passwords (1).

- Compress into a ZIP before sending (SQL Server sources only): the agent puts each .bak into a ZIP once SQL Server has written it and, with Verify each copy after writing it, once RESTORE VERIFYONLY has checked it. To restore such a copy you take the .bak out of the ZIP first. Unticking the box of a source that has a saved password removes that password when you save, and the editor says so; tick it again to keep the password. Unticking it also drops a password you typed and had not saved yet; a Compression level you chose comes back if you tick the box again before saving, but a source saved without the ZIP is saved with Normal.

- Compression level: None (store only), Fast, Normal (default) or Maximum. Normal is what folder ZIPs have always used. Maximum is slower and uses more processor time on the computer, for little gain on data that is already compressed: in a folder ZIP, archives such as ZIP, 7z and RAR, photos, audio and video files and the newer Office and OpenDocument files (.docx, .xlsx, .pptx, .odt, .ods, .odp) are stored as they are whatever the level, and a backup that SQL Server has already compressed only gets a little smaller.
- Password (optional): 8 to 128 characters, using only the letters A to Z without accents, digits, spaces and simple symbols such as ! @ # $ % & * - _ . (not € £ § º ª « »), so that every unzip tool reads it the same way. With a password every file in the ZIP is encrypted with AES-256, and the copies open with 7-Zip, WinRAR or WinZip but not with Windows Explorer. Show, at the right of the field, shows what you are typing. A saved password is not shown in the field, which reads •••• (kept — type to replace): leave it empty to keep the password, or type a new one to replace it when you save.
- Reveal shows the saved password, with Copy, and turns into Hide saved password; every reveal is recorded in the Audit log. Remove password removes the saved password when you save, so that the next copies are not encrypted, and turns into Undo removal until then. Both appear only for a source that has a saved password, in a job that has been saved. Replacing or removing a password does not touch the copies already made: they keep the old one, which is why the editor asks you to reveal it and note it down first, and why the source keeps it among its previous passwords.
- Previous passwords (n), with the note Showing them is recorded in the Audit log., appears once the source has had a password replaced or removed. Click it to list them, newest first, each with Replaced on or Removed on and the date, up to job revision N (the last revision of the job that still had it), the name the source had then when it was different, and the password with Copy. Each one opens the copies of the runs that used those revisions; Job revision is in the run detail, where Password of this copy finds the right one for you (see The password of a copy). Click it again, close the source or close the editor, and the passwords are hidden again; showing them again is a new entry in the Audit log.

A ZIP without a password is written exactly as before, so any tool opens it, Windows Explorer included. With a password the agent uses the AES-256 encryption of the WinZip format. A ZIP with a password, and every ZIP of a SQL backup, is opened again once it is written, with its password, and every file in it is read back and compared with what went in; a ZIP that fails this check is not sent. The ZIP of a SQL backup is built in the staging folder, one database at a time, while its .bak still exists, so that folder needs room for about the size of the largest .bak; the .bak is deleted once its ZIP has passed the check. The copy is named like the .bak with .zip added, for example ERP__SalesDB__20261012-020000.bak.zip, and the .bak and .bak.zip copies of a database form one series, so switching the option on or off does not start the retention count again. The password is kept encrypted on the server and in the agent's copy of the job definitions, and never appears in a log, a run report or an error text.

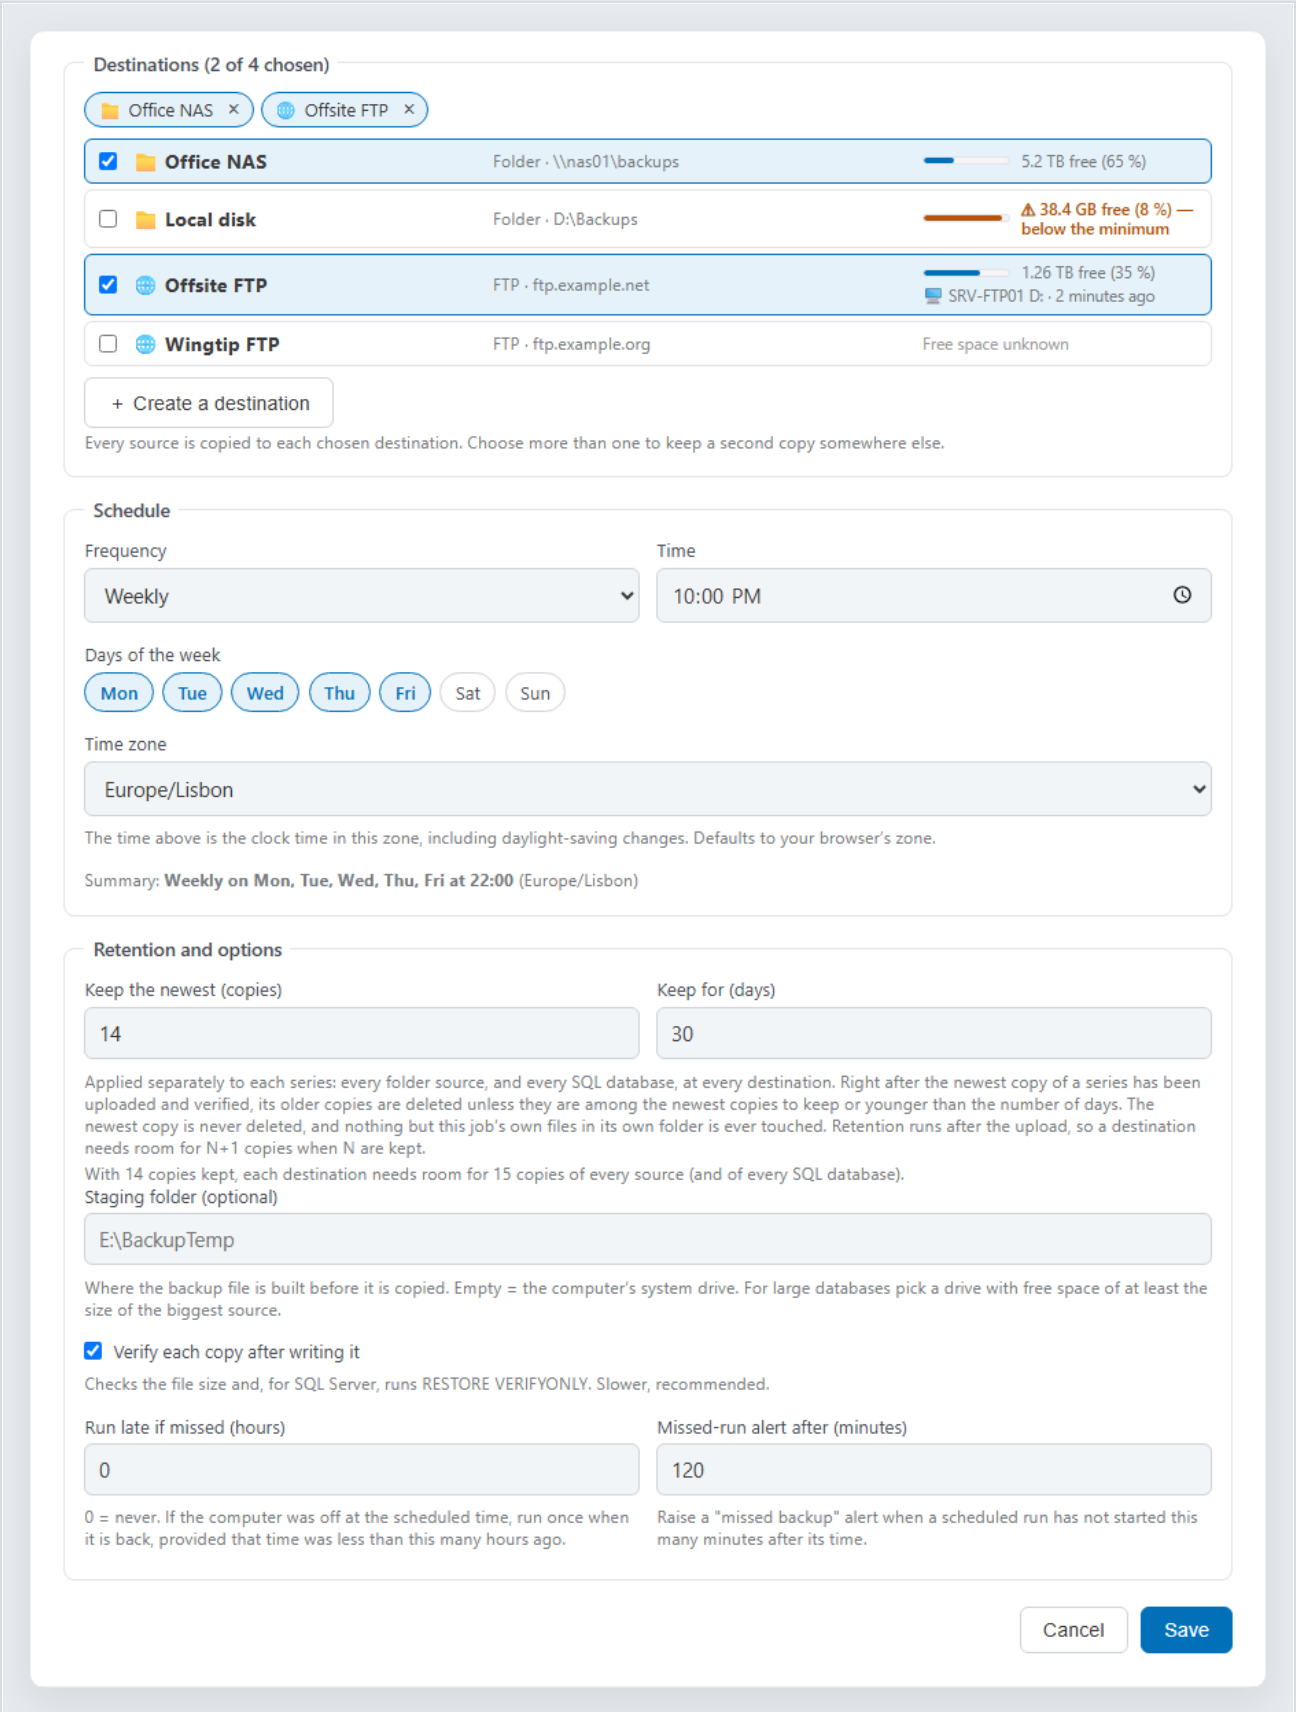
The archive options need agent 1.2.0 or later on the computer that runs the job. When a source uses one of them, that is a ZIP of a SQL backup, a level other than Normal or a password, and the computer has not reported agent 1.2.0 or later, the editor shows “This computer's agent is too old for the archive options (ZIP, compression level, password): it needs agent 1.2.0 or later.”, followed by the advice to use Force update, and Save does not save; the server refuses such a save as well. The message appears under Computer (agent), in the open source that uses the options and after Cannot save yet. Agents look for updates by themselves 2 minutes after their service starts and then every 6 hours; to update one at once, use Force update in the Remote actions tab of the computer's page. While a backup is running on that computer, Force update answers “A backup is running - the update is postponed until it finishes (at most 6 h).” and does not try again by itself, so use it again once the backup has ended. When the agent has restarted, click Save again: the editor reads the versions again and saves at once if the agent is now recent enough, or adds Checked again at …: the computer still does not report

agent 1.2.0 or later. If the agent that runs the backups turns out to be older after all, the server never hands it such a job with its sources and destinations, so it makes no copy at all rather than a copy without the password (see Troubleshooting).

> **i** Keep each password somewhere safe outside SLApp Helpdesk as well: without it the copies cannot be restored. Inside the app an administrator can find it again with Reveal and Previous passwords in the job editor, for a source that is still in the job, and with Password of this copy in the run detail, also for a deleted job or a source removed from the job. For a deleted job or a removed source the run detail is the only way, and only while a run that made such a copy is still in the Runs tab, which keeps runs for 180 days (Remove takes a run out sooner); retention never deletes their copies, so note the password down before you delete the job or remove the source. A password that the server can no longer decrypt is not shown at all (see Troubleshooting). Every reveal, of the current password or of previous ones, is recorded in the Audit log, and the log entry of a saved job says when a password was set, replaced or removed, never its value.

The job editor: Destinations

The second half of the editor, shown below, holds Destinations, Schedule, and Retention and options; Schedule and Retention and options have their own sections after this one. Destinations lists the destinations of the library, one row each: a tick box, the name with the icon of its type, the type and the location, for example Folder · \\nas01\backups, and a short free-space bar with a line such as 5.2 TB free (65 %), the full text in its tooltip, or Free space unknown; a destination that reads its free space from a computer adds a line such as SRV-FTP01 D: · 2 minutes ago. Tick the ones the job writes to, up to 20; the title counts them, as in Destinations (2 of 4 chosen). Every source is copied to every ticked destination, so tick more than one to keep a second copy somewhere else. Above the list the chosen destinations are repeated as chips, whatever the list shows, each with a × that unticks it, or No destination chosen. With more than 6 destinations the list gets a filter box, Filter by name, type or location…, and a Show only the chosen box, and it scrolls inside the window; No destination matches the filter. says when nothing matches. + Create a destination opens the destination editor without leaving the job; the destination is saved as soon as you click Save in that window, and it is ticked for you, with the filter cleared so that you can see it. That window has no Test button: to test the new destination, finish or cancel the job editor and use Test in the Destinations tab. When the library is empty the section says The library has no destinations yet. Create the first one. If a ticked destination was deleted in the meantime the editor shows Destination no longer available (untick it), at the top of the list and among the chips.



The second half of the job editor: Destinations (2 of 4 chosen), with the chips of the chosen destinations and one row per destination with its free space, then Schedule, and Retention and options, with Cancel and Save.

The job editor: Schedule, and when the computer is off

- Frequency set to Daily: every day at the Time you set; a new job starts at 02:00. The Time field is your browser's own time control, so it may display 11:30 PM where the Summary line shows 23:30;

the schedule itself is always kept as HH:mm (24 hours).

- Frequency set to Weekly: on the Days of the week you tick, at the Time. The day buttons start on Monday and a new job has Monday to Friday ticked.
- Frequency set to Monthly: on the Day of the month (1 to 31, starting at 1) at the Time. In months that are shorter the job runs on the last day of the month.
- Frequency set to Every N minutes: every Every (minutes) minutes (a new job starts at 60), with a minimum of 15 and a maximum of 525,600 (one year). The agent counts the interval from the end of the previous run, so runs never overlap and never pile up. A new job of this kind runs as soon as the agent has picked it up. The Next run column is only an estimate for such a job: before the first run it counts from the moment you saved, so it can show a later time than the real first run; afterwards it is the start of the last run plus the interval (or the moment you last changed the job, if that is later), while the agent counts from the end of that run, so the real start is a little later, by the duration of the last run.

Time zone sets the zone in which the time is read, including its daylight-saving changes, and starts as the zone of your browser. A job set to 02:00 in Europe/Lisbon runs at 02:00 Lisbon time all year. If a time does not exist on the day the clocks go forward, the job runs the same amount of time later as the jump (a job set to 01:30 in Lisbon runs at 02:30 that day); if a time happens twice when the clocks go back, it runs once, the first time. If the computer does not know the chosen time zone, the agent uses the computer's own time zone instead, without any warning in the app, so check the Next run and the first run times after choosing an unusual zone. The line under the schedule, Summary, repeats the schedule in words with its time zone. The time zone also applies to the Next run column and to the missed-backup alerts, but not to file names, which are always in UTC.

The schedule is kept by the agent, on the computer, so the app does not need to be reachable for a backup to start, and the agent works from its saved copy of the job definitions whenever SLApp Helpdesk is unreachable. The agent starts a scheduled run when it sees the scheduled time, within five minutes of it. What happens when the computer was off or the agent was not running at that time depends on two fields in Retention and options.

- Run late if missed (hours): the default 0 means never. With a number N, a job whose scheduled time was missed runs once, as soon as the agent is back, if the scheduled time was less than N hours ago. That run shows as Catch-up in the Runs tab. It is one run, never one per missed day. If more than N hours have passed, that occurrence is skipped and the job simply waits for the next one. The maximum is 720 hours. It has no effect on Every N minutes jobs: an overdue one simply runs once as soon as the agent is back.
- Missed-run alert after (minutes): the default is 120, and the allowed range is 15 to 10,080. If a scheduled time has passed and no run has started within this many minutes, SLApp Helpdesk adds a Missed entry to the Runs tab and opens a Backup missed alert, once for each missed time. This does not depend on the catch-up window: a job that cannot catch up still reports the miss, and a job that does catch up closes the alert when the late run finishes. A Missed entry is judged by what the server has heard, so a computer that did run the backup but could not reach SLApp Helpdesk for longer than this time can show a Missed entry and a Backup missed alert for a backup that did run. When the delayed report arrives the alert closes, but the Missed entry stays in the history next to the real run; Remove in the Runs tab clears it (see Removing a run from the history). A paused job does not pile up missed times, and a job you save after the scheduled time of the day waits for the next time. For an Every N minutes job the server expects the next run N minutes after the start of the previous run plus this delay, while the agent waits N minutes after the end of the run, so set the delay longer than the longest run, or the job can be reported as Missed while it is running normally.

A job never overlaps itself. A scheduled time that arrives while the same job is still running (for example after a Run now that took long) is not started at that moment. Once the running backup ends, that time

is run only if it is still within five minutes of the scheduled time or inside the Run late if missed window; otherwise it is skipped, and if no run starts within Missed-run alert after minutes of it, SLApp Helpdesk reports it as Missed.

The job editor: Retention and options

Retention decides how many copies of each series are kept. Keep the newest (copies) starts at 7 and accepts 1 to 100,000; Keep for (days) starts empty and accepts 1 to 36,500. The agent applies them exactly like this. The section also holds the staging folder and the verification option, described after the rules.

- Retention is applied separately to each series: every folder source, and every database of a SQL Server source, at every destination.
- It runs right after the newest copy of a series has been uploaded and verified at a destination. The agent then lists the job's folder at that destination and deletes the older copies of that series that are neither among the newest N nor younger than D days. When both fields are filled in a copy is kept if either condition holds, so setting both keeps more copies, not fewer. The age is taken from the date in the file name.
- Every run makes a new copy of each series, including a run started with Run now and a Catch-up run, and every copy counts toward the newest copies to keep. With Keep the newest (copies) set to 7, seven manual runs in one afternoon replace the whole previous week of copies; fill in Keep for (days) as well if you want older copies to survive.
- The newest copy is never deleted, and neither is the copy that was just made. Nothing is deleted when the series did not produce a good copy in this run, so a bad night never costs you earlier good copies.
- Only files named exactly as the module names its own copies, inside the job's own folder, are ever touched. Any other file in that folder, and anything in other folders, is left alone. Unfinished uploads of the series left behind by an interrupted run (names ending in .part) are removed once they are old: after 12 hours on an FTP server (when the next upload to that folder starts) and after 24 hours on a folder or share (when retention runs for that series, so only if a retention rule is set).
- Because old copies are deleted after the new one has arrived, a destination needs room for N+1 copies of every source when N copies are kept. The editor writes this under the fields, for example that each destination needs room for 8 copies of every source and of every SQL database when 7 are kept.
- With both fields empty nothing is ever deleted and the destination will eventually fill up; the editor warns you.
- If retention cannot do its work, for instance because it cannot delete a file, the item of that copy becomes a Warning that starts with “Retention was skipped” or “Retention could not delete old copies”.
- Staging folder (optional): where the ZIP of a folder source is built before it is copied. Empty means a work folder on the system drive of the computer, under C:\ProgramData\SLApp-Plugin\backup\staging. It must be a full path. Before building a ZIP the agent checks that this disk has room for the source and refuses with “Not enough free space in the staging folder” if not, so for large folders choose a staging folder on a big disk. The work folder is removed when the run ends, however it ends. A SQL Server source uses it only with Compress into a ZIP before sending: its .bak is written by SQL Server into the instance's own backup folder, whose volume the agent checks itself, and the ZIP of each database is then built here, one database at a time, once the agent has checked that this disk has room for about the size of that .bak.
- Verify each copy after writing it: ticked by default. After each upload the agent compares the size of the copy at the destination with the original and removes a copy that does not match. For SQL

Server sources it first runs RESTORE VERIFYONLY on the .bak and does not upload a file that fails. A ZIP with a password, and the ZIP of a SQL backup, are also read back in full before they are sent, whether this option is ticked or not (see The job editor: archive options). It is slower and recommended. Run late if missed (hours) and Missed-run alert after (minutes), the last two fields of the section, are described under the schedule.

Where the files land and how they are named

Every copy is written to the destination as <destination path>/<job folder>/<file>. For example, a job called Server files with a folder source called Documents and a SQL Server source called ERP, written to the share \\nas\backups, produces \\nas\backups\Server_files_1a2b3c4d\Documents__20261012-020000.zip for the folder and \\nas\backups\Server_files_1a2b3c4d\ERP__SalesDB__20261012-020000.bak for the database SalesDB. With Compress into a ZIP before sending ticked, the copy of SalesDB is ERP__SalesDB__20261012-020000.bak.zip instead.

- The job folder is the job name with every character other than letters, digits, dot, underscore and hyphen replaced by an underscore, followed by an underscore and eight characters taken from the job's identifier. Two jobs therefore never share a folder, even with the same name or on different computers.
- The file name starts with the File name prefix of the source, which is made from the source name in the same way, then two underscores, then for SQL Server the name of the database, then two underscores and the date and time. A database name that contains other characters, such as a space or an accent, is changed in the same way and gets a short code of eight characters appended, so two databases never end up with the same file name. Two sources of the same job never share a prefix: the module adds -2, -3 and so on when needed.
- The date and time in the file name are UTC, even though the job runs on its own time zone. A copy made at 02:00 in Lisbon in summer is stamped 010000.
- The job folder and the file prefix are fixed when the job and the source are created, and the editor shows them read-only (Destination folder and File name prefix). Renaming a job or a source changes only its label, so old copies are never stranded in a folder that retention no longer looks at. If you need a different folder name, create a new job and delete the old one; the old files stay where they are.
- While a copy is being written its name ends in .part. It is renamed to the final name only when the transfer is complete and its size matches, so a half-written file is never mistaken for a backup. On FTP the same layout is used with / as the separator, and the .part is also what an upload whose connection broke continues from (see When an FTP upload is cut).

When an FTP upload is cut

From agent 1.2.0, an FTP upload whose connection breaks does not start again from zero. The agent connects again, asks the server how many bytes of the .part file it already has and sends only the rest, appending it to the .part with the FTP command APPE. It carries on for as long as each attempt adds at least 1 MB at the server, with pauses that grow from 2 seconds to one minute, and gives up after three attempts in a row without progress or after 50 reconnections. A transfer that has taken no data from the file for two minutes counts as broken too. When continuing is impossible or unsafe, because the server refuses APPE or does not append correctly, or the .part has vanished, is larger than the file or has no known size, the agent deletes the .part and sends the file again from the beginning, as older agents do. An upload that never breaks does none of this and takes no longer.

A connection that broke can still deliver some of its data late, and the server then writes it into the .part. So after a cut the agent sends nothing for at least 60 seconds, and not before the size of the .part has stopped changing. Before giving the copy its final name it checks that the size is exactly that of the

original for a few seconds, compares the whole file with the original when the server offers a checksum command, and reads the size once more after the rename. If anything differs, it deletes the copy and sends that file again from the beginning, with the text “a late write from the dropped session was detected”; after three such restarts in a row the upload fails with that text. These checks catch any late data that reaches the server within 60 seconds of the cut. When the server has a checksum command, the comparison also catches any difference that is already in the file when it is made, just before the rename. Data that arrives later still, after those final checks, or more than 60 seconds after the cut on a server without a checksum command, can go unnoticed, and a server that reports no size cannot be checked at all; the warning of the copy says which checks were made, and a test restore remains the surest check.

From agent 1.2.9, a long upload is also sent in parts of at most 10 minutes: at the end of each part the agent closes the data connection normally, the server confirms it, and the next part is appended with APPE from the size the server reports. Some firewalls stop passing the data of a connection that stays open too long (one did so after about 14 minutes); with parts no connection lives that long, so such an upload is no longer cut. Parts that go well are not mentioned in the copy's warning, and a server that refuses APPE gets the whole file in one transfer, as before. From agent 1.2.6, when after a cut the server refuses the .part as in use (“being used by another process”: Microsoft IIS keeps the file of a dropped connection open for several minutes), the agent waits up to 20 minutes for it to be released instead of giving up, and the warning says how long it waited.

A copy that needed a new connection has passed the checks that its warning lists, but its card is a Warning and the run reads Completed with warnings. The text says what interrupted the upload and what followed, for example “The connection to ftp.example.net:21 dropped 2 times during the upload; it was resumed each time from the bytes the server had (total 48.2 GB, 1 h 52 min).”, and ends with the checks made against late data, which start with Checked for late writes of the dropped session, or Not checked for late writes of the dropped session when the server tells neither the size nor a checksum. An upload that fails in the end says how far it got, for example “FTP upload of '…' on ftp.example.net:21 failed after 31.5 GB of 48.2 GB (51 min, 4 resumes): …”, and its .part is deleted. Long transfers also keep their connections open with regular keep-alive messages, so that routers and firewalls are less likely to cut them for looking idle.

The Calendar tab

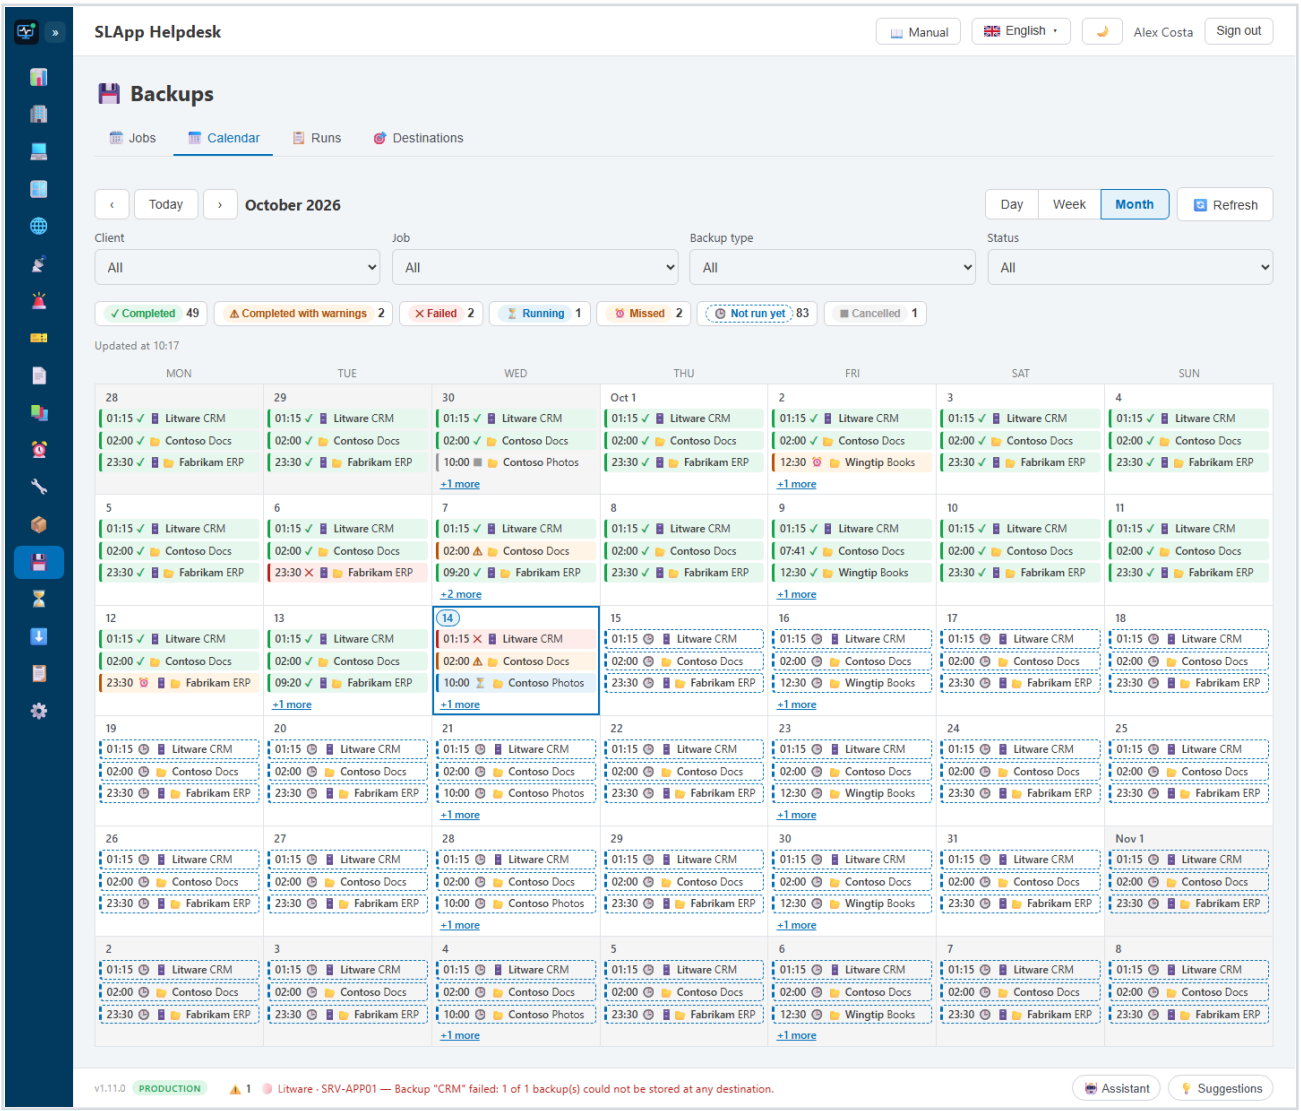
The Calendar tab, the second one, lays every backup of every job out on a calendar, so that a single look tells you what ran and how it ended, what is running now, what did not run and what is still to come. The Runs tab lists the runs that exist. The Calendar also draws the ones that do not exist yet, or that never happened, because it works out from the schedule of each job when the job is due and marks the times that no run covers. It only reads: nothing you do on the calendar itself starts, changes or deletes a job, or sends anything to a computer. The only actions come from the windows it opens, such as Cancel run in the detail of a backup that is running, and they work exactly as in the Runs tab.

The top of the tab has the buttons that move through time, with the tooltips Previous month, Next month, Previous week, Next week, Previous day and Next day according to the view, and Today, which brings you back to the present; then the title of the period on screen, the Day, Week and Month buttons, which choose the view, and Refresh. Weeks start on Monday. The first time you open the tab it shows the Month view, and from then on it opens in the last view you chose, which your browser remembers. The period, the filters and the counters are not remembered: every time you open the tab it starts on today with nothing filtered. Under the filters and the counters, Updated at gives the time of the last answer from the server. The clock of the server at that moment, not the one of your browser, decides which entries are Due now and which are Missed.

Times are shown in the time zone of your browser, and an entry sits on the day, in that zone, of its time: the start of a run, or the scheduled time of a backup that did not start. The time zone of the job itself is

in the tooltip of every entry and in the window of a backup that has not run yet, so you can compare it with the schedule you wrote in the job.

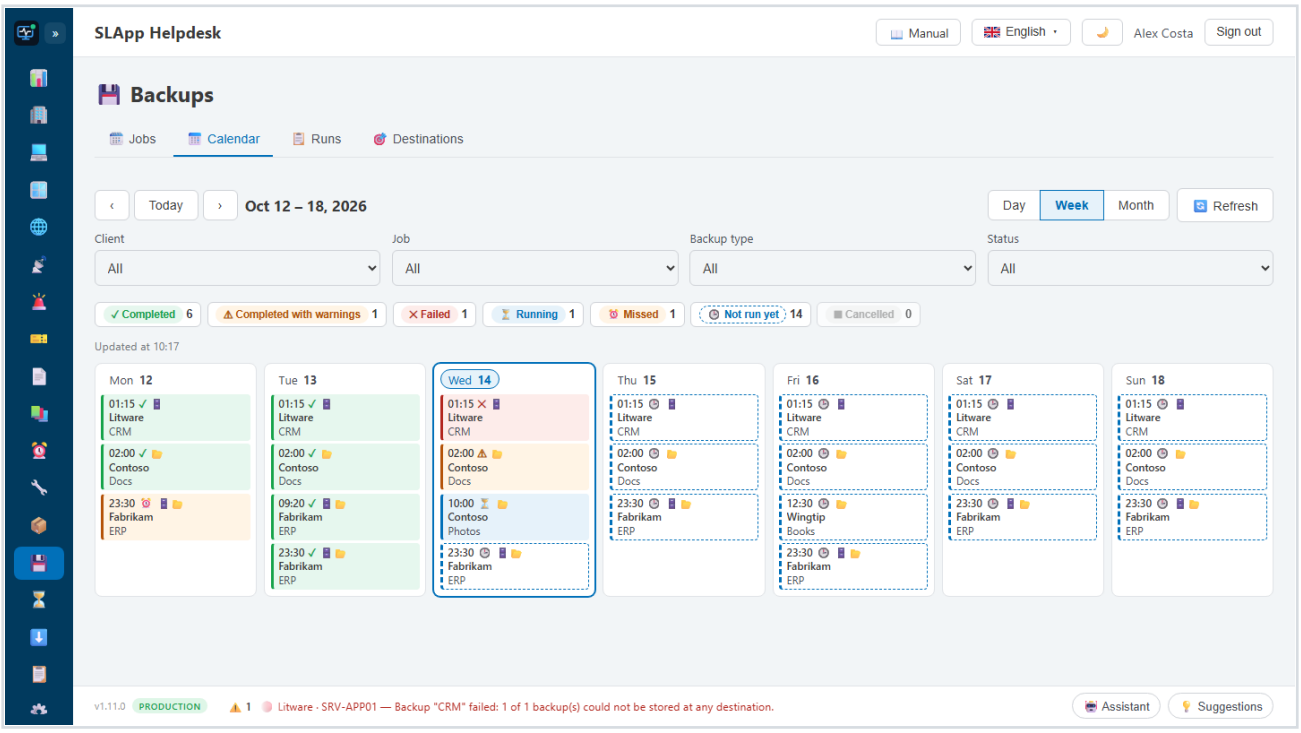
The Month view shows six whole weeks, from the Monday on or before the 1st of the month, so the first and last rows can hold days of the neighbouring months; those days are shaded and their entries are drawn like any other. Today is outlined and the 1st of every month carries the name of the month. Each day shows up to three entries as small chips with the time, the icon of the state, an icon for each type of source (SQL Server and folders have one each), the client in bold and the job. The chip of a backup that has not run yet has a dashed outline, and one that is due now has the same shape, filled. The three chips are not the first three of the day but the most worrying ones: Failed first, then Missed, Completed with warnings, Due now, Running, Cancelled, Not run yet and, last, Completed, drawn in time order, so that calmer entries never push a failure out of sight. When the day holds more, a +N more button appears under the chips. It also shows, as an icon and a number, each state that needs attention among the entries it hides, and its tooltip spells them out, for example “+2 more (Failed ×1, Completed ×1). Open this day.” Click +N more, or the number of the day, to open that day in the Day view; this does not change the view that is remembered. The tooltip of the number of a day gives its date, how many backups it holds and in which states, and ends with “Open this day.” Hover over a chip to read its tooltip: the time and the state, the type and the client, the job and the computer, the time in the job's time zone and, when there is one, the error.



The Calendar tab in the Month view: the filters, one counter for each state, and in each day up to three entries (the most worrying ones when the day has more) with +N more for the rest.

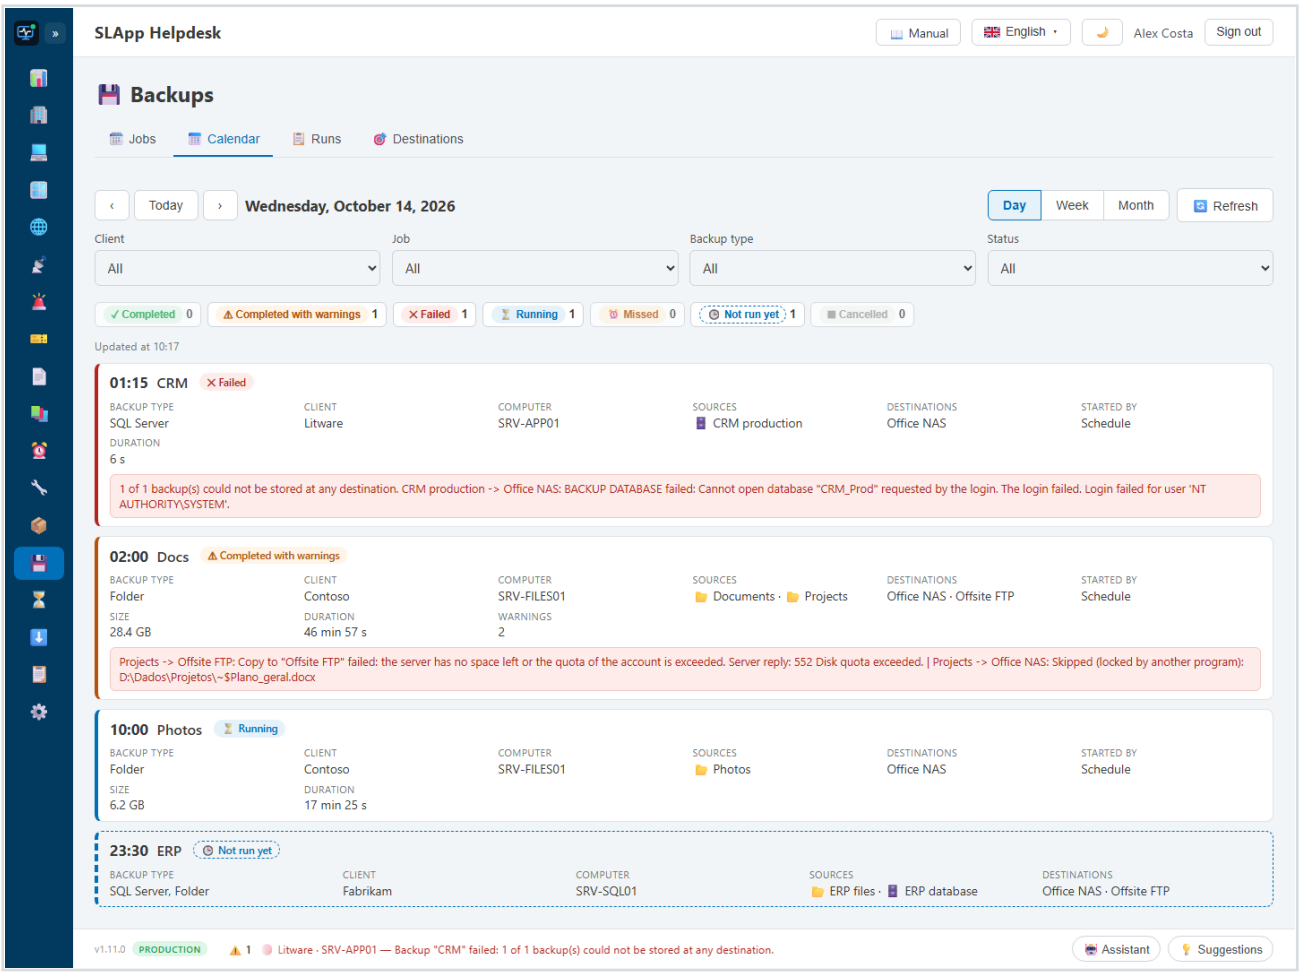
The Week view shows the seven days, Monday to Sunday, side by side. The weekday and the number

at the top of each day are a button that opens that day in the Day view, with the same tooltip as the number of a day in the Month view. Every entry of the day is drawn, with no limit and no +N more, as a block of three lines: the time with the icons of the state and of the type, then the client, then the job. A day with nothing in it shows a dash. On a narrow window the seven days are stacked one under the other.



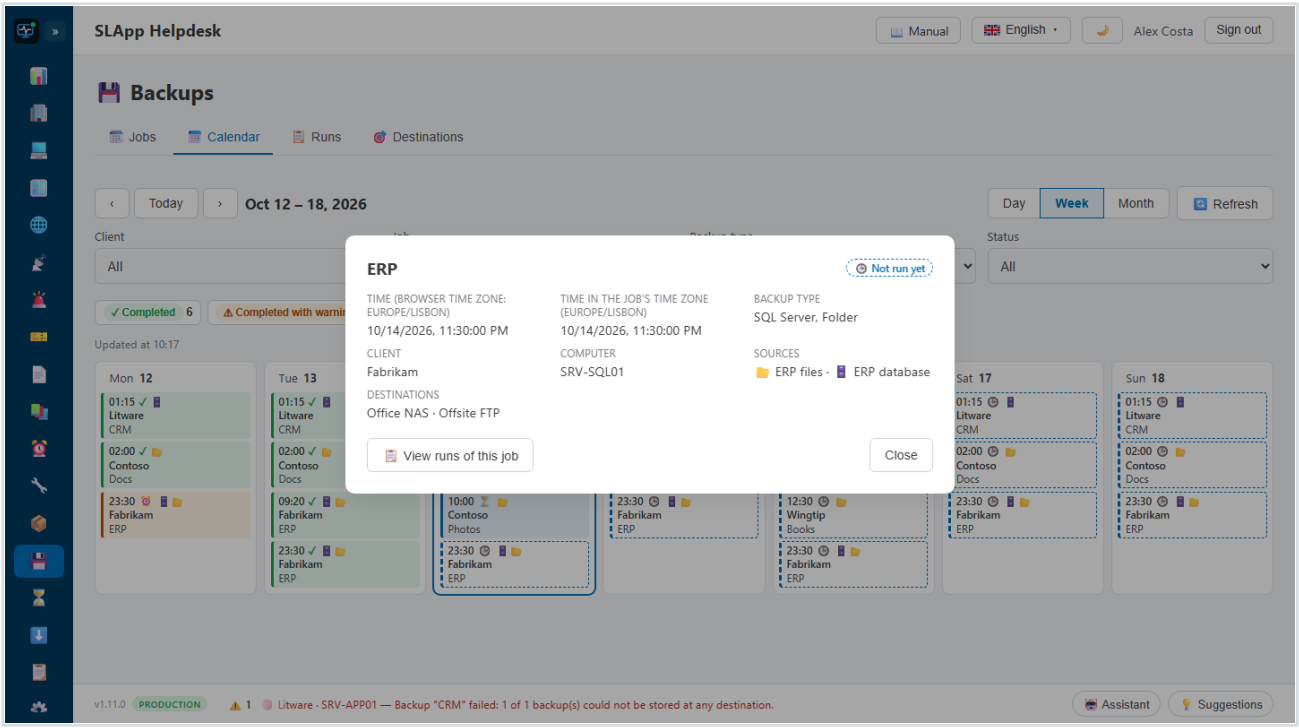
The Week view: every entry of the seven days, with today outlined.

The Day view lists the backups of one day in time order, one row for each, with everything the chips only hint at. The title of the row is its time and the job. Next to it are the state, with its icon and its words, and the word estimated for an estimated time. Below come Backup type (SQL Server, Folder or both), Client, Computer (its name is a link that opens the computer's page), Sources (each with the icon of its type), Destinations and, when they apply, Started by (Schedule, Manual or Catch-up, for a run), Size (for a run that produced files), Duration (for a backup that has started) and Warnings (for a run with warnings). If the run has an error, its summary is written under the row, in a box whose text you can select and copy: it is the text of the error as a single paragraph, cut at 300 characters (with … at the end when it is longer); the run detail has the whole text. A Due now row adds “The scheduled time has come. Waiting for the agent to start the backup (the computer may be offline or busy).” and a Missed row adds “The scheduled time passed without a run starting (for example the computer was off or offline). No backup was made for it.”



The Day view: one row for each backup of the day, with its state, client, computer, sources, destinations, duration and error.

Click a chip, a block or the title of a row to open the entry. A backup that has a run, whatever its state, opens the run detail, the same window as in the Runs tab, with the cards of every file and destination and, while the run is running, Cancel run. A time with no run, which is a Not run yet, Due now or Missed entry that the Calendar worked out itself, has no run to open, so it opens a small window instead. At the top of the window are the job and its state. Then come the time in your browser's time zone, Time (browser time zone: …), and in the job's, Time in the job's time zone (…), followed by Backup type, Client, Computer (a link to the computer's page, as in the Day view), Sources and Destinations. An estimated time adds the note “Estimated time: this job repeats at a fixed interval counted from its last start, so the exact time can shift.” A Due now or Missed window adds the same explanation as the Day view. View runs of this job takes you to the Runs tab, already filtered on that job. Close, or the Esc key, closes the window.



The window of a backup that has not run yet: the time in your browser's time zone and in the job's, the type, client, computer, sources and destinations, and View runs of this job.

Reading the Calendar: states, schedules and estimates

Every entry is either a run, which is also in the Runs tab, or a time that the Calendar plans from the schedule of an enabled job and that no run covers. The state is never given by colour alone: every entry has the icon of its state, the Day view and the windows add its name in words, and the tooltip of a chip spells it out. The entries that have not run yet are dashed. The states are these.

- Completed (a tick), Completed with warnings (a warning triangle), Failed (a cross), Running (an hourglass), Cancelled (a stop square) and Missed (an alarm clock) are the statuses of a run, described under The Runs tab.
- Not run yet (a clock, dashed outline): a time that is still to come and that no run has covered. For a job that repeats every N minutes it is an estimate, as described below.
- Due now (the same clock, filled): the time has come and nothing has started. The Day view and the window say that the agent is being waited for (the computer may be offline or busy). For a Daily, Weekly or Monthly job this lasts for the delay set in Missed-run alert after: the entry becomes a run when the agent starts the backup, and Missed when the delay is over. For an Every N minutes job see below. The counters count it under Not run yet.
- Missed: a time that passed more than the Missed-run alert after delay ago with no run to show for it. The server records a missed time as a run when its own check, which runs every five minutes, finds it: the Runs tab then lists it and the Backup missed alert opens. A time that the check has not recorded, such as an older one that passed while the server was stopped, is worked out by the Calendar itself and has no run to open. A time whose run you removed in the Runs tab is drawn the same way, as Missed with no run to open and no alert (see Removing a run from the history).

For a Daily, Weekly or Monthly job the Calendar goes through the schedule one time after another, in the time zone of the job and with its daylight-saving changes, which is the same calculation as the Next run column. A run covers a time when its Scheduled for is within a minute of it, or when it started between five minutes before the time and as many minutes after it as Missed-run alert after says (120 unless you changed it). This is the rule used to decide that a backup was missed, so a Run now or a Catch-up run that starts inside that window also counts as the run of that time. A run covers one time only. A covered time is drawn once, as the run, so a daily job has one entry for each scheduled time

and not a run plus a planned one (plus one more entry for every run you start by hand with Run now, unless that run starts close enough to a scheduled time to cover it). A time that nobody covers is Not run yet when it is in the future, Due now when it passed less than the Missed-run alert after delay ago, and Missed when it passed more. A time from before the job was created is not drawn, and a time from before the job was last saved or resumed is not drawn as Missed, because the job did not exist in that form then; this is the same rule as the check that raises the Backup missed alert. One difference: such a time that is still less than the Missed-run alert after delay old is drawn as Due now until the delay is over, and then it disappears.

A paused job, which is also what happens to the jobs of a computer that is removed in Computers, has no planned entries at all, neither upcoming nor Missed. That is why the future of a paused job is empty, and why the Job filter marks it with (paused). Its runs in the period are still drawn. When you resume the job its upcoming times come back, and the times that passed while it was paused are not drawn as Missed. A time that is less than the Missed-run alert after delay old when you resume the job is still drawn as Due now until that delay is over, and then it disappears, because the agent does not catch up the times of a paused job. A job that you delete disappears from the Calendar, its past runs included, although the Runs tab keeps them.

A job that repeats Every N minutes starts each run N minutes after the end of the previous run, so its times cannot be known in advance, and the Calendar does not invent any for the past: it draws only its real runs, with the Missed ones that the server recorded. For the future it draws estimates: the first is the Next run of the Jobs tab and the others follow one interval apart. An estimate has ~ in front of its time (for example ~14:30), the word estimated in the Day view and a note in its window. The Calendar counts the interval from the start of the last run, while the agent counts it from the end, so a real run starts a little later than its estimate, by the duration of the previous run, and the estimates move after every run. For the same reason an estimate whose time has come is drawn as Due now at the current time, and it can read Due now for as long as the last run lasted (or while a run longer than the interval is still going) although the agent is not yet due. It stays Due now until the job starts and is never turned into Missed by the Calendar. If the job really is late, the server records one Missed run once the Missed-run alert after delay has passed, and that entry is drawn at the time N minutes after the last start, next to the Due now one.

> **i** The type, the sources and the destinations of an entry are those the job has today, also for a run from months ago: if you add a source to a job later, its old entries show it too. The run detail shows what that run really did. The computer and the client, on the other hand, are those where the run really ran, so moving a job to another computer does not rewrite its old entries.

Calendar filters, counters and limits

Four lists under the title narrow what is drawn. Client lists the clients that have a job, with No client for the jobs of computers that belong to no client. Job lists the jobs, only those of the chosen client when one is chosen, with (paused) after the name of a paused job. Backup type is All, Folders or SQL Server, and a job with both kinds of source belongs to both. Status is All or one of the seven states; Due now is not a separate item, those entries are found under Not run yet. If you choose a client that does not own the job you had selected, the Job filter is cleared. Clear filters appears as soon as one filter is on.

Between the filters and the calendar there are seven counters, one for each state, in this order: Completed, Completed with warnings, Failed, Running, Missed, Not run yet and Cancelled. Each shows how many entries of that state the period holds. The period is what is on screen, so for the Month view it is the six weeks of the grid, shaded days included. The counters follow Client, Job and Backup type, but not Status, and Not run yet includes the Due now entries. A counter is also a button: click it to show only that state (the Status list follows) and click it again to remove the filter. Its tooltip reads, for example, “Show only: Failed”. A counter at zero is a little paler.

To keep the page fast, the server sends at most 200 planned entries for each job and 4000 entries in all, runs and planned ones together. When a period holds more, a notice above the calendar says “There are too many backups in this period, so some are not shown. The ones furthest from today go first, and failed or missed backups are kept whenever possible. Choose a shorter period or filter by client or job to see them all.” What stays is chosen in this order: first every entry that needs attention (failed, completed with warnings, running, cancelled, missed and due now), then the rest, the ones nearest to today first, so that a busy month still shows today's failure, the missed night and what comes next. A job that repeats every 15 minutes, with 96 planned entries a day, reaches the limit quickly. Client and Job narrow what the server sends, so they help (the No client choice does not: it only hides what has already been loaded); Backup type and Status only hide what has already been loaded. The same notice appears when the period, with the day before it and the Missed-run alert after delay after it, holds more than 20,000 runs: the Calendar then leaves out the oldest part of the period (its first days can look empty), because it cannot tell whether an earlier time was covered.

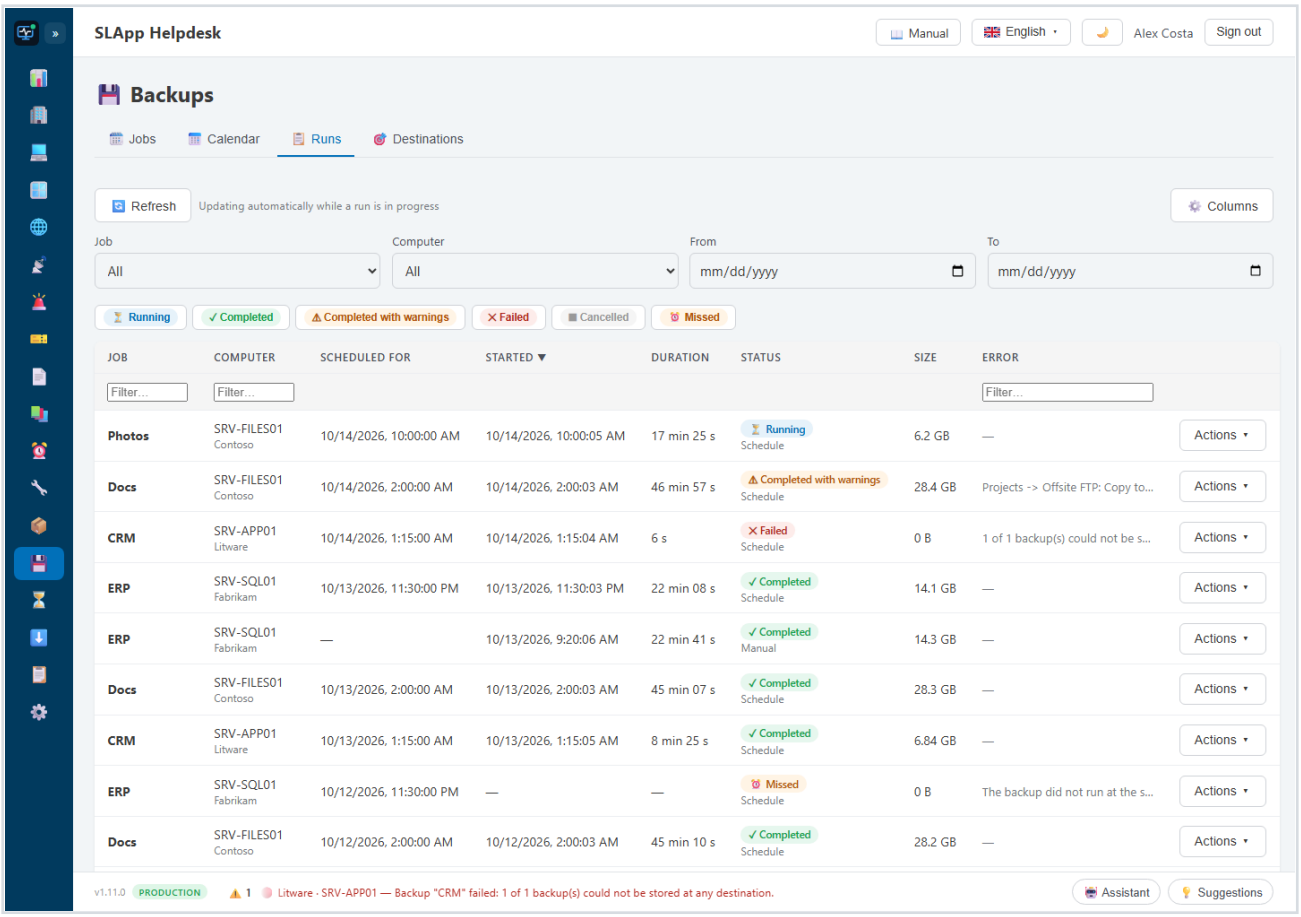
While the answer is on its way the counters show a dash, the calendar is dimmed and a line says Loading…. If the server cannot be reached, a box shows the error with a Try again button. When a period has nothing to show it says “No backups scheduled in this period.”, with “Create a job in the Jobs tab and its schedule appears here.” underneath when your company has no job at all, or “This job is paused, so it has no upcoming backups. Resume it in the Jobs tab.” when the Job filter is on a paused job. If there are backups but the filters hide all of them, it says “No backups match the filters.” and offers Clear filters.

The Calendar asks the server again by itself every 30 seconds while the page is visible, and again when you come back to the page after a while. For two minutes after you click Run now in the Jobs tab it asks every 10 seconds, so the new run appears quickly. It does nothing while the page is hidden. The Refresh button asks at once.

> **i** On a phone or a narrow window the Week view stacks its seven days one under the other (below about 860 pixels), and the Month view (below about 640 pixels) replaces the chips by the icons of the states present in each day, the worst first and at most four, with the total number of backups of the day. Tap a day to open it in the Day view. The state is still given by the icon and not only by the colour.

The Runs tab

The Runs tab is the history of every run of every job. The top of the tab has Refresh, Clear filters when a filter is on, and Columns. Above the table you can filter by Job, Computer and a range of dates (From and To, whole days in your own time zone), and the row of status buttons shows one status at a time; click a status again to remove it. While a run is running, and for two minutes after you click Run now, the list refreshes itself every 10 seconds; while a run shown in the list is Running, the note Updating automatically while a run is in progress appears next to Refresh. The list shows 100 runs at a time; Load more brings the next ones. Run history is kept for 180 days; to take a single run out sooner, see Removing a run from the history.



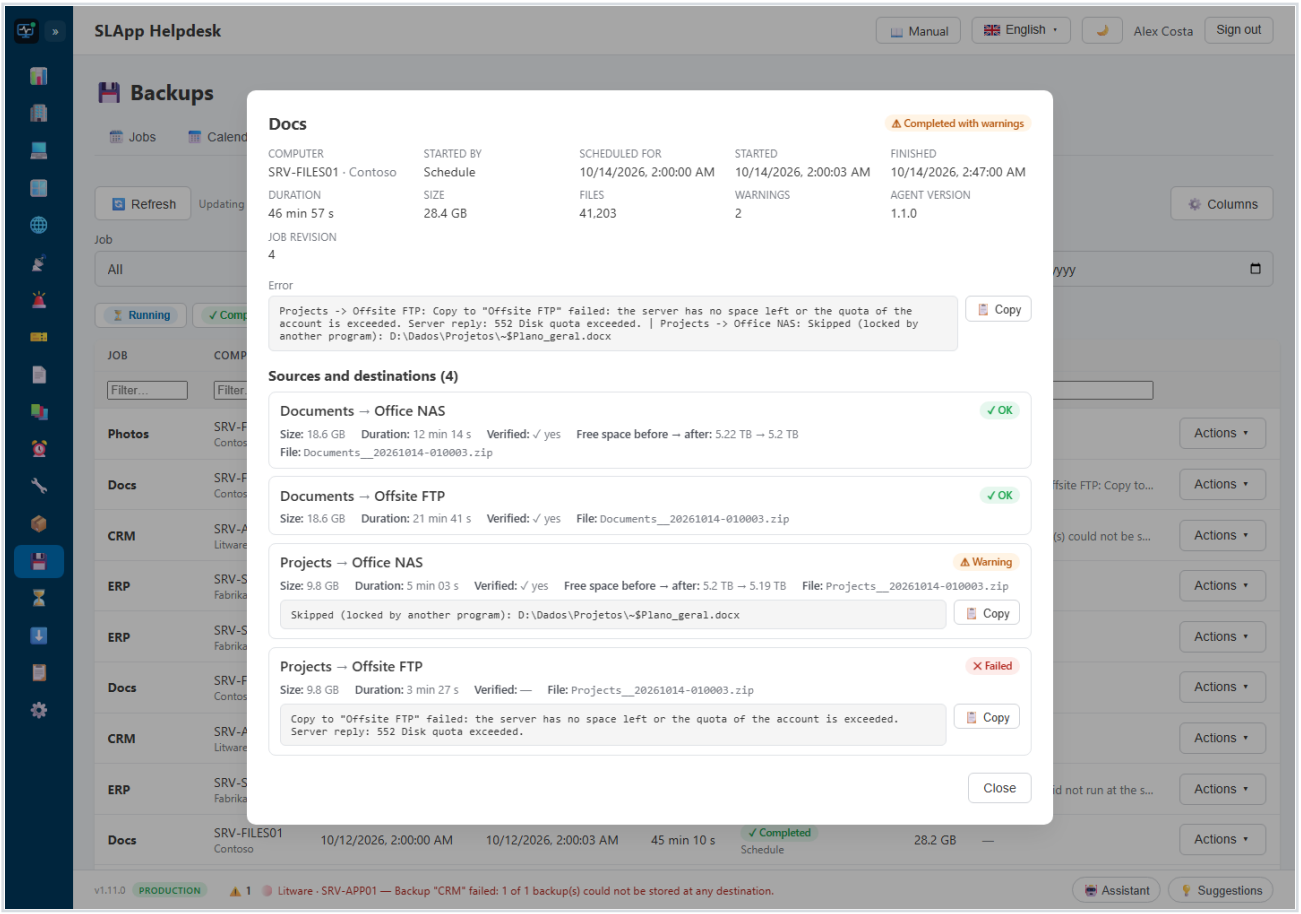
The Runs tab: every run with the time it was due, its start, duration, status, size and error, and a button for each status above the table.

The table shows Job and Computer (with the client underneath; a click on the computer's name opens the computer's page rather than the run), Scheduled for (the time the run was due, empty for a run started with Run now), Started, Duration, Status with Schedule, Manual or Catch-up underneath to say what started the run, Size (the total size of the files the run produced; a ZIP or a .bak is counted once, however many destinations it was copied to) and Error (the summary of what went wrong, cut short). Finished, Started by, Files, Warnings and Agent version are available through Columns. Started by says how the run began (Schedule, Manual or Catch-up), not who started it: to see which technician clicked Run now, Cancel run or Remove, open the Audit log and filter the entity type Backup job (the actions are run_requested, run_cancelled and run_deleted). A run has one of these statuses.

- Running: in progress. The agent reports about every 30 seconds. If the computer is switched off or loses power during a run, the run stays Running until the agent service starts again (it then becomes Failed with the error interrupted) or until six hours pass without a report (the server then closes it as Failed with the error “No report from the agent for 6 hours while the backup was running; it is assumed lost.”).
- Completed: every folder and every database was copied to every destination and nothing needed a warning.
- Completed with warnings: every folder and every database has at least one good copy somewhere, but something went wrong on the way. Typical causes are a destination that failed while another one worked, files that were locked and skipped, links that were not followed, a retention step that could not delete, or an FTP upload that had to be resumed after its connection broke (see When an FTP upload is cut). The Error column and the run detail say which.
- Failed: at least one folder or database has no good copy at any destination, or the run could not start or finish. Examples are a destination that is unreachable and the only one, a source that produced nothing, a staging folder without room, and a run that was cut short, which reads interrupted.

- Cancelled: somebody used Cancel run, which is in the Actions menu of a Running run and in its detail window. It asks you to confirm (Cancel the running backup “name”? The agent stops as soon as it can.) and then answers Cancellation requested. The agent stops as soon as it can. Until the run ends it shows Cancellation requested… under its status, and Cancel run is no longer offered for it. The server marks the run and, when the agent is connected, tells it at once; otherwise the agent learns it from the answer to its next report, within about 30 seconds. A copy that was already complete stays and an unfinished one is dropped. A cancelled run does not open, change or close Backup failed or Backup missed alerts.
- Missed: the scheduled time passed and no run started. It is created by the server, as described under the schedule, and no backup was made for it.

Click a run, or choose View details in its Actions menu, to open the run detail. The title is the job name with the status chip. Under it come the computer, whose name is a link that opens the computer's page, and the client, Started by (how the run began), Scheduled for, Started, Finished (or in progress), Duration, Size, Files, Warnings, Agent version and Job revision, which is the version of the job's settings that this run used and goes up with every change to the job, Pause and Resume included, and with every change to one of its destinations. If the run has an error it is shown in a box with a Copy button. While the run is Running the window also has a Cancel run button.



The run detail: the summary of the run and one card for every file and destination with its status, size, verification and file name.

Below the summary there is one card for every file the run produced and every destination, titled with the source, an arrow and the destination (a SQL Server source has one card per database, and the File name tells them apart), and with its own status chip: OK, Warning, Failed or Skipped. Skipped means that nothing was attempted, because the source could not be read or the destination could not be reached earlier in the same run; a source that could not be read has one Failed card, with no destination, and one Skipped card for each destination. Each card shows Size, Duration, Verified (yes when the copy was checked; no when the copy is there but the destination could not report the size of the file, so the check could not be made; and a dash when Verify each copy after writing it is off, or

when the copy failed or was skipped and so was never checked), Free space before › after for the destination, and the File name. An error is shown in a copyable box under the card. While the run is Running the window refreshes itself every few seconds.

The Copies tab

The Copies tab shows what each destination really holds. At the end of every run the agent lists the job's folder at each destination and reports the copies it finds, after retention has deleted the old ones; the tab shows that list. There is one row per series: a folder source at one destination, or one database of a SQL Server source at one destination. You can filter by Job, Computer, Destination and State, and the buttons above the table count the rows in each state.

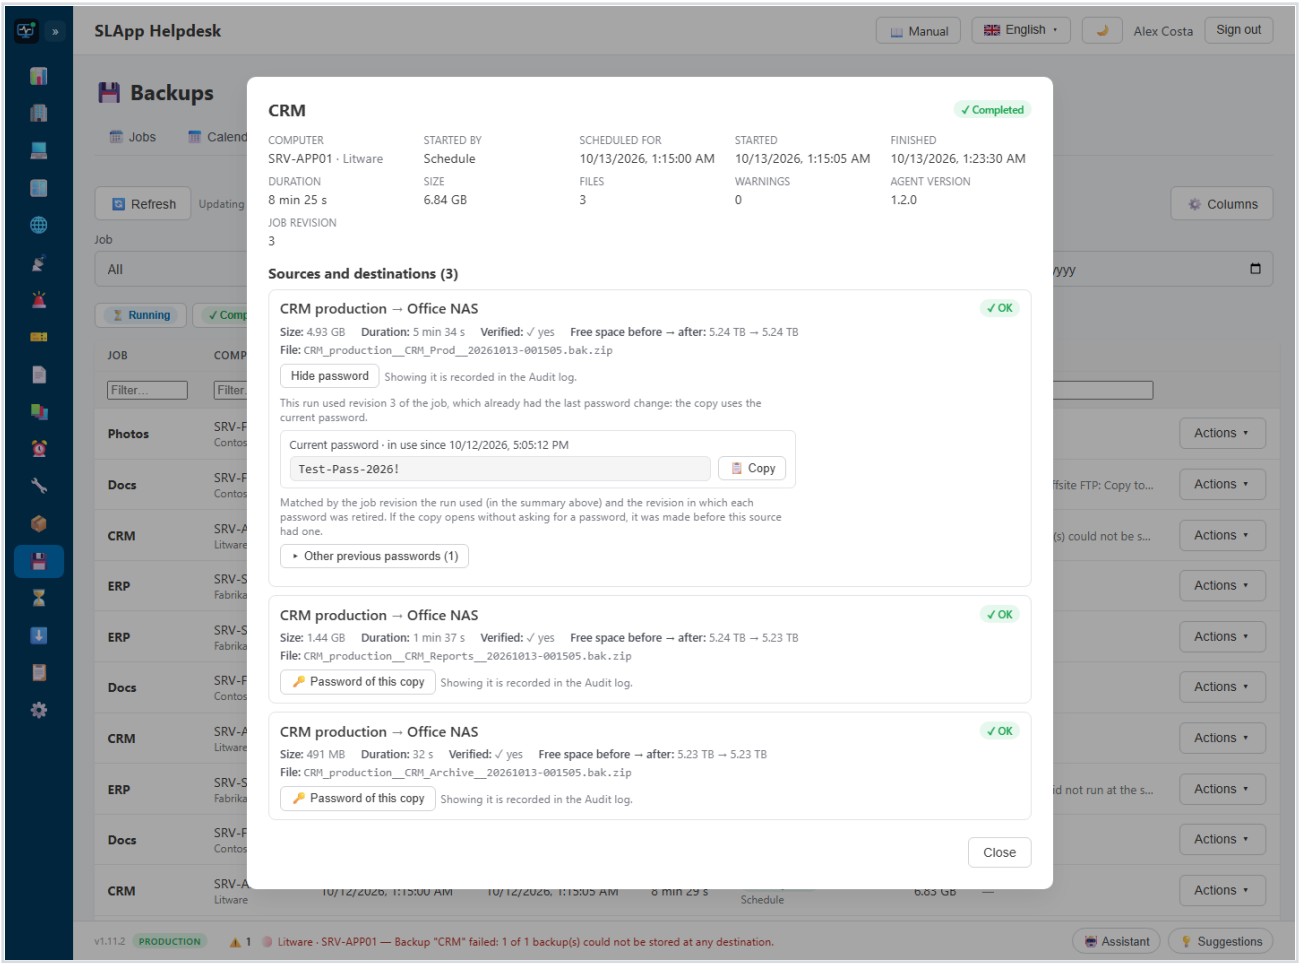
- Columns: Job, Source / DB, Destination, Latest copy (date and how long ago), Copies (how many complete copies, and the retention rule), Total size, Inventory (when the destination was last listed) and State.
- OK: the latest copy is there and nothing is wrong.
- Out of date: the latest copy is older than the job's schedule implies (for example a daily job whose newest copy is several days old); the row says what was expected.
- Needs attention: copies that retention should already have deleted are still there, or there are unfinished uploads (.part files).
- No copies: the destination holds no complete copy of that series.

Click a row to see its copies, newest first: file name, date, size, whether it was verified, its state (Available, Should have been deleted or Unfinished upload (.part)) and View run, which opens the run that produced it. A series with very many files lists only the newest ones and says how many there are in total.

> **i** The inventory comes from agent 1.2.14 or newer and appears after the next run of each job. Until then, a notice above the table lists the jobs without an inventory and why (agent too old, or waiting for its next run). On FTP destinations unfinished uploads are not listed.

The password of a copy

When the file of a card is a ZIP (the copy of a folder source, or the .bak.zip of a database) of a source that has or had an archive password, the card has a Password of this copy button, with the note Showing it is recorded in the Audit log. It is there for the runs of a deleted job and for a source removed from the job too. Click it and the window works out which password that copy was made with and shows it, with Copy; Hide password, or closing the run detail, hides it again.



The run detail of a run whose copies have a password: Password of this copy, open on the first card, matched the current password by the job revision the run used and shows it with Copy, with Other previous passwords (1) underneath.

- By job revision first. A run reports the revision of the job's settings it used (Job revision, in the summary), and every previous password records the revision of the save that retired it, so the answer is exact, even for a backup that waited hours to start. The window then says, for example, “This run used revision 3 of the job, which already had the last password change: the copy uses the current password.”
- By time when the revisions cannot tell, for example for a password retired before revisions were recorded: the window compares the start of the run with the dates the passwords were retired and says that the copy most likely uses the password it shows. It also says why this is not certain: a backup that waited for a free slot (a computer runs 2 at a time) or ran on a computer that was offline used the settings it had when it was due to start, possibly hours before it started.
- Each password shown says what it is: Current password · in use since a date, Replaced on a date · up to job revision N, or No password · removed on a date. When there are several they are listed in the order to try them, with Try them in this order. A copy that opens without asking for a password has none.
- Other previous passwords (n), under them, lists the other previous passwords of the source, the ones that do not match this copy, to try only if it opens with none of the passwords above. It is the only place where the other passwords of a deleted job or of a removed source can be seen, and only while that run is kept in the Runs tab (180 days).
- A password set on a source that already had copies does not reach those copies: they have no password. That is why the window can end with If the copy opens without asking for a password, it was made before this source had one.

Removing a run from the history

Remove takes one run out of the history, for example a test run, or a Missed row that sits next to a real run of the same time (see Missed-run alert after in the job editor). A Missed row for a time that really had no run leaves the Runs tab, but the Calendar can go on drawing that time as Missed (see below). It is in the Actions menu of the run and in its right-click menu, in the Runs tab only (the run detail and the Calendar do not have it), and it is offered for every run that is not Running: Completed, Completed with warnings, Failed, Cancelled and Missed. A Running run has no Remove, not even one that already shows Cancellation requested…: use Cancel run, while it is still offered, and remove the run once it has ended. A run whose computer went silent is closed by the server as Failed after six hours without a report (see Running above) and can be removed from then on. Runs are removed one at a time: there is no multiple selection, and the only other clean-up is the automatic one after 180 days. Remove also works on the runs of a deleted job, which stay in the Runs tab; for those the window has no line about the Calendar.

Remove asks first. The window is titled Remove this run? and its first line names the run: the job, the date and time of the run (its start, or the time it was due for a run that never started) and its status. Under it the window says “The record of this run and its details for each source and destination will be erased. The backup files already written to the destinations are NOT deleted, and the alerts this run raised stay open until you close them in Alerts.” For an enabled job that does not repeat every N minutes (see below) it adds a last line, “If this run covered a scheduled time, the Calendar may now show that time as missed, because no run covers it any more (no alert is raised).” The buttons are Cancel, which has the focus when the window opens so that pressing Enter by mistake removes nothing, and Confirm; Esc closes the window too, unless the server is already working on the removal.

- Erased: the record of the run and the cards of its files and destinations. It disappears from the Runs tab and the Calendar stops drawing the run itself; the scheduled time it covered may be drawn as Missed afterwards (see below). In the Jobs tab, Last run then shows the most recent run that is left, whatever its status, or Never run when there is none.
- Not touched: the copies. The ZIP and .bak files the run already wrote stay at the destinations. Nothing is sent to the agent or to a destination, and the job, its settings and its schedule do not change.
- Not touched: the alerts. The Backup failed, Backup missed or Backup destination space alerts that the run raised stay open, and removing a run neither opens, closes nor changes any alert. Close them in Alerts, with Resolve, as usual.
- Recorded: the removal is written to the Audit log as run_deleted, under the entity type Backup job. The details name the job and give, in brackets, the status and the trigger of the run as the server stores them (for example completed_with_warnings, manual) and the time it started in UTC, or a dash when the run never started, as with a Missed row.

Removing a run does not by itself make the app raise a Backup missed alert, send an e-mail or record a Missed run. This needs some care, because the server checks every five minutes whether each job missed its last scheduled time, and it calls a time missed when no run covers it: removing the run that covered the latest time would have made that check report, a few minutes later, a backup that did run. So Remove settles it in the same step as the erasing: when the removal is what leaves a time uncovered, that time is marked as already dealt with and the check leaves it alone. A time that really was missed and is waiting to be reported is not affected and is reported as usual, and removing a Missed run does not make the server record it again. One rare case needs a word. A manual run that started up to five minutes before a scheduled time that has not come yet counts as covering that time. If you remove such a run, that coming time is marked as dealt with in advance, so no Backup missed alert follows if the agent then fails to run it; the exception is when the latest scheduled time is itself a real miss that has not been reported yet: the real miss comes first, the coming time is not marked, and

it is judged without the removed run.

The Calendar does not read that mark. It works out which times have a run from the runs that are left (see Reading the Calendar: states, schedules and estimates), so after you remove the run that covered a scheduled time of a Daily, Weekly or Monthly job, that time can be drawn as Missed, with no run to open and no alert behind it, because the history no longer says that it ran. This is expected, and it includes a Missed run that you remove: the row leaves the Runs tab, but if no other run covers that time, the Calendar goes on drawing it as Missed, with no run to open, for as long as the job is enabled, does not repeat every N minutes and has not been saved or resumed since that time. A time that is still within the Missed-run alert after delay is drawn as Due now until the delay is over. A job that repeats Every N minutes, a paused job and a deleted job get no such entries: for them the removed run simply disappears from the Calendar, and that is why the window adds its extra line about the Calendar only for an enabled job that does not repeat every N minutes.

One caveat for a job that repeats Every N minutes. The server counts its interval from the later of two moments: the start of the latest real run (a run that started, not a Missed row) and the last time the job was saved. If you remove that latest run and the job has not been saved since, the missed-backup check goes back to the run before it, treats that time as already dealt with and has nothing newer to judge: it stays quiet until the job has a new real run, or until you edit and save the job or pause and resume it, which starts the count again from that moment. If the computer stops running backups before that, no Backup missed alert is raised for it in the meantime (the alerts about a computer being offline are not affected). Removing any other run of such a job, or the latest run of a job that was saved after it, changes nothing for that check. The Next run estimate of the job, in the Jobs tab and in the Calendar, is counted in the same way, so it moves when you remove its latest real run; the schedule that the agent follows does not change.

While the server works, both buttons are disabled and Confirm reads Removing…; then the window closes and the list shows the notice “Run removed.” and loads again. If another technician removed the run first, the notice is “That run had already been removed.”, which is just as good. If the run started in the meantime (the list was out of date), the server refuses and its message, “A running backup cannot be removed. Cancel it first.”, appears inside the window, which stays open while the list reloads so that the row shows its real state. Any other failure, such as a lost connection, is shown in the same place.

> **i** In the rare case that the agent sends again the final report of a run because it never received the answer of the server, and you removed that run in between, the server stores the report like any new one: the run is back in the list and, when it is the latest run of the job, its alerts are evaluated as for any new report (a Failed run opens or refreshes a Backup failed alert, with an e-mail if the alert is new or more severe; a Completed one closes the open Backup failed and Backup missed alerts of the job). Remove the run again.

Alerts and e-mails

The module raises three types of alert of its own. They do not come from rules, so there is nothing to set up in the Rules tab of Alerts; they appear in the Alerts feed, in the footer ticker and in the Alerts tab of the client, linked to the computer that runs the job (for a destination that reads its free space from a computer, the Backup destination space alert is linked to that computer instead, see below). Their type, which you can read in the tooltip of the icon in the first column and pick in the Type filter (the filter lists only the types that are in the feed at that moment), is Backup failed, Backup missed or Backup destination space, shown with a floppy disk, an alarm clock and a file cabinet icon respectively. You can Resolve them and create a ticket from them like any other alert. The backup behind a Backup failed or Backup missed alert is also drawn in the Calendar tab, on the day it ran or should have run, as a Failed, Completed with warnings or Missed entry.

- Backup failed opens when the final report of the latest run of a job is Failed. It is Critical and reads Backup “name” failed followed by the error. It also opens, as a Warning, when a run is Completed with warnings but at least one item failed, for instance when one of two destinations is unreachable; the message then starts Backup “name” completed with errors and names the failed sources and destinations. It closes when a later run is Completed, or Completed with warnings without any failed item. A Cancelled run does not change it. A run that is Running and stops reporting for six hours is closed by the server as Failed, and raises this alert too.
- Backup missed is a Warning. It opens once for each missed time, after the delay set in Missed-run alert after, and reads Backup “name” did not run at the date and time in the job's time zone, adding that the computer has been offline since a given time when that is the case. It closes when a later run of the job finishes as Completed or Completed with warnings. It is only raised for enabled jobs of computers that are still in the app.
- Backup destination space is about a destination, on a computer that has an active job using it. It is a Warning when the free space is below Minimum free (GB) or Minimum free (%), and Critical when the agent finds that the next copy does not fit, which also fails the run. The Value column of the feed, which you turn on with Columns, shows the free space in GB. It closes when a check or a successful run sees healthy figures, for every computer when the destination is an FTP server or a network share; a Critical one is not lowered by a later reading that only says the space is low. Unknown space never raises it: a destination that cannot be reached at all shows as Failed items in the run and as Failed in Last check and, when another destination worked, as a Warning-level Backup failed alert. For a destination that reads its free space from a computer the alert follows that computer's reading instead: every five minutes SLApp Helpdesk compares the reading with the same minimums and, while an active job uses the destination, opens or updates the alert on that computer, and it closes the alert when the reading is healthy. Such a reading raises a Warning at most and never lowers a Critical that is already open; a reading older than 30 minutes, or one that does not list the disk, changes nothing, and what the agents report about that destination no longer opens or closes the alert. The check before each copy still fails a run whose file does not fit, and Backup failed says so, but it no longer raises this alert for that destination.

An e-mail is sent when a new alert opens and when an open alert becomes more severe (a Warning that turns Critical). It is not sent when an alert is refreshed or resolved. The recipients are the active users of your company who have an e-mail address and access to Backups or Alerts: the administrators, a technician whose modules include Backups or Alerts, and a technician with no module rule (who has Alerts, like every regular module); a technician whose modules include neither gets none (see Who receives the alert emails in the Alerts section). There is no per-job setting and no list of outside addresses for backup alerts. The e-mail is written in the e-mail language of your company, and since version 1.14.81 so is the alert's message (the names in it, and an error text the app does not know, stay as they were written). If you Resolve an alert by hand and the next run fails again, a new alert opens and a new e-mail goes out. Each e-mail that was sent, or that failed, is listed in Settings › Mails. Deleting a job, deleting a destination, or removing the computer in Computers closes the alerts that belong to them. An administrator can stop the three backup alerts for your whole company with the Generate alerts checkbox next to the module's title (see Turning alert generation on or off per module in the Alerts section); the runs, missed ones included, are still recorded.

Restoring from a backup

The module does not restore by itself. There is no Restore button, and nothing is ever written back to a computer: you fetch the file from the destination and restore it yourself, wherever you need it. Look at the date and time in the file name, which are UTC, to choose the copy. To restore a folder from a ZIP:

1. Open the destination and the job's folder and copy the ZIP you want to a disk with enough room, since the ZIP expands to the size of the original files.

2. Extract it into a new, empty folder, not over the live data: in Windows Explorer right-click the ZIP and choose Extract All, or use any ZIP tool. A ZIP with a password needs 7-Zip, WinRAR or WinZip, which ask for the password (see below): Windows Explorer cannot open it. Use a short folder such as D:\Restore: Windows Explorer cannot always extract very long paths.
3. Compare, then copy back the files you need. The ZIP keeps file names, folders and modified dates but not NTFS permissions, so restored files take the permissions of the folder you copy them into.

To restore a database from a .bak:

1. If the copy is a .bak.zip, take the .bak out of the ZIP first, with 7-Zip, WinRAR or WinZip when it has a password. Copy the .bak to a folder on the SQL Server computer that the account of the SQL Server service can read. It does not have to be the original computer: it is a standard SQL Server backup.
2. If you want to check it first, run RESTORE VERIFYONLY FROM DISK = N'D:\Restore\ERP__SalesDB__20261012-020000.bak' WITH CHECKSUM in SQL Server Management Studio or sqlcmd.
3. Run RESTORE DATABASE [SalesDB] FROM DISK = N'D:\Restore\ERP__SalesDB__20261012-020000.bak' WITH REPLACE, RECOVERY. REPLACE overwrites an existing database of the same name, and nobody may be connected to it: run ALTER DATABASE [SalesDB] SET SINGLE_USER WITH ROLLBACK IMMEDIATE first and SET MULTI_USER afterwards. On another server, or next to the original under another database name, add a MOVE clause for each data and log file whenever the original folders do not exist or are in use there; RESTORE FILELISTONLY lists them. A .bak can only be restored on the same SQL Server version as the one that made it or a newer one.

To open a ZIP that has a password, with 7-Zip for example:

1. Find its password. In the Runs tab open the run that made the copy (the date and time in the file name are its start, in UTC), click Password of this copy on the card of that file and copy the password. If several are listed, try them in that order (see The password of a copy). Runs are kept for 180 days: for an older copy of a source that is still in the job, use Reveal and Previous passwords in the job editor and pick the password by the date of the copy.
2. On the computer where you restore, install 7-Zip if it is not there; WinRAR and WinZip work as well.
3. Open the ZIP with 7-Zip, extract it into a new, empty folder and type the password when 7-Zip asks for it.
4. For a .bak.zip, restore the .bak that came out as described above for a database; for a folder copy, compare and copy back the files you need.

> **i** A COPY_ONLY backup restores like any other full backup; it is simply not a base for the differential backups of another tool. The module encrypts only the ZIPs of a source that has a password; apart from that, if a database uses Transparent Data Encryption, SQL Server encrypts its backup and the server that restores it needs the same certificate. And a backup you have never restored is a hope, not a backup: restore a ZIP and a .bak to a test location now and then, and open the result.

Limits and good practice

- Windows computers only, and only SQL Server instances installed on the same computer as the agent. Destinations are folders (local or network), FTP and explicit FTPS; SFTP, implicit FTPS and cloud storage are not supported.
- Every run is a full copy: there are no incremental or differential file backups, and the transaction log of a database is not backed up.

- The copies are encrypted only when their source has an archive password (AES-256 ZIPs, agent 1.2.0 or later). Without one, a ZIP or a .bak can be opened by anyone who can read the destination, so use Explicit FTPS rather than plain FTP and keep the destination protected. With one, nobody can extract the contents without the password, but the ZIP format still shows the names, sizes and dates of the files inside to anyone who can read the copy, and the name of the copy itself tells which job, source and database it is.
- From agent 1.2.0 an FTP upload whose connection breaks continues from the bytes the server already has, for as long as it makes progress (see When an FTP upload is cut). With older agents, or with a server that cannot append to a file, a transfer that breaks starts again from the beginning. Wrong passwords, certificate problems, permission problems and a full account are never retried.
- At most two jobs run at the same time on one computer, and a job never runs twice at once; the others wait their turn.
- A running backup holds back the agent's automatic update for up to six hours. Reinstalling the agent by hand, with the installer or the command from Download agent, does not wait and cuts the backup short: the run is reported as Failed with the error interrupted, and when the job allows late runs and its time is still inside the window, the agent tries it once more.
- Removing a computer in Computers uninstalls its agent and pauses its jobs, keeping their history. The Remove window tells an administrator how many backup jobs the computer has (“This computer has 2 backup jobs. They will be paused when the computer is removed (their run history stays).”; the number includes jobs that were already paused). The jobs stay in the Jobs tab as Paused, with no Next run. To carry on, open each job, choose the new computer in Computer (agent) (the removed one is listed as unavailable), click Save, then open the job's Actions menu and choose Resume.
- Keep at least one destination that is not the computer being backed up, and ideally one in another place. A copy on the same disk does not protect against that disk failing.
- The short commands the module sends to the agent appear in the Command queue under their own names, without an icon: backup-sync (sent by Run now) and backup-cancel (sent by Cancel run), both only when the computer is connected at that moment (an offline computer finds the request by itself when its agent reconnects), and backup-test-destination (sent by Test on a destination). They carry only identifiers. No antivirus exclusion is required for the module to work; if a real-time scanner slows large backups down, the folders the agent works in are the staging folder and, for SQL Server, the SLApp folder inside the instance's backup folder. The Calendar tab sends no commands of its own; the Cancel run button of a run opened from it sends backup-cancel, as in the Runs tab.
- Creating, changing, pausing, resuming, running (Run now), cancelling a run, removing a run from the history and deleting jobs, creating, changing and deleting destinations, and revealing archive passwords, the current one or the previous ones, are recorded in the Audit log under the entity types Backup job and Backup destination.

Troubleshooting

Open the failed run in the Runs tab (or click it in the Calendar: both open the same run detail), find the card that is Failed or Warning, and read its error. The messages below are the real ones, as the agent and the server write them: they are always in English, whatever language the app is set to. Each entry gives the cause and what to do.

- “Network path not found” or “Network name not found” for a share. Cause: the computer cannot resolve or reach the server, or the share name is wrong. What to do: open the share from that computer with its \\server\share path, check the name and the network, then run Test again from the same computer.

- “Access denied to \\server\share”, “Logon failure” or “not allowed to sign in”. Cause: the user and password in the destination are wrong, or the account may not write to the share. What to do: fix User and Password in the destination and make sure the account has write permission on the share. For a NAS or a workgroup, the computer's own access is not enough.
- “The drive of the destination does not exist”, “The destination path does not exist or its drive is not available” (destination) or “The source folder was not found or cannot be read by the agent” (source). Cause: a mapped drive letter such as Z: was used, or the path is wrong, or the disk or share is not connected; the agent service runs as SYSTEM and does not see mapped drives. What to do: use the network path \\server\share instead of the drive letter, open that path from the computer to check it, and make sure SYSTEM may read the folder.
- FTP “login refused - wrong user name or password”. Cause: the credentials are wrong or the account is disabled. What to do: correct User and Password, then Test.
- FTP “the server's TLS certificate was rejected”. Cause: the server has a self-signed or expired certificate, or a certificate for another name. What to do: for a server you control, tick Accept an untrusted / self-signed certificate (the error message calls this option “Trust certificate”); otherwise install a certificate from a trusted authority and use the name in the certificate as Server.
- FTP “the server does not support explicit FTPS (AUTH TLS)”, or “the TLS (FTPS) negotiation did not succeed”. Cause: the server does not offer FTPS, or only an older TLS. What to do: enable FTPS with TLS 1.2 or later on the server, or, on a trusted network, choose None (plain FTP).
- FTP “Connection refused”, “Timed out connecting” or “Cannot resolve the host name”. Cause: wrong address or port, a firewall, or the service is down. What to do: check Server and Port and the firewall, from the computer that runs the job.
- FTP “the data connection could not be established”. Cause: in passive mode the server's firewall must allow its passive port range. What to do: open that range on the server, or switch Passive mode.
- FTP “the server requires TLS session reuse”. Cause: the server insists on a feature this client cannot provide. What to do: turn that requirement off on the server (for vsftpd, require_ssl_reuse=NO), or use another server.
- FTP “permission denied or path not found”, or “Cannot create the folder ... on the FTP server”. Cause: the account cannot write to the Remote folder, or the folder is wrong. What to do: fix the Remote folder or the rights of the account on the server, then Test.
- FTP “no space left or the quota of the account is exceeded”, or “Not enough free space at the destination”. Cause: the destination is full. What to do: free space or enlarge it. Remember that retention deletes old copies after the new one is stored, so with the newest N kept the destination needs room for N+1 copies.
- An FTP copy is a Warning whose text starts with “The connection to … dropped N times during the upload” or “The FTP server at … did not answer within 120 s when asked the size of the temporary file right after a transfer it had not confirmed”, and goes on with “it was resumed each time from the bytes the server had”, “the server already had every byte: nothing was sent again” or “each time the upload started again from the beginning (…)”. Cause: the connection to the server broke, or the server stopped answering, during the transfer, and the agent finished the upload on a new connection; the end of the text says which checks against late data the copy passed (Checked for late writes of the dropped session …) or that the server allowed none (Not checked for late writes of the dropped session …). What to do: nothing if it is rare. If it keeps happening, look for what cuts long connections between the computer and the server, such as an unstable line or a router or firewall that drops connections after some time.
- FTP “a late write from the dropped session was detected”. Cause: after a connection broke, data from it still reached the server and was written into the copy, so the agent deleted that copy and sent the file again from the beginning. In a warning, the copy that stayed passed those checks again;

after three such restarts in a row the upload fails with this text. What to do: run the backup again and, if it repeats, check the network between the computer and the server as above.

- FTP “… failed after X of Y (…, N resumes): …”. Cause: the upload could not be finished, although it was resumed; the figures say how much the server had, how long it took and how often it was resumed, and the text after the colon says why it stopped, for example a connection that kept breaking without progress or that could not be opened again. What to do: act on that reason; the unfinished file is deleted, and the next run sends the file from the beginning.
- “Not enough free space in the staging folder”. Cause: the disk used to build the ZIP is too small; for a SQL Server source with Compress into a ZIP before sending, it must hold the ZIP of one database, about the size of its .bak. What to do: choose a Staging folder on a bigger disk in the job.
- “The staging folder ran out of space while the zip file was being written”. Cause: the size of a very large source could not be estimated in advance. What to do: choose a Staging folder on a bigger disk in the job.
- “nothing to back up”. Cause: the folder is empty, the Include only and Exclude patterns match nothing, or the agent cannot read the folder. What to do: check the path and the patterns, and that SYSTEM may read the folder or share.
- SQL Server “refused the login of the Windows account the agent runs as”, or “login failed for user”. Cause: SYSTEM has no login in the instance, or the SQL login or password is wrong. What to do: add a login for SYSTEM with enough rights, or switch the source to SQL Server login and enter a valid user and password.
- SQL Server “has no permission to back up database”, “Could not find out where the SQL Server instance stores its backups”, or “The SQL Server service account cannot write to the backup folder”. Cause: the login lacks rights, or the service account of the instance has no Modify permission on its default backup folder. What to do: use a login with the sysadmin role, or at least permission to back up the database and to see the backup folder, and give the service account Modify permission on that folder.
- SQL Server “was created by SQL Server, but the agent ... cannot read”, or “is not visible to the agent”. Cause: SYSTEM and Administrators have no read and delete permission on the SLApp folder inside the instance backup folder, or that folder is on a network share. What to do: grant the permission, or move the instance's default backup folder to a local disk.
- SQL Server “Not enough free disk space on the SQL Server backup volume”. Cause: the volume of the instance backup folder cannot hold the .bak files, which all stay there until the source is finished. What to do: free space, or move the instance's default backup folder to a larger volume.
- SQL Server “database 'name' NOT BACKED UP: ...”. Cause: this prefix is used for every database that could not be backed up; the text after the colon says why: the database is not online (state SUSPECT, RECOVERY_PENDING or EMERGENCY, or, for a database you listed by name, any state other than ONLINE), the login has no permission, there is not enough free space on the SQL Server backup volume, or the backup found damaged pages. What to do: act on the reason: bring the database online or repair it, and for “damaged pages” run DBCC CHECKDB on the database. The other databases of the source are still backed up, but the run is Failed because this one has no copy.
- SQL Server “Database(s) not found on SQL Server instance” or “No user database to back up was found”. Cause: a name in the list is misspelled, or the instance has no user databases that are online. What to do: check the spelling of the names (see the SQL Server source fields) and that the databases are online.
- SQL Server “points to another computer” or “runs on ..., not on this computer”. Cause: the instance is not installed on the computer that runs the job. What to do: only instances installed on the same computer as the agent are supported, so run the job on the computer that hosts the instance.

- “Verification failed”. Cause: the copy at the destination has a different size from the original (“has N bytes instead of M”), or SQL Server did not confirm the .bak (RESTORE VERIFYONLY). A copy that fails the check is removed, or never uploaded. What to do: run the backup again and, if it repeats, check the destination and the disk.
- “interrupted”. Cause: the agent service was stopped, updated or the computer restarted during the run. What to do: nothing, unless it repeats; check that the computer is not restarting at the backup time.
- “No report from the agent for 6 hours while the backup was running; it is assumed lost”. Cause: the computer or the agent stopped during the run and never reported again. What to do: nothing, unless it repeats; check that the computer is on and the agent is online (see Running under the run statuses).
- “Configuration error: A stored credential could not be decrypted”. Cause: the server could not read a saved password. What to do: open the job and the destination and type the password again.
- In the job editor, Reveal answers “The stored archive password could not be decrypted on the server (check the encryption key); edit the job and save the password again.”, or a password under Previous passwords or Password of this copy reads “This password could not be decrypted on the server (check the encryption key).” Cause: the server can no longer read the stored password, for example because its encryption key was changed. What to do: the copies made with that password open only with the password you noted outside SLApp Helpdesk, or once whoever runs the server restores the encryption key it had. For the next copies, type the password again in the job editor and save; until then the runs of the job fail with “Configuration error: A stored credential could not be decrypted”.
- In the job editor, “This computer's agent is too old for the archive options (ZIP, compression level, password): it needs agent 1.2.0 or later”. Cause: a source uses a ZIP of the SQL backup, a compression level other than Normal or a password, and the computer has not reported agent 1.2.0 or later yet. What to do: use Force update in the Remote actions tab of the computer's page (or Actions on N in the Computers list), wait for the agent to restart and click Save again, which checks the version again; or put the options back to their defaults to save without them. If Force update answers that a backup is running and the update is postponed, use it again once that backup has ended. Left alone, agents look for updates every 6 hours.
- “Configuration error: This job uses archive options (ZIP of the SQL backup, compression level or password) that need agent 1.2.0 or later”. Cause: the agent service that runs the backups on that computer is older than 1.2.0, for example because an older installer was run again. It received the job without its sources and destinations, so it made no copy rather than a copy without the password, and until it reaches the server it runs no backup at all, because it cannot read the job definitions that agent 1.2.0 saved on the computer. What to do: update the agent with Force update, or reinstall it with the command from Download agent; the next run works.
- A ZIP does not open, or its password is refused. Cause: Windows Explorer cannot open a ZIP that has a password, or the copy was made with another password than the one you tried (an earlier one, or none). What to do: open it with 7-Zip, WinRAR or WinZip and take the password from Password of this copy in the run detail, trying the passwords in the order shown and then Other previous passwords (see The password of a copy).
- A job does not start. Cause: it is Paused, the computer is off or its agent is older than 1.1.0, the Next run is still in the future, you saved the job after today's time, the same job is still running from an earlier time, or two other jobs are already running on that computer and this one is waiting for its turn. What to do: check Status and Next run in the Jobs tab, the computer's dot, the Runs tab and the Calendar (a time that has come and has not started is drawn there as Due now, and as Missed later), and use Run now to try.

28. Third-party backups

The Third-party backups module watches the backups that other software makes on your clients' Windows computers: Veeam Backup & Replication, Veeam Agent for Windows, Windows Server Backup and Acronis, plus any program that writes to a Windows event log, through custom events. It reads their results from the Windows event logs with the same read-only check that the Active Directory module uses, once a day on every Windows computer, lists every backup job it finds with its last result and last success, and opens an alert when a job fails or has had no successful run for longer than expected. Nothing is installed on the computers and the backup software is never touched. Open it from Infrastructure › Third-party backups in the sidebar. Every technician can use it, run a check, set how often a job is expected to succeed and manage the custom events; forgetting a job, deleting a custom event and closing an alert by hand are for administrators. The backups that SLApp Helpdesk makes itself are in the Backups module.

How jobs are found

The daily check (see The daily check in the Active Directory section) reads the events of the last 48 hours of each supported product, plus its latest successes however old they are, up to ten per job. A job appears in the list after the first check that sees it, with no setting on your side. A job is identified by its product and its name, ignoring capitals, and the module keeps its last 20 runs. The event logs keep old results for months, so a job that only shows in that history, with no run in the last 48 hours and no success in the last 8 days, is not added: a job removed or renamed long ago does not come back as a failing one.

- Veeam Backup & Replication: the Veeam Backup event log. Each event that reports the end of a job session gives a run, its result taken from the job result Veeam writes in it (success, warning or failed) or, when there is none, from the level of the event; the duration comes from the event that started the same job.
- Veeam Agent for Windows: the Veeam Agent event log, in the same way.
- Windows Server Backup: the Microsoft-Windows-Backup event log. Event 4, a backup that finished successfully, is a success, and every error or critical event of that log is a failed run. All of them form one job per computer, named Windows Server Backup, which appears once a backup has run.
- Acronis, on a best-effort basis: Acronis documents neither its event sources nor its event ids, and writes to the Windows event log only when that option is switched on in its settings, which it is not by default. The check reads the Application log entries of the sources whose name contains Acronis: errors are failed runs, warnings are warnings, and an information entry counts as a success only when its text says the backup succeeded. When Acronis is installed but no backup event is found, the computer's page says so.
- Synology Active Backup for Business: its results stay on the Synology NAS, so the check only notes that it is installed, and the computer's page says that its results cannot be read here.
- Custom events: for any other software, you name the event log, optionally its source, and the event ids of a success, a failure and a warning (see Custom events).

The list

The module lists every job of every active Windows computer, sorted by client, computer, product and job. The top of the page has Custom events, which opens the custom events of every computer, and Columns. Under the introduction come a client list (All clients, — no client — or a client that has jobs), a search box that looks in the job, the computer and the product, a product list (All products or one of them) and the Show list: Show all, With problems, Overdue, Ignored and No result yet. With problems keeps the jobs that are not ignored and whose last run failed or had warnings, that are overdue or that

have an open alert. When filters hide jobs, a counter such as 8 of 31 says how many are left. Until the first job is found the page explains how the jobs are read. The list refreshes itself every minute, and every 15 seconds while a check of an online computer waits for its result.

- Client, and Computer with its online dot; the name is a link that opens the computer's page.
- Product and Job; hover the job for the message of its last run.
- Last result: Success, Warning, Failed, or No result while no run has been seen.
- Last success, or Never, with a red Overdue badge when the job has gone too long without one; hover the badge for the expected period.
- Last run, and Alerts, the number of open alerts of the job.
- Expected: the expected frequency, which you can change right in the list (see Expected frequency and overdue jobs).
- Duration of the last run, First seen, Last seen and Message are available through Columns.

Click a header to sort by that column and again to reverse it. The row under the headers filters Client, Computer, Job and Message by text, Product, Last result (Success, Warning, Failed, No result), Expected and Alerts (With alerts, Without alerts) by their values, Last success by Overdue, Last 24 hours, Last 7 days, Older or Never, Last run, First seen and Last seen by the same periods, and Duration by Under 1 hour, 1 hour or more or Unknown. Click a row to open the job window. At the end of the row, the refresh button checks the computer now and, for administrators, the bin button forgets the job. The right-click menu offers Open details, Check now, one Expected: entry per frequency (the current one is ticked), Open computer, Custom events of this computer…, Copy job name and Copy computer name, and Forget job for administrators.

Expected frequency and overdue jobs

Each job has an expected frequency, how often it should succeed: Daily, Weekdays (Mon-Fri), Weekly or Ignore. When a job is first seen, the module works it out from the dates of its successes, and afterwards it never changes it by itself. Any technician can change it in the Expected column, in the job window or with the right-click menu; a notice confirms it, for example Expected frequency of Nightly NAS: Weekly, the job's alerts are evaluated again at once, and the change is written to the Audit log under the entity type Third-party backup.

- How it is worked out: with fewer than three successes, Daily. Otherwise from the typical gap between successes: up to a day and a half gives Weekdays (Mon-Fri) when every success fell on a weekday and no gap was longer than three and a half days, and Daily otherwise; up to 8 days gives Weekly; anything longer gives Daily.
- When a job is overdue: when it has had no success for its expected period plus one day, that is 48 hours for Daily, 48 hours counted on weekdays only for Weekdays (Mon-Fri), so a weekend does not count, and 8 days for Weekly. The time counts from the last success or, for a job that never succeeded, from when it was first seen.
- Ignore: the job stays in the list but never raises an alert, any open alert of it closes, and it is left out of the client report. Use it for a job that still runs but needs no watching.

The job window

The title is the job's name with its product. Under it come the computer, with its online dot and a link to its page, the client, the time of the computer's last check and Check in progress while a check waits for its result, with the buttons Check now, Custom events for a job that comes from a custom event, and Forget job for administrators (not for a custom event's job). Last result shows the last result, the last success (with Overdue when it applies), the last run, its duration and its message. Monitoring has the Expected list, First seen, Last seen and a reminder of when an alert opens. When the backup part of

the computer's check has been failing, a red line says how many times and why. Recent runs lists up to 20 runs with When, Result, Duration, Event ID and Message, and Open alerts lists the open alerts of the job, and the computer's Backup check failing alert when there is one, with their Severity, Type, Message and Opened at; administrators also have a Resolve button. While a check of an online computer waits for its result, the window looks again every 10 seconds.

Custom events

A custom event turns entries of any Windows event log into a backup job of one computer, for backup software that the check does not know. Custom events at the top of the module opens the custom events of every computer, and Custom events of this computer… in the right-click menu of a job, or the Custom events button of the computer's page, opens those of one computer. The window lists them with the Computer and Client (in the window of every computer), Name, Event log, Source, the success, failure and warning ids, Last read, which says Not read yet until the next check, and Last error, for example Event log 'Backup' not found on this computer. The pencil button, or Edit in the right-click menu, changes one; Delete is for administrators.

1. Click + Add custom event.
2. In the New custom event form, choose the Computer when you opened the window of every computer; the list holds the Windows computers, grouped by client.
3. Type a Name, for example Nightly backup to the NAS, and the Event log as Event Viewer names it, for example Application or a log of its own. Source is optional and limits the event to that source.
4. Fill in Success event ids (required), and Failure event ids and Warning event ids if the software has them, separated by commas.
5. Click Save. A notice says that the event is read on the computer's next check; Check now on the computer reads it at once. After that check the custom event is a job of product Custom event, listed even while it has no run yet.

- Each check reads the events with these ids from the last 48 hours, and the latest successes however old, and counts each one as a success, a failure or a warning according to the list its id is in.
- A computer can have up to 10 custom events. The ids are whole numbers from 0 to 65535, at most 20 in total, and an id can be in one list only. The name has up to 100 characters, and the event log name may hold letters, digits, spaces and - _ / . ( ).
- Renaming a custom event renames its job, and its open alerts follow the new name. The job of a custom event cannot be forgotten: delete the custom event instead.
- Deleting a custom event (administrators) asks Delete the custom event … of …? Its job leaves the list and its open alerts are resolved.
- Adding, changing and deleting custom events are written to the Audit log under the entity type Third-party backup.

Forgetting a job (administrators)

Forget job, in the job window, at the end of the row and in the right-click menu, is for a job that was removed or renamed in the backup software. It asks Forget … on …? It leaves the list and its open alerts are resolved. It only comes back if the job runs again — its old events in the computer's logs never bring it back. After Confirm, the job leaves the list, its alerts close and the forgetting is written to the Audit log. If the job runs or succeeds again later, it comes back as the same job with its monitoring started afresh. For a job that still runs but needs no alerts, set its expected frequency to Ignore instead.

Check now

Check now, in the job window, with the refresh button at the end of a row, in the right-click menu or on

the computer's page, runs the whole daily check of that computer at once: every backup job of the computer and, on a domain controller, its Active Directory part. A notice says that the check is queued, or that the computer is offline and the check runs as soon as it is back online. The result arrives when the computer has run the check, usually within a minute or two. A computer can be checked at most once every 5 minutes, and never while a check of it is still waiting for its result.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. Their source is Third-party backup, and they appear in the Alerts feed, on the Dashboard, in the footer ticker and in the Alerts tab of the client, with the computer and the job in front of the message, as in SRV01 / Nightly NAS; you can create a ticket from them like from any other alert. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- Third-party backup result: Critical when the last run of the job failed and a Warning when it finished with warnings, with a message such as Veeam Backup & Replication job "Nightly NAS" failed at 2026-10-30 02:14 UTC followed by the message of the event. It closes when a later run succeeds.
- Third-party backup overdue: Critical when the job is overdue, with a message such as No successful run of Windows Server Backup job "Windows Server Backup" since 2026-10-27 (expected daily). It is evaluated at every check and also every few minutes between checks, so it opens on time even when nothing new is logged, and it closes once a success is read. One that an administrator resolved by hand is not opened again until the computer's next check.
- Backup check failing: a Warning about the computer when the backup part of its check failed three times in a row and the computer has at least one job that is not ignored. It closes when the check reads the event logs again, or when nothing is left to watch.
- On the first evaluation of a job, a Warning opens without an email; Critical alerts always send one. Emails go to the active administrators of your company who have an email address, when an alert opens. The email is written in the email language of your company, and since version 1.14.81 so is the alert's message (the job's name and the message of the event stay as they were written).
- The alerts of a job close when it is set to Ignore, forgotten or deleted, or when its computer is removed, and follow the job when the computer or the custom event is renamed or the computer moves to another client. Only an administrator can close an alert by hand, with Resolve, and that is written to the Audit log under the entity type Third-party backup.

> **i** Each Windows computer's Overview lists its own jobs, the backup software found and its custom events under Active Directory & third-party backups (see the Computers section). The client report has a Third-party backups section (PDF and Excel) with every job that is not ignored, its last result, last success, expected frequency and whether it is overdue; the full report leaves it out for a client without such jobs.

29. Syslog

The Syslog module is a "Log Center" for the network gear that has no agent of its own — routers, switches, NAS units, access points. It works through a computer you already monitor: turn on its collector, point the device's own syslog setting at that computer's IP and a UDP port, and every line the device sends is parsed, stored and searchable here, with the search covering both the reporting device's name and the message itself so a device never goes silently unfound. Open it from Infrastructure › Syslog in the sidebar.

The list

The table shows Received, Severity, Device (the router, switch, access point or NAS that sent the line,

shown by the name it reports with its IP under the name; a line that carries no name shows the IP alone), Collected by (the computer whose agent received it, with its client underneath), Category (the syslog facility, e.g. kern, auth, local0), Program (the program that wrote the line or, for gear that sends no program name, a code taken from the message itself — the category it starts with, such as Arp, DOS, AAA or Port, or the message number or ID the device gives — or a dash when there is none) and Message. Click a column header to sort by it; click again to reverse. Filters above the table narrow the search on the server: free text (device name, its IP, program and message), Client, Computer (only computers actually receiving are offered here — see Enabling a collector below), Severity (Critical or worse / Error or worse / Warning or worse) and a From/To date range. A second, lighter row of column filters refines what is already loaded, the same way as the Audit log.

- Choosing a Client narrows Computer to that client's collectors and clears any previous computer choice.
- Choosing a Computer is enough on its own — the server does not need the client as well.
- Entries are kept for 14 days; older ones are removed automatically.
- Received is the moment the line reached the collecting computer. From agent 1.0.45 each line is stamped when it arrives, so lines that reach the server later (for example when a batch could not be delivered at once and was sent again) keep their real time instead of all showing the moment of upload. With an older agent, or when the stamp is not believable (more than 5 minutes ahead of the server clock, or more than 14 days old), the time the server received the batch is used instead. From agent 1.2.0 the collector also keeps on disk the lines it could not deliver, up to about 50 MB or 14 days, so an outage of the server or an orderly restart of the agent does not lose them (a crash or a power cut can still lose the last few seconds of lines): they arrive later, each one once, with their real time.
- By default the list is ordered by Received, newest first, and lines with the same time always keep the same order. Load more continues right after the last line shown, so no line repeats or goes missing between pages, even while new lines keep arriving.

What the receiver understands

Network gear does not agree on a log format, so the receiver on the collecting computer recognises the shapes that real routers, switches and access points send and fills in Device and Program as well as each line allows. Nothing is dropped because it does not fit: a line the receiver cannot read is kept exactly as it arrived. The reading is done by the agent on the collecting computer, so it applies once that computer runs agent 1.0.44 or later.

- Standard syslog lines (RFC 3164 and RFC 5424) carry the device name and the program name themselves, and are read from there.
- Routers, switches and access points of the Vigor family send the device's name, then its MAC address, then the message, with no host name. They appear under the device's own name, and Program shows the category the message starts with (Arp, DOS, AAA, Port...) instead of the MAC address.
- Access points that send a priority value above the standard range (the Vigor access point is one) keep their severity and show a dash under Category, because the facility they send is not one of the standard ones.
- TP-Link switches, whose lines start with a date such as 2026-10-05 13:09:43, appear under the switch's name, and Program shows the message number.
- Vigor switches that put the year in the date (for example Oct 05 2026 13:01:08) appear under their name, and Program shows the code the message carries, such as AAA for a %AAA-5 message.
- Dell iDRAC sends no date and no name, only Severity, Category, MessageID and Message. Its lines

appear under the iDRAC's IP, Program shows the message ID (for example RAC1195), and the category is kept at the start of the message in square brackets.

- Anything the receiver still does not recognise is kept as received: the whole line is the message and the IP of the sender is the device.

> ℹ Lines received before the update keep the names and programs they were stored with until they expire after 14 days. For a while the same device can therefore appear under two names — an older one that may end in a colon and show a MAC address under Program, and its real name for the lines received since.

Enabling a collector

Nothing listens on the network until you turn a collector on for a specific computer. Click Collectors in the module header to open the Collectors panel, which — unlike the main list — shows every Windows computer of the tenant, since this is where you turn a new one on.

1. Find the computer that is already on the client's network (the one you want the router/switch/NAS to send its logs to). A client and folder filter at the top help narrow a long list.
2. Tick Receiving to turn the collector on, and set the UDP port it should listen on (514 by default).
3. Save. Point the network device's own syslog client setting at that computer's IP address and the port you chose.
4. The first lines usually arrive within a few seconds of the device sending them.

- A computer whose agent is too old to run the receiver is marked warning: agent too old in the panel — turning the switch on would do nothing until the agent is updated.
- Entries column shows how many lines a collector has received in total, so you can tell at a glance whether a device is actually reaching it.
- Alert on critical/error (see Syslog alerts below) only becomes available once Receiving is on, and switches itself off if you turn receiving back off.

> ℹ A busy collector on one client can dominate a free-text search across the whole list if you do not pick a specific client or computer first — narrow the search when you are chasing a particular device.

Reading a message

Click a row's message to open it in full, untruncated, in a small popup. From there, Ask the AI sends the AI Assistant the whole entry — device, collector, severity, category, program and message, plus the collecting computer's internal ID, so it may use its own tools on that machine if that helps make sense of what happened. Every row also has a Create ticket button, which opens Tickets with a new ticket already filled in from that entry (client, computer, subject and description).

Refreshing the list

Refresh in the header lets the list refresh itself automatically, every 5, 10, 20 (the default), 30 or 60 seconds — a quiet background reload that never resets your scroll position or flashes back to a loading screen. Your choice is remembered in this browser for next time.

Ignoring known harmless ARP lines

A router's ARP-spoofing defense logs "addr mismatch" whenever a MAC address and an IP address do not line up, which is usually a useful security signal. A few setups cause it harmlessly, most commonly the multicast address of a Microsoft NLB cluster, whose MAC deliberately encodes its own IP. Click ARP exceptions in the module header to list the exact pairs to ignore.

1. In the ARP "addr mismatch" exceptions window, click + New exception next to the title. In the New ARP exception window, type the IP address (for example the cluster's virtual IP) and the MAC address the router reports for it, written like aa:bb:cc:dd:ee:ff. A note, such as the name of the cluster, is optional.
2. Click Create. The pair appears in the table; click Remove on a row to delete an exception again.

Only an ARP "addr mismatch" line that contains both the IP and the MAC of a pair you listed is ignored. It is discarded before it is stored, so it neither appears in the list nor raises an alert. Every other "addr mismatch" line is kept, so real spoofing detection is not affected. Lines that were stored earlier are not removed, and the exceptions apply to your whole company, whichever collector receives the line.

> **i** Before hiding these lines, make sure the address is really reachable through the router: the router's spoofing defense may be refusing it, and the line can then be the only visible sign of a real problem. Removing the exception makes the lines appear again from that moment.

Hiding noisy lines with line filters

Some devices write the same harmless line over and over, for example a router that logs every accepted Telnet login from your own PC. A line filter tells SLApp what to do with such lines without touching the rest of that device's log. It is more general than the ARP exceptions above, which stay exactly as they are for the ARP "addr mismatch" case. Click Line filters in the module header, next to ARP exceptions, to open the Syslog line filters window. Anyone can open it and read the filters; creating, turning on or off and removing a filter needs an administrator.

A filter is deliberately narrow. It applies to one device, given by its exact IP address, and only to the lines from that device that contain all of the one to three texts you give. The texts are not case-sensitive, each one needs at least 3 characters (spaces do not count) and at most 100, and a line that lacks any one of them is never touched. Because a filter always needs a device and at least one real text, it can never cover a whole device or all lines. Each filter does one of two things:

- Do not alert: the line is still stored and searchable, but it never opens a Syslog alert and never refreshes one that is already open. An open alert still closes by itself after 24 hours without a line that counts, as described under Syslog alerts.
- Discard: the line is thrown away the moment it arrives. It is never stored, so it cannot be searched, counted or recovered later, and it never alerts. Lines stored before the filter existed are not removed.

1. Open Line filters and click + New filter, next to the title of the window; the filter is filled in a window of its own. Alternatively, right-click a line in the list and choose Filter lines like this…: the window then opens with that line's device and a suggested text already filled in.
2. Under Device, pick one of the devices that sent lines in the last 3 days (the list offers up to the 100 busiest) or type its exact IP address.
3. Under The line must contain all of, type one to three texts: a recognisable piece of the line, such as a source address. At least one is required.
4. Choose the action: Do not alert (the default) or Discard. You can add a Notes entry, up to 200 characters, saying why the filter exists.
5. Read the message under the form: when you started from a line it says whether the filter matches that line, and it says how many of the lines from that device currently loaded in the list would match. This is a quick check on what is loaded, not a count over the whole history. Then click Create. For Discard, a confirmation (Discard matching lines for good?) appears first.

The table in the window shows, for every filter, the Device, the Text (all must appear), the Action, Lines handled (how many lines the filter has handled, with the date and time of the last one, or never) and an

On box. Untick On (or right-click the row and choose Turn off) to pause a filter without losing it. Remove deletes a filter after a confirmation: matching lines are then treated normally again, but lines that were already discarded do not come back. The device, texts and action of a filter cannot be edited; remove it and create a new one instead. A company can have at most 200 filters, and the same device with the same texts cannot be added twice, even typed in another order.

The server applies the filters when each batch of lines arrives, so they work with any agent version, only affect lines that arrive after the filter is saved, and, like the ARP exceptions, apply to your whole company whichever collector receives the line. When several filters match one line, Discard wins over Do not alert, and the line is counted on one filter only.

> **i** Prefer Do not alert: the line stays searchable, and you can always see what was hidden. Use Discard only when the volume itself is the problem, and check the Lines handled counter after a while: a filter that handles far more lines than you expected is catching too much.

Syslog alerts

Each collector has its own Alert on critical/error switch, off by default, in the same Collectors panel. Turning it on makes the collector raise a normal alert, visible in the Alerts feed with the type Syslog, whenever a device sending to it reports a line of Error severity or worse — Error, Critical, Alert or Emergency, the same threshold as Error or worse in the Severity filter. The alert itself is always Critical, and no email is sent for it. A line that reaches the server more than 24 hours after the collector received it (a backlog delivered after a long outage) is stored and searchable but does not open or refresh an alert. An administrator can stop the Syslog alerts of every collector at once with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section); the lines are still received, stored and searchable.

- There is one alert per device, not one per collector. Each router, switch or access point sending to the collector (told apart by its IP address) gets its own alert, so a noisy device does not hide a real problem on another.
- The alert shows the most severe matching line of the latest batch of lines the device sent (the collector forwards lines every few seconds), and its message starts with the device's name.
- While the alert stays open, further matching lines refresh that message and add to a counter. The counter is the Value column of the Alerts feed, which is hidden by default — turn it on with the Columns button.
- The alert closes by itself after 24 hours without a new matching line from that device. The check runs every ten minutes, so it can take a little longer. You can also click Resolve on it earlier.
- Once an alert is closed, a new matching line from the device opens a new alert.

Some lines never count towards alerts, although they are stored and searchable like any other:

- The bare --SendMailAlert-- marker that the Vigor routers write just before they send their own mail alert. It carries no information.
- DrayTek DoS Defense blocks of the scanner or traceroute kind (the types trace_route and tcp_flag) coming from a public Internet address. They are constant background noise on a router that faces the Internet.
- Lines that arrive with no severity at all.

Every other line of Error severity or worse still alerts, including DoS Defense floods and DoS Defense blocks coming from a local or private address, unless you have asked for it not to. An ARP "addr mismatch" pair that you have listed under ARP exceptions in the module header is dropped before it is stored, so it neither appears in the list nor alerts. A line covered by a line filter either stays out of the

alerts (Do not alert) or is dropped before it is stored (Discard); see Hiding noisy lines with line filters above.

> **i** Some devices label routine lines with a severity of Error or worse — a mesh Wi-Fi node logging its everyday kernel messages as Error is a typical example. The collector can only go by the severity the device gives, so such a device keeps raising alerts until its log level is lowered on the device itself.

30. Network configs

Network configs keeps a copy of the configuration of your clients' routers, firewalls and switches. Once a day, the agent of a Windows computer on the same network as the device, the collector, signs in to the device over SSH with a login from the client's password vault, reads its configuration and sends it to SLApp Helpdesk, which stores it as a new version whenever it differs from the last one. You can compare any two versions, download one as a text file and announce a planned change so that it raises no alert, and you are alerted when a configuration changes without warning, when the backups keep failing or when a device presents a different SSH key. The agent only reads: it never changes anything on a device. Open the module from Infrastructure › Network configs. Every technician sees the devices and their versions, starts a backup with Back up now and announces a change with Expect a change; adding, editing, testing and removing devices, trusting a new SSH key and opening the versions that contain passwords or keys are for administrators.

What it needs

- A collector: a Windows computer of the same client, with the agent at version 1.2.29 or later, that can reach the device's SSH port. The backups run from that computer, so it has to be online when they run. A Linux or macOS computer cannot be a collector. A Windows computer with an older agent can be chosen, but its backups fail until the agent is updated, and the list and the device's window say so.
- A login in the client's password vault (the Passwords tab of the client) with a user name. The form only lets you pick it, so the password is never typed there. The collector receives the user name and the password over HTTPS for the duration of one backup; the agent never writes them to its logs nor keeps them, and masks the password with asterisks in its error messages. The systems that have an enable mode can use a second vault entry for the enable password.
- A login the device allows to show its configuration over SSH from the collector.

Supported systems and what is copied

Each device has a System. The agent signs in, runs the commands below, answers the device's --More-- pages by itself and keeps what they print:

- Fortinet FortiGate (FortiOS): show, the whole configuration, as in the web interface's backup. The values FortiOS encrypts again on every export are left out of the comparison.
- MikroTik RouterOS: /export terse. Passwords and keys are left out unless the device is set to export them (see Passwords and keys below).
- Cisco IOS / IOS-XE: terminal length 0 and show running-config, entering enable mode first when the login lands at the > prompt.
- Cisco Small Business (SG / CBS): terminal datadump and show running-config, entering enable mode when needed.
- HPE Aruba ArubaOS-Switch / ProCurve: no page and show running-config, entering enable (manager) mode when needed.

- HPE Aruba CX (AOS-CX): no page and show running-config.
- Ubiquiti EdgeOS / UniFi gateway: terminal length 0 and show configuration commands.
- pfSense and OPNsense: the file /conf/config.xml, read over SSH. The login needs shell access (on pfSense, admin or a user with the User - System: Shell account access privilege; on OPNsense, root or a member of the admins group with a login shell), and the file always contains private keys and password hashes.
- Juniper Junos: show configuration | display set | no-more.
- Linux / other (custom command): the one-line Command you type, for example cat /etc/network/interfaces; what it prints is the configuration kept.
- DrayTek Vigor (DrayOS): sys version, show lan, show dns and srv nat showall. This is a snapshot of the main settings, marked settings snapshot in the list, that shows what changed but cannot be restored: DrayOS has no text export of its configuration. For a copy you can restore, use the router's own System Maintenance › Configuration Backup, an encrypted .cfg file.

> **i** Each device gets 60 seconds for the connection and all its commands, and a collector backs up at most four devices at once. A configuration larger than 20 MB is refused. On a DrayTek the agent never sends a command starting with sys cfg, which would reset the router to its factory settings.

Adding a device (administrators)

Click + Add device. The Add network device window asks for:

- Client *: the client the device belongs to. It cannot change later: to move a device to another client, remove it and add it again.
- Address and collector: Linked to an SNMP device or Standalone. Linked uses an SNMP device the client already has on a computer's Devices tab (printers cannot be linked; a device already backed up is marked so): the device's address and the computer that polls it become the address and the collector of the backups, and follow it when they change. Standalone asks for a Collector computer, chosen among the client's computers (the ones that cannot run the backups yet are marked not Windows or with their agent version), and the Address (IP or host name).
- SSH port (22 by default) and Name, optional, for example Office firewall. Without a name the device is shown by the SNMP device's label or by its address.
- System *: see Supported systems and what is copied above; the hint under it says what is run. Linux / other also asks for its Command *.
- Login (vault entry) *: an entry of the client's password vault with a user name, listed with its user name and, for an entry of a computer, the computer's name. For the systems with an enable mode, Enable password (vault entry) is only used when the login lands at the > prompt and enable mode asks for a password; Same as the login is the default.
- The passwords and keys box: see Passwords and keys below.
- Alert when the configuration changes: Off, Warning (the default) or Critical, for a change nobody announced with Expect a change.
- Scheduled daily backups: ticked by default. Unticked, the device is paused: no daily backup and no alerts, while Back up now still works.

Save adds the device. Each device gets its own random time of day for its daily backup, so that a client's devices do not all run at once, but the first backup does not wait for it: it is taken within a minute and becomes the device's first version, with no change alert. Test connection saves the device and asks the collector to connect and read the configuration right away, then shows the result in the window: Connected and the number of lines read, with the device's SSH host key, or the reason it failed. A test

that succeeds stores the configuration like any other backup. Each device can be added only once per client: an SNMP device has at most one backup, and so has an address and port.

The device list

The table has a row per device: Device (its name, with the address and port underneath, a padlock when its versions contain passwords or keys, a satellite-dish icon when it is linked to an SNMP device, and a Paused badge), Client, System, Collector (marked offline, with the agent version when it is too old, not Windows or removed when it cannot run the backups), Last backup (when, Succeeded or Failed — hover a failure to read the reason — and Waiting or Running while a backup is under way), Versions, Last change (when the configuration last changed), Schedule (Daily at a time of day, in your own time zone, with the next run on hover, or Paused) and Alerts (the open alerts). Host key, Login and Added can be shown with Columns.

Above the table, a client picker, a search box and a Show picker narrow the list: Show all, With open alerts, Last backup failed, Changed in the last 7 days, Never backed up, Paused and Change expected. The row under the headings filters by device, client, system, collector (Online, Offline or Needs attention), last backup (Succeeded, Failed or Never) and alerts. Each row ends with a refresh button for Back up now and, for administrators, a pencil button for Edit and Remove. Click a row to open the device's window, or right-click it for Open, Back up now, Expect a change…, Edit, Test connection, Trust the new key (when the device presents a new key), Copy address, Copy host key and Remove; the ones for administrators are only shown to them.

Back up now

Back up now, open to every technician, asks the collector for a backup at once: the button shows Backing up… and a message then tells you how it went, First version stored, New version followed by its number or No change, with the number of lines read, or why it failed. It needs the collector online: when it is offline you are told straight away and nothing is queued, and only one backup of a device runs at a time. A scheduled backup, on the other hand, waits up to six hours for an offline collector to come online before it counts as failed. Every backup reaches the collector through the command queue, so it also appears in Automation & updates › Command queue, as a Network configuration backup that carries nothing but an identifier, never the login.

The device's window

Click a device to open its window. The top line shows the address and port, the system, the client and the collector, with Back up now and Expect a change… for everyone, and Edit and Test connection for administrators. Below it:

- Backups: Last backup (with the reason of a failure), Last successful, Last change, Schedule, Failures in a row (the scheduled backups that failed one after the other) and Change alert.
- Access: Login (vault), Enable password when the system has one, Address and collector (From the SNMP device followed by its name, or Its own (standalone)), Passwords and keys (Included — administrators only, or Not included) and SSH host key (the fingerprint pinned and since when).
- Versions: every stored version, newest first, with Captured (when, and whether it came from a Scheduled, Manual, Test or Automatic backup), Lines (and the size) and Changes (the lines added and removed, or First version). A version captured while a change was expected carries an expected badge. Each version has View, Download and Compare with previous; tick two versions and click Compare selected to compare any pair.
- Last runs: the last 20 backups, with When, Kind, Status, Result, Duration, Version and By (who asked for it, or System for the scheduled and automatic ones).
- Open alerts: the device's open alerts, with Resolve for administrators.

A new version is only stored when the configuration really changed: before comparing, the lines that change on every backup (timestamps, values the device encrypts again on each export, pager prompts) are left out. Every version of the last 30 days is kept, and at least the newest 50. Download saves a version as a .txt file named after the device, the version and its date. Saving new settings for a device cancels a backup still running with the old ones (Cancelled: the device's settings changed before the backup finished), which does not count as a failure.

Comparing versions

Compare with previous and Compare selected open Changes from version X to version Y. Unified shows the removed and added lines in one column and Side by side puts the two versions next to each other (your choice is remembered), while Lines of context shows 3 lines or 10 lines around each change, or the Whole file. The lines that change on every backup are left out here too, so two versions that only differ in them show No difference. When the versions differ in too many places, or the comparison is too long to show in full, the window says so: download both versions to compare them elsewhere. The email of a configuration change alert has a See the changes button that opens this comparison.

Passwords and keys

A configuration can contain passwords and private keys, and the box in the device's form says whether its versions do. For MikroTik it reads Export with passwords and keys (show-sensitive), off by default: unticked, the export leaves the passwords out; ticked, it includes them. For pfSense and OPNsense it is always ticked: Contains passwords and keys — always, for this system. For the other systems, This configuration contains passwords or keys in clear marks the versions as containing them, without changing what is read. Those versions are marked with a padlock: only administrators can view, download or compare them (the buttons are disabled for other technicians), and every view, download or comparison of one is recorded in the Audit log under Network config backup.

The SSH host key

The first time the collector connects to a device, the device's SSH host key is pinned: its fingerprint appears under Access as SSH host key, with Pinned on and the date. Every later backup checks it. When the device presents a different key, the backup fails before the login is sent, the device's window shows The device presents a different SSH host key with both fingerprints, and a critical alert opens. The device's backups stay stopped until an administrator clicks Trust the new key and confirms in the Trust the new SSH host key window; that key is then pinned instead and the alert closes. Only trust it when you know why the key changed (the device was replaced or reset, or its firmware was reinstalled) and after checking the fingerprint on the device itself: otherwise someone may be intercepting the connection. If the device shows its pinned key again, there is nothing left to trust and the alert closes by itself.

> **i** Changing a device's address or SSH port, or its link to an SNMP device, in its form starts afresh: the key is pinned again at the next connection, and an open SSH host key changed alert closes. Changing the address of a linked SNMP device on the Devices tab, which any technician can do, keeps the pinned key and expects it at the new address, so a different device answering there is refused until an administrator trusts it; the move is recorded in the Audit log.

Expect a change

About to change a device yourself? Click Expect a change… (any technician), choose For 1, 2, 4, 8, 24, 48 or 72 hours and click Expect the change; doing it again replaces the window already set. Until then, a backup that finds a different configuration stores it as an expected version, without a change alert, and the scheduled backups and Back up now go on as usual. When the time is up, one closing backup

is taken automatically (Automatic in Kind) and becomes the new reference, also without an alert. The device's window shows Change expected until the date and time, with who set it, and Stop expecting ends the window early, after which no closing backup is taken and the next change alerts as usual.

Changing a device's system, command or passwords option, its address or its link has the same effect on its next backup, which becomes the new reference without a change alert: the window then shows The next backup becomes the new reference — it raises no change alert. A new device's first backup is that reference too.

Network config alerts

The module raises client alerts that appear in Alerts (source Network config), on the Dashboard and in the footer ticker under the device's name; right-clicking one in Alerts offers Open the network device and, for a change, Show the changes. Every active administrator with an email address is emailed when one opens and when it becomes critical; the first evaluation of a device opens its warnings without an email, while a critical alert always sends one. Only administrators can resolve them by hand.

- Network configuration changed: a backup found a configuration different from the previous version while no change was expected. Its severity is the device's Alert when the configuration changes, Warning by default or Critical (Off raises none). The message gives the version and the lines added and removed. Each change opens a new alert, closing the previous one, and each closes by itself after 7 days.
- Network config backup failing (warning): two scheduled backups failed in a row; the message says how many and why. It closes at the next successful backup. A failed Back up now or Test connection is not counted.
- SSH host key changed (critical): the device presented a different host key, and its backups are stopped (see The SSH host key). It closes when an administrator trusts the new key, or when the device shows the pinned key again.

Pausing a device by unticking Scheduled daily backups closes its open alerts, and a paused device raises none. Removing a device (administrators: Remove, then Confirm in the Remove device window) deletes it with all its stored versions and closes its alerts. Adding, editing and removing devices, trusting a key, announcing a change and opening a version with passwords or keys are recorded in the Audit log. To stop the alerts of every device at once, an administrator unticks Generate alerts at the top of the page (see Turning alert generation on or off per module in the Alerts section); the daily backups and their versions go on.

31. Printers

Printers lists the network printers of all your clients with the level of their toner, ink and other supplies, read over SNMP with the standard Printer MIB by the agent of a computer on the printer's network, and opens an alert when a supply runs low. A printer is an SNMP device of type Printer: it is added and set up on the Devices tab of that computer, like a NAS or a router, and this module gathers every printer of your company in one list. Printers can also be found with a scan: the agent of a Windows computer looks for the printers on its network and the ones you pick are added in one go. Open it from Infrastructure › Printers; every technician can use it.

Adding a printer

A printer is added where every SNMP device is: on the Devices sub-tab of a computer on the same network as the printer (see Custom fields, Passwords, Attachments, Network and SNMP devices in the Computers chapter). Any technician can add one.

1. In Computers, open a Windows computer on the printer's network and go to its Devices sub-tab.
2. Click + Add device, give the printer a Label and choose Printer as its Type. A note under the type explains that the agent reads the supplies with the standard Printer MIB and that most network printers answer SNMP v2c with the community public; it also warns you when the computer's agent is older than 1.2.29.
3. Fill in Host / IP (the printer's address), Port (161), Community (often public) and Interval (seconds), as for any SNMP device.
4. Leave Raise alerts for this device ticked and set Supply — warning (%) (15 by default) and Supply — critical (%) (5 by default); leave a field empty to turn that level off. Tick Email the team when an alert opens if you want emails. The volume, temperature, CPU, interface and firmware settings of the other device types are not shown: they do not apply to a printer.
5. Click Save. The supplies appear after the next poll: use Poll now not to wait for the interval.

> **i** Reading the supplies needs the agent 1.2.29 or later on a Windows computer. An older Windows agent still polls the printer, which then shows as online or offline but without its supplies, and the Linux and macOS agents do not poll SNMP devices at all. In the Printers list, the Supplies column then says what to do: update the agent of that computer to 1.2.29, or add the printer on a Windows computer with the agent on its network. It shows This device does not report its supplies over SNMP when the printer itself does not answer the Printer MIB.

Finding printers with a scan

Instead of adding the printers of a client one by one, you can let the agent of one of its Windows computers look for them. The agent scans its local network (the same sweep as the Scan network button of the Network tab, plus an SNMP v2c question to every device that answers), the printers it finds are listed in a window, and the ones you tick become SNMP devices of type Printer polled by that computer, exactly as if you had added them by hand. Every technician can do it, from the Find printers button at the top of Printers, or under the message shown while the list is still empty.

1. Click Find printers. The Find printers window opens.
2. Optionally choose the client in Client (optional) to shorten the list of computers; All clients lists them all, and No client the computers without a client.
3. In Computer that scans, choose a Windows computer on the printers' network. Each computer shows its client and Online or Offline. An offline computer can be chosen: the scan waits in the Command queue and starts as soon as the computer comes back online, for up to 6 hours. A note also says when the agent's service is not connected, and the Scan button stays disabled when the computer's agent is older than 1.2.31. When the computer already has a scan from the last 24 hours, the window shows it straight away (still running, or with its results); New scan goes back to the form.
4. Check SNMP community (public by default) and SNMP port (161). The same community is used afterwards to poll the printers you add.
5. Click Scan. The window shows the scan waiting for the computer and then Scanning the network for printers — this can take a few minutes, with the time elapsed. It keeps waiting however long it takes; you can close the window and open Find printers again later: choosing the same computer shows the scan again.
6. When the scan finishes, a line says which subnet was scanned, how many hosts and how many printers were found. The table lists them with Address, Name / model, Serial number, MAC / manufacturer, SNMP, Ports and Status; click a heading to sort, use the search box to narrow the list, and right-click a row to copy its address or serial number. A printer whose address is already

one of this computer's devices, or that has the same serial number as a device of the same client, cannot be ticked: Status says Already in Printers when that device is an active printer, Already a device of this computer when it is another type of device or a stopped one (find it on the computer's Devices tab), or Same serial number as followed by its name. The new printers that answered SNMP are ticked for you; the ones marked Not confirmed (no SNMP response), found only by an open printer port, are left unticked.

7. Click Add N printers. The printers are added to Printers (and to the computer's Devices tab) with the default supply thresholds, 15% and 5%, and the agent starts polling them on its next SNMP cycle, in about a minute. The rows you added now show Already in Printers. Close closes the window.

> **i** Finding printers needs agent 1.2.31 or later on a Windows computer, and only SNMP v2c is supported. The scan covers only the computer's own /24 (254 addresses) of its main network, slowly and only on the printing ports and SNMP, so that it does not look like an attack to an antivirus; on a wider network the list is marked as possibly incomplete, so scan again from a computer on the other part of the network if a printer is missing. The community never appears in the Command queue (the command only carries the scan's identifier): it goes to the agent over HTTPS and is erased 24 hours after the scan, after which that scan's printers can no longer be added. A printer added while Not confirmed shows as offline with the No SNMP response alert until SNMP is enabled on it with that community. Each printer added is recorded in the Audit log.

The printer list

The table has one row per printer: Client, Printer (its label, with the address underneath, and an Inactive badge when it is switched off in its settings), Model, Supplies (a small bar for each supply), Pages (the page count the printer reports), Status, Alerts (the number of open alerts), Last poll (how long ago, with the last error under it) and Polled by (the computer, a link to its page). Serial number can be added with Columns. Above the table, a client picker (only the clients that have printers), a search box (printer, address, model, serial number, client or computer) and a Show picker narrow the list: Show all, Low supplies (a supply under the printer's warning level, or empty), Offline, With open alerts and Reporting errors. The row under the headings filters by client, printer, model, status (Online, Offline or Inactive) and alerts. The list reloads itself every minute.

Status shows Online or Offline, the printer's own state (Ready, Printing or Warming up), its device status when it is not running normally (Warning, Testing or Down) and the problems it reports, such as Paper jam, Door open, Out of paper, Low toner or Service requested; the ones that stop the printer are shown in red. While a printer is offline, for example switched off for the night, its last known supplies and state stay on screen, the state faded.

Click a row to open its detail underneath: every supply as a bar with its name and level, Pages printed, the printer's state and its alert thresholds, for example Alert thresholds of this printer: warning under 15%, critical under 5% (changed on the Devices tab of the computer that polls it). The refresh button at the end of the row asks the agent for a poll now; right-click a row for Poll now, Open the computer, and Copy address, Copy model or Copy serial number.

Reading the supply bars

- Each supply is named as the printer reports it, for example Black Cartridge HP CF259A, or else by its type and colour, with a colour swatch for black, cyan, magenta and yellow.
- The bar shows the level in percent. It turns amber under the printer's warning level, and red under its critical level or when the supply is Empty.
- Some printers only say whether a supply is empty: Some left then stands for any level above zero, and Level unknown for a supply the printer does not measure. Neither raises an alert.

- Waste toner boxes and the other containers that fill up instead of running out show how full they are: a percentage followed by full, and Full in red when there is no room left. They never raise an alert and do not count as low supplies.
- Pages is the printer's own page counter.

The same bars appear on the Devices tab of the computer that polls the printer, where clicking the printer's name shows its supplies and detail, and on the client's Devices tab.

Printer alerts

A printer's alerts are SNMP device alerts: they appear in Alerts with the type Printer supply, linked to the computer that polls the printer and with the printer's name at the start of the message, and you resolve them or create tickets from them like any device alert. With Raise alerts for this device ticked:

- A supply under Supply — warning (%) opens a Warning, for example Black Toner low: 12% (threshold 15%), and under Supply — critical (%) it becomes Critical, Black Toner very low: 3% (threshold 5%). An empty supply is Critical, or a Warning when only the warning level is set. Each supply has its own alert.
- A printer that stops answering opens the Critical No SNMP response alert, as every SNMP device does.
- Unlike the other device types, a printer raises no reboot or interface alert: printers are switched off at night and their virtual interfaces come and go.
- With Email the team when an alert opens ticked, the active users of your team with an email address who have access to Printers or Alerts, and the administrators, are emailed when a new alert opens (see Who receives the alert emails in the Alerts section).

A supply alert closes by itself at the first poll after the supply is replaced, and the No SNMP response alert at the first poll that gets an answer. New thresholds, or a level left empty, apply from the next poll, and unticking Raise alerts for this device closes every open alert of the printer. An administrator can stop the alerts of every printer at once with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section); the other SNMP devices have their own switch, on the Devices tab.

32. Lifecycle

The Lifecycle module shows, for every computer of every client, how old it is, when its operating system stops receiving security updates and when its warranty ends, and from that what to do about it and by when: replace the computer, upgrade its operating system or renew its warranty. It also exports the replacement plan of the next three years to Excel, for one client or for all of them, and opens alerts when an operating system or a warranty is about to run out. Open it from Infrastructure › Lifecycle in the sidebar; the page itself is titled Hardware lifecycle. Every technician can use it, look a warranty up and type the dates of a computer. The replacement ages, the warning days and the manufacturers' warranty API keys are for administrators.

The top of the page and the summary

The top of the page has Replacement plan, which exports the plan (see The replacement plan), Settings, shown to administrators only, which opens the lifecycle settings in a window, and Columns. Under the introduction come a client list (All clients, — no client — or one client that has computers), a search box that looks in the computer, client, model, serial number and operating system, and the Show list: Show all, To replace this year (overdue included), To replace next year, Unsupported operating system, Operating system support ending soon, Warranty expiring, No warranty date and With open alerts. When filters hide computers, a counter such as 12 of 48 says how many are left. The

list is read when you open the module and again after each of your actions; it does not refresh by itself.

Four cards sum up the client chosen, or all clients: To replace this year, with N overdue in red for computers whose replacement date has already passed; To replace next year; Unsupported OS, with +N ending soon underneath; and Warranties expiring, with N expired underneath. Click a card to list only those computers and click it again to list them all. With All clients chosen and at least two clients, By client (N) opens a table with the same counts per client and the number of computers; click a row, or Show this client's computers in its right-click menu, to choose that client in the client list.

The list

- Client and Computer: the computer's name is a link that opens its page; a virtual machine carries a VM badge.
- Type: Workstation or Server, followed by · VM for a virtual machine. A computer is a server when its operating system is a server edition or its asset type is Server, and a workstation otherwise.
- Hardware (manufacturer and model) and Serial number, as the agent reports them.
- Age: in years with one decimal, followed by where it comes from: purchase, warranty, BIOS, OS install or agent. A ~ in front marks an approximate age, and the figure turns red once the computer is due for replacement. Hover it for the source in full, the date it counts from and the replacement date (see Age and when to replace).
- Operating system: the product, a chip (Supported, Ending soon, Unsupported or Unknown) and the end-of-support date. Hover it for the reason, the version to move to, whether paid Extended Security Updates are still possible and whether the edition had to be guessed.
- Warranty: a chip (Active, Expiring, Expired or No date), the end date and its source, Dell, Lenovo or HP when a lookup found it and a writing hand when it was typed by hand. Hover it for the days left, why the warranty can or cannot be looked up and the result of the last lookup. A virtual machine shows a dash.
- Recommendation: Replace, Upgrade OS, Renew warranty or OK; a computer can have several. Hover it for the next step and its date, for example Next step: Replace by 01/03/2027, the earlier of the replacement date and the end of support of the operating system.
- Purchase price, in euros. Purchase date, Replace by and Alerts (the number of open lifecycle alerts) are available through Columns.

Click a header to sort by that column and again to reverse it. The row under the headers filters Client, Computer, Hardware and Serial number by text, and Type, Operating system, Warranty and Recommendation by a list of their values. At the end of each row, the refresh button looks the warranty up (see Warranty) and the pencil button opens the dates of the computer (see Typing the dates of a computer); each appears only when it can be used. The right-click menu of a row offers Open computer, Check warranty, Edit dates, Copy computer name and Copy serial number; an entry that cannot be used is greyed out and its tooltip says why. When no manufacturer keys are set and some computers have no warranty date, a line above the list explains how to fill them in. The last line of the page says when the operating system support dates were last checked against Microsoft's, Apple's, Canonical's and Debian's own pages.

Age and when to replace

The age counts from the best date the computer has. The purchase date typed by a person is the only exact one. Without it, the module takes the start of the warranty that a Dell, Lenovo or HP lookup returned, which is close to the date the computer was shipped. Without either, it takes the oldest of three dates: the BIOS date, the installation of the operating system and the installation of the SLApp agent. None of the three is the computer's birth date, and all of them only ever move forward, because every BIOS update and every Windows feature update writes a newer one; the oldest is the closest

guess. Every age that does not come from a purchase date is marked ~ as approximate. A computer whose dates are all unknown shows a dash.

The replacement date is the start of the age plus the replacement age set for its type, 5 years for workstations and 7 for servers unless an administrator changed them. Once that date has passed the computer is due, its age turns red and its recommendation is Replace. Virtual machines have no hardware of their own: they get no age, no replacement date and no warranty, only the support of their operating system.

Operating system support

The end of support of an operating system is the last day it receives security updates; after it, nothing fixes the next vulnerability. The dates are kept in SLApp Helpdesk itself, from the vendors' published pages. Supported means more than 180 days to go, Ending soon 180 days or less, Unsupported that the date has passed, and Unknown that the system was not reported yet, is not tracked or is a version that the table does not know yet.

- Windows 10 and 11 by version, with the dates of the Home and Pro editions or of the Enterprise and Education editions, and the LTSC, LTSB and IoT LTSC editions with their own dates. When the agent did not report the edition, the Home and Pro dates are used and the tooltip says so. For Windows 10 22H2, whose support ended, the tooltip says until when paid Extended Security Updates can still be bought. Windows Vista, 7, 8 and 8.1 are long out of support.
- Windows Server 2003 to 2025 by the end of their extended support, and the Windows Server versions without a year in their name, such as version 23H2, by the end of their own support.
- macOS: Apple publishes no end dates, so the three most recent versions count as supported and an older one as unsupported, with the advice to update if the Mac supports it.
- Ubuntu LTS releases by the end of their standard support (interim releases are not tracked) and Debian by the end of its LTS support. Other Linux distributions are not tracked.

An operating system that is Ending soon or Unsupported gives the computer the recommendation Upgrade OS, and the tooltip names the version to move to when there is one.

Warranty

The warranty end date comes from one of two places. A date typed by a person always wins: in Edit dates, in the computer's asset or by a client contact in the Portal. A lookup never replaces it; clear it if you want the lookup to fill it in again. Otherwise, for Dell, Lenovo and HP computers (HP Enterprise excluded), SLApp Helpdesk can look the warranty up by serial number with your own API keys (see Warranty API keys). The chip reads Expiring when the end date is within the warning days, 60 unless an administrator changed them, and Expired once it has passed; both give the computer the recommendation Renew warranty.

- Lookups run by themselves in the background: a computer is looked up again about once a week after an answer, found or not, and a day after a lookup that failed, which keeps the date it had. Each manufacturer is asked a limited number of times every 15 minutes, so a large fleet is looked up over several hours.
- When a manufacturer asks to slow down, its lookups pause for the time it asks for (30 minutes when it does not say, at most two hours). When it rejects your keys, the lookups of your company pause for 6 hours, or until you save new keys.
- Check warranty (the refresh button, the right-click menu, or the card on the computer's Overview) looks the warranty up at once and says what it found, for example PC-042: warranty until 31/05/2027 (Dell), or that the manufacturer has no warranty for this serial number. It works at most once a minute per computer.

- A computer cannot be looked up, and Check warranty is not offered, when it is a virtual machine, when it has no asset yet (for example because it has no client), when its date was typed by hand, when it is not a Dell, Lenovo or HP, when it reports no usable serial number, or when the keys of its manufacturer are not set. The tooltips say which applies.
- When the manufacturer also returns the start of the warranty, the age of a computer that has no purchase date counts from it.

Typing the dates of a computer

1. Click the pencil button at the end of the row, choose Edit dates in its right-click menu, or click Edit dates on the computer's Overview card.
2. The window Lifecycle dates — PC-042 shows the hardware, serial number and client of the computer. Fill in Purchase date (from 1990 up to tomorrow), Purchase price (€) and Warranty expires. Under the fields, the window says where the current warranty date came from and the result of the last lookup.
3. Click Save. A notice says Dates of PC-042 saved, the row is updated and the computer's alerts are checked again at once. Cancel closes the window without saving.

> **i** The dates are stored on the computer's asset, the one generated from the computer, which is why a virtual machine, or a computer that has no asset yet because it has no client, cannot hold them. A changed warranty date becomes a date typed by hand, while a date left as it was keeps its source. Every change is written to the Audit log of the computer with the old and new values.

The replacement plan

Replacement plan downloads an Excel workbook with the plan of the client chosen in the client list, of all clients, or of the computers without a client when — no client — is chosen; the other filters of the page do not apply. It is written in the language of the app and named, for example, replacement-plan-contoso-2026-10-30.xlsx. Every technician can export it.

- The Replacement plan sheet groups the work by quarter, from today to the end of the eleventh quarter after the current one (the date after Plan until at the top of the sheet), and lists only the quarters that have something in them: Overdue first, then the others in date order. Each computer gives at most two lines: Replace on the day its age reaches the replacement age, and Upgrade OS on the day its operating system support ends, unless the computer is due to be replaced in that quarter or before it anyway. An outdated macOS, which has no end date, is listed as Overdue. A plan with nothing in it says Nothing to replace or upgrade within this plan.
- The heading of each quarter counts its computers, each computer once, and adds up their known purchase prices as a rough budget, saying how many computers have no price; a Total line closes the sheet.
- The Computers sheet lists every computer of the plan's scope with its client, type, hardware, serial number, age and its source, operating system and support, warranty and its source, recommendation, replacement date, purchase date and purchase price, with filters on the header row.

Lifecycle settings

Only administrators can change the lifecycle settings. They are in Settings › General › Lifecycle, and the Settings button of the module opens the same form in a window. Replace workstations after (years) accepts 1 to 15 and starts at 5, Replace servers after (years) accepts 1 to 20 and starts at 7, and Warn before a warranty expires (days) accepts 0 to 365 and starts at 60, where 0 means no warranty alerts. The form also explains how the age is worked out and lists the alert rules. Click Save; every change is

written to the Audit log under the entity type Lifecycle. A new number of warning days is applied to every computer by the next passes of the lifecycle check, which runs every 15 minutes, and the warranty alerts that it opens send no email.

Warranty API keys

The keys of the warranty lookups are set by an administrator in Settings › General › Connectors › Warranty, one block per manufacturer with a Configured or Not configured badge and a short explanation of how to obtain them.

- Dell (TechDirect): Client ID and Client secret, received once Dell approves your request for the Warranty API in Dell TechDirect.
- Lenovo: the ClientID that Lenovo gives to its partners, through your Lenovo account representative or partner programme.
- HP: API key and API secret of an app created after HP approves your enrolment in the HP Warranty API. HP keys expire every 90 days, so save the new ones when they do.
- Type the keys and click Save. They are stored encrypted and never shown again: once a manufacturer is configured, its fields say Stored — leave empty to keep it. Test asks the manufacturer whether it accepts the stored keys, so save new keys before testing them. Remove deletes them after a confirmation; the dates already found stay, and new lookups stop until keys are set again.
- Saving and removing keys are written to the Audit log under the entity type Lifecycle, never the keys themselves. Without any keys, warranty dates are simply typed by hand.

Alerts and emails

The module raises its own alerts, so there is nothing to set up in the Rules tab of Alerts. They appear in the Alerts feed and on the Dashboard linked to the computer, with the source Lifecycle, in the Alerts tab of the client, and in the footer ticker as a single item, and you can create a ticket from them like from any other alert. Each computer is checked once a day, and again at once after Check warranty or a change of its dates. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

- OS support ending: a Warning when the support of the computer's operating system ends within 90 days.
- OS support ended: Critical once that date has passed, and for a macOS that is no longer one of the three latest versions. It replaces the OS support ending alert, which closes.
- Warranty expiring: a Warning when the warranty ends within the warning days, with a message such as Warranty expires on 2026-12-01 (Dell). It closes once the warranty is renewed, or once it has expired: an expired warranty and the age of a computer are statuses shown in the list, the plan and the client report, never alerts.
- The operating system alerts close by themselves once the computer runs a system whose support ends more than 90 days away, and all the alerts of a computer close when it is removed. An operating system that becomes Unknown leaves its alerts as they are.
- Emails go to the active administrators of your company who have an email address; the email is written in the email language of your company, and since version 1.14.81 so is the alert's message. The first check of each computer opens its alerts without an email, so turning the feature on does not email you about every old computer, and so does a warranty alert that opens only because a date was just found or typed, or because the warning days changed.
- When one pass of the check opens several alerts, each administrator receives one email, Hardware

lifecycle alert: N new, that lists them (up to 50), instead of one email per computer.

- Only an administrator can close an alert by hand, with Resolve; if the condition is still true, the next daily check opens it again, with an email.

> **i** Lifecycle alerts stay open for months. So that they never push an operational alert out of sight, the alerts card of the Dashboard lists them after every other alert, and the footer ticker sums them up in a single item such as 14 hardware lifecycle alerts. The same information is on each computer's Overview, in the Hardware lifecycle card (see the Computers section), and in the Hardware lifecycle section of the client report, which leaves out the purchase prices.

33. Websites

The Websites module watches the sites you host from the outside — no agent involved. The server itself probes each site on its own schedule and records whether it answers, how fast, and how long its SSL certificate has left; a one-line tracking snippet adds cookieless visitor analytics on top. Open it from Cloud & web › Websites in the sidebar. Alerts raised here appear in the normal Alerts feed, prefixed with a globe icon. To also watch the registration and the email records (SPF, DKIM, DMARC) of a site's domain, add it in Cloud & web › Domains, whose Suggest button proposes the domains of the monitored websites.

The websites list

The list shows one row per monitored site with the columns Name, URL, Client, Status (Online or Down), Response (the last check, in milliseconds), SSL (days until the certificate expires, orange when close, red when past), Visits today and Alerts, plus a Remove button. Next to the name a small badge shows the tracking snippet state: a green Tracking snippet badge when the monitor found the site’s snippet in the page, a grey no snippet badge when it did not, and nothing while the site has not yet been checked successfully. The list refreshes itself every 60 seconds.

- The bar above the table has a client selector (including — no client —), a search box that matches name or URL, and a snippet filter (all / installed / without snippet / not checked yet). When filters are active a counter shows N of M.
- A filter row under the header adds free-text filters on Name and URL and dropdowns on Status (Online/Down), SSL (<= 30 days / Expired) and Alerts (with / without).
- Every column header sorts on click (click again to reverse): name, URL, client, status, response time, SSL days, visits and open alerts.

Adding a website

1. Click + Add website, or — when a Plesk panel is connected in Settings › General › Connectors — click Import from Plesk to pick domains straight from your hosting panel (sites already monitored are marked, Select all new ticks the rest, and an optional client is applied to all).
2. Fill in Name and URL (https://…). Optionally pick the Client the site belongs to.
3. Choose the Check interval (how often the server probes the site).
4. Review the alert options: Alerts enabled for this website, Alert when response exceeds (ms, empty = off), Alert when SSL expires within (days, 0 = off), Email the team when an alert opens, and Also notify these emails for external recipients such as the site’s webmaster.
5. Click Save. The first check runs within 30 seconds.

The website detail page

Click a row to open the site’s page. Tiles at the top summarise Uptime over 24 h, 7 and 30 days, the

average response time, the SSL expiry (with issuer), visits and unique visitors, and the average time on page. Below them, a chart plots the response time of the last 24 hours — red dots are failed checks — and a second chart shows visits per day over the last 30 days, followed by Top pages, Referrers, Browsers, Operating systems and Devices, and a table of recent visits with the full detail of each one. Check now forces an immediate probe. Removing a site stops monitoring and hides it everywhere; its history is kept until it ages out (checks for 90 days, visits for 400 days).

The tracking snippet

Each site has a unique tracking code, and the detail page shows the ready-made snippet with a Copy button — one script line to paste before </body> on every page of the site (or in its shared template). It counts the page, title, referrer, UTM campaign fields, browser, device, language, screen size and time on page, using no cookies and no fingerprinting. It does keep a random visitor id in the browser's storage, which EU rules only allow after the visitor's consent: tick Wait for the visitor's consent under the snippet and it stores and sends nothing until the site's cookie banner calls slappTrack.consent() when the visitor accepts (on every page), while slappTrack.revoke() forgets the visitor. The line pasted on the site stays the same, so the option can be turned on or off at any time. The monitor also looks for the snippet in the pages it downloads: the badge in the list and in the snippet box turns to Tracking snippet once it is found, so you can see at a glance which sites are actually counting.

> **i** If the same build of a site is deployed to more than one environment (production, staging, developers' machines), wrap the tag in a small hostname check so that only the real domain loads it — otherwise test traffic is counted as production visits.

Website alerts

Three alert types come from this module, all raised per site and visible in the Alerts feed with the site name: Website down (Critical) opens only after two consecutive failed checks, so a single blip never pages anyone, and sends a recovery email when the site answers again; Slow response (Warning) opens while the response time stays above the site's threshold; SSL certificate (Warning, Critical once expired) opens when the certificate is within the configured number of days of expiry. Turning Alerts enabled off for a site closes its open alerts at the next pass. With Email the team when an alert opens ticked, emails go to the active users of your team who have access to Websites or Alerts and to the administrators (see Who receives the alert emails in the Alerts section), plus the addresses in Also notify these emails, which receive them as typed; with it unticked, the site's alerts email nobody, not even those addresses. An administrator can stop the alerts of every site at once with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section); the checks go on and the list keeps showing each site's status.

The hosting subscription (Plesk)

Sites hosted on a connected Plesk panel show two extra columns in the list (and cards on the detail page): Setup and Expiration — the same dates the panel shows on its Subscriptions screen, refreshed automatically every 12 hours. A subdomain inherits the dates of its parent subscription. The Expiration value turns amber when 30 days or less are left, red once it has passed, and is replaced by a red Suspended badge when the subscription is suspended in the panel.

Each website also has a Hosting expiry alert (days) setting (edit form, default 30, 0 turns it off): when the subscription is that close to expiring, a site_hosting alert opens — a warning while there is still time, critical once expired — and it resolves itself when the renewal shows up in the panel. Like every website alert, it honours the site's Alerts enabled switch and the e-mail setting.

Editing and assigning from the list

Every row has a pencil button that opens the edit form directly, without going through the detail page.

With rows selected, the selection bar also offers Assign to client: pick a client (or “no client” to clear) and it is applied to every selected site at once. The websites assigned to a client are listed on that client’s sheet, in its Websites tab, with status, SSL, hosting dates and today’s visits.

34. Domains

The Domains module watches your clients' Internet domains from the outside, with no agent involved. Once a day the server itself reads when each domain's registration expires, which nameservers and mail servers (MX) it points to, and its email authentication records — SPF, DKIM and DMARC — as the rest of the Internet sees them. It warns you before a domain expires, when its email setup breaks and when its mail servers or nameservers change, which may mean that someone has taken the domain over. Open it from Cloud & web › Domains in the sidebar; the page is titled Domains & email. Every technician can add, edit, check and remove domains; only an administrator can resolve one of their alerts by hand. Additions, edits and removals are recorded in the Audit log under the entity Domain.

The domains list

Each row is one monitored domain, with the columns Domain (a notepad icon after the name when it has notes; hover it to read them), Client, Expires, Registrar, SPF, DKIM, DMARC, Last check (Not yet until the first check) and Alerts (the number of open alerts, red when one of them is critical, amber otherwise). Notes and Added, the date the domain was added, can be shown with Columns. Click a row to open the domain's detail window. The list refreshes itself every 60 seconds. Until the first domain is added, the page suggests adding one or using Suggest.

- The bar under the title has a client selector (All clients, — no client — and each client that has domains here), a search box (Search domain, registrar or notes…) and a Show selector: Show all, With open alerts, Expiring or expired, Email authentication problems (an SPF, DKIM or DMARC chip in Warning or Error) and Not checked yet. While filters hide rows, a counter shows how many are shown out of the total (N of M).
- A filter row under the headers narrows each column: free text for Domain and Registrar; Expires (All, Expiring, Expired, No date); SPF, DKIM and DMARC (All, OK, Warning, Error, Unknown — Unknown also takes the domains never checked); Alerts (All, With alerts, Without alerts).
- A click on a header sorts by that column and a second click reverses it; sorting by SPF, DKIM or DMARC puts Error first, then Warning, Unknown and OK. Columns opens the List columns window, where you choose the columns and drag them into your order; the choice is saved for your user.
- Each row ends with the refresh button (Check now), the pencil button (Edit) and Remove. A right-click on a row offers Open details, Edit, Check now, Copy domain name, Copy registrar (when one is known) and Remove.

Reading the Expires column and the SPF, DKIM and DMARC chips

- Expires shows the date the registration expires and a badge: N days left in green; amber once the domain is inside its own warning window (Warn before expiry (days), 30 by default); red during the last 7 days, on the last day (Expires today) and after it (Expired N days ago). A dash means there is no date; hover it to see why: the domain was not checked yet, the lookup failed (with the error), the registry does not publish the expiry date, or the domain's extension has no registry service that can be read.
- SPF, DKIM and DMARC each show a chip: OK in green, ! Warning in amber, × Error in red, a grey ? Unknown when the record could not be read, and a grey ? alone while the domain was never checked. Hover a chip to read what it means in one sentence, for example Valid SPF — 4/10 DNS lookups, ends with -all, or DMARC in monitoring only (p=none): nothing is quarantined or rejected.

Adding a domain

1. Click + Add domain.
2. Type the Domain *, for example example.com. You can paste a URL or an email address instead: only the domain is kept, in lowercase and without www. A name with accented letters is stored in its Internet form and shown as you wrote it.
3. Choose the Client the domain belongs to, or leave — no client —.
4. Check Warn before expiry (days): 30 by default, from 0 to 365; 0 turns the expiry alerts off. Whatever the value, the alert turns critical 7 days before the expiry.
5. If the client's email provider signs with a DKIM selector that is not among those always tried (selector1, selector2, google, default, k1, s1, s2, dkim, mail), type it in Extra DKIM selectors; separate several with commas (up to 10).
6. Leave Alert on SPF and DMARC problems ticked, unless the domain sends no email, and add Notes if useful.
7. Click Save. A toast confirms Domain example.com added; its first check starts right away and the row fills in shortly after.

> **i** You can monitor a subdomain too (shop.example.com): its registration is that of example.com, named as Registered domain in the detail window, and when it has no DMARC record of its own the record of example.com is used, as receiving mail servers do. A domain removed earlier can be added again: it comes back with its change history (the toast says it was added again, with its previous history), but its monitoring starts afresh, like a new domain's. A domain that is already monitored cannot be added twice.

Suggest: adding the domains SLApp already knows

Suggest opens the Suggested domains window with the domains your company already has on record for its clients and does not monitor yet: the Email domains of the client records, the domains of the sites in Cloud & web › Websites and the verified domains of the clients connected to Microsoft 365 (their .onmicrosoft.com domain left out). Each name is reduced to its registered domain (www.shop.example.com becomes example.com) and listed with its Client, when one is known, and its Source: Client email domain, Website: followed by the site names, or Microsoft 365. Domains with a client come first.

1. Click Suggest. Every suggestion comes ticked, except those marked Removed before: you monitored them once and removed them, so they are left for you to decide; tick one to monitor it again and its history comes back.
2. Untick what you do not want. The checkbox in the header ticks or unticks everything, and a right-click on a row selects it or copies the domain name.
3. Click Add N domains. A toast says how many domains were added and, when some were not, how many were skipped (invalid, repeated, already monitored or client not found).

Domains added through Suggest are not checked on the spot: the server takes them a batch at a time, every five minutes, so that a large import does not hit the registries all at once, and they show their first results within the next minutes. When there is nothing left to suggest, the window says so.

Editing and removing a domain

Edit (the pencil button on the row, Edit in the right-click menu or Edit in the detail window) opens the same form, titled Edit domain. The domain name itself cannot be changed: remove the domain and add the new one. A new Warn before expiry (days), or a change to Alert on SPF and DMARC problems, is applied to the alerts as soon as you save, from the last check, without waiting for the next one; an alert

that opens because of it is emailed as usual. Giving the domain another client moves its open alerts to that client.

Remove asks for confirmation (Remove domain): the checks stop and the domain's open alerts are resolved. Its history is kept, so adding the same name again later brings it back.

Check now

Check now (the refresh button on the row, in the right-click menu or in the detail window) runs a full check at once; when it ends, a toast says Check finished: example.com and the new result is shown. A domain can be checked at most once a minute and never twice at the same time; too soon, the toast says Too soon to check this domain again — try again in N s. Registries limit how often one address may query them, and all your technicians share the server's address, so use Check now after fixing something rather than to refresh the page.

The domain detail window

A click on a row opens the detail window. Its header shows the domain (and the Internet form of an accented name), its client, the time of the last check and the buttons Check now and Edit. A domain not checked yet shows Not checked yet — the first check runs within a few minutes, or use Check now, and the window looks again every 10 seconds until the first result arrives. The window has these parts:

- Registration: Expires with its badge, Registrar, Registered on, Registry status (the registry's codes, such as clientTransferProhibited) and Source (RDAP or WHOIS); for a subdomain, Registered domain names the domain whose registration it is. A note says when the domain's extension has no registry service that can be read, when the registry does not publish the expiry date, or when the lookup failed, with the error and how many checks in a row have failed.
- DNS: the Nameservers; At the registry, when the nameservers recorded at the registry differ from those the domain's DNS answers with; and the Mail servers (MX), each with its priority.
- SPF: the chip and its explanation, the Record, the DNS lookups it needs out of 10 (in red above 10), the Final rule (its all term), the Includes and the Redirect.
- DMARC: the chip and its explanation, the Record, the Policy, the Subdomain policy, the Percentage (pct) and Aggregate reports to (rua); for a subdomain without a record of its own, a note names the domain whose record applies. A crossed-out bell next to the SPF and DMARC titles means that Alert on SPF and DMARC problems is off for this domain.
- DKIM: the chip and its explanation, the Keys found — each selector with TXT key, or CNAME › the name it points to, marked no key published yet or revoked when that is the case — and the Selectors tried.
- Open alerts: the Severity, Type, Message and Opened at of each open alert of the domain, with a Resolve button for administrators.
- Raw records (click to expand): the domain's TXT records, its _dmarc record, MX and NS, exactly as read.
- Change history: the changes seen by the checks, newest first (up to 100), with When, Record (MX, Nameservers, SPF, DMARC or DKIM), Before and After.

What is checked

- Registration: read with RDAP, the registries' modern lookup service (the server of each extension comes from IANA's list), or with WHOIS when an extension has no RDAP service or it fails — .pt domains, for example, are read from DNS.pt's WHOIS. While the registry cannot be read, the last expiry date known keeps the expiry alerts working.
- Nameservers (NS) and mail servers (MX): compared with the previous check. A change is written to the Change history and raises an alert; the same mail servers with other priorities are not a change.

- SPF: the domain must have exactly one v=spf1 record. SLApp follows its include and redirect terms and counts the DNS lookups the record needs: receivers allow 10 at most, and past that SPF fails for every message — the classic outage after one more email service is added to the record. The count follows RFC 7208 for include, a, mx, ptr, exists and redirect; the extra lookups behind an mx term are not counted, so another checker may find a slightly higher number. The final all term decides what happens to the other servers: -all and ~all are OK; ?all, or no all term, is a warning because it does not tell receivers to refuse them; +all, which allows any server, is an error. An unknown term, an include loop, or an included domain without exactly one SPF record makes the record invalid (permerror), also an error.
- DMARC: the record published at _dmarc.example.com for example.com. p=quarantine and p=reject are OK. p=none, monitoring only, is a warning but never raises an alert, because it is the normal first step of a DMARC rollout. No record is a warning; more than one record, or a record without a valid p= policy, is an error. A subdomain without a record of its own is covered by its parent domain's record (its sp= policy, else p=), as receiving mail servers do.
- DKIM: keys cannot be listed, only asked for by name, so SLApp tries the selectors the large email platforms use — selector1 and selector2 (Microsoft 365), google (Google Workspace), default, k1, s1, s2, dkim and mail — plus the domain's Extra DKIM selectors. A key found is OK. A selector whose CNAME leads to no key (what a Microsoft 365 domain shows when its DKIM signing was never switched on), only revoked keys, or no key on any selector tried is a warning. DKIM never raises an alert: a selector nobody guessed is not a problem.

The DNS questions go to public DNS servers (1.1.1.1 and 8.8.8.8, unless your server was set up otherwise), never to the server's own DNS, so a hosting or Active Directory server that answers for its own zones cannot hide what the rest of the Internet sees. A lookup that fails shows as unknown and leaves the alerts as they are: a record that could not be read is not a record that was fixed.

When the checks run

The server checks each domain once a day: about 24 hours after its previous check, plus a random delay of up to an hour, so that domains added together spread out. Every five minutes it looks for the domains that are due and checks them one after the other. A domain added with + Add domain is checked straight away, while the domains added through Suggest wait for the scheduler, which takes them within minutes. Nothing depends on the clients' computers being on, and Check now adds a check whenever you need one.

Domain alerts

- Domain expiring: the expiry date is within the domain's Warn before expiry (days); a warning, critical in the last 7 days. Domain expired (critical): the date has passed. Both close by themselves when the renewal shows up at the registry, and neither opens when the warning days are 0.
- SPF record (warning): no SPF record, more than one, more than 10 DNS lookups, an invalid record or +all. DMARC missing (warning): no DMARC record, more than one, or none with a valid policy. Both need Alert on SPF and DMARC problems ticked on the domain, and both close by themselves once the record is fixed; ?all and p=none show as warnings on the chips but raise no alert.
- MX changed and Nameservers changed (critical): the mail servers or the nameservers differ from those of the previous check, with the old and the new ones in the message. It may be a planned move, or someone taking over the client's mail or DNS. Each change opens an alert of its own (the previous one of the same kind closes); it closes by itself 7 days after the change, or when an administrator resolves it.
- Registration lookup failing and DNS lookups failing (warning): the registry, or the domain's NS, MX or TXT records, could not be read on 3 checks in a row, that is three days. They close at the next check that reads them.

When an alert opens, every active administrator of your company with an email address receives an email headed Domain alert, naming the domain, the client, the message and the time, with a subject such as Warning alert: contoso.com — DMARC record missing (Critical alert for a critical one). If the alert later turns from warning to critical (an expiring domain entering its last week), a second email follows, headed Alert now critical, with a subject starting with Now Critical. Closing an alert sends nothing. The first check of a domain, including one added again after a removal, opens its warnings without email, so that a Suggest import does not send dozens of DMARC emails; a critical alert is always emailed. The alerts show in the domain's detail window, in the Alerts module (the message starts with the link icon and the domain name, and their Source is Domain), in the client's Alerts tab, on the Dashboard and in the footer ticker. Only an administrator can resolve one by hand, and the Audit log records it on the domain. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

> **i** When the DKIM chip says No DKIM key on the usual selectors — add the provider's own selector in Edit, open an email the client sent and look at its headers: the s= value of the DKIM-Signature line is the selector to add in Extra DKIM selectors.

> **i** SLApp only reads public information: adding, checking or removing a domain never changes anything at the registrar or in the domain's DNS.

The Domains tab of a client

Each client's sheet has a Domains tab that lists the client's domains with the columns Domain, Expires, Registrar, SPF, DKIM, DMARC, Last check and Alerts, the same filters under the headers, sorting and Columns. It is read-only: a click (or Open details in the right-click menu) opens the detail window without Check now, Edit or Resolve, and domains are added, changed and checked in Cloud & web › Domains. A client without domains shows No domains for this client yet — add them in Cloud & web › Domains.

Domains in the client report

Client reports (Reports) have a Domains section: one row per monitored domain of the client with Domain, Expires, Days left, Registrar, SPF, DKIM, DMARC and Last check, plus an Email authentication chart that counts, for each record, the domains that are OK, in Warning, in Error or Not checked. DKIM shows the first selector with a key, as in OK (selector1), and DMARC shows the policy, as in Warning (none). The section shows the domains as they are on the day the report is made, not over the report period. The full report leaves it out when the client has no monitored domains; a custom report that includes it always shows it.

35. Microsoft 365

The Microsoft 365 module shows each client's Microsoft 365 at a glance — its users, licenses, multi-factor authentication (MFA), mailboxes, administrators, domains and Microsoft Secure Score — and lets your technicians do the everyday user work without opening the Microsoft 365 admin center: reset a password, block a sign-in, sign a user out everywhere, assign or remove a license, set automatic replies and create a user. It works through an app that SLApp creates in each client's own tenant when an administrator connects the client, once. From then on the server reads the tenant every 6 hours and raises alerts, for example when an administrator has no MFA. Nothing is changed in a client's tenant unless a technician clicks and confirms it, apart from the renewal of the app's own secret. Open the module from Cloud & web › Microsoft 365 in the sidebar; the same data is on the Microsoft 365 tab of

each client's sheet.

Who can do what

- Every technician: opens the module and the clients' Microsoft 365 tabs, runs Sync now and uses the user actions (each one confirmed first); on the client's Data tab, imports users as contacts and works with the BitLocker keys.
- Administrators only: Connect Microsoft 365, Reconnect and Disconnect, and Resolve on a Microsoft 365 alert. The server refuses these to anyone else, and a sign-in is only accepted if the person who started it is still an active administrator when it comes back. Only administrators are also shown the list of what went wrong in the last connection.
- In the client's tenant: a Global Administrator of the client's Microsoft 365 signs in once to connect it, and again for each Reconnect. That sign-in is used only during the connection and never stored.
- Every connection, disconnection and user action is recorded in the Audit log under the entity Microsoft 365; passwords never are.

Connecting a client

A client is connected from the Microsoft 365 card at the bottom of its Data tab (Clients › the client › Data). When an administrator clicks Connect Microsoft 365, a Global Administrator of the client signs in with Microsoft. With that sign-in SLApp creates an app named SLApp Helpdesk - Microsoft 365 in the client's tenant, grants it the permissions it needs and gives it a secret of its own; every later sync and action signs in as that app. The powerful permissions the setup needed are taken back at the end. When you have no Global Administrator account of the client's own, its primary contact can connect it from the client portal instead (My company > Connect Microsoft 365, with the same Global Administrator sign-in); the connection then works exactly the same, the audit log names the contact and the last connection's problems say it came from the portal.

1. Open the client on its Data tab and find the Microsoft 365 card, which reads Not connected. A technician who is not an administrator sees Ask an administrator to connect instead of the button.
2. Click Connect Microsoft 365. A Microsoft window opens and the button reads Awaiting approval…. If the browser blocks the window, a message asks you to allow pop-ups for this site; allow them and click again.
3. In the Microsoft window, pick or type the account of a Global Administrator of the client's Microsoft 365 and accept the permissions Microsoft lists. If the role is assigned through Privileged Identity Management (PIM), activate it first. The sign-in can be used only once and must be completed within 10 minutes.
4. SLApp sets the app up, and the window shows the outcome at the top (Microsoft 365 connected, Microsoft 365 connected, with warnings, or Microsoft 365 connection failed), the Tenant, Signed in as, and every step with OK, Warning, Failed or Skipped. When all the steps are OK the window closes by itself and a toast says Microsoft 365 connected; with warnings or a failure it stays open, so that you can read what happened and what to do.
5. The card now reads Connected, with the tenant's name and default domain, Connected by followed by the account and the date, and when the App secret expires. The first sync runs within a few minutes; Import users and the BitLocker keys work at once.

What each step of the connection does

The steps appear in this order in the Microsoft window:

- Sign-in with Microsoft: the Global Administrator's sign-in. A personal Microsoft account is refused; sign in with a work account of the client's Microsoft 365.

- Global Administrator check: the account must hold the Global Administrator role, directly or through a group. Nothing is written in the client's tenant before this check passes.
- Microsoft 365 tenant: a client already connected to one tenant cannot be connected to another until it is disconnected. When another of your clients already uses the same tenant, the step says so.
- App "SLApp Helpdesk - Microsoft 365": the app is created, or found again on a Reconnect, and its list of permissions is written again, so a Reconnect also repairs an app someone edited in the Microsoft Entra admin center. Its notes say that SLApp Helpdesk created it for your company and that deleting it disconnects SLApp.
- App in the tenant (enterprise application): the app's instance in the tenant, which receives the permissions and the role.
- API permissions: the administrator's consent to each Microsoft Graph application permission the app needs (see the tip below). A permission that cannot be granted is a warning, and whatever needs it fails until a Reconnect.
- User Administrator role: given to the app, because resetting passwords needs it. If it cannot be given, password resets fail (a warning).
- App owner (automatic secret renewal): the app becomes the owner of its own registration, so that it can renew its secret by itself. If this fails, the renewal will fail too (a warning).
- App secret: a new secret valid for 24 months (12 or 6 months when the tenant's policy does not allow that long), stored encrypted.
- Saving the connection: from this step on, the connection works.
- Removing old secrets: any other secret of the app is deleted, so the only one left is the one SLApp holds.
- Removing the setup permissions: the permissions the Global Administrator's sign-in gave SLApp's own sign-in application, needed only to set the app up, are taken back, even when the connection failed. If that fails, the window says how to revoke them by hand: Microsoft Entra admin center > Enterprise applications > the SLApp Helpdesk application > Permissions.

The sign-in, the Global Administrator check, the tenant, the app, the enterprise application, the secret and the saving are required: when one of them fails, the connection fails, the steps after it are Skipped and the window says why, for example The account that signed in is not a Global Administrator of this Microsoft 365, or This client is already connected to another Microsoft 365 tenant. Disconnect it first to connect a different one. The other steps only leave warnings, which a Reconnect repairs. Administrators also find the steps that went wrong under the card, after The last connection reported problems:, each with what to do about it.

> **i** The app receives these Microsoft Graph application permissions: User.ReadWrite.All, Directory.Read.All, Organization.Read.All, AuditLog.Read.All, UserAuthenticationMethod.Read.All, Reports.Read.All, MailboxSettings.ReadWrite, BitlockerKey.Read.All, Device.Read.All, Domain.Read.All, SecurityEvents.Read.All and Application.ReadWrite.OwnedBy (to renew its own secret), plus the User Administrator directory role. Give the client's IT contact this list if they ask what is being granted.

The connection's state

The card on the Data tab, the client's Microsoft 365 page and the module show the state of the connection as a badge. Hover it to read the reasons, which are also listed under the card and at the top of the client's Microsoft 365 page:

- Connected (green): everything works.
- Needs reconnect (amber): the last connection did not finish; some permissions are missing (they are

named); the app is missing the User Administrator role, so password resets fail; the app does not own its registration, so its secret cannot renew itself; or the app secret expires soon and could not be renewed, or has expired. A Reconnect fixes all of these.

- Failing (red): the last 3 syncs failed; the error is shown with the last sync.
- Old connection (amber): a client connected before SLApp created an app for each client; see Old connections below.
- Not connected (grey).

Reconnect

Reconnect, for administrators, is on the Data tab card, at the top of the client's Microsoft 365 page and in the module's right-click menu. It runs the same sign-in as Connect, with a Global Administrator of the same tenant, and every step is safe to repeat: the same app is found again, whatever is missing is granted, the role and the ownership are set again and the secret is replaced. Use it when the state is Needs reconnect or Failing, when an error message ends with reconnect, and to upgrade an old connection. A successful Reconnect also closes the Microsoft 365 connection failing and Microsoft 365 app secret expiring alerts.

Old connections

Clients connected with the earlier Connect Office 365 button use the old multi-tenant consent instead of an app of their own. Their badge reads Old connection and the card says Old connection — reconnect to enable Microsoft 365 management (users, licenses, MFA, mailboxes). The contacts import and the BitLocker keys keep working meanwhile. Until an administrator clicks Reconnect, the client is not synced and raises no Microsoft 365 alerts, its Microsoft 365 page shows only this notice, and no user action is possible. Reconnect creates the client's own app and upgrades the connection; there is no need to disconnect first.

Disconnecting a client

Disconnect, for administrators, asks for confirmation in the Disconnect Microsoft 365 window. SLApp first revokes the app's secret in the client's tenant when it can, so that the app left behind can no longer sign in; then it forgets the connection and its secret, stops syncing the client, drops what it had read from Microsoft 365 and closes the client's Microsoft 365 alerts. The contacts and BitLocker keys already imported stay.

The app itself stays in the client's tenant: deleting it there is the client's decision. A card titled The app stays in the client's tenant shows where to do it — Microsoft Entra admin center > Entra ID > App registrations > All applications > "SLApp Helpdesk - Microsoft 365" (with its App ID) > Delete — with a Copy App ID button, and adds Its secret was revoked, so it no longer works. when the revocation succeeded. If the same tenant is connected again later, the app is reused when it still exists.

The Microsoft 365 module

The module lists, one row each, the clients with a Microsoft 365 connection, including old connections and connections left unfinished. The columns are Client, Tenant (its name and default domain), State, Users (enabled / total), Licenses (in use / bought, with a badge such as 1 license type full), MFA (the share of users with MFA, in red below 100%, and a red badge such as 2 admins without MFA), Secure Score (current / maximum), Last sync (pending while a sync runs, never before the first one, and a warning whose tooltip gives the error when the last syncs failed), App secret (the expiry date with its badge) and Alerts (the open alerts, red when one is critical). When no client is connected, the module says No client is connected to Microsoft 365 yet.

- Above the list: a search box (Search client, tenant or domain…), a state selector (All states or one state) and a Show selector: Show: all, With open alerts, MFA gaps (an administrator without MFA, or

MFA below 100%), Licenses exhausted (a license type with every unit in use) and App secret expiring (within 30 days). While filters hide rows, Showing N of M says how many are shown.

- The filter row under the headers adds text filters on Client and Tenant, Last sync (All, Failing, Never synced) and Alerts (All, With alerts, Without alerts). A click on a header sorts by it (State puts Failing first), and Columns chooses and orders the columns.
- A click on a row opens that client's Microsoft 365 page inside the module; the refresh button at the end of the row is Sync now. The right-click menu offers Open, Open client, Sync now, Reconnect (administrators), Copy tenant ID and Copy domain.
- The list refreshes itself every minute, and every 5 seconds while a sync runs.
- Under the list, N clients without Microsoft 365 (click to expand) names your other active clients, with a reminder of how to connect one; a click on a name opens that client's Data tab.

A client's Microsoft 365 page

This page opens from a row of the module, with ‹ Microsoft 365 to go back and Open client to open the client's sheet on its Microsoft 365 tab, and it is also the Microsoft 365 tab of the client's sheet. Its top card shows the state, the tenant's name and default domain, Sync now and, for administrators, Reconnect, followed by these facts: Tenant ID, Connected (the date and who signed in), Last sync, Users (N enabled of M), Licenses (N of M in use), MFA (the percentage, how many of how many, and a badge for administrators without MFA), Secure Score (current / maximum and percentage), App secret expires, App (the app's name; hover it for its App ID) and Open alerts.

Notes under the facts say what could not be read and why, for example Last sign-in dates need Microsoft Entra ID P1 in this tenant, MFA is read user by user (no Microsoft Entra ID P1 for the registration report), Secure Score is not available for this tenant, Only the first 5,000 users are synced, or that Microsoft 365 hides user names in its usage reports, so mailbox sizes cannot be matched to users — with the setting to untick in the Microsoft 365 admin center: Settings > Org settings > Reports > Display concealed user, group, and site names in all reports. When a part could not be refreshed by the last sync, it keeps its previous values and is listed with its error. Then come the Open alerts of the client (Severity, Type, Message, Opened at), with Resolve for administrators.

Below the card, four sub-tabs show what the last sync read, each with its count: Users, Licenses, Admins and Domains. They only appear for a client connected through its own app.

Users

The Users sub-tab lists the tenant's users with the columns Name (a Guest badge for guests, a shield icon for an administrator — hover it for the roles — and a building icon for a user synchronised from the on-premises Active Directory), UPN, Status (Enabled or Blocked), Licenses (a warning when Microsoft reports a license assignment error), Last sign-in, MFA (registered — hover it for the methods —, × no MFA method registered, ? not known yet), Mailbox (the share of the quota in use, in red from 90%; hover it for the sizes) and Created. Department, Job title, Usage location and Administrator roles can be added with Columns.

- The toolbar filters by a search on name, UPN or email, by status (All statuses, Enabled, Blocked), by license (All licenses, With a license, Without a license, or one license), by MFA (MFA: all, MFA: yes, MFA: no, MFA: unknown) and by role (All roles, Administrators, Guests); the row under the headers adds a text filter per column, and a click on a header sorts.
- Click a user to select it: the action buttons of the toolbar act on the selected user — + New user, Reset password, Block sign-in (Unblock for a blocked user), Sign out everywhere, Licenses and Out of office. Until a user is selected they are disabled (Select a user in the list first).
- A right-click on a user offers Reset password, Block sign-in (or Unblock), Sign out everywhere, Assign license, one Remove followed by the license name for each license the user holds, and Out

of office, plus Copy UPN and Copy name.

- At most 500 rows are shown at once; above that, a note asks you to refine the search.
- Last sign-in needs Microsoft Entra ID P1 in the client's tenant; without it the column shows a dash (hover it for the reason). Never means that the user never signed in.

Licenses, Admins and Domains

- Licenses lists the tenant's subscriptions: License (the friendly name and the part number), Status (Active, Warning, Suspended, Locked out or Deleted), Total, In use, Free (a red badge when more are in use than bought, an amber 0 when none is free), Suspended / warning and Waiting (users waiting for a license of this type: Microsoft could not assign them one because none was free). Paid licenses come first; free and trial subscriptions are greyed out, marked with an information icon and left out of every total. Right-click a license and choose Show users to see who holds it in the Users sub-tab, or copy its part number. The friendly names come from Microsoft's official list of licence names, which SLApp downloads again about once a week, so new licences get their name without an update of the app; a licence Microsoft has not named yet shows its part number.
- Admins lists the users who hold an active administrator role in Microsoft Entra ID, with Name, UPN, Administrator roles, MFA, Last sign-in and Status, and the same right-click actions as Users. Reader roles, such as Directory Readers, Reports Reader or Message Center Reader, do not make a user an administrator.
- Domains lists the tenant's domains with Domain, Default, Initial (the .onmicrosoft.com one), Verified, Services and Authentication (Managed or Federated). The verified domains are also offered by Suggest in Cloud & web › Domains.

Acting on a user

Every action runs in the client's tenant, as the client's own app, only after you confirm it, and any technician can use them. A toast says how it went, the user's row is read again from Microsoft 365 at once, and the action is recorded in the Audit log under Microsoft 365. When Microsoft refuses, the message says why in plain words, for example A User Administrator cannot change an administrator — do it in the Microsoft 365 admin center, This user is synchronised from the on-premises Active Directory — change it there, No free license of this type — buy more or free one up, or The SLApp app is missing the permission User.ReadWrite.All: reconnect. The reset, block, unblock and sign-out windows also warn you beforehand when the user is synchronised from the on-premises Active Directory (Microsoft may refuse the change, or the next synchronisation may undo it) or is an administrator (the app's User Administrator role cannot change most administrators).

Resetting a password

1. Select the user and click Reset password, or right-click the user and choose Reset password.
2. In Reset the password of ana@example.com, leave Require a password change at next sign-in ticked (the default) unless the user should keep the new password, and click Reset password.
3. The Temporary password window shows the new password with a Copy button and says whether the user has to change it at the next sign-in. It is shown only once and stored nowhere: copy it now, give it to the user safely, and click Close.

SLApp generates the password: 16 characters with lowercase and uppercase letters, digits and symbols, leaving out the characters that are easy to confuse (0 and O, 1, l and I), because a temporary password is often read out over the phone. The password of an administrator usually has to be reset in the Microsoft 365 admin center, since the app's User Administrator role cannot change most administrators.

Blocking, unblocking and signing out

- Block sign-in asks Block the sign-in of ana@example.com? Every session is closed too, and signing in stays impossible until unblocked. SLApp disables the account and then revokes all of its sessions, so a stolen session ends with it.
- Unblock, on a blocked user, asks Let ana@example.com sign in again? and enables the account.
- Sign out everywhere asks Close every session of ana@example.com, on every device? The user has to sign in again; some apps can take up to an hour to notice.

A user's licenses

Licenses (Assign license in the right-click menu) opens the user's Licenses window. Assigned licenses lists each license the user holds with a Remove button, and shows the assignment error in red when Microsoft reported one. Under Assign license, the — choose a license — list offers the licenses that still have free units, each shown with how many are free; choose one and click Assign. Microsoft needs a usage location before it assigns a license: when the user has none, a Usage location field appears for a two-letter country code (PT by default). Remove asks, for example, Remove Microsoft 365 Business Standard from ana@example.com? The services included in it stop for this user. A license that comes from a group cannot be removed here: change the group membership instead.

Automatic replies (Out of office)

Out of office opens the user's automatic replies with their current state: on, scheduled between two dates, or off. Tick Only during this period to give a Start and an End, in your local time. Write the Reply inside the organisation (required), choose who outside the organisation receives a reply in Outside the organisation, reply to — No one, The user's contacts only or Everyone — and write the Reply outside the organisation, or leave it empty to send the same reply as inside. Each reply can have up to 4,000 characters. Save and turn on switches the replies on; Turn off, shown while they are on, switches them off. The user needs an Exchange Online mailbox.

Creating a user

1. In the Users sub-tab, click + New user.
2. Fill in First name and Last name. Display name * and Alias * follow them (for example Ana Silva and ana.silva) until you type in them yourself; the alias takes letters, digits, dots, hyphens and underscores.
3. Choose the Domain * among the tenant's verified domains, the default one first. When none is listed, sync the client first.
4. Check the Usage location *, a two-letter country code (PT by default): Microsoft needs it to assign licenses.
5. Tick the Licenses to give; only those with free units are offered.
6. Leave Require a password change at next sign-in ticked unless the user should keep the first password, and click Create user.
7. The Temporary password window shows the new user's password once, as for a reset. If the user was created but the licenses could not be assigned, the window says so; assign them later with Licenses.

The sync

The server reads each connected client every 6 hours, plus up to half an hour so that clients spread out, and reads a new connection within a few minutes. A sync reads the organization and its domains, the users, the licenses, who holds an administrator role, MFA, mailbox usage and the Secure Score. If the organization or the users cannot be read, the sync fails: it is tried again after 15 minutes, until the

third failure in a row, which raises an alert and goes back to the 6-hour rhythm. Any other part that cannot be read keeps its previous values and is listed as a warning on the client's page. A sync never changes anything in the client's tenant, apart from renewing the app's secret when it is due.

Sync now, open to every technician, reads the client at once, at most once a minute and never twice at the same time. The page waits up to 25 seconds; a longer sync goes on in the background (the toast says, for example, Sync of Contoso started — it goes on in the background) and the page tells you when it ends.

- Users: up to 5,000 per tenant. Last sign-in dates need Microsoft Entra ID P1 in the tenant.
- MFA: from Microsoft's registration report when the tenant has Microsoft Entra ID P1; otherwise user by user, up to 300 users per sync, administrators first, the following syncs completing the rest. The MFA figures count the enabled members who have a license or an administrator role; guests are left out.
- Mailboxes: from Microsoft 365's usage report, which can only be matched to the users when the tenant does not hide user names in its reports (Microsoft 365 admin center > Settings > Org settings > Reports).
- Licenses: only the licenses actually bought count in the totals and the alerts; free and trial subscriptions, which come with thousands of units nobody paid for, are listed but never counted.

Microsoft 365 alerts

- Administrator without MFA (critical): one for each enabled administrator with no MFA method registered, naming the roles. It closes when the administrator registers MFA or stops being an enabled administrator.
- Users without MFA (warning): one per client, with how many enabled members with a license, or administrators, have no MFA method registered. It closes when none is left.
- Licenses exhausted (warning): one for each paid license type with more licenses in use than bought, or with every unit in use while users wait for one. A license type that is simply full, with nobody waiting, is normal and raises nothing. It closes once the license type is back in order.
- Mailbox almost full (warning): one for each mailbox at 90% of its quota or more, only when the usage report shows the user names. It closes when the mailbox drops below 90%.
- Microsoft 365 connection failing (warning): the last 3 syncs failed, with the error. It closes at the next sync that works, or with a Reconnect.
- Microsoft 365 app secret expiring (warning; critical in the last 7 days and once expired): the app secret expires within 30 days and its automatic renewal did not succeed. It closes once the secret is renewed, or with a Reconnect.

When an alert opens, every active administrator of your company with an email address receives an email headed Microsoft 365 alert, with its target (a user's UPN, a license or the client), the client, the message and the time, and a subject such as Warning alert: Contoso — 3 user(s) without MFA (the target, then the message; Critical alert for a critical one). If the alert turns from warning to critical, another email follows, headed Alert now critical, with a subject starting with Now Critical. Closing an alert sends nothing. The first sync of a connection opens its warnings without email, except the two about the connection and its secret; critical alerts are always emailed. While a part cannot be read, such as MFA or the mailboxes, its alerts stay as they are. The alerts show on the client's Microsoft 365 page and in the module's Alerts column, in the Alerts module (the message starts with the people icon, and their Source is Microsoft 365), in the client's Alerts tab, on the Dashboard and in the footer ticker. Disconnecting or deactivating a client closes its Microsoft 365 alerts. An administrator can stop this module's alerts for your whole company with the Generate alerts checkbox at the top of the page (see Turning alert generation on or off per module in the Alerts section).

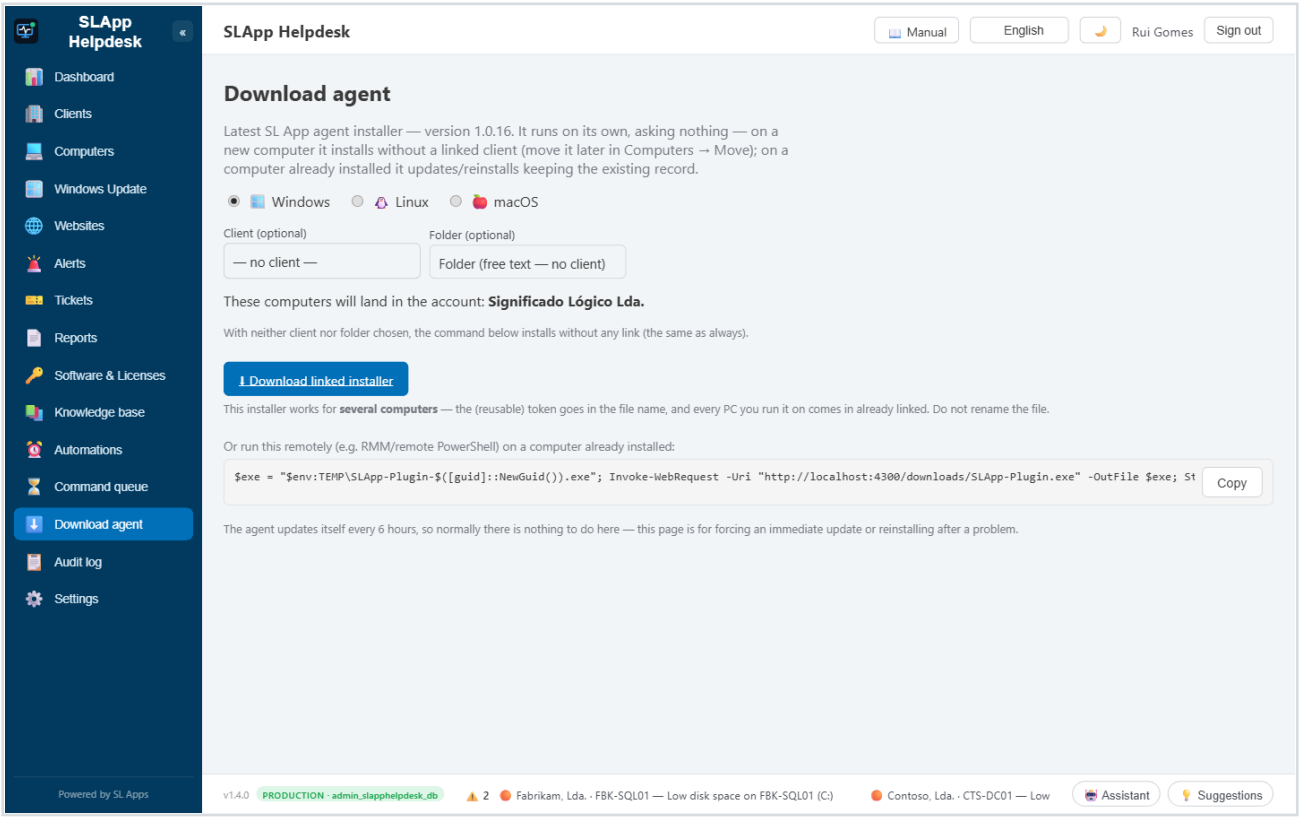
The app's secret

The app's secret is valid for up to 24 months and renews itself: within 30 days of its expiry, the next sync adds a new secret, checks that it works and only then removes the old one, trying again every 6 hours while it fails. The renewal needs the App owner step of the connection to have succeeded. The expiry date appears as App secret expires, in amber in the last 30 days and in red in the last 7. If the renewal fails, the client's page shows The automatic renewal of the app secret failed: followed by the error, and the Microsoft 365 app secret expiring alert opens; a Reconnect gives the app a new secret.

> **i** What you see is the picture of the last sync: a change made in the Microsoft 365 admin center appears in SLApp after the next sync, or at once with Sync now, while an action made in SLApp updates its user straight away.

36. Download Agent and installing on computers

The SL App agent is the small program that runs on every computer you monitor and support. Once installed, it keeps the computer's record up to date (hardware, software, disks, network, patches), sends heartbeats so you can see whether the machine is online, and lets you run remote commands, scripts and remote sessions from the app. On Windows, from version 1.1.0, it also runs the backup jobs you define in the Backups module. This section explains how to get the agent onto Windows, Linux and macOS computers, how to make each install land in the right client and folder automatically, and what the agent looks like on the computer afterwards.



The Download agent page, with the OS selector, optional client and folder, and the linked installer.

The Download agent page

Open Administration › Download agent in the sidebar (a technician who sees no other module of that group finds Download agent directly in the sidebar, without the group). The page shows the version of the latest available installer and lets you build an installer (or a command) that is already tied to the destination you choose.

- OS selector (Windows, Linux, macOS): picks which installer or command is shown below. Windows gets a downloadable installer plus a remote command; Linux and macOS get a single terminal one-liner.
- Client (optional): the client the computers should be linked to. Leave it on — no client — to install without a link; the computer then appears in the unassigned Computers list and you can move it to a client later with Move.
- Folder (optional): with a client chosen, this is a dropdown of that client's folders (— no folder — to skip). With no client chosen it becomes a free-text field (Folder (free text — no client)).
- These computers will land in the account: … — always tells you which account the installs will end up in. With no client chosen it shows your own company name (from Settings › Company), so an unlinked install is never ambiguous.
- A short status line confirms what the command/installer does: either that it installs already linked to the chosen client and folder, or that it installs without any link.

Behind the scenes the page prepares a reusable install token for the chosen client and folder (reusing an existing one if it already exists). While that happens you see Preparing the link…; if it fails, a message tells you to check your connection, with a Try again button.

Installing on Windows

With Windows selected you get two options. The first is the Download linked installer button (labelled Download installer — client name when a client is chosen). The file you download carries the install token in its file name, so the same file works on as many computers as you like, and every PC you run it on comes in already linked to the chosen client and folder. Do not rename the file, or the link is lost.

1. Open Administration › Download agent in the sidebar and leave Windows selected.
2. Choose the Client (optional) and, if you want, the Folder (optional) the computers should land in. Check the account shown after These computers will land in the account.
3. Click Download linked installer and copy the file to the computer (USB stick, shared folder, or send it to the user).
4. Run the file on the computer. It installs on its own, without asking anything, and the computer appears in Computers a few moments later, already linked.
5. Repeat on every other computer with the same file.

The second option, under Or run this remotely (e.g. RMM/remote PowerShell) on a computer already installed, is a PowerShell command you can copy with the Copy button (it briefly reads Copied! to confirm). It downloads the latest agent and runs it silently, including the same link to client and folder. It is meant for machines you already reach remotely (for example through another RMM tool or a remote PowerShell session) and is also the way to force an immediate update or to reinstall after a problem: on a computer that already has the agent, running it updates or reinstalls while keeping the existing computer record (on a machine whose computer was removed in the app it creates a new computer instead; see Removing a computer and its agent in the Computers section).

> **i** Exclude the agent from the antivirus on the computer to guarantee uninterrupted remote access and inventory. Some security products otherwise quarantine a freshly downloaded, self-installing program.

Installing on Linux and macOS

Select Linux or macOS and the page shows a single command to copy. Run it in a terminal on the target computer as root (for example through sudo). It downloads the install script, fetches the agent and sets it up as a background service; there is no window and nothing else to click through. The

command already contains the install token, so the computer arrives linked to the chosen client and folder, and the same command can be run on every other computer. The only requirements are curl and, on Linux, systemd. If someone is signed in to a desktop session on the machine, the same command also installs the small menu-bar icon described below; on a headless server it is simply skipped and the background agent is unaffected.

The Install agent wizard in Computers

If you are already in the Computers list, the + Install agent button opens a short wizard called Install the SL App agent with three steps: Select OS, Assign agent and Install agent. Back takes you to the previous step and Cancel closes the wizard at any point.

1. Select an OS: pick Windows (.exe), Linux or macOS and click Next.
2. Assign the agent to a client: choose the Client, or leave — no client (I'll link it later via Move) — to keep the computer unassigned for now. If the client is already known (for instance when the wizard is opened from a client's page) the field is shown read-only.
3. The wizard checks whether a reusable token already exists for that client (or for unassigned installs when no client is chosen). If it does, it tells you so and reuses it. If not, you can fill in Folder (optional) and Description (optional), e.g. "Reception PC", and the token is generated when you click Next (the button shows Generating… meanwhile).
4. Install the agent: the Install token is shown with a Copy button. When an existing token was reused, the label adds (reusable — use the same one on every computer); a freshly generated token is reusable too. Below it, for Windows, you have three tabs: Install via command line (a PowerShell command to run on the computer), Download installer (the Download linked installer button — or Download installer for client name — with the same do-not-rename rule and the antivirus reminder) and Share link (a URL you can send to the client; whoever opens it downloads the installer that links itself, and the link works for every computer). For Linux and macOS there is just the terminal command.
5. Click Done when you have copied what you need.

> **i** Install tokens are reusable on purpose: generate one per client (or per client and folder) and use it on every upcoming computer. You do not need a new token per PC. Existing tokens can be reviewed from the client's sheet.

What the agent does on the computer

On Windows the agent runs as a background service and also places a small icon in the notification area (tray). When it starts, a notification SL App — Agent started appears and the icon's tooltip shows the agent version and the client it is linked to (SL App v… · client). Right-clicking the icon opens a menu with Open support ticket, Status, Password manager (only on some computers, see below) and Exit. Open support ticket opens the client portal in the default browser, already signed in as the contact linked to the computer and with the computer pre-selected on the new support request (if the computer has no linked contact who is allowed to use the portal, the normal Microsoft sign-in appears first). If the server cannot be reached, the agent falls back to a local form (Open a support request) with Subject, Type, Impact, Priority and Description, plus an Attach files… button; the computer and the signed-in user are linked to the request automatically and the ticket is created directly in your Tickets module. Double-clicking the icon (or choosing Status) opens the Status — SL App Agent window, which shows the agent version, the server it talks to, the computer name, whether it is currently connected and the time of the last heartbeat. The window has Sync (sends a fresh inventory immediately), Update (checks for a newer agent; if one is available the agent restarts itself) and Open ticket buttons, a usage chart (CPU, RAM and system disk) with a selectable period (Hour, Day, Week, Month), and two lists: Disks

(Drive, Total, Free, Usage) and Connected network adapters (Adapter, IP, Speed).

From version 1.2.13 the Windows service starts as Automatic (Delayed Start): Windows starts it about two minutes after the computer boots, so that on a busy server it does not compete with everything else that starts at that moment. If the agent stops unexpectedly, Windows starts it again after one minute (after five minutes from the third failure in the same day). On top of that, a scheduled task called SLApp-Plugin Watchdog, run by the system account, checks every 10 minutes (the first time 10 minutes after the boot) that the service is running and starts it if it is not; this also covers a service that did not manage to start at all after a restart. A service that an administrator set to Disabled is left alone. Agents installed before 1.2.13 get these settings at their next automatic update, the agent puts them back by itself if someone removes them, and uninstalling the agent or removing the computer also removes the task.

The tray menu also offers Password manager (agent 1.0.47 or later), but only on computers where the password manager has been switched on. It is the same switch as the SLApp Vault extension: tick Install the SLApp Vault extension on this computer in the SLApp Vault extension card of the computer's OS and security sub-tab, or use Install SLApp Vault extension in the Actions on N dialog of the Computers list, and the computer gets both the browser extension and the tray entry; Remove SLApp Vault extension switches both off. The entry does not appear at once: it shows up after the agent's next heartbeat, and goes away (closing the window if it is open) after the heartbeat that follows you switching it off.

Choosing Password manager opens the personal password vault, the same Vault page as in the client portal, in a window of its own, so neither a browser nor the extension is needed. The window opens already signed in as the contact linked to the computer; if there is no such contact (none is linked, or the linked contact is inactive, not allowed to use the portal or belongs to a different client), it shows the normal Sign in with Microsoft page. The vault still asks for its master password, which the agent never stores, and it stays end-to-end encrypted: nobody, technicians included, can read it without that password. Only one such window is open at a time. The server refuses the request on any computer where the switch is off. If the Microsoft Edge WebView2 Runtime is missing from the PC, or fails to start, the agent offers to open the same vault in the default browser instead. The vault itself is described in the Client Portal Guide.

On Windows the agent can also show the user three prompts. When a technician starts a remote session, a Remote access request window asks the user to Authorize or Don't authorize; if nobody answers before the countdown runs out, the organisation's policy decides (Settings › General › Remote control): access is refused by default (Refuse automatically). Since version 1.14.10 this is also true for organisations created earlier: their setting was switched once to Refuse automatically, and the Audit log shows a System entry for each switch. An administrator can choose Authorize automatically again. A locked or sign-in screen is not affected, and a computer set to Connect without asking for authorization shows no popup at all. When a restart is requested, a Computer restart notice tells the user to save their work and lets them choose Restart now, Cancel, or Postpone by 1, 2, 3 or 4 hours; here too, the restart goes ahead on its own when the countdown ends.

The third one (agent 1.2.41 or later) only appears where the organisation switched on Ask the user to restart (see Asking the user to restart in the Settings section): when a Windows workstation is waiting for a restart, the signed-in user gets a window titled SL App — Restart needed. It says that updates or changes on the computer only finish after a restart, asks the user to save their work first (open programs are closed 30 seconds after Restart now) and says that if they postpone they will be asked again in N hours, the number the organisation chose. It has two buttons, Postpone and Restart now, and no countdown: it stays on top of the other windows until the user answers, but it never takes the keyboard from the program being typed in, and even after a click on the window Enter, Space and Esc all mean Postpone — only a click on Restart now restarts the computer. Closing it with × or Alt+F4 also

counts as Postpone. Restart now has Windows restart the computer 30 seconds later, with the notice SL App: the computer restarts in 30 seconds to finish updates. If Windows does not let the user restart (a group policy that took that right away from users) or another Windows user signed in meanwhile, nothing restarts: a message asks the user to restart the computer from the Start menu as soon as they can, and the window comes back after the usual number of hours. The window closes by itself, with nothing recorded, once the computer no longer needs a restart or the option is switched off. It is shown in English, Portuguese or French: the language chosen for the computer's contact or, without one, the language of Windows. Each window and each answer appears in the History card of the computer's Remote actions tab.

During a remote session the computer shows a small panel, SL App · Remote connection active. Under Connected to it shows the computer's own name in bold and, beneath it, who is connected: the technician and their organization (technician · organization). Then come the elapsed time (Connected for …), a line when the session is being recorded, and End connection, which ends the session after a confirmation. The panel can be minimized (Minimize / Show details). Before agent 1.2.39 the bold line was the technician's name.

The tray menu also has Agent activity (agent 1.2.1 or later). After asking for the password set in Settings › General › Remote control › Agent activity window, it opens a window with three tabs: Commands and scripts (the remote commands, scripts and maintenance steps waiting in the queue and the last 100 done, with their text and result — double-click a row for the full detail), Backups (the backup jobs of the computer with their next run, and the last runs) and Updates (the agent's version and whether a newer one is on its way, the Windows updates still to install and the software with an update available). Refresh reloads it without asking for the password again. While no password is set the window says so; a wrong password is refused, the attempts are limited, and every opening is recorded in the Audit log.

When a technician starts a conversation (see Chatting with the person at the computer in the Computers section), the agent opens a Conversation with … window by itself, on top of the other windows and flashing in the taskbar — also when the message was sent while nobody was signed in. The person answers in the box at the bottom (Enter sends, Shift+Enter starts a new line), attaches files with the paperclip button or pastes a screenshot or copied files with Ctrl+V, and sees their own messages on the right. The files of the conversation are listed under it: an image opens in the default viewer, anything else is only saved where the person chooses. Closing the window only hides it; the tray menu item Conversation with the technician, shown while a conversation is open, brings it back. When the technician ends the conversation the window says so and stops accepting replies. The person can never start a conversation.

On Windows, the agent's installation folder, C:\Program Files\SLApp-Plugin, holds a single file: SLApp-Plugin.exe. Installing or updating the agent copies nothing else into it. Agents older than 1.0.46 copied other files into it by mistake and could leave hundreds of unrelated files behind (temporary files, scripts, even other programs' libraries); from 1.0.46 the agent tidies the folder by itself every time its service starts, removing any other file at the top of the folder that is more than an hour old. Subfolders are never touched, so there is nothing to clean by hand: a machine that still has the leftovers loses them when its service starts on 1.0.46 or later (after the next automatic update, for example). The cleanup is best-effort: a very large folder or a file in use may leave some behind until the next service start.

From version 1.1.0 the Windows agent can also run scheduled backups (see the Backups section). What it needs for that is kept outside the installation folder, in C:\ProgramData\SLApp-Plugin\backup: an encrypted copy of the job definitions, the state of each job, the reports that are waiting to be delivered and a temporary work folder where ZIP files are built (a job can choose another staging folder). The work folder is emptied when each run ends, and leftovers of a run that was cut short are removed when the agent service starts again.

From version 1.2.0 the Windows agent also keeps on disk what it streams to SLApp Helpdesk while the server cannot be reached: the usage samples (CPU, memory, disks and network) behind the charts and the usage alerts, the Syslog lines of a collector (see the Syslog section) and the Windows events watched by event log alert rules. They wait in C:\ProgramData\SLApp-Plugin\outbox, a folder that only SYSTEM and administrators can open (the notification-area program keeps its own in the user's profile), so an outage of the server, a lost network or an orderly restart of the agent or of the computer no longer loses them, within the limits below. What is still only in memory can be lost on a crash or a power cut: the samples taken since the agent last sent them (one sending interval, 60 seconds by default) and the last few seconds of Syslog lines. When the server answers again, the current data goes first and what waited follows, oldest first, in batches spread out with short random pauses so that a whole fleet coming back does not overload the server; whatever is sent twice is stored once. Each kind has a limit, about 20 MB or 7 days of samples, 50 MB or 14 days of Syslog lines and 10 MB or 7 days of events, beyond which the oldest are dropped; an entry that the server keeps refusing is dropped too, and when the folder cannot be used the agent sends as it did before 1.2.0 and loses what does not get through. Samples that arrive late fill the charts but never open, update or close a usage alert, which always follows the current samples. A watched event found during the outage is reported when it arrives, once, and the agent moves on in the event log only after the events it found are safely stored; after a long gap it looks back at most 7 days. Removing the computer in Computers deletes these folders together with whatever is still waiting in them. The Linux and macOS agents do not have this store.

On Linux and macOS the agent itself is a background service with no window. If the install found someone signed in to a desktop, it also adds a menu-bar icon with the same three choices, Open support ticket, Status and Exit (the Password manager entry exists only on Windows). Its Agent status window shows the computer name, server, connection state, last heartbeat and version, with Sync now, Check for update and Open support ticket buttons (this last one opens your support portal in the browser, as on Windows), a CPU and Memory usage chart with a period selector, and Disks and Network sections. There are no remote-access or restart prompts on these systems; everything else is managed from the app.

The agent also uninstalls itself when its computer is removed in Computers (see Removing a computer and its agent in the Computers section). From agent 1.0.48 on Windows and 1.0.14 on Linux and macOS, removing a computer in the app removes the agent from the machine within a couple of minutes, or a couple of minutes after the machine next connects if it was off. To manage the machine again you install the agent again, and it comes in as a new computer.

Automatic updates

The agent updates itself every 6 hours on every platform, so normally there is nothing to do: when you publish a newer version it reaches the whole fleet on its own. The Download agent page and the remote command are there for the exceptions, such as forcing an immediate update on one machine, reinstalling after a problem, or re-linking a computer. A user can also trigger an update check from the Status window: the Update button on Windows, Check for update on Linux and macOS. On Windows, the first check is made 2 minutes after the agent service starts, and the automatic update waits while a backup is running, for at most six hours, so that it does not cut the backup short. Force update waits as well: while a backup is running it answers “A backup is running - the update is postponed until it finishes (at most 6 h).” and does not try again by itself, so use it again once the backup has ended, or leave it to the next automatic check.

> ℹ Always use the command or installer produced by the app rather than a copy you saved earlier: the command includes the server address and token the agent needs, and without them the agent would register against the wrong place and never show up in Computers.

Testing on the dev server next to production

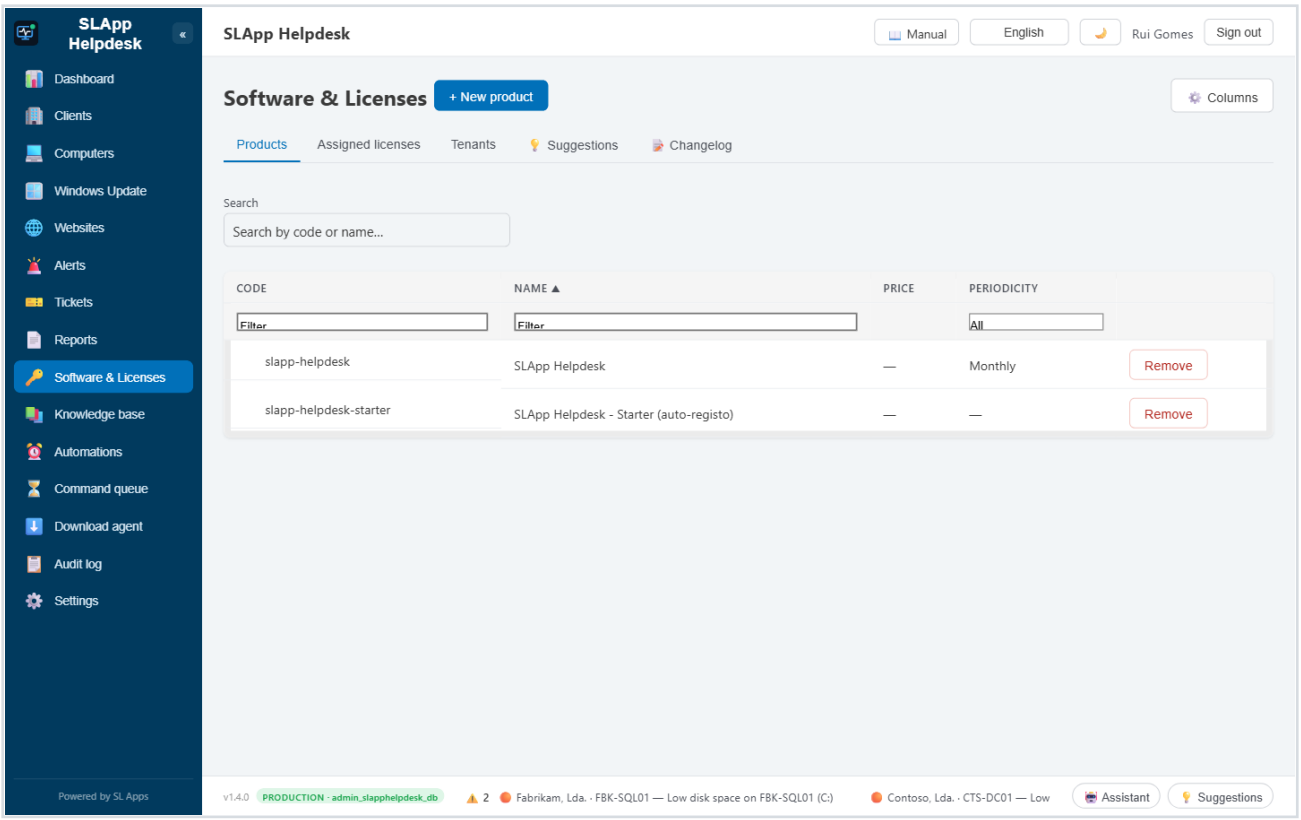
From agent 1.2.38 (Windows only), a computer already managed from production can also be managed by an organisation on SL Apps' own development server (slapp-api.ill-dev.com), so that a change can be tried on real computers before it is published. Only servers on a fixed list built into the agent are accepted, and today that list holds just the dev server. On a computer that already has the agent, an installer pointing at any other server stops with a message saying that server is not allowed (in unattended mode it exits with code 8 and writes the reason to install_error.txt). The production organisation is not asked for its co-management consent, because the dev server cannot reach it; between organisations on the same server every usual rule stays the same.

To add the dev organisation, run the dev installer with the dev server address and an install token from the dev Download agent page, for example: SLApp-Plugin.exe --api=https://slapp-api.ill-dev.com/api --token=<dev token>. The --api argument is required: without it the installer asks production about the token and simply reinstalls the agent. The agent is not reinstalled (only an agent older than 1.2.38 is updated first); the link is recorded and the agent service registers the computer on the dev server within about a minute. The tray then shows the dev organisation with "(dev)" after its name. Automatic updates, the interface language and the other settings of the computer as a whole keep following the production organisation only, so the dev organisation can never update or downgrade the agent. Each organisation's heartbeats, inventory, chat, remote access, backups and scripts go to its own server. When the production organisation removes the other organisations from the computer, the dev organisation is removed too (on the computer only; the dev server keeps showing it offline).

> ℹ Publish agent 1.2.38 in production before running the dev installer on production computers: on an older agent the dev installer first updates it with the dev build.

37. Software & Licenses (operator only)

Software & Licenses is the operator's console for the SL Apps catalogue: the products you sell, the licenses you have assigned to your clients, the SLApp Helpdesk tenants those licenses created, and the inbox of improvement suggestions sent by users. It only appears in the sidebar, as Administration › Software & Licenses, for the operator tenant (the SL Apps company itself); customers who run their own SLApp Helpdesk instance never see it. The screen has four tabs: Products, Assigned licenses, Tenants and Suggestions.



Software & Licenses — the Products tab

Products tab

A product is anything you license to a client: SLApp Helpdesk itself, ConsultisMed, or any other software. The table lists each product's Code, Name, Price and Periodicity. Click any cell to edit the product, or use Remove to hide it (after a Remove product confirmation; removed products are deactivated, not destroyed, and existing licenses keep working). The Search box filters by code or name.

The product form contains:

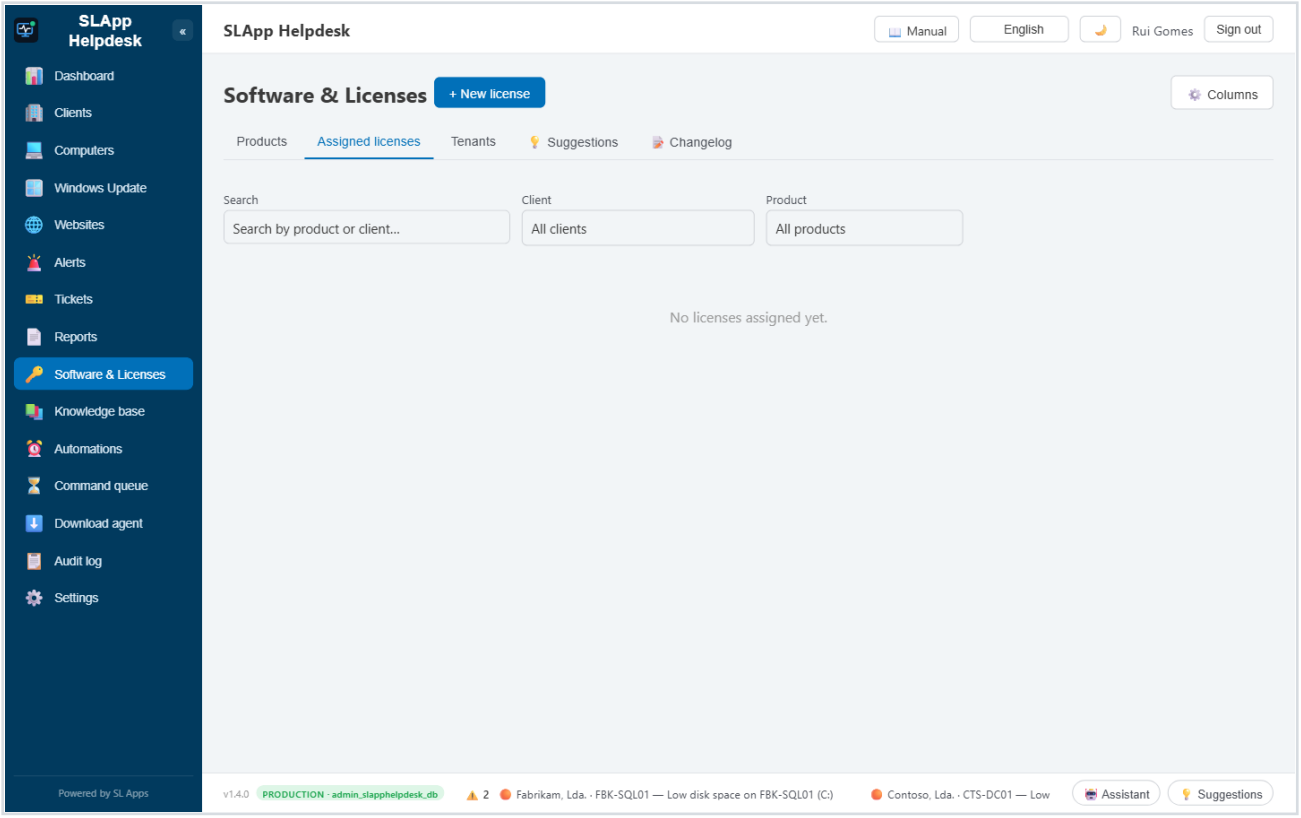
- Code — a short identifier. The codes of the two SLApp Helpdesk products are what triggers automatic tenant provisioning (see below), and the ConsultisMed product's code is what the ConsultisMed integration looks up, so do not change them.
- Name (required) and Description.
- Price (€) and Periodicity — One-time, Monthly or Yearly. This is the product's base price.
- Price per seat (€) and Price per user (€) — default prices used when assigning a license of this product to a client; they can be changed per license.
- Welcome email — a rich-text template sent to the new tenant's administrator (the client's primary contact, or the person who signed up for a self-service trial) when a license of this product is created. Use the placeholders {{client}} (client name), {{product}} (this product's name) and {{logo}} (the SL Apps logo); they are replaced at send time.

1. Open Administration › Software & Licenses and stay on the Products tab.
2. Click + New product.
3. Fill in at least the Name; add Code, Description, Price (€) and Periodicity.
4. Set the default Price per seat (€) and Price per user (€) if the product is billed per computer or per user.

5. Adjust the Welcome email if you want a custom message for this product (a sensible default is already filled in for new products).
6. Click Save.

> A new product is never saved with an empty Welcome email: if you leave it blank when creating the product, the default template is stored instead. If you later clear the Welcome email of an existing product and save, the app falls back to a fixed generic welcome message at send time instead of sending an empty email.

Assigned licenses tab



Assigned licenses with expiration badges

This tab lists every active license, sorted by expiration date, with the columns Product, Client, ConsultisMed Tenant ID, Seats, Users and Expiration. Licenses whose expiration date has passed get an Expired badge; those expiring within the next 30 days get an Expiring badge. Filter with the Search box (product or client) and the Client drop-down (All clients by default). Click a row to edit it, or Remove to deactivate it — the Remove license confirmation reminds you that any tenant provisioned from that license will be suspended automatically.

> Last access shows, for an SLApp Helpdesk license, the last time a technician of its tenant used the app (web or mobile), in your own date and time format; hover it to see how long ago that was. It is recorded at most every 5 minutes, so it can be a few minutes behind, and the client contacts who sign in to the portal do not count. A dash means it is unknown — hover it to see why: nobody has signed in to the tenant yet, the license has no provisioned tenant yet, or the product does not report its access to SLApp Helpdesk (ConsultisMed, Migration TT, Home Automation, Contrata, TemosTudo, BrincaAprendendo and the other products only verify their license key here). The filter under the column header shows Last 7 days, Over 30 days ago or No access recorded. If you had already customized the columns of this list, Last access is added at the end; move it with Columns.

The license form holds: Product and Client (both required, and locked once the license exists), Plan

(only for SLApp Helpdesk), License key, Seats, Users, Activation date, Expiration date, Periodicity (Monthly, Semiannual or Yearly), Allowed domain / endpoint, Price per seat (€), Price per user (€) and Notes. When the license belongs to a ConsultisMed instance, a read-only ConsultisMed Tenant ID field is shown at the top; it is filled automatically when ConsultisMed is provisioned for that client from the client's file, or when ConsultisMed itself reports the link back — never typed by hand.

> **i** Plan decides how an SLApp Helpdesk license is charged and which modules its tenant has: Essential or Complete are charged per technician (Users, or the active users when Users is empty) plus VAT, with no monthly minimum unless the product sets one; no plan keeps the product's flat price and every module. The prices themselves are set on the SLApp Helpdesk product, in its Plans and demo section: the monthly and yearly price per technician of each plan, an optional monthly minimum (0 = none), and what a new license gets — the plan, the billing period and how many days the demo lasts (used by self sign-up and when you pick the product in New license). Changing the plan takes effect at once; moving to Essential also resolves the open alerts of the modules the tenant loses and switches the password manager extension off. The Plan column of the list shows each license's plan, and a self-service trial starts on Complete.

> **i** Billing on the license form sets how SLApp bills an SLApp Helpdesk license: — (not billed by SLApp), Trial, Active (billed by invoice, with Billing period Monthly or Yearly) or Declined. Switching a license to Active removes its expiration date — from then on the invoices decide access — and starts billing today unless a start date was already set. For an Active license the form shows Paid until / end of the demo instead of the expiration date: change it to lengthen or shorten the demo (billing starts the next day) as long as the license has no invoice yet. When billing starts in the middle of a month, the days left in that month are invoiced on the start day (prorated, due 7 days later); the monthly invoices follow from the 1st. For a Trial license the expiration date is labelled End of the demo. The activation date is the day the license was created.

> **i** Dedicated database: the price of the option the customer can tick on its Plan card is set on the SLApp Helpdesk product, in Dedicated database (€ / month, without VAT) (empty = 50 €; a yearly license pays it × 12). The Dedicated DB column of Assigned licenses shows whether each license has it (Yes / No) and a Tonight, Moving or Failed badge while a move is pending; the Tenants tab shows the same badges next to the database, plus Paid, not moved or Own DB, not billed when the billed option does not match where the data is (nothing is moved automatically in that case).

> **i** The prices on the website (www.slapp-helpdesk.com, Pricing section, in English, Portuguese and French) come from this product: the Essential and Complete prices per technician (paying yearly and paying monthly) and the Dedicated database price. A price saved here shows on the site within about 5 minutes, without a deploy; if the site cannot reach the API, it shows the prices written in the page.

The Dedicated DB column, or Dedicated database… in a license's right-click menu, opens a window with the option (billed), where the data is, the customer's last change, the last move and its error. An administrator can switch the option there without the customer's once-a-month limit; the move still waits for the next night (01:00-05:00, Lisbon). When a move fails, the customer only sees that the support team is handling it and the operator receives an email; the window then offers Retry tonight and Retry now.

> **i** Retry now runs the move at once, outside the night window. Moving a company to its own database blocks it for a moment at the end of the copy, but moving it back to the shared database runs in a single transaction and can block the whole API for several minutes (about 10 in production) — only do that out of hours.

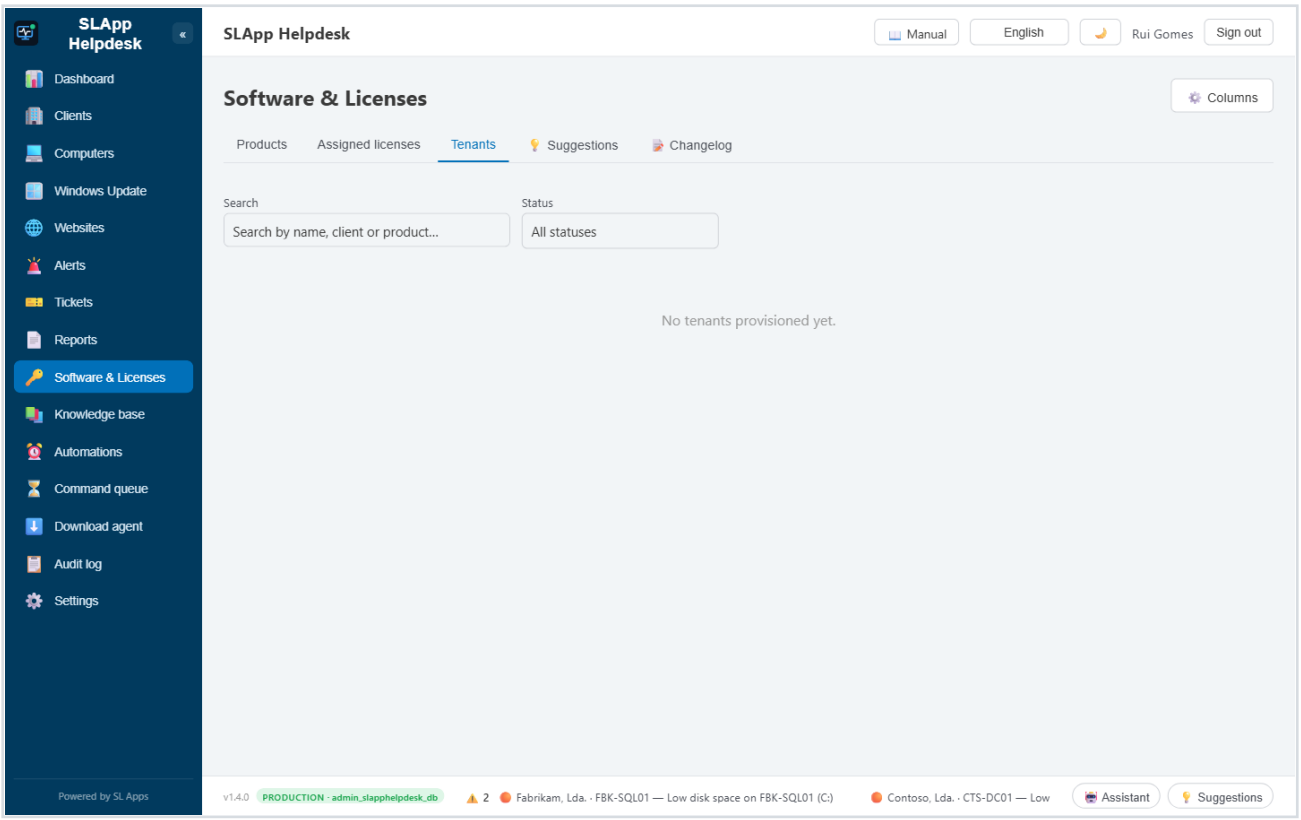
1. Switch to Assigned licenses and click + New license.
2. Choose the Product and the Client. These cannot be changed later, so double-check them.
3. Enter the License key, Seats and Users, the Activation date and Expiration date, and the Periodicity.
4. Adjust Price per seat (€) and Price per user (€) if they differ from the product defaults; add Notes if needed.
5. Click Save. The license is recorded, and a 'New license created' email with the product, client, seats, users, periodicity and expiration is sent to the operator's notification address.

> **i** When the product is SLApp Helpdesk (or its Starter edition), saving the license also provisions a tenant for that client (see the next section). Only one tenant is ever created per client: a second SLApp Helpdesk license for the same client does not create another one. For any other product nothing else happens beyond the notification email.

> **i** Provisioning is re-checked every time a license is saved, so if the tenant failed to appear when the license was created (for example because the email could not be sent), simply open the license and click Save again.

> **i** You cannot lower Seats on a license that backs a provisioned tenant below the number of computers that tenant already has active: remove some computers first, then lower the limit. Users can be lowered below the users the tenant already has: the save warns you, and only that many users can sign in (administrators first, then who joined first) until the company deactivates users; the others see that their company has no license available for their account. Activating or approving a user when every seat is in use is refused with 1 more license is needed (the company itself cannot lower its number of technicians below its active users). Removing a computer also uninstalls its agent from the machine (see Removing a computer and its agent in the Computers section).

Tenants tab



Provisioned tenants and their status

A tenant is a separate SLApp Helpdesk instance for one of your clients. Tenants are created automatically, never by hand, in two ways: when you assign an SLApp Helpdesk license to a client, or when someone signs in with a Work/School Microsoft account from a domain that no tenant authorizes yet (a self-service trial, described below).

Provisioning copies the client's company details (name, tax ID, address, postal code, city, phone, email, authorized domains) into the new tenant, seeds the same starter catalogue every tenant gets, and turns the client's primary contact (or the first active contact with an email) into an approved administrator — for a self-service trial, the person who signed in instead. That person then receives the product's Welcome email with a sign-in button; they simply sign in with their Microsoft account. If the client has no domains configured, the domains of its active contacts' emails (the primary contact first) and then of the client's own email are used instead — never a free-mail domain such as gmail.com or outlook.com, nor a domain another tenant already authorizes; if none is left, the tenant is created but nobody can sign in until its authorized domains are set by hand. The domains the tenant is created with are locked, for a license and a self-service trial alike: they show a padlock in the tenant's Settings › Company and nobody in the tenant, administrators included, can remove them. The one exception is a typed Domain that another tenant already authorizes: it is authorized but not locked, so either tenant can still remove it and solve the clash.

Self-service trials: the sign-up only goes ahead for a member of their own Work/School Microsoft directory — never a personal Microsoft account, a guest account (invited into another organization's directory) or a free email domain such as gmail.com. Only your own clients are looked at; the clients of the other tenants (an MSP's own customers) never block a sign-up:

- None of your clients has the domain in its Domains: a new client named after the domain is created, with the Self-service badge in Clients and the person as its only, primary contact. If a client of another tenant already has that domain, the new client gets no Domains (the portal sends a new contact to the client that has their domain, across tenants); the tenant still authorizes the domain.
- Exactly one of your clients has the domain, is active and has no tenant yet — an existing customer such as Contoso: that client is reused instead of creating a duplicate. The person becomes an active

contact of it: a deleted contact with the same email (in any capitalization) comes back as an ordinary contact — not primary, without the View as, staff-change, billing or credit permissions of the portal, and with its portal sign-in link cleared unless it is the same Microsoft account — or else a contact is created with the name from Microsoft, and the Audit log records it. The other contacts are not touched, the person (not the primary contact) becomes the tenant's administrator and the tenant authorizes all of the client's Domains.

- Either way the client gets an SLApp Helpdesk trial license: 5 computers and 1 user, Billing Trial, with the plan, billing period and demo length of the product's Plans and demo section (Complete, Monthly and 30 days by default). The person is an approved administrator at once and receives the Welcome email; your notification address receives a New self-service tenant email, which names the client when an existing one was reused and says so when the new client was left without Domains.

> **i** No tenant is created when the one client with the domain already has a tenant (add the domain to that tenant's authorized domains instead), already has an SLApp Helpdesk license with no tenant (save that license again to provision its tenant), is inactive, has another Domain that a tenant already authorizes or that is a free email domain, or has your own company's tax ID; nor when several of your clients, or several tenants, have the domain. The reason is written to the server log, and the person sees Your email domain is not authorized to sign in to this application. Colleagues signing in at the same moment create a single tenant and a single license; on the 1-user trial only the first one gets in.

> **i** Locked domains of existing tenants: the server decides them once per tenant, at its first start on version 1.14.85, and every hour after that for a tenant a pass could not reach. A tenant created from a client locks the domains that are both in that client's Domains and in its current authorized list, except one another tenant also authorizes; nothing is ever added to its list, and nothing is locked when the two have no domain in common. The operator tenant (Significado Lógico) locks significado-logico.pt: the domain is added to its list when the list has other domains (recorded in the Audit log), while an empty list — anyone may sign in — stays empty and the lock applies from the next time it is saved with a domain. Tenants not created from a client lock nothing. Locked domains cannot be changed anywhere in the application, not even here: for now only Significado Lógico can change them, directly in the database (company.locked_domains, in the tenant's own database when it has a dedicated one).

> **i** To end a trial for good, click Remove on its license in Assigned licenses: the confirmation is Remove tenant, and the tenant, everything in it and its license are deleted. The client stays, with its contacts and its other licenses, unless the sign-up created it for that trial and it has no other license; the Audit log entry says when the client record was kept. A contact the sign-up reactivated or created stays active. A client that is kept — also when the tenant of an ended subscription is deleted automatically — can start a new trial the next time someone signs in from its domain.

The table shows Name, Client, Product (or No license), Status, Expiration and Created, with a Search box (name, client or product) and a Status filter (All statuses, Active, Suspended, Cancelled). The action buttons depend on the current status:

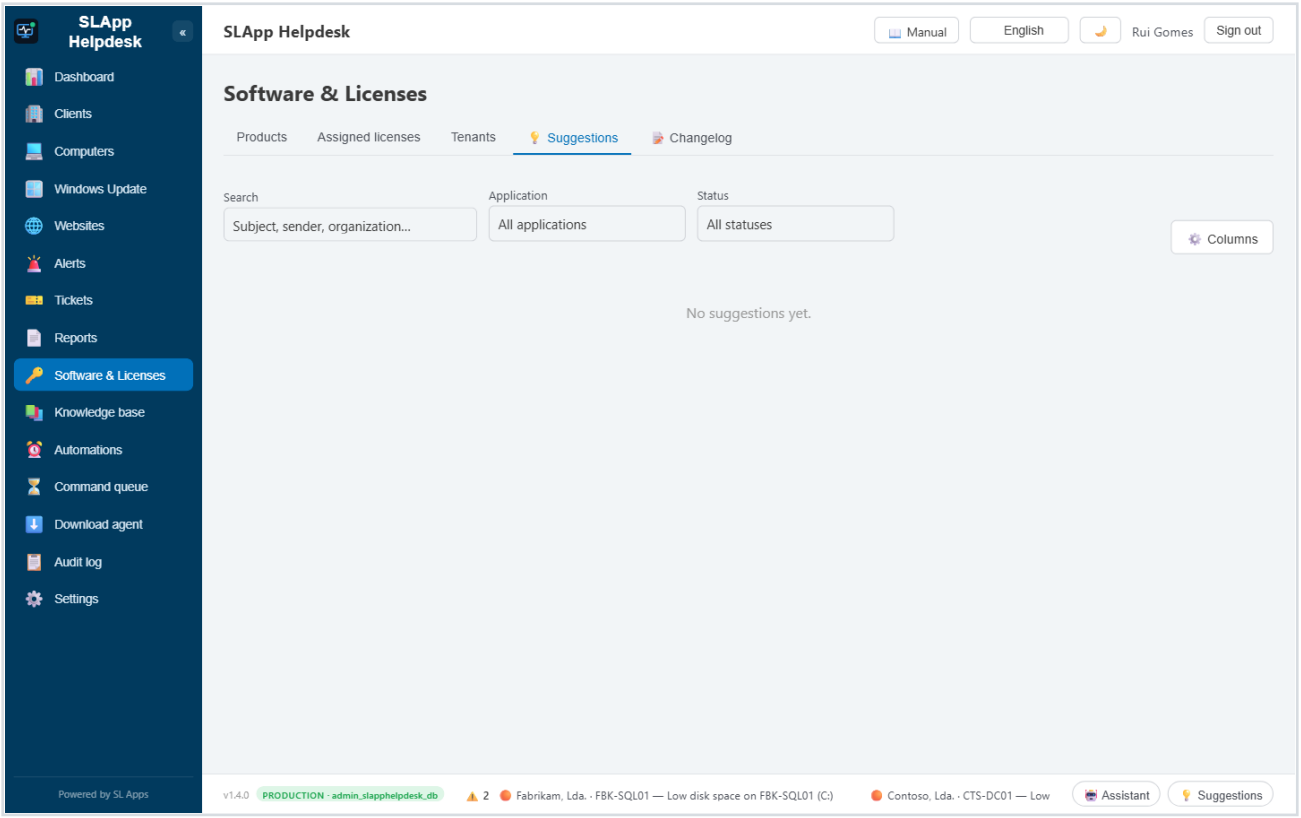
- Suspend — shown for Active tenants. Users of that tenant are refused on their very next request; no data is deleted. Note that the automatic status sweep (see below) puts a Suspended tenant back to Active within about a minute if its license is still active and unexpired, so a manual Suspend only sticks once the license has been removed or has expired; to block access while the license is still valid, use Cancel tenant.
- Reactivate — shown for Suspended or Cancelled tenants.
- Cancel tenant — shown unless already cancelled. Asks for confirmation; data is kept, but nobody can sign in until you reactivate it.

> Changing a tenant's status has an immediate effect on a paying customer, so only administrators of the operator tenant can do it; other operator staff see the buttons but the change is refused.

> Support for a tenant that locked itself out with Require multi-factor authentication for technicians: right-click the tenant and choose Turn off MFA requirement. After the confirmation Turn off the MFA requirement?, its technicians can sign in without multi-factor authentication again; the notice reads MFA requirement turned off for … or … did not require MFA. This entry only ever turns the option off, and the change is recorded in the tenant's own Audit log. Its administrator can turn it back on in Settings › Users › Security.

> Status also follows the license automatically: roughly every minute the app suspends Active tenants whose license was removed or has expired, and reactivates Suspended ones whose license is active and unexpired. A Cancel tenant is never undone automatically, only by Reactivate.

Suggestions tab



Suggestions received from this app and other SL Apps products

Every idea submitted through the Suggestions button in the app footer (any logged-in user can send one by choosing a Type — Change suggestion or New feature — filling in a Subject and Description and clicking Send suggestion) lands here, together with suggestions forwarded by other SL Apps products such as ConsultisMed. Each new suggestion also triggers a notification email to the operator's notification address, so nothing is lost.

Columns: Date (date and time), Application (the product code it came from), Type (Change or New feature), Subject, Sent by (name and email), Organization (the sender's tenant or company, with the environment underneath) and Status. Filter with Search (subject, sender, organization), the Application drop-down (All applications by default) and the Status drop-down (All statuses by default).

- Show / Hide — expands the full description inline under the row.

- Mark reviewed — available while the suggestion is New.
- Mark done and Reject — available while it is New or Reviewed.
- Reopen — available on Done or Rejected suggestions; puts them back to Reviewed.

> **i** Statuses are New, Reviewed, Done and Rejected. Use Mark reviewed as your 'seen it' acknowledgement and keep the list filtered on New to triage only what has not been looked at yet.

Billing tab

Billing lists every invoice of the licenses billed by invoice: client, product, plan, period, technicians, amount without and with VAT, due date, status (Pending, Paid, Cancelled; an overdue one is highlighted) and who registered the payment. The chips at the top show the pending total, how many invoices are overdue and what was paid this month; filter by status, month or client, or tick Overdue only. Invoices are issued automatically on the 1st of each month; Issue invoices now runs the same check at once (it never issues an invoice twice).

1. Open Administration › Software & Licenses and pick the Billing tab.
2. Find the invoice (filter by client or tick Overdue only).
3. Click Mark as paid, add a note if you like (for example the transfer reference) and confirm.
4. The tenant is unblocked within a minute if this invoice was what blocked it.

> **i** Set back to pending undoes a payment registered by mistake (the tenant is blocked again if the invoice is overdue). Cancel removes an invoice from billing; for the current month a new one is issued within the hour with the current number of technicians, so to waive a month mark it as paid with a note instead. Only administrators see these buttons.

When a dedicated database move succeeds, a separate prorated invoice, Dedicated database (prorated dd/mm-dd/mm), is created for that license from the move day to the end of the period already invoiced (monthly or yearly), so you can issue the invoice; the regular invoices after it include the option, shown under the amount as + dedicated DB and the value, but only while the data really is on the dedicated database. Turning the option off creates no credit.

Changelog tab

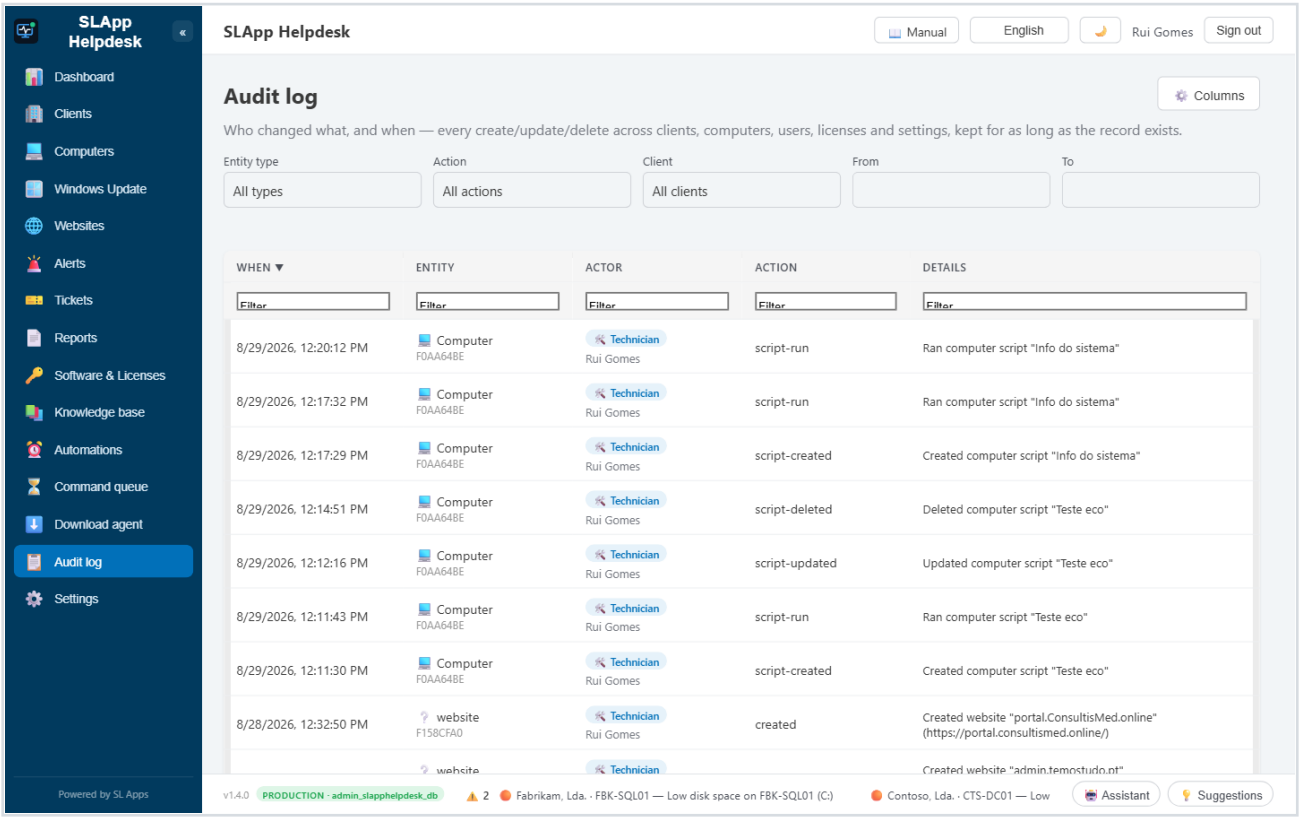
The Changelog tab lists what shipped in each product (new features and fixed bugs), typed here or sent by the product's own release process; it feeds the public What's new page. Each entry has a Stage: Roadmap means the version is published to DEV but not yet in production — the entry appears under the Roadmap tab of the public page, without a date, and never changes the product's current version — while Released means it is in production and appears under What's new (rows still in Roadmap carry a Roadmap badge). When a version reaches production its Roadmap entries with that version or lower move to Released automatically, dated that day: SLApp Helpdesk does this itself when its production backend starts, and other products through POST /api/changelog/external/release. Changing an entry from Roadmap to Released by hand dates it today if it had no date. Since version 1.14.81 the title and the description of each entry are shown in the language of the app, or in English when that translation is missing.

38. Audit Log (administrators, or by explicit permission)

The Audit log answers the question "who changed what, and when". Every time a technician creates, updates, deletes or deactivates a client, contact, staff user, license, alert rule, automation, backup job or backup destination, custom report or company-wide setting, a line is written here automatically.

Entries are also written by the system itself, for example when a scheduled monthly report is emailed to a client. The log is append-only: nothing in the app lets anyone edit or remove an entry, and the name of the person who acted is stored as a snapshot, so the line stays readable even if that user account is later deleted.

The module is only visible to administrators. It appears in the left sidebar as Administration › Audit log; technicians without admin rights do not see the entry at all, and it is never exposed to the client portal. Since version 1.14.15 the server also refuses the log to anyone who is not an administrator.



The Audit log with its filter bar and the event table.

What gets recorded

Each entry belongs to one entity type. The types you will find in the log are:

- Client: created, updated and deactivated (deactivating is what happens when a client is removed from the Clients module).
- Contact: created, updated and deleted, including bulk removals from a client's contact list. When the primary-contact flag changes, an extra line records the old and new value. A change of Can use "View as" in the portal is recorded as updated with the field portal_can_view_as, and a change of Can see the billing in the portal or Can buy credit packages in the portal with the field can_view_billing or can_buy_credits, with the old and the new value. A contact with that permission who views the portal as someone else adds view_as_started and view_as_ended, view_as_action for each change made meanwhile (with the method and the address, never the content) and view_as for what was only looked at (at most once every 10 minutes per contact viewed).
- Staff user: approved, rejected, set back to pending, deactivated, reactivated, deleted, and admin rights granted or revoked. Also tunnels_granted and tunnels_revoked (the RDP/HTTP tunnel permission, with the old and the new value), tunnel_denied (a tunnel refused to a technician without that permission) and mfa_rejected (a request refused because the sign-in did not use multi-factor authentication, at most once every 15 minutes per person).
- License: created, updated and deactivated.

- Alert rule: created, updated and deleted (the rule type is included in the details).
- Automation: created, updated, enabled, disabled and deleted.
- Report definition and Report subscription: custom report created, updated, deleted or generated on demand; monthly report subscriptions set up (subscribed), paused, resumed and removed (unsubscribed), plus the system entry sent each time a scheduled report goes out.
- Billing: exported, each time an administrator exports a month from Reports › Hours & billing (the details give the month, the format and the number of rows).
- Company: company information and global monitoring intervals updated. Also the Require multi-factor authentication for technicians option (field require_mfa), whether changed by an administrator or turned off by SLApp support.
- Email settings: email settings updated or cleared.
- Connector: Claude API key and payment settings updated or removed. This includes the Eupago webhook key (only that it was changed or removed, never the key), the Assistant actions option of the AI connector, and rejected, a System entry when an Eupago notification arrives without a valid signature while a webhook key is set (at most one every 10 minutes).
- remote_settings: the Remote control settings, including If nobody answers in time (field consent_timeout_action, with the old and the new value), plus one System entry per organisation when version 1.14.10 switched it from Authorize automatically to Refuse automatically.
- Audit log: retention_changed (the retention period, old and new value), exported (each CSV export, with its date range) and purged (a System entry of the daily clean-up, with how many entries were deleted).
- rate_limited, on the technician, portal contact or computer concerned, when the application had to slow down its requests again and again (Rate limit exceeded repeatedly …).

The Backups module adds two more entity types. Backup job: created, updated, deleted, enabled and disabled (what Resume and Pause do, and also what happens to the jobs of a computer that is removed in Computers), run_requested (Run now), run_cancelled (Cancel run), run_deleted (Remove, in the Runs tab, which erases one run from the history), archive_password_revealed (Reveal on a source, or Password of this copy in the run detail, showing the current archive password) and archive_password_history_revealed (Previous passwords, or Password of this copy, showing the previous ones). Backup destination: created, updated and deleted. The details name the job or destination and never contain a password; when an archive password is set, replaced or removed, the created or updated entry of the job says so, for which source, without the value, and the entry of a destination says when its free space starts or stops being read from a computer.

Nine more entity types come from the Domains, Microsoft 365, Vulnerabilities, Lifecycle, Agenda, Network configs, Active Directory, Third-party backups and Printers modules. Where a module raises client alerts, an alert an administrator resolves by hand is recorded too, with the action alert_resolved, on the item the alert was about. No password, secret or API key is ever written to the details.

- Domain: created, updated, deleted (Remove) and restored (a removed domain added again, which keeps its history), plus alert_resolved.
- Microsoft 365: connected, reconnected, connect_failed and disconnected (the Connect Microsoft 365, Reconnect and Disconnect buttons); secret_rotated, a System entry each time the app's secret renews itself; the user actions password_reset, sign_in_blocked, sign_in_unblocked, sessions_revoked, user_created, licence_assigned, licence_removed, out_of_office_set and out_of_office_cleared; plus alert_resolved. A reset password never appears in the log.
- Vulnerability: bulk_update, each time an administrator uses Update on all N computers; the details name the application, the client when there is one, how many computers it was queued on and how

many of them were offline.

- Lifecycle: updated when an administrator changes the settings in Settings › General › Lifecycle or sets a manufacturer's warranty API keys, removed when those keys are removed; the keys themselves are never written. A change of a computer's lifecycle dates (purchase date, warranty end, purchase price) is written on its Computer instead.
- On-site visit: created (scheduled), updated, rescheduled, confirmed (Mark as confirmed), client_confirmed (the client in the Client Portal, with the contact as the author), started, finished, cancelled, completion_edited (an administrator correcting the time, travel, kilometres or signature of a finished visit), confirmation_requested (the confirmation request to the client, by Notify client or sent by itself when the visit was scheduled or rescheduled, which the details then say), client_notified (Notify client, the visit's details only) and report_emailed (Email the report to the client). The private link of the Agenda's My calendar (iCal) is recorded on the technician's Staff user instead: calendar_feed_created (Create my link or New link) and calendar_feed_revoked (Switch off).
- Network config backup: created, updated, deleted (the device and all its stored versions), host_key_trusted (Trust the new key), expect_change and expect_change_cleared (Expect a change…), viewed_version and downloaded_version (only for versions that contain secrets, comparisons included), plus alert_resolved.
- Active Directory: dc_marked, dc_unmarked and dc_auto, when a computer is set to always or never be monitored as a domain controller, or back to automatic detection; plus alert_resolved.
- Third-party backup: frequency_changed (the expected frequency of a job), job_forgotten (Forget job), custom_event_created, custom_event_updated and custom_event_deleted, plus alert_resolved.
- SNMP device: created, when a technician adds an SNMP device on a computer's Devices tab (the details name the device, its type, its address and the computer that polls it) or when printers are added from a printer scan (Added printer "…" (…) polled by … - found by a printer scan); deleted, when a device is removed. The SNMP community is never written.

The Computer entity type collects what is done on one computer, for example its lifecycle dates edited in Infrastructure › Lifecycle (updated), a vulnerability or hardware lifecycle alert resolved by hand (alert_resolved), a chat with its user started or ended (chat_started, chat_closed), its own scripts created, changed, run or deleted, the recording of a remote session started, stopped, watched or deleted, the Agent activity window opened on it, and co-management changes. It also holds rdp_tunnel_opened and http_tunnel_opened (an RDP or web tunnel opened through the agent) and ai_action_approved / ai_action_rejected (an AI Assistant action approved or rejected, described as AI assistant: <action> on <computer>, with the command).

Staff changes (see the Staff changes chapter) add the entity type Staff change request: created, when a contact sends a request from the Portal or a technician records one (the details give the kind, the person, the start date or the last working day and the ticket number); action_done, each time a step of a ticket's Staff change tab succeeds, saying what was done, for example Created the Microsoft 365 user ana@example.com with Microsoft 365 Business Standard (ticket #128); marked_done (Mark as done); and cancelled (Cancel request, by a technician or, from the Portal, by the client), with the reason when one was given. Each step also writes its usual entries: Microsoft 365 for the user actions, Contact for the contact created or deactivated. A change of a contact's Can request staff changes, by a technician or by the primary contact in the Portal, is recorded on the Contact as an update showing the old and the new value.

> **i** Ticket changes are not in the Audit log. Each ticket keeps its own activity history, which you can see in the ticket itself. Likewise, Session appears in the Entity type filter but no events are currently written for it, so selecting it returns an empty list.

Columns

- When: the date and time of the event, shown in your local time zone and formatted according to the language you use in the app. Newest events are at the top.
- Entity: the entity type (with its icon) and, underneath in small grey text, the first eight characters of the record's identifier, which helps you tell two records with the same name apart.
- Actor: who did it. A coloured badge shows the actor type, Technician, Portal client or System, and the actor's name (or email, when no name was stored) appears below it. System entries have no name.
- Action: what happened, by its name in the language of the app, for example Created, Updated, Deleted, Deactivated, Approved, Admin rights granted, Subscribed or Sent. Pointing at it shows its code (such as admin_granted), and an action that has no name yet shows the code itself.
- Details: a readable summary such as Updated client "Acme" or Deactivated license: Office 365 (Acme). When the entry records a single field change instead, you see the field name followed by the old value, an arrow and the new value. A dash means there is nothing more to show.

Filters

The filter bar sits above the table and applies immediately; there is no search button. Every change reloads the list from the start.

- Entity type: restrict the list to one type; All types shows everything.
- Action: a list of the actions that have actually been used in your company so far, by name and in alphabetical order, so it only contains values that will return results. All actions is the default.
- Client: show only the events that concern one client; All clients is the default. This matches the client record itself, its contacts, its computers, its custom reports and report subscriptions, and the entries of its domains, its Microsoft 365, its on-site visits, its network config backups, its Active Directory (its domain controllers and AD domains), its third-party backups (its computers' jobs and custom events), its staff change requests and the SNMP devices its computers poll, printers included (the entries of a device that has since been removed only appear with All clients). Nothing else is matched: company-wide events (staff users, company settings, email settings, connectors, billing exports, the lifecycle settings and warranty keys, the bulk updates of the Vulnerabilities module) are not about any particular client, and License, Alert rule and Automation events are not linked to the client either, so none of these appear while a client is selected. To review a client's licenses, alert rules or automations, leave Client on All clients and use the Entity type filter instead.
- From and To: a date range. From starts at midnight of the chosen day and To runs to the very end of its day, so both days are fully included. Leave either one empty for an open-ended range.
- Clear filters: appears at the top right only when at least one filter is set, and resets all five filters at once.

Reviewing recent changes

1. Click Administration › Audit log in the sidebar. The most recent 200 events are loaded, newest first.
2. Narrow the list with the Entity type, Action, Client, From and To filters as needed.
3. Read the Actor and Details columns to see who did what. Compare the short identifier under the entity name when two records share a name.
4. If the list ends with a Load more button, click it to fetch the next 200 events with the same filters; keep clicking until the button disappears, which means you have reached the end.
5. Click Clear filters to go back to the full list.

Finding out who changed a specific client

1. Open the Audit log and choose the client in the Client filter.
2. Optionally set Entity type to Client to hide contact and report events, or to Contact to focus on the contact list.
3. Set From and To to the period you are investigating.
4. Read the Actor column: the badge tells you whether a technician or the system made the change, and the name below it identifies the person.

> **i** If No audit events match these filters is shown, check the date range first: the From and To fields are interpreted in your local time zone, so an event recorded late in the evening may fall on the next day for a colleague in another time zone.

> **i** Writing the audit entry never blocks the action it describes. If the database is briefly unavailable, the change still goes through and the entry is simply missing, so treat an absence in the log as strong evidence rather than absolute proof that something did not happen.

Exporting the log (CSV)

Export CSV, in the top bar before the columns button, opens Export the audit log (since version 1.14.15). From and To are filled in from the module's date filters or, when those are empty, with the first day of the current month and today. When a type, action or client filter is set, Apply the type, action and client filters limits the file to the same entries. Export downloads audit-log-<from>-to-<to>.csv: one CSV file (UTF-8, ';' separated, UTC times), oldest entry first, that opens directly in Excel; the notice reads Exported …. The export itself is recorded in the log (exported). A range with more than 100 000 entries is refused rather than cut short, with Could not export the audit log: …; choose a shorter range.

How long entries are kept

Retention opens Audit log retention. Choose Keep entries for N days (between 90 and 3650; Default: 365 days.) or Keep entries forever (while the subscription lasts); the window also shows the Oldest entry. Save confirms Audit log retention saved. Entries older than the period are deleted automatically once a day, and each clean-up leaves a System entry (purged). The line under the module's introduction shows the policy in force, for example Entries are kept for 365 days. (default) or Entries are kept forever.

> **i** When a shorter period would remove existing entries, Shorten the retention? asks first and says how many entries older than the new period the next daily clean-up will delete. This cannot be undone, so export them first if you need a copy. Every change of the retention is recorded (retention_changed).

39. Transactions (administrators, or by explicit permission)

Transactions lists, in one place, every credit package of every client of your organisation: the packages requested in the client portal, paid by invoice or online, and the packages registered by the team in a contract's Credits window, plus the transactions created by the automatic billing of contracts. It is the page to check what was sold, what has been paid and what is still awaiting payment, without opening the contracts one by one. Amounts include VAT.

The page appears in the sidebar as Administration › Transactions. It is only shown to administrators, and only when your plan includes the Contracts module (Complete plan), like the other credit screens; the server also refuses the list to anyone who is not an administrator, because it shows money.

What is listed

There is one row per transaction, the most recent first:

- A request made in the portal and paid by invoice (origin Portal, method Invoice): Awaiting payment until an administrator marks it as paid, then Paid, or Cancelled if it was cancelled.
- An order paid online in the portal (origin Portal, method Multibanco, MB WAY or Card): the state of its payment, that is Pending, Paid, Failed or Refunded; a pending order whose time to pay has run out shows Expired.
- A package registered by the team (origin Team, method Registered by the team): Paid, or Voided if the package was voided afterwards.
- A portal request that was paid and whose package was later voided shows Voided, so the list never counts it as paid money.
- The package that a portal request created is not listed again: the row of the request already stands for it.

Licence subscriptions (service operator only)

In the service operator's own organisation, the one that sells the SLApp Helpdesk licences, Transactions also lists the licence subscription invoices of every customer, next to the credit packages, and the introduction at the top says so. Every other organisation sees only its credit packages, exactly as before.

- Date is the day the invoice was issued and Client the customer that holds the licence (with the name of its organisation under it when it differs). Contract shows the product, the plan and the period, for example SLApp Helpdesk · Complete · Nov 2026, or the first and last day followed by (yearly) for a yearly licence; the prorated charge of a dedicated database shows its days. Package and Credits show —.
- Status follows the licence billing: Awaiting payment while the invoice is pending, Overdue once its due date has passed unpaid, Paid or Cancelled. Paid on and Marked / cancelled by come from the billing as well.
- Mark as paid and Cancel, on an invoice awaiting payment or overdue, are the same actions as in the Billing tab of Administration › Software & Licenses, with the same rules and the same effect on the customer's licence; each opens its own window to confirm.

Contract transactions

Each contract with Automatic billing ticked (see Automatic billing of a contract in the Clients chapter) gets one transaction at the start of each of its periods. These rows are listed in every organisation, with origin Automatic and method Invoice, and Contract shows the contract's name followed by the period: a whole month, for example December 2026, or else its first and last day.

- Amount (VAT incl.) is the total: the contract's value, plus the previous month's extras on a monthly contract with an hours bank, plus VAT at the rate in force when the transaction was created. Changing the rate in Settings never changes the transactions already created.
- Status is Awaiting payment when the transaction is created, then Paid or Cancelled; a contract transaction is never Overdue.
- Details, on the row or in its right-click menu, opens the transaction's window: the Period (with extras only (contract ended) for the extras of a contract's last month), the status and the invoice number, the lines with their Description, Quantity, Unit price and Amount (the contract itself, then any Extra hours, Travel billed separately and Kilometres, with their month), then Amount excl. VAT, VAT (with its rate) and Total. While the transaction is awaiting payment the window offers Mark as paid and Cancel, and in every state Attach invoice (PDF) or View invoice.

- Mark as paid opens Mark the transaction as paid, with Invoice no. (optional) and Invoice (PDF, optional), to attach the invoice in the same step. Cancel opens Cancel the transaction?, with an optional internal Reason: the client no longer sees it and that period is never billed again; a paid transaction cannot be cancelled. If someone else paid or cancelled it in the meantime, a message says so and the list is refreshed.
- The same actions are on the row and in its right-click menu (Open is not offered for these rows). Each one shows a message and is recorded in the Audit log.

Columns

Date (when the request was made or the package registered), Client, Contract, Package, Credits, Amount (VAT incl.), Method, Status, Invoice no., Paid on, Marked / cancelled by (who marked the request as paid, cancelled it or voided the package) and Origin are shown by default. Reason (why a request was cancelled or a package voided) and Requested / registered by can be added with Columns, which also changes their order; the choice is kept for you. Payment references, Multibanco entities and MB WAY phone numbers are never shown here.

Two more columns are shown by default: Invoice (PDF), with the name of the invoice file attached to the row (click it to open the file), and Type, which says Credit package, Contract or, for the service operator only, Subscription. Amount excl. VAT, VAT rate and VAT can be added with Columns; they are filled in on contract transactions and subscriptions.

Each row starts with an icon for the type of transaction: shopping cart a credit package requested in the client portal, ticket a credit package registered by the team, document a contract transaction created by the automatic billing and, for the service operator only, key a licence subscription. Hover over it to see the type and the origin.

Filters and totals

- Period: This month, Last 3 months (the default), Last 12 months or Custom, which shows From and To (both days included). It filters on the date the request was made or the package registered.
- Status: All statuses, or one of Awaiting payment, Pending, Paid, Voided, Cancelled, Expired, Failed and Refunded.
- Client: All clients, or one client.
- Method / origin: All, Team, Portal (all), Portal - Invoice, Portal - paid online, Portal - Multibanco, Portal - MB WAY, Portal - Card or Automatic (contracts).
- Search: a text of up to 100 characters looked for in the client, contract and package names, the invoice number, the names of the people involved and the reasons.
- Refresh reloads the list, and Clear filters, shown while a filter differs from the default, puts them all back.

Under the filters, the totals always cover every row that matches the filters: the number of transactions, the Total of their amounts, how much is Paid, and Awaiting payment, with how many and how much (requests by invoice awaiting payment and online payments still pending). At most 2000 rows are shown, the most recent; when there are more, a notice asks you to narrow the period or the filters, and the totals still count all of them.

The Type filter offers All types, Credit packages, Contracts and, for the service operator, Subscriptions, and when more than one type has rows a line under the totals splits the total between them. The service operator also gets the status Overdue, for subscription invoices past their due date, and its totals add Overdue, with how many and how much, when there is any. Export CSV includes the contract and subscription rows like the others, and always has the Amount excl. VAT, VAT rate, VAT and total of every row (empty where a row has none).

Actions

- Mark as paid and Cancel, on a request by invoice that is awaiting payment, open the same Mark the request as paid and Cancel request windows as the contract's Credits window and do exactly the same (see Credit contracts in the Clients chapter): marking as paid adds the package and emails the receipt to the client; cancelling requires a reason and adds no credits. The list is refreshed afterwards.
- Open, on a row whose contract is still active, opens that contract's Credits window without leaving the page.
- The right-click menu of a row offers the same actions, plus copying the invoice number or the client's name.
- Export CSV downloads the rows currently shown, with the visible columns in their order (separator ;, UTF-8, decimal comma in Portuguese and French), like the billing export.

Invoice (PDF)

Any row, a portal request, a package registered by the team or a subscription invoice, can have its invoice attached as a PDF file, so that the client downloads it whenever it needs it. The file is kept in the database, never on the server's disk. Only PDF files are accepted, up to 10 MB.

- Attach invoice (PDF), on a row without one, picks the file on your computer; while it is being sent the row shows Uploading….
- View invoice opens the attached PDF. The restart button replaces it with another file, and the bin button asks for confirmation and removes it, after which the client can no longer download it. The right-click menu of the row offers the same actions.
- The window to mark a request as paid, here and in the contract's Credits window, and the window to mark a subscription invoice as paid have an Invoice (PDF, optional) field, to attach the invoice in the same step.
- Every attach, replacement and removal shows a message and is recorded in the audit log.

Where the client finds it: for a credit package, the contacts with Can see the billing in the portal or Can buy credit packages in the portal (see Contacts in the Clients chapter) download it in the portal, under Purchases in the Credits card (Invoice (PDF)), in My company for the primary contact and in Billing for the others. The first time an invoice is attached to a portal request that is awaiting payment or paid, the client is also emailed, to the contact who made the request and the active contacts with Can see the billing in the portal (or else the client's email), with the PDF attached; the other emails of the request go to the contact who made it, or else to the contacts with Can see the billing in the portal; a PDF over 2.5 MB is not attached and the email points to the portal instead. Replacing the file sends no new email. For a subscription, the customer's administrators download it in Settings › License, in the Invoice column of the list of invoices.

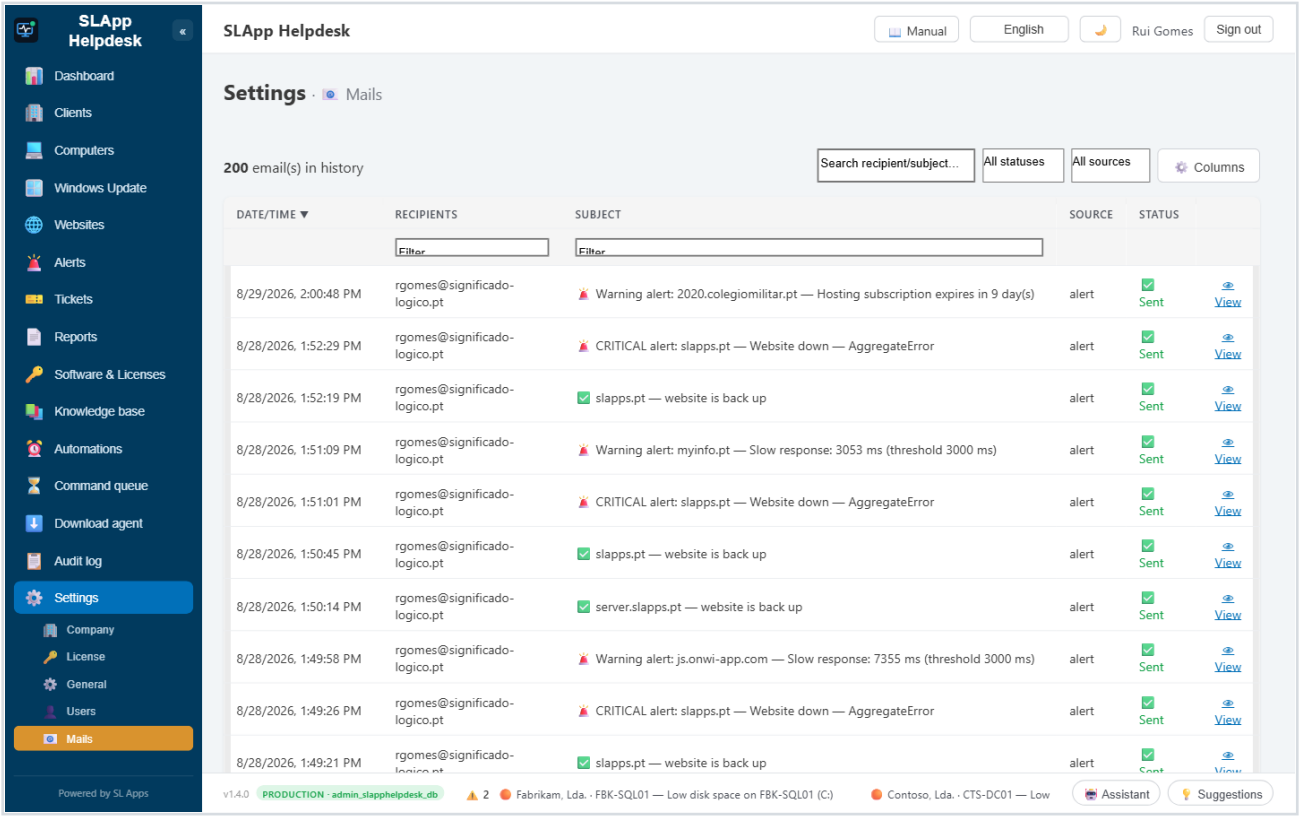
For a contract transaction, the client's contacts with Can see the billing in the portal, primary or not, download the PDF in the portal, in the Invoices card (My company for the primary contact, Billing for the others). The first time a PDF is attached to a contract transaction that is not cancelled, they are also emailed, in the client's language, with the contract and the period, as for a credit package; when the client has no active contact with that permission, the email goes to the client's own email address.

> **i** To see only what still has to be invoiced or collected, choose Status Awaiting payment and the period Last 12 months: the Awaiting payment total at the top is the amount still to receive.

40. Settings (administrators only)

Settings is where you shape the application for your company: company details and sign-in rules, the value lists that feed dropdowns, SLAs, patch windows, integrations, user access and the email log. The Settings entry is the last one in the sidebar and only appears for administrators; when it is active, a submenu shows five groups: Company, License, General, Users and Mails. The script library is not one of them: it is the Scripts tab of Automation & updates › Automations (see Scripts below).

Company



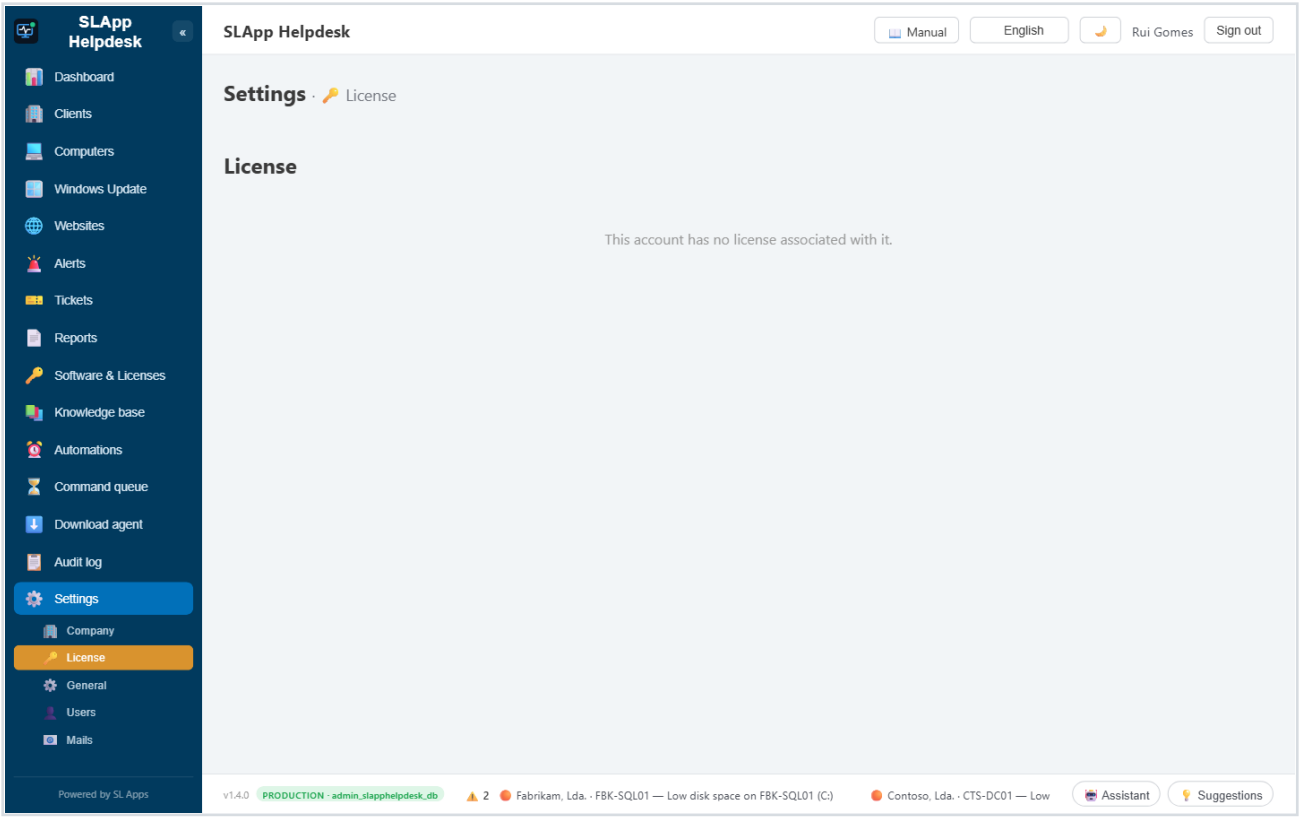
Company details, logo, icon and authorized sign-in domains

This group holds the details of the company using the application, used for example in the footer of emails and documents. Fill in Company name, Tax ID, Phone, Email, Website, Address, Postal code and City. Below the fields you can upload a Logo (scaled to at most 1200 px, so it prints sharp on the cover of the client reports) and a square Icon (256 px). The Logo (or, without one, the Icon) and the Company name also replace "SLApp Helpdesk" at the top of the app for every user of the company; on the dark theme the logo sits on a light background so it stays readable.

Domains authorized to sign in controls who can log into the application: only emails from these domains are accepted. Type a domain and press Enter, a comma or the Add button to add it as a chip (shown as @domain); the × on a chip removes it. A chip with a padlock and no × is a locked domain: the company's own domain, set when the company was created (hovering it, and the line below the chips, shows Company domain set when the company was created; it cannot be removed). It always stays in the list and nobody in the company can remove it, administrators included; to change it, contact the service operator. Leaving the list empty means no restriction (not possible while the list has a locked domain). A new company's list starts filled in: with the domains of the client it was created from, or else with the domains of its contacts' emails (never a free-mail domain such as gmail.com, nor one another company already uses), and those starting domains are locked — except one that another company also authorizes. You cannot add a free email domain either, nor one that another company or another customer of the service operator already uses: the message appears above Save and nothing is saved.

A new person from an authorized domain does not get in straight away — they wait for approval in Settings › Users. Press Save to store everything; a Saved confirmation appears next to the button, while an error is shown above Save and as a notification.

License



Your license: plan, seats and payment

License shows your own subscription: Product, Plan, Status (Active, Suspended or Cancelled), Billing period (Monthly, Semiannual or Yearly, shown only when the license has one), Expiration date (or No expiration date; a license billed by invoice shows Paid until — the end of the last paid period — instead, and Demo period until while it is in its demo or its billing has not started yet), and two seat counters, Computers and Users, shown as used / limit. A counter without a limit shows how many are active followed by "no limit" (for example "20 active · no limit" — the plans have no limit on computers); a counter turns red once it reaches the limit. If nothing is linked yet you see "This account has no license associated with it." Seats are managed by the service operator on your behalf.

SLApp Helpdesk is sold in two plans, both charged per technician (every person with access to the console; customer portal users are free) with unlimited devices: Essential (79 € per technician a month paid yearly, 95 € paid monthly) and Complete (119 € / 139 €). VAT (23%) is added to what is charged; there is no monthly minimum, so a single technician pays only the plan price. A license from before the plans keeps the price agreed for it and every module; its Plan shows "Agreed price (no plan)".

- Essential — devices (Windows, macOS, Linux), inventory and vulnerabilities, alerts, remote access and RDP, scripts, maintenances and automations, Windows and software updates, tickets with SLAs, calendar, knowledge base, the customer portal and the mobile app.
- Complete — everything in Essential, plus backups (files and SQL Server), third-party backup monitoring, Active Directory, Exchange, Hyper-V, network devices (SNMP, Syslog, network configurations, printers), the credentials vault, contracts with the hours bank or with credits and billing export, and the AI assistant.

Under the details, administrators see the Plan card: pick Essential or Complete and set the Number of technicians (users with access to the console — never fewer than the active users; to lower it, first

deactivate the users who no longer need access in Settings › Users, where the Admin and Active columns are shown in colour) and, for a license billed by invoice, the Billing period (Monthly or Yearly — it applies from the next invoice; a year already invoiced runs to its end), then press Save changes. Under the choice, the card works out the amount as you change it — technicians × price per technician, monthly and yearly, net and with VAT. A new plan changes the modules at once and the page reloads; the new amount applies from the next invoice.

Next to them, the Plan comparison card lists what each plan includes, with a in the Essential and Complete columns: first what both plans have (devices and inventory, remote access, tickets, SLAs and time tracking, scripts and automation, updates, monitoring and alerts, the customer portal and the mobile app), then what only Complete has (backups, third-party backups monitoring, Active Directory, Exchange, SQL Server, storage, failover clusters, DHCP, IIS, Hyper-V, network, the credentials vault, contracts and billing, and the AI assistant). The column of your current plan is highlighted and marked "your plan".

> **i** A module that is not in your plan disappears from the menu, from the computer and client tabs and from reports, and any attempt to open it shows "This module is not included in your plan". Nothing is deleted: moving back to Complete brings every module back with its data.

Dedicated database

At the bottom of the Plan card, the Dedicated database (+50 € / month + VAT) option moves your company's data to a database of its own, for better isolation and performance; it costs 600 € a year on a yearly license. Tick it, check the new monthly and yearly totals under the choice and press Save changes. The move is done that night, between 01:00 and 05:00 (Lisbon time), and nothing is lost; until then the card says the move is done tonight, and once it is done it shows Active since and the date. Unticking it moves the data back to the shared database the same way.

When the move is done, a separate prorated invoice (Dedicated database (prorated …)) charges the option from that day to the end of the period already invoiced; from then on every invoice includes it (shown as incl. dedicated database). Turning it off gives no credit for the rest of the period.

> **i** The option can be changed once a month; the card tells you from which date you can change it again. A change made today can still be undone before tonight's move, and undoing does not count toward that limit. It cannot be changed while a subscription is active (cancel it first in Manage subscription) or a Multibanco / MB WAY payment is pending, and if a move fails the card says so — the support team is notified and handles it.

Free month and billing

A company that signed up on its own gets one free month on the Complete plan. Seven days before it ends, administrators see a card at the bottom left of the screen: "Your trial ends on …". Continue opens a window to pick the plan (Essential or Complete) and the billing period (monthly or yearly); it shows the price per technician, the total with and without VAT and an estimate of the first payment. Do not continue asks for confirmation and hides the card; the × hides it until the next day. If nobody confirms, the application is blocked when the free month ends — only Settings › License stays open, and Continue is still there to come back.

After Continue, the license is billed by invoice: the first invoice covers the rest of the month after the free month (prorated by days), then one is issued on the 1st of every month (once a year for a yearly license) with the plan price for the technicians at that moment. Settings › License lists the invoices — period, plan, technicians, amount without and with VAT, due date and status. An invoice has to be paid within the first 7 days; if it is still unpaid on the 8th, the application is blocked until the payment is

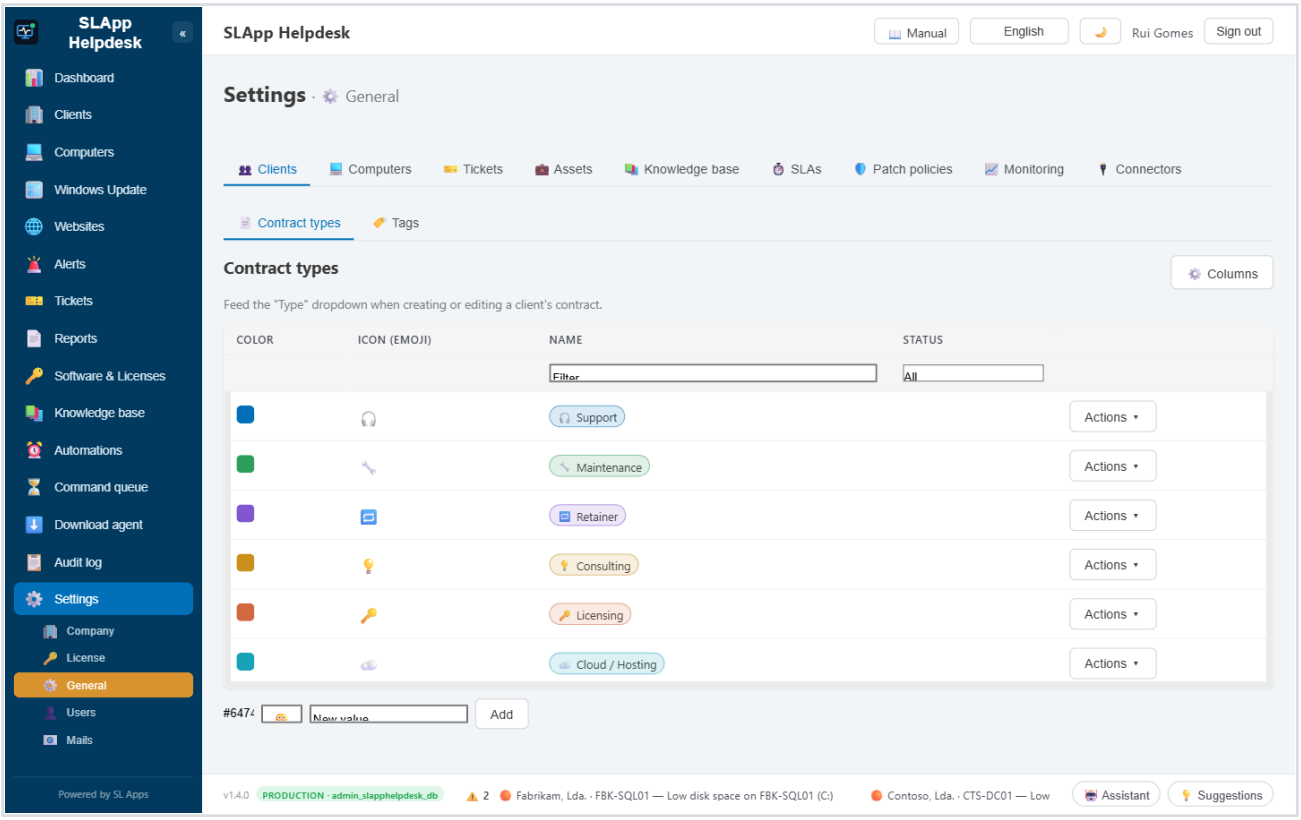
registered by the service operator, and Settings › License explains why.

When the service operator has attached the invoice as a PDF, the Invoice column of that list shows Invoice (PDF): click it to download the file. Only administrators see it.

> ℹ Payments are registered by the service operator (bank transfer or another agreed method): there is no card or Multibanco payment inside the application for now.

General

General gathers everything else behind internal tabs: Clients, Computers, Tickets, Assets, Lifecycle, Knowledge base, SLAs, Patch policies, Remote control, Alerts, Retention, Monitoring (only for the operator's own account), Backend instances (only for the administrators of the operator's own account) and Connectors.



Settings › General: one tab per area (Clients, Tickets, Assets, Knowledge base, SLAs, Patch policies, Monitoring, Connectors).

- Clients — Contract types: the values offered in the Type dropdown of a client's contract.
- Computers — two tabs: Tags, the colored tags your team puts on computers (a maintenance flag, a priority, anything; they are assigned from a computer's own page), and Restart prompt, described under Asking the user to restart below.
- Tickets — Ticket types, Ticket priorities, Ticket statuses, Impact levels, Product families, Products and Templates (Product families and Products are described under Product families and products below, Templates under Ticket templates). Each ticket type also has a Responsible group selector (No group by default; only active groups are offered, and a deactivated group still linked shows as Inactive group): tickets of that type are assigned to that group on creation. Each status also has a Behavior, set in its Edit dialog: Starting status for new tickets, Counts as resolved (sets the resolved date) and Locks editing and stops timers (closes the ticket).
- Assets — Asset custom fields (a name plus a type: Text, Number or Date) that appear in every asset form, including the client portal; plus Asset types, Asset statuses and Asset event types (used in each asset's history). An asset's type cannot be changed after creation, and deleting a field does not

erase values already stored.

- Lifecycle — when a computer is due for replacement and how early a warranty about to end raises an alert, used by Infrastructure › Lifecycle, its replacement plan and the client report: Replace workstations after (years) (1 to 15, 5 by default), Replace servers after (years) (1 to 20, 7 by default) and Warn before a warranty expires (days) (0 to 365, 60 by default; 0 turns the warranty alerts off). Under the fields, a note explains how the age of a computer is counted and an Alerts box sums up when lifecycle alerts open. Press Save (it briefly shows Saved). Administrators also reach these settings from the Settings button of the Lifecycle module. After a change of the warning days every computer's warranty alert is evaluated again shortly afterwards, and an alert that opens only because of the new value sends no email.
- Knowledge base — the category › section tree. Click + New category next to the title to create a category, or + Section next to a category to add a section to it; both open a small window that asks for the name. Use the up/down arrows to reorder, pencil (or a double-click on the name) to rename in the same window, bin to delete and people to set who sees an item (All clients or a chosen set; a section also inherits its category's audience). Items marked with a padlock were created by a client in the portal and cannot be changed by the team.
- SLAs and Patch policies — described below.
- Remote control — four cards. Remote-access authorization popup: how long the popup on the user's screen waits, and whether access is authorized or refused when nobody answers. Command timeout: how long the agent lets a command run and how long the server waits for its reply. Agent activity window: Set password / Replace password / Remove password for the Agent activity window of the agent — the same password for every computer; only a hash of it is kept, so it can be replaced or removed but never shown again. Files in conversations: the largest file (1 to 20 MB) and the accepted file extensions — by default images, PDF, text, Office documents and archives, but no programs or scripts; Restore the default list puts that list back, and Accept any type of file lifts the restriction (with a warning). Since version 1.14.10, If nobody answers in time is Refuse automatically for every organisation unless an administrator chose Authorize automatically again; a change is recorded in the Audit log with the old and the new value.
- Alerts — which modules generate alerts (one switch per module, the same as the Generate alerts box on each module's page); described under Alert generation below.
- Retention — how long your company keeps the history its computers and monitors report (performance charts, Windows events, website checks, backup runs…); described under Data retention below.
- Monitoring (operator only) — Local sampling (seconds), Sending to the server (seconds) and Heartbeat interval (seconds) for every agent; the heartbeat decides how quickly Online/Offline gets updated. Press Save (it briefly shows Saved). Agents pick the new values up by themselves; no reinstall or restart is needed. The values can only be raised from the ones the fleet runs at today: at least 10 seconds for Local sampling, 60 seconds for Sending to the server (and never less than the sampling) and 30 seconds for the Heartbeat interval. Each field shows its minimum under it (Minimum N seconds.), and a lower value is refused with that message.
- Backend instances (operator administrators only) — the server processes of the application and whether their background jobs are running; described under Backend instances below.
- Connectors — Email, Payments, AI (Claude), Plesk, Warranty and TeamViewer; all but TeamViewer are described below.

All value lists share the same editor: a table with the columns Color, Icon, Name and Status (the list of product families has no color and no icon; ticket types add a Responsible group column and ticket statuses a Behavior column) and, at the end of each row, an Actions menu. Click a header to sort, use the row under the headers to filter (type in the Name cell; in the Status cell pick All, Active or Inactive),

and use Columns to choose which columns are shown. Actions offers Edit, Deactivate / Reactivate and Delete (right-clicking a row opens the same menu without Delete); Edit (or a double-click on the row) opens the Edit value window, where you change the name, color and icon. An inactive value stays in the list, dimmed, but is no longer offered for new records. To add a value, click + New value next to the list's title: the New value window asks for the same fields as Edit (Name, Color and Icon), Create adds the value and Cancel closes the window without changes; if the value cannot be saved, the reason stays inside the window. Every icon field in the application, here and in the other settings editors, has a smiley button next to it (tooltip Choose an emoji) that opens a list of suggested emojis grouped by theme, with a search box and a Clear option; you can still type any emoji in the field yourself (Win + . on Windows). Deleting a value never changes records that already use it.

> **i** Impact levels drive two things at once: the Impact dropdown on tickets and the rows of every SLA grid. Create the levels first, then define the SLAs.

Intervention types and credit packages

With the Contracts module (Complete plan), the Clients tab of General has two more tabs next to Contract types: Intervention types and Credit packages. They feed the credit contracts (see Credit contracts in the Clients chapter) and start empty, so create them before the first credit contract. Every technician sees both lists; only administrators add, edit, deactivate and reactivate.

- Intervention types (+ New type): the Name, the Credits per hour (from 0 to 10000, up to two decimals; for example On site 1.5, Remote 1, Urgent 2), a colour and an icon, and Use for travel, which suggests the type for the travel of on-site visits. Changing the rate only affects new time entries: each entry keeps the rate its type had when the type was chosen.
- Credit packages (+ New package): the Name, the Credits, the Price (the total the client pays, VAT included) and the Validity in months, from 1 to 120, or No expiry. The validity counts from the date of each purchase.

> **i** Removing a type or a package only deactivates it: the time entries and purchases that use it keep it, the row stays in the list, dimmed, and Reactivate in its Actions menu (or on right-click) brings it back. Two active types, or two active packages, cannot have the same name.

Asking the user to restart

Settings › General › Computers › Restart prompt has one card, Ask the user to restart. With it on, when a Windows workstation of your company is waiting for a restart (updates or changes that only finish after restarting), the agent shows the signed-in user a window with Restart now and Postpone, and shows it again every N hours (Ask every) until the computer has been restarted. It needs agent 1.2.41 or later. Tick Ask the user to restart when the computer is waiting for a restart (off by default), set Ask every (hours) — a whole number from 1 to 24, 2 by default; the field is greyed out while the box is not ticked — and press Save: Restart prompt settings saved confirms it. A number outside 1 to 24 is refused with The number of hours must be a whole number between 1 and 24. Only administrators can change this setting: a technician who opens the card sees that same sentence and no Save button. Each change is recorded in the Audit log with the old and the new value, and the agents pick it up at their next heartbeat (every 30 seconds by default), with no reinstall or restart.

The agent shows the window only when all of the following are true, and it checks them again at every heartbeat:

- The computer is waiting for a restart: Windows Update or a Windows component update asked for one — the same checks as the restart pending badge of the computer. Files that an installer replaces at the next start (PendingFileRenameOperations) count only until the computer has

restarted once: some programs, such as antivirus products and updaters, write them again at every start, and without this exception the window would come back forever. The badge of the computer is not affected by this exception.

- It is a Windows workstation. Servers are never asked (a restart there would cut off everyone who works on them), and neither are Linux and macOS computers.
- Someone is signed in and can see the desktop: not on a locked screen or in a disconnected session, not while a technician is connected to the computer through remote access or while the agent's own Computer restart notice is open, and never while another Windows user is also signed in (Switch user) — a restart would close that person's programs without warning them.
- At least the chosen number of hours has passed since the window was last shown or postponed on that computer, and at least 10 minutes since the user signed in or the agent started. The time of the last window is kept both on the computer and in SLApp Helpdesk, so it survives restarts and also applies when another person signs in.

What the user sees and what each button does is described under What the agent does on the computer in the Download Agent and installing on computers section. Nothing restarts without a click on Restart now, and after that click Windows restarts the computer 30 seconds later. Every window shown and every answer appears in the History card of the computer's Remote actions tab as Restart prompt shown, Restart postponed by the user or Restart accepted by the user, with agent in the Technician column and the Windows user (for example DOMAIN\user) in the Detail column; a Restart accepted by the user that Windows refused is shown in red.

> **i** On a computer managed by more than one organisation, the window appears as soon as one of them has the option on, every N hours with the smallest N among those that have it on, and each organisation that has it on sees the answers in its own History.

Product families and products

The Product families and Products tabs under Tickets feed the Product family and Product fields of the ticket forms, both in this app and in the Client Portal's New support request form. Until at least one active family exists, those fields are not shown. Only administrators can change them.

Product families is the same table as the other lists, without color and icon. Type the name under New value… (up to 100 characters) and press Add; a new family goes to the end of the list. The up and down arrows at the start of each row move it one place up or down, and the order of this list is the order in which the ticket forms offer the families. The arrows are disabled while a column filter or a sort is active, because the neighbours you see would not be the real ones: clear them first. The order always includes the inactive families, which stay listed, dimmed. Edit changes the name, and renaming a family renames its products with it. Deactivate hides the family and its products from new tickets while keeping them for later; Reactivate brings them back. Delete asks for confirmation (Delete product family), says how many products go with it and removes the family and all of its products for good.

Products lists the products under each family. The Product family drop-down above the table shows All or one family (an inactive family is marked Inactive). Each row has the arrows, the Product family, the Product, its Status and an Actions menu with Edit, Deactivate / Reactivate and Delete. To add a product, click + New product next to the title: the New product window asks for the family (only active families are offered) and the name (up to 100 characters), and Create adds it. It goes to the end of its family. The same product name can exist under different families, but not twice in the same one. Edit (or a double-click on the row) opens the same window as Edit product, where you rename the product; the family is shown but cannot be changed. The arrows move a product within its own family only. Delete removes a product at once, without asking. If there are no families yet, the tab tells you to add one in Product families first.

> ℹ Tickets keep the text they were created with: renaming, deactivating or deleting a family or a product never changes an existing ticket. In the ticket forms, Product only appears once you choose a family that has active products.

Ticket templates

The Templates tab under Tickets holds your company's ticket templates. A template is a ticket prepared in advance: a Subject, a Description, a Ticket type, a Priority and an Impact, and an ordered list of tasks that is copied into every ticket created from it. Its kind is chosen when it is created and cannot be changed afterwards. General templates are offered in From template when a technician creates a ticket (see Creating a ticket in the Tickets chapter). New employee and Employee leaving templates are the ones the staff change requests build their tickets from (see the Staff changes chapter): of each of these two kinds exactly one template is active, marked In use by requests, and every request of that kind uses it.

The table lists the templates with the columns Name, Kind, Ticket type (Default when the template leaves it to the usual default), Tasks (how many) and Status; Subject and Updated can be added with Columns. Inactive templates stay listed, dimmed. The row under the headers filters by name, kind and status, and a click on a header sorts. The Actions menu at the end of each row offers Edit, Duplicate (it asks for the copy's Name, the original's followed by (copy) to begin with, and its Kind, the original's unless you choose another, and creates an inactive copy with the same fields and tasks: as the kind cannot be changed on an edit, this is how a template created with the wrong kind becomes the right one, for example a General copy of a New employee template so it is offered in From template), Deactivate or Activate, and Delete, after a confirmation; tickets already created from a template keep their tasks. Right-clicking a row offers the same entries except Delete, plus Copy name.

- Activating a New employee or Employee leaving template first asks Activate the template? and says which template it replaces; from then on the requests of that kind use it, and the one used until then is deactivated.
- The template in use can be neither deactivated nor deleted: those entries are disabled, and hovering them says why. Activate another template of the same kind first.
- A General template is activated and deactivated freely; an inactive one is no longer offered in From template.

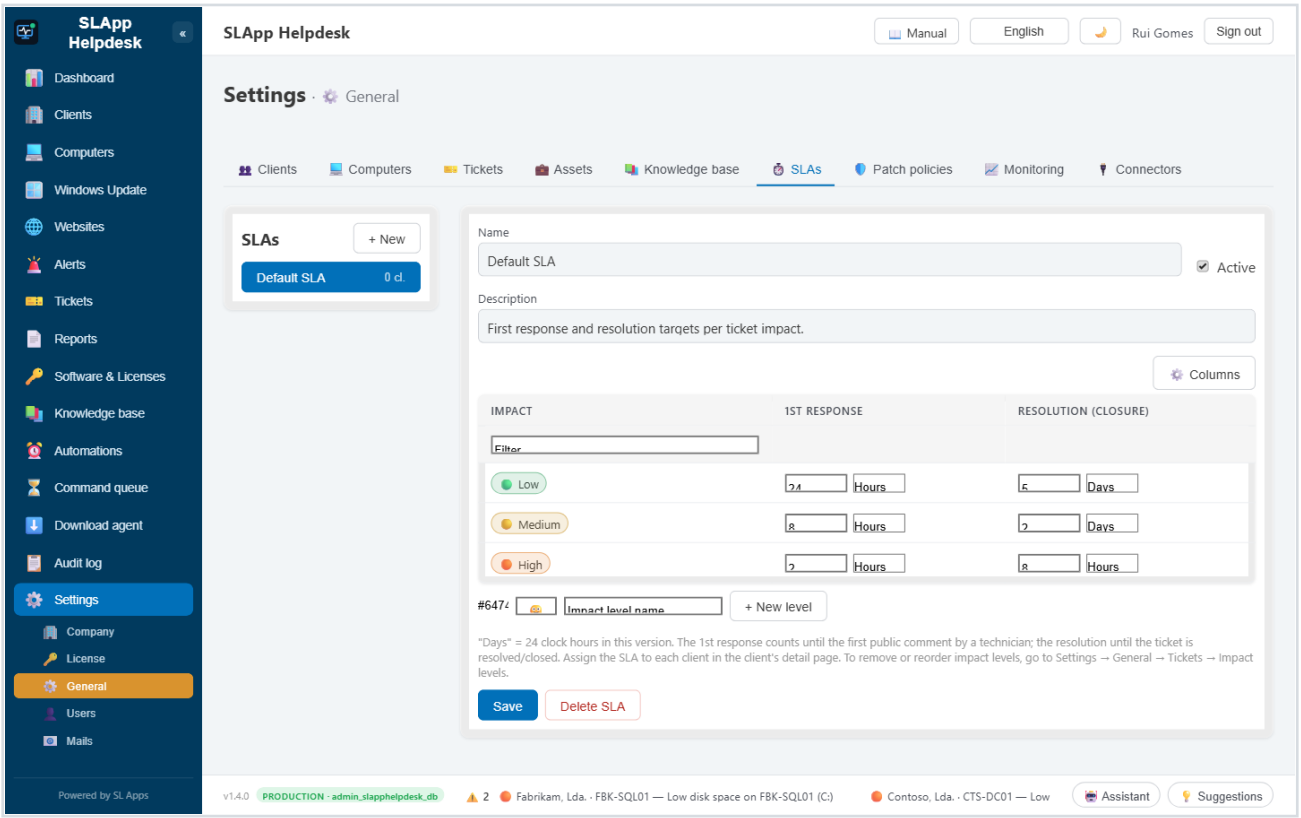
1. Click + New template, or Edit in a row's Actions menu.
2. Type the Name * and, for a new template, choose its Kind: General, New employee or Employee leaving.
3. Optionally choose a Ticket type, a Priority and an Impact among the active values; — Default — leaves the value a new ticket gets anyway.
4. Write the Subject * and, if useful, the Description. {{name}}, {{client}} and {{date}} in them are replaced when a ticket is created (see the tip below).
5. For a new template, choose whether it is Active. A General one starts ticked. A New employee or Employee leaving one starts ticked only when its kind has no template in use yet; ticked otherwise, it replaces the one in use when you save, as the line under the box says.
6. Under Tasks, click + Add task for each task (up to 100) and give it a title *, an optional note, a Group and, when it matches a step of a ticket's Staff change tab, an Action (for example Microsoft 365: create the user): the task is then ticked by itself when that step is run. The up and down arrows change the order, and × removes a task.
7. Click Create, or Save. The button stays disabled until the template has a name and a subject and every task has a title.

> **i** {{name}} is the person a staff change request is about (the new employee or the person leaving), or the contact of a ticket created with From template; {{client}} is the client's name; {{date}} is the start date or the last working day of a request, or the day a ticket is created with From template, written the way your company's email language writes dates. An Action only ticks a task in the ticket of a staff change request whose kind offers that step: the form warns when a General template, or a template of the other kind, carries one that would never be ticked.

Each company starts with one active New employee template and one active Employee leaving template, in its email language, with the checklists listed under The templates behind the requests in the Staff changes chapter. They are only created for a company that has no template of that kind, so they can be adapted freely: they are never put back.

SLAs and patch policies

The SLAs tab shows your SLAs on the left (each with its client count, e.g. "3 cl.", and an (inactive) suffix when switched off) and the editor on the right: Name, an Active checkbox, Description, and a grid with one row per impact level and two targets, 1st Response and Resolution (closure), each a number plus Hours or Days. + New level, under the grid, adds an impact level in the same New value window as the Impact levels list. Days means 24 clock hours. The first response counts until the first public comment by a technician; resolution until the ticket is resolved or closed. You assign the SLA to each client in the client's detail page, and the editor lists Clients with this SLA. Delete SLA (after confirmation) leaves those clients without an SLA, so their tickets stop being monitored.



The SLAs tab.

1. Open Settings › General › SLAs and press + New; the New SLA window suggests the name New SLA (change it if you like) and Create creates the SLA and selects it.
2. Give it a Name and, optionally, a Description; keep Active ticked.
3. For each impact row, set the 1st Response and Resolution (closure) values and units.
4. Press Save — the button briefly shows Saved.

5. Open each client's detail page and pick this SLA there.

Patch policies define maintenance windows during which approved Windows Updates install automatically (manual install or removal per computer is always possible regardless). The table shows Name, Target, Days (Every day when all seven are selected), Window, Status and Last run, with pause/tick to deactivate or activate, pencil to edit and bin to delete. + New policy asks for a Name, a Client (All clients by default) and, once a client is chosen, a Folder; the Days buttons (Mon–Sun) toggle individually; Window start and Window end may span midnight (e.g. 22:00–02:00); a Status dropdown sets the policy Active or Inactive. Save is enabled only with a Name and at least one day. Saving a policy with no client triggers a confirmation, because it will apply to every computer of every client.

Alert generation

Settings › General › Alerts chooses which modules generate alerts. Its card, Alert generation, opens with the line Choose which modules generate alerts. Alert rules (Alerts › Rules) keep their own enable/disable. and lists every module of your plan that raises alerts of its own in a table with the columns Module, Open alerts (how many of its alerts are open now), Generate alerts (a checkbox followed by On or Off) and Last change (when the switch was last changed, and by whom). Each checkbox is the same switch as the Generate alerts box at the top of the module's own page, so a change made in one place shows in the other: unticking asks Stop generating <module> alerts? and resolves the module's open alerts at once, while ticking turns it back on with no question. A note under the table sums up what each state does; the details, including which alerts come back without an e-mail, are under Turning alert generation on or off per module in the Alerts section.

> **i** Network devices (SNMP) covers every SNMP device except printers, which have their own row (Printers). The alert rules are not listed: each alert rule keeps its own Enabled / Paused switch in Alerts › Rules.

Data retention

Settings › General › Retention decides how long your company keeps the history that its computers and monitors report. Each category starts at the platform maximum — what SLApp Helpdesk keeps for every company — and you can only keep it for fewer days, never more. Every day the clean-up deletes, for your company only, whatever is older than the number of days chosen here. The table has three columns: Data (the category, with a short explanation), Keep for (the number of days, marked (maximum) when it is the maximum) and Allowed (the range accepted).

- Performance charts — CPU, memory, disk and network usage of each computer: 7 to 30 days (30 by default). Not below 7 days, because the Week chart and the free space of backup destinations read the latest samples of a computer, which must still be there after a computer was switched off for a week.
- Windows events — the critical, error and warning events of the Application and System logs: 1 to 7 days (7 by default).
- Agent errors — the errors the agent itself reported (the Agent errors channel of Events): 1 to 30 days (30 by default).
- Network log (Syslog) — the messages received from routers, firewalls and other network devices: 1 to 14 days (14 by default). The row is shown only when your plan includes the Network logs module.
- Website checks (uptime) — the result of every availability check of a website: 62 to 90 days (90 by default). Not below 62 days, because the monthly client report covers the whole previous calendar month, which on the last day of a month starts about 61 days back.
- Website visits — the page views counted by the tracking snippet: 30 to 400 days (400 by default). Not below 30 days, because the Visits tab shows the last 30 days.

- Backup run history — every run of a backup job and its items: 62 to 180 days (180 by default). Not below 62 days, because the backup calendar covers 62 days. Whatever the value, the latest real run of each job is never deleted — the missed-run check of a job that runs at long intervals (monthly, or up to once a year) and the job status are based on it — and neither is a run still in progress. The row is shown only when your plan includes the Backups module.

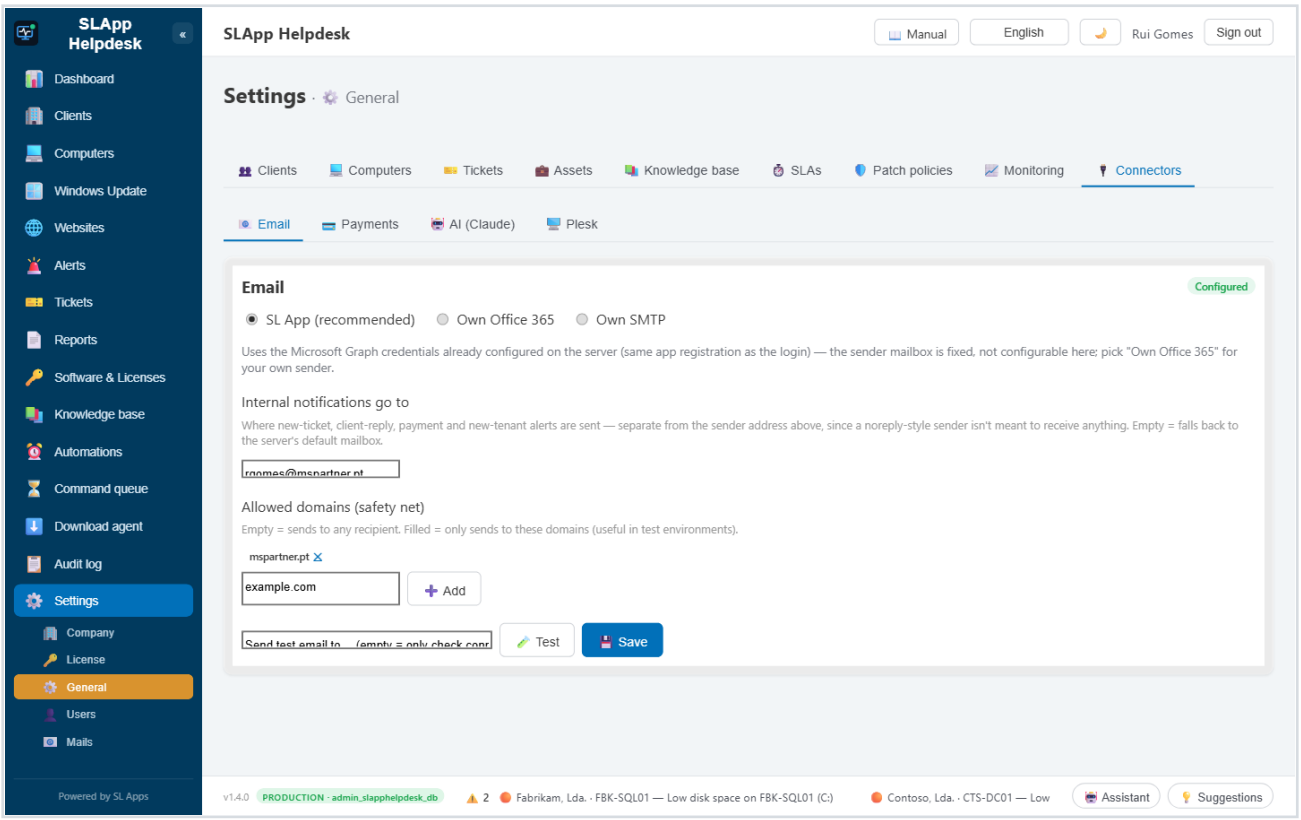
Change the numbers and press Save. Discard changes puts the saved values back, and Use the maximums fills in every maximum (press Save to apply it). A value outside the allowed range is refused with a message naming the category. Only administrators can change the values; other technicians see them read-only, with Only an administrator can change these values. Each change is recorded in the Audit log with the old and the new number of days. A row hidden because the plan does not include its module keeps any shorter value saved earlier, and it is still applied.

> **i** A shorter value takes effect at the next daily clean-up: the older history is then deleted for good and cannot be recovered. The audit log is not in this table: it has its own retention (365 days by default, from 90 days up to 10 years or forever), and the Audit log retention… button at the bottom of the card opens the same window as in Audit.

Connectors

Email decides how the application sends mail; a Configured / To configure badge in the card header shows where you stand. Choose SL App (recommended), which uses the server's fixed sender and needs no fields, Own Office 365 (Tenant ID (Directory), Client ID (Application), Client secret and Sender mailbox — the Entra app registration needs the Mail.Send application permission consented by an admin) or Own SMTP (SMTP server, Port, User, Password, Sender (From:) and a Secure connection checkbox: 465 is direct TLS, 587 does STARTTLS). Internal notifications go to sets where new-ticket, client-reply, payment and new-tenant alerts land; empty falls back to the server mailbox. Allowed domains (safety net) restricts sending to listed domains — handy in test environments; type a domain and press Enter or Add, and × on a chip removes it; empty sends to anyone. The box Send test email to… plus Test checks the connection only when empty, or really sends a message when an address is given; Save stores the card.

Plesk links one or more Plesk hosting panels over their REST API so the Websites module can import domains. Add a server with a name and the panel URL (https://panel.example.com:8443 — an explicit :443 is respected for panels that publish the API there), then either an API key (recommended; Plesk › Tools & Settings › API keys) or an administrator login, and tick Accept self-signed certificate if the panel has no valid certificate on its port. Test checks the connection and shows the panel’s hostname and version on the row. Credentials are stored encrypted and are never sent back to the browser; when editing, an empty key or password field keeps the stored one. Once at least one server is connected, the Websites module shows the Import from Plesk button.



The Connectors tab.

> **i** Stored passwords and secrets are never shown again; the field reads "set — fill in to replace". Leaving it empty on Save keeps the current value, so you can change the port or sender without re-entering the password.

Payments selects the online gateway used for client payment requests: None, Eupago, or SIBS, Stripe and PayPal (coming soon). Only Eupago has a working flow today; for the others the credentials are stored but a notice reminds you the charging flow is not implemented, and Test stays greyed out. For Eupago enter the API key (reference creation), decide whether Sandbox (test environment) stays on, and optionally the Management Client ID and Client secret, needed only to refund from the application. Test validates the management credentials (the reference key itself is only confirmed on the first real payment); Remove, shown whenever a provider other than None is selected, clears every stored credential after confirmation. Right after the Management credentials, Webhook key (Webhooks 2.0) is optional: paste the key from the Eupago backoffice and, in the backoffice's Webhooks 2.0, set the address shown under Webhook URL to set in the Eupago backoffice (Webhooks 2.0, method POST): (it ends with your organisation's own identifier). With a key, Notifications without a valid signature are rejected. and each refusal is recorded in the Audit log; without one, notifications are accepted unsigned. In both cases a payment is only marked paid after the application has checked it with Eupago. Tick Remove the webhook key on save to go back to unsigned notifications; if the saved key can no longer be read, a warning says so and every notification is refused until you type the key again and save.

With the Contracts module, the Payments card ends with Sell credit packages in the client portal, off by default and saved with the card. Ticked, the contacts of a client with a credit contract who have Can buy credit packages in the portal can get credit packages in the portal (see Credit contracts and Contacts in the Clients chapter), and How clients pay for credit packages decides how. Invoice (no online payment), the default, needs no payment provider: the client requests the package in the portal, the administrators get an email to issue the invoice, and the credits are added when an administrator marks the request as paid. Online payment (Multibanco, MB WAY, card) keeps the payment in the portal: by Multibanco or MB WAY through your Eupago account, or by card through your Stripe account, so the money goes to

your account. Each method is only offered while its provider is active and configured (with none, a warning says clients cannot buy online), and a package is only added once the payment is confirmed, so the webhooks must be set: for Eupago the Webhook URL above, for Stripe an endpoint at the address shown in its block that receives the events checkout.session.completed and checkout.session.expired. The card payment opens in a new tab of the client's browser. Changing the mode only affects new requests.

Below it, Send the package requests to takes one or more email addresses, separated by ; or commas (up to 10, each checked when you save). The email of each package a client requests in the portal to be paid by invoice (the purchase order the invoice is issued from) goes to those addresses instead of to the administrators; left empty, the administrators get it, as before. The field can only be edited while Sell credit packages in the client portal is ticked, and it is saved with the card.

With the Contracts module, a card titled Automatic contract billing comes next, with VAT rate (%), 23 by default: the VAT added to the transactions that the automatic billing of contracts creates (see Automatic billing of a contract in the Clients chapter). It takes a number from 0 to 100 with at most 2 decimals (a comma also works); anything else shows The VAT rate must be between 0 and 100, with at most 2 decimals. and nothing is saved, and an empty field means 23. It is saved with the Payments card and only applies to the transactions created from then on: each existing transaction keeps the rate it was created with.

AI (Claude) stores the API key that powers AI features such as Generate with AI in Scripts. Paste the key, press Test connection (a free check — it consumes no tokens), then Save; the badge switches from Not configured to Connected. Any administrator can see the key in clear text, and Remove deletes it after confirmation.

Under the AI card, Assistant actions has one checkbox, Technicians approve every action that changes something, ticked by default since version 1.14.14. While it is ticked, every command that changes something, every restart and every shutdown the AI Assistant wants to perform waits until the technician approves it (see The AI Assistant in the Help tools chapter); read-only diagnostics still run straight away, and the Run without asking box is no longer offered. Ticking it takes effect at once (The assistant now waits for approval before every action with effect). Unticking it asks first, Turn off the approval step?, and brings back the previous behaviour: the assistant asks in the chat, and each technician can turn on Run without asking. Only administrators can change it, and each change is recorded in the Audit log.

Warranty (its card is titled Warranty lookup (Dell, Lenovo, HP)) holds your own keys for the manufacturers' warranty APIs, so that Infrastructure › Lifecycle looks up the warranty end date of Dell, Lenovo and HP computers by serial number, in the background (each computer about once a week). There is one block per manufacturer, with a Configured or Not configured badge, a line saying where to get the keys, and its fields: Dell (TechDirect) needs a Client ID and a Client secret (the Warranty API requested in Dell TechDirect), Lenovo a ClientID (which Lenovo gives to its partners) and HP an API key and an API secret (the HP Warranty API at developers.hp.com; HP keys expire every 90 days). The first time, fill in every field of the block and press Save; afterwards the fields stay empty (they read Stored — leave empty to keep it) and only what you type replaces the stored value. Test asks the manufacturer whether it accepts the stored keys (save new ones before testing them), and Remove deletes them after confirmation: the dates already found stay, new lookups stop. When a manufacturer rejects the keys, its lookups pause for 6 hours; a successful Test lifts the pause. The keys are stored encrypted, never shown again and never written to the Audit log, which only records that they were set or removed. Without keys, warranty dates are typed by hand in Infrastructure › Lifecycle (right-click › Edit dates), and a date typed by hand always wins over a lookup.

Backend instances (operator only)

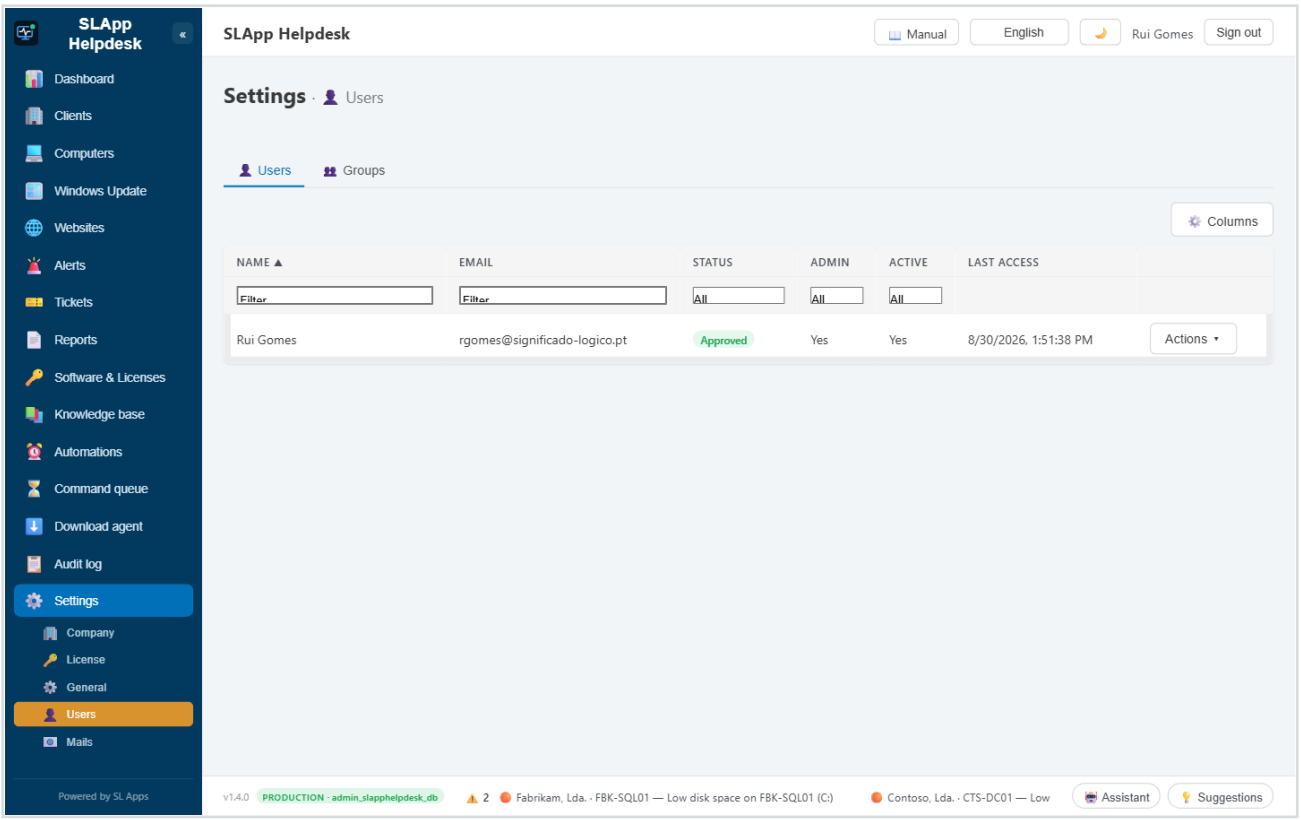
This tab exists only for the administrators of the operator's own account: it is about the servers that run SLApp Helpdesk itself, not about any client. Every server process of the application registers itself here and reports in every minute. The table shows them with Version, Host, PID, Role, Started, Last seen and State, and a line above it says what the process that answered your screen has decided: Schedulers running, when it is the newest version, or Schedulers PAUSED, with the newer version it has seen. The list refreshes by itself every 30 seconds; Refresh reloads it at once. Processes unseen for more than 24 hours no longer count and are hidden: Show the N older than 24 hours lists them again, and the server deletes them by itself after 7 days.

Why it matters: the background jobs (maintenances, offline and other alerts, patch windows, reports, backup checks and synchronisations) must be run by the current version of the application only. An old process left running next to a newer one, for example after an update that did not stop it, would otherwise carry on running them with outdated rules. So a process is out of date as soon as a newer version has run at any time in the last 24 hours, and an out-of-date process pauses all its background jobs by itself; it keeps serving the screens normally. While a process is paused, the administrators of the operator's account see the banner Background tasks are paused on this server. at the top of the application, with a button that opens Settings › General. The State column uses these tags:

- newest — the highest version seen in the last 24 hours.
- live — seen in the last 3 minutes: the process is running now.
- stale — not seen for more than 3 minutes but within the last 24 hours, so its version still counts.
- older than 24 h — no longer counts when the processes decide whether they are out of date.
- this process — the process that answered your screen.

The normal answer to an out-of-date process is to stop it. Forget is for a deliberate rollback: if you go back to an older version on purpose, the newer process is gone but would keep the older one paused for up to 24 hours. Press Forget on the newer row and confirm in Forget this backend instance?; the older version resumes its background jobs within half a minute. Forget is not offered on the row of the process that answered you. Only a process unseen for 3 minutes can be forgotten: on a live row Forget is greyed out (hover it for the reason), because a process that is really running would register itself again within a minute - stop it first. Every Forget is recorded in the Audit log.

Users and groups



Technicians waiting for approval are listed first

The Users tab lists everyone who can sign in to the application, with the columns Name, Email, Status (Approved, Pending or Rejected), Admin, Active and Last access. Pending accounts are sorted to the top. The buttons on each row depend on the account: Approve and Reject for pending users, Reauthorize for rejected ones, Make admin / Remove admin, Deactivate / Reactivate, and Delete (after confirmation). The Tunnels column (since version 1.14.12) shows Always (admin) for administrators and Yes or No for everyone else, and can be filtered (All, Yes, No) and sorted. On the row of a technician who is not an administrator, the Actions menu (also on right-click) offers Block RDP/HTTP tunnels or Allow RDP/HTTP tunnels. Two remote session permissions (since version 1.14.71) are set in the user's profile (Actions › Edit profile), in the Account card under Remote access: Can connect without asking for authorization (the computer's user is not asked to allow the remote session) and Can turn off session recording (may untick Record session when the company records every session automatically). Both are off for every user, administrators included, until an administrator ticks them; each box is saved at once with a notice (Permission granted: … / Permission removed: …), and the Audit log records every change. The server checks them again each time a session starts.

1. Make sure the person's email domain is listed under Company › Domains authorized to sign in, and ask them to sign in once — their account appears as Pending.
2. Open Settings › Users and press Approve on their row (or Reject).
3. If they should manage settings, press Make admin.

> **i** Deactivate blocks sign-in temporarily and is reversible. Delete is permanent: the person can no longer sign in and their tickets, sessions and actions stay in the history without the name. You cannot delete your own account.

> ℹ Every technician can open RDP and HTTP tunnels through the agent until an administrator blocks them. A blocked technician sees Remote Desktop (through the agent) in the Remote access menu (and in the Actions menu of the Computers list) and the web connections of the computer's Network tab greyed out, with You are not allowed to open RDP/HTTP tunnels. Ask an administrator (Settings › Users).; tunnels already open stop working within about a minute. A notice confirms each change ("…" can now open RDP/HTTP tunnels / "…" can no longer open RDP/HTTP tunnels), and the Audit log records every tunnel opened or refused and every change of the permission.

The Groups tab defines teams of technicians. + New group asks for a Color, an Icon, a Name, a Description and the Members (tick boxes — a user can belong to several groups; inactive users are flagged). Each group card offers Edit, Deactivate / Reactivate and Delete; deactivated groups are dimmed with an Inactive badge and no longer offered when linking ticket types. Link a group to a ticket type in General › Tickets › Ticket types: tickets of that type are assigned to the group and the members are notified by email. Deleting a group leaves those ticket types without a responsible group and removes the group from tickets already assigned — prefer Deactivate when in doubt.

Module permissions (groups and users)

Since version 1.14.60 you can choose which sidebar modules each technician can open. In the group editor (+ New group or Edit), the Modules section lists the modules included in your plan, grouped as in the sidebar, with the buttons All and None. Tick the modules the group grants; with none ticked the editor shows None chosen = every module (this group sets no rule). Each group card shows All modules or the number of modules chosen (for example 4 modules).

To change it for one person, open the Actions menu on their row of the Users tab (also on right-click) and choose Modules — it is not offered for administrators. The Modules of … dialog shows the person's groups with what each one grants, and for every module a choice of Inherit (the default: follow the groups), Allow or Deny, followed by the effective result, tick or no-entry sign, and where it comes from: from group …, allowed for this user, denied for this user, none of the user's groups grants it, or no rules - every module. Press Save; a notice confirms Modules of "…" saved.

- Administrators always see every module: the permissions never apply to them.
- Dashboard and My profile are always available.
- A technician with no rule — in no group, or only in groups with no module chosen, and with no Allow or Deny of their own — sees every module, as before version 1.14.60. Deactivated groups do not count.
- Otherwise a technician sees a module when at least one of their active groups grants it.
- The person's own Allow or Deny wins over their groups for that module.
- The plan still decides first: a module your plan does not include stays hidden whatever the permission. Settings and Software packs stay for administrators only.
- Alert emails follow these permissions too (since version 1.14.80): a technician never gets the email of an alert of a module they have no access to, Alerts counting as every module, unless an administrator typed their address in Also notify these emails of a rule or a website. A note with an envelope under the modules, in the group editor and in the Modules of … dialog, says so; the details are under Who receives the alert emails in the Alerts section.

Since version 1.14.63 four modules that used to be for administrators only can also be given to a technician: Transactions, Hours & billing (the tab of Reports), Backups and Audit log. They are marked with a padlock only by explicit permission in the group editor and in the Modules of … dialog, and they follow a stricter rule: a technician gets one only when one of their active groups ticks it or their own row says Allow. A technician with no rule does not get them, so nothing changes until you grant them, and

Deny always wins; when a technician does not have one, the dialog shows administrators only - needs an explicit permission. Ticking only these modules in a group does not restrict the group's other modules (the summary reads No regular module chosen (every module) + 1 by explicit permission), and All ticks the regular modules only. Hours & billing still needs the contracts of your plan, and Backups the backups of your plan. A few actions inside them stay for administrators only: in Backups, adding, editing, testing and deleting a destination (it holds the FTP or share credentials) and showing an archive password; in Audit log, the retention settings; in Transactions, Mark as paid and Cancel of a licence invoice (they can unblock a suspended tenant). A technician with Backups who changes the SQL Server instance, login or computer of a source saved with a SQL password must type the password again.

> **i** A forbidden module disappears from the sidebar, and reaching it another way (a link, a remembered page or the start page chosen in My profile) shows No access to this module with a Go to the dashboard button; the server refuses that module's requests as well. What other modules need stays available: a technician with only Tickets still picks the client, contact and computer of a ticket, but cannot open the Computers module. A computer's sheet belongs to both Computers and Clients (the client's sheet shows its computers), so to hide computers completely deny both. Every change to a group's or a user's modules is recorded in the Audit log.

Requiring multi-factor authentication (Security tab)

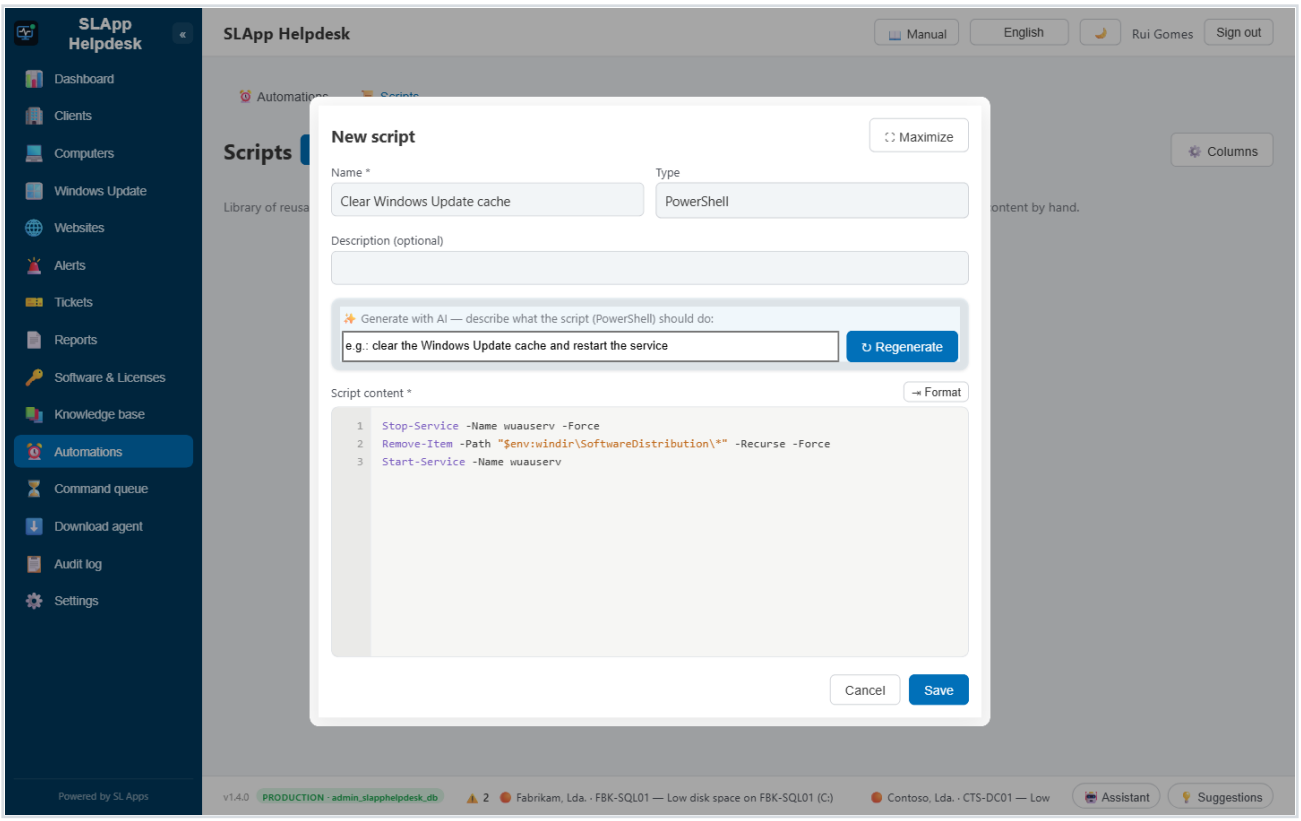
The third tab of Settings › Users, Security (since version 1.14.13), has one card, Multi-factor authentication. While Require multi-factor authentication for technicians is ticked, technicians can only use SLApp Helpdesk if their Microsoft sign-in used multi-factor authentication (MFA), and this is checked on every request, not only at sign-in. The option does not set MFA up: your Microsoft 365 administrator must require it (Security defaults or Conditional Access). The card first shows how your own current sign-in authenticated: Your current sign-in used multi-factor authentication. (tick), Your current sign-in did NOT use multi-factor authentication. (warning) or Microsoft did not report how you signed in (question mark).

1. Sign in with MFA; if the card says your current sign-in did not use it, sign out and sign in again with MFA.
2. Open Settings › Users › Security and tick Require multi-factor authentication for technicians. The box can only be ticked when your own sign-in used MFA, so that you cannot lock yourself out; otherwise the reason is shown under it.
3. Confirm in the Require MFA? window: technicians who sign in without MFA are blocked immediately, including those already signed in. A notice confirms Multi-factor authentication is now required.

> **i** A blocked technician sees a full-page Multi-factor authentication required screen with Sign in again (a fresh Microsoft sign-in, where they complete the second step) and Sign out. Personal Microsoft accounts, and sign-ins through another identity provider that does not pass the MFA result on to Microsoft, are refused while the option is on. Unticking it is always possible and takes effect at once (Multi-factor authentication is no longer required). Every change of the option and every refused request (mfa_rejected, at most once every 15 minutes per person) is recorded in the Audit log.

> **i** If the card says This cannot be turned on yet: SLApp cannot see how anyone signed in. Contact SLApp support., Microsoft is not yet telling SLApp how people sign in, and the option stays off. If an organisation ever locks itself out, SLApp support can turn the requirement off for it.

Scripts

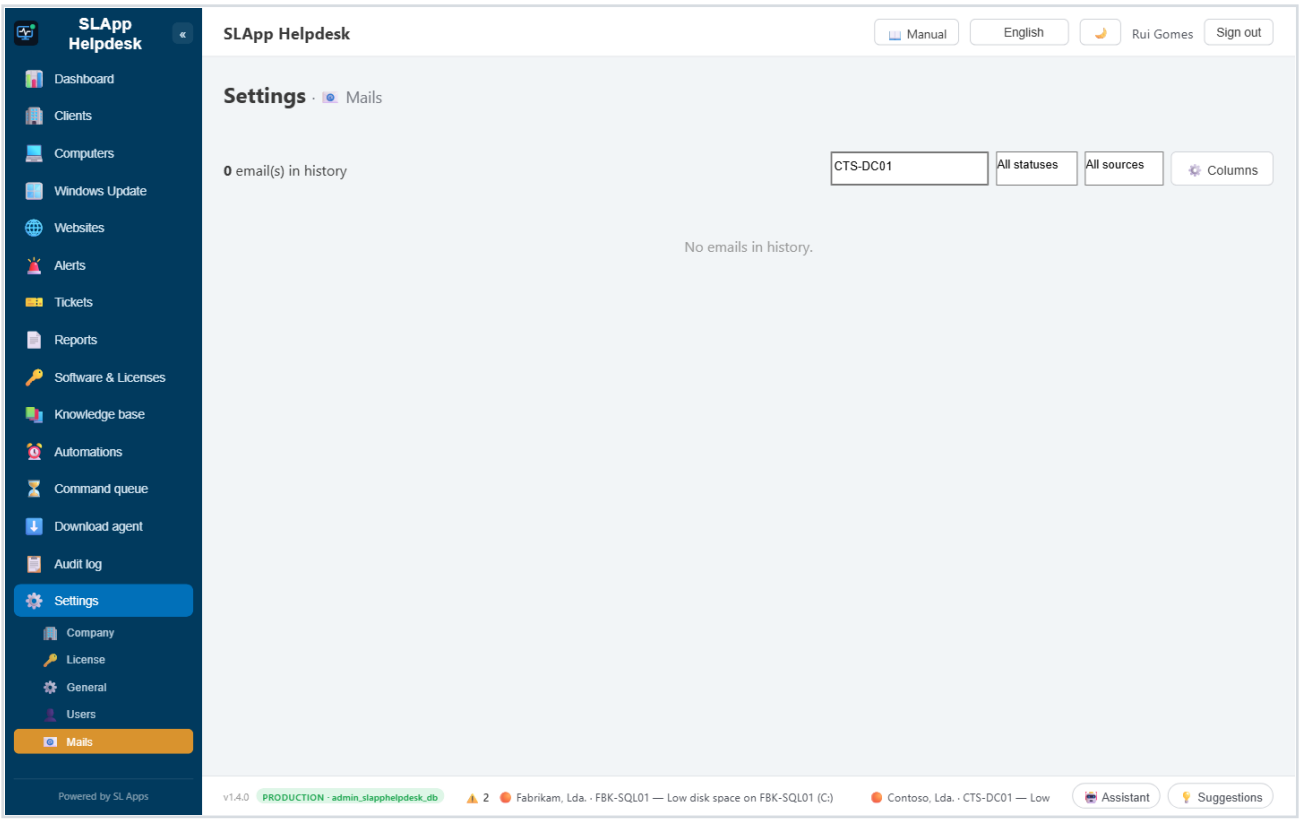


Script editor with syntax highlighting and AI generation

Scripts is a library of reusable PowerShell and CMD scripts that technicians pick in Computers › Remote actions instead of pasting code by hand. It is not a Settings group: open it from Automation & updates › Automations, Scripts tab. The table shows Name, Type and Description; click a row to edit it, or press Remove to delete it after confirmation.

1. Press + New script and enter a Name; pick the Type (PowerShell or CMD) and, if you like, a Description (optional).
2. Either paste or write the code in Script content (the editor shows line numbers and syntax highlighting), or describe what it should do in the Generate with AI box and press Generate or Enter (Regenerate rewrites existing content). This needs the AI (Claude) connector.
3. Use Format to re-indent (only leading whitespace changes) and Maximize (Restore to go back) to get a full-screen editor; Tab inserts two spaces.
4. Press Save — it is enabled only when both a name and content are present.

Mails



History of every email sent by the application

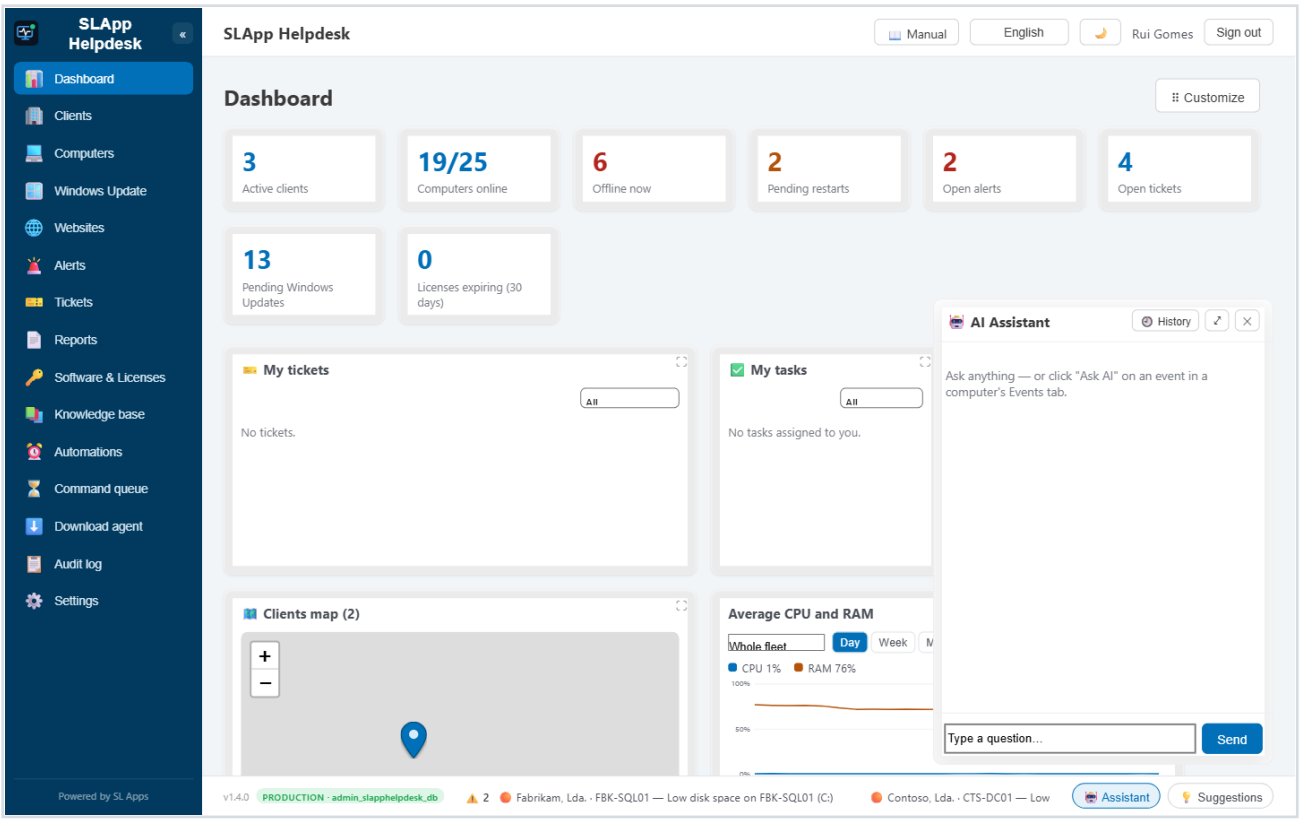
Mails is the sending log: the header shows how many email(s) in history there are (the most recent 200) and, in red, how many failed. Filter with Search recipient/subject…, the status list (All statuses, Sent, Failed) and the source list (All sources, then each source that appears in the log, by its name, such as Test e-mail for messages sent from the Email connector's test box; in the table, pointing at a source shows its code). Columns are Date/Time, Recipients, Subject, Source and Status. Press a Failed status to read the Failure reason, and View to open the Sent email body when it was stored.

> **i** If nothing seems to reach clients, check this log first: a failed entry with its reason usually points straight at the Email connector (wrong credentials, or a recipient blocked by the Allowed domains safety net).

41. Help tools: AI Assistant, Suggestions and this manual

Three small helpers are available from every screen of SLApp Helpdesk, regardless of which module you are in: the AI Assistant panel and the Suggestions form sit in the bottom bar (footer) of the app, and the Manual link sits in the top header next to the language switcher. Opening any of them changes nothing by itself, so feel free to explore; just keep in mind that the AI Assistant can act on your managed computers when you ask it to (see below).

The AI Assistant



The AI Assistant panel, opened from the Assistant button in the footer

Click the Assistant button (robot icon) in the footer to open a chat panel in the lower-right corner of the screen, just above the footer. The button stays highlighted while the panel is open. The panel floats on top of whatever module you are using and keeps its conversation when you switch modules, so you can ask a question, go look at a computer, and come back to the answer. Click the Assistant button again, or the X (Close panel) button, to hide it without losing the thread. An empty panel shows a hint reminding you that you can also start from an event in a computer's Events tab.

The assistant is an IT support helper. It is built to explain Windows events and remote-action results, but you can ask it anything. Beyond answering questions, it has a small set of tools it can use on your managed fleet:

- List your managed computers (with client, online/offline status, operating system and the last signed-in user), optionally filtered by a name, hostname or client you mention. The list returns at most 50 computers, so name the client or computer when your fleet is large.
- Run a CMD or PowerShell command (PowerShell by default) on a computer through the installed agent and read the output back. By default it runs elevated (as the system service, which is always connected); it only switches to the signed-in user's session when the question genuinely needs the user context.
- Restart a computer or shut it down. The assistant is instructed to do this only when you explicitly ask for it, never on its own initiative to fix a problem.
- Before running a command whose effect is not obvious, it is instructed to explain in one sentence what it is about to do.

> ℹ Any command, restart or shutdown the assistant performs goes through exactly the same channel as the buttons in a computer's Remote actions tab, and is recorded in that tab's History ("History — everything done from the application") with your name in the Technician column. There is no distinction between an action you clicked and one you asked the AI to do, so treat its requests with the same care.

Since version 1.14.14, while Technicians approve every action that changes something is on (Settings ›

General › Connectors › AI › Assistant actions, on by default), nothing with an effect runs until you approve it. When the assistant wants to restart or shut down a computer, or to run a command that changes something, the Approve the assistant's action window opens: The assistant wants to do the following. Nothing runs until you approve it., then what the assistant said and the list of actions (Run a command on …, with the command, the Interpreter and as SYSTEM (service) or in the signed-in user's session; Restart …; Shut down …) and Expires at … (30 minutes later). The window does not close when you click outside it.

- Approve and run (red when a restart or a shutdown is included) runs exactly the actions shown, and the assistant carries on (Approved — the assistant is running it). It may ask again for its next step; while it runs, the button reads Running….
- Reject tells the assistant not to run them (Rejected — the assistant was told not to run it).
- Decide later closes the window and leaves a Waiting for your approval card in the conversation, with a Review button that reopens it. Until you decide, the question box and Send are disabled (Approve or reject the pending action first).
- A request older than 30 minutes, or one already decided in another window, is not run: the notice says This approval request expired — ask the assistant again or This action was already decided (in another window?).

> **i** Read-only diagnostics, for example Get-… commands, ipconfig /all, dcdiag, repadmin /showrepl, w32tm /query or systemctl status, still run straight away; anything the application cannot be sure is read-only asks for approval. While the approval step is on, the Run without asking box is not shown. Every approval and rejection is recorded in the Audit log on the computer (ai_action_approved, ai_action_rejected).

The assistant needs a Claude API key configured by an administrator under Settings, Connectors, in the section AI (Claude). Without it, sending a message shows an error in the panel telling you the Claude connector is not configured. The same key also powers the Generate with AI button in the script editor.

Controls in the panel header:

- History: lists your previous conversations (title and last-updated date), most recent first, up to the 100 most recent. Click one to reopen it and continue where you left off. Only your own conversations are shown; other technicians cannot see them, and you cannot see theirs. Use Back to return to the current chat.
- Close conversation (shown once the chat has messages): ends the current conversation and starts a fresh, empty one. The old conversation stays in History.
- Maximize / Restore: enlarges the panel to a larger window, or back to the compact size.
- Close panel (X): hides the panel only. The conversation is kept and reappears next time you open the assistant.

1. Click Assistant in the footer (bottom-right of the screen).
2. Type your question in the box at the bottom (Type a question…) and press Enter or click Send. While the answer is being prepared you will see Thinking… and the input is temporarily disabled.
3. Read the reply. If the assistant needs to look something up or act on a machine, it does so in the background and includes the outcome in its answer. If something goes wrong (connector not configured, AI service unreachable), the error appears in the panel under the messages.
4. Continue the conversation with follow-up questions; the whole thread is sent each time, so the assistant remembers the context.

5. When you are done with a topic, click Close conversation to start a new thread, or simply close the panel.

> There is a limit on how many tool actions the assistant can chain inside a single reply. If it reaches that limit, it answers that it could not complete the request (too many chained actions); split the task into smaller questions and ask again.

Ask the AI from an event or an action result

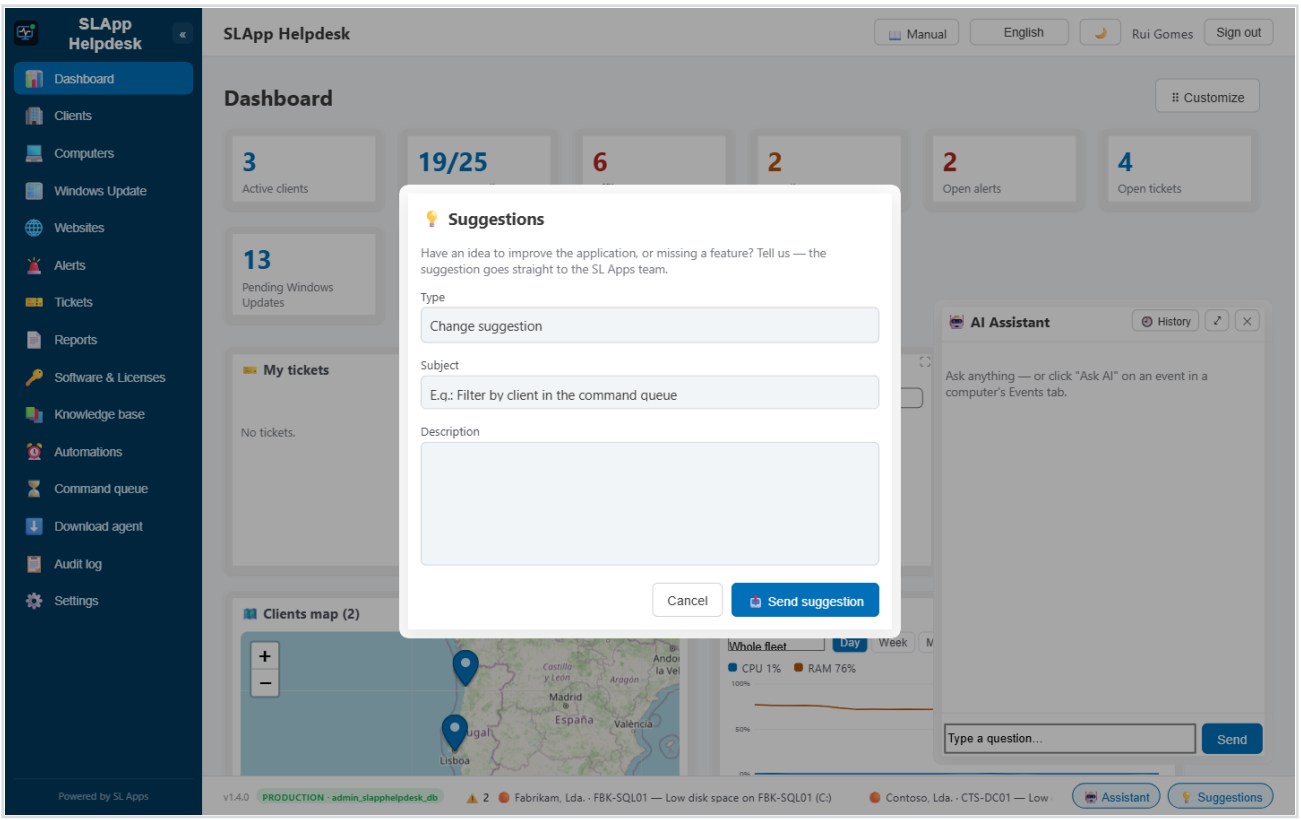
The fastest way to use the assistant is from a computer's detail sheet. In the Events tab, every Windows event row has an Ask the AI button: clicking it opens the panel (if closed) and immediately sends a ready-made question containing the computer name, date, level, channel, source, event ID and message, asking the assistant to identify the problem, the probable cause and a practical solution. Likewise, in the Remote actions tab, click the command text (Detail column) or the result text (Result column) of a row in the History to open it in full; that window has an Ask the AI button (next to Copy and Close) which sends the action and its full text for the assistant to explain whether it went well and what each relevant line means, and closes the window.

> Because these prompts include the computer's identifier, the assistant is able to investigate or fix things directly on that machine using its tools. For commands whose effect is not obvious it explains what it intends to run first; if you only want an explanation and no action, say so in a follow-up message. With the approval step on, those actions wait for you in the Approve the assistant's action window.

> The first message of a conversation becomes its title in History (shortened to 80 characters), so a short, descriptive opening question makes old threads easier to find later. Conversations started from an "Ask the AI" button are titled with the beginning of the ready-made prompt ("Explain this Windows event…" or "Explain the result of this remote action…"), so they all look alike in History; use the date to tell them apart.

> If a command the assistant ran comes back as "scheduled", the machine is not necessarily offline: it may be online but without an active command channel at that instant, and the command will run once the agent reconnects. Also note that the assistant is currently set up to answer in European Portuguese regardless of the interface language you have selected.

Suggestions



The Suggestions form, opened from the footer

Any signed-in user can propose an improvement or request a feature. Click the Suggestions button (light-bulb icon) in the footer to open the form; the button stays highlighted while the form is open, and clicking it again, clicking Cancel or clicking outside the form closes it. Your suggestion is stored and sent straight to the SL Apps product team, together with your name, your e-mail address, your organisation and the environment you are using, so they can get back to you if they need more details.

- Type: choose Change suggestion (improve something that already exists, the default) or New feature (something that does not exist yet).
- Subject: a short one-line title, up to 150 characters. Example: Filter by client in the command queue.
- Description: the idea itself, up to 5000 characters. Say what you would like to change or have, and on which screen or in which situation it is needed.

1. Click Suggestions in the footer.
2. Pick the Type, enter a Subject and write the Description. The Send suggestion button stays disabled until both text fields are filled in.
3. Click Send suggestion. While it is being submitted the button reads Sending… and Cancel is disabled. If the submission fails, the error is shown in the form and you can try again.
4. A confirmation (Suggestion sent) appears. Click Close to return to the app.

> ℹ Suggestions are not support tickets. If something is broken and you need help now, click Report a problem (bug icon) in the footer instead: it opens a support request with the SLApp Helpdesk team, with the subject, what happened, a priority and up to three screenshots (paste them with Ctrl+V), plus the screen you are on, the app version and the browser; the same window lists the problems your company reported and their status. Click one to read the support team's replies (never our internal notes) and answer it with text and a screenshot; a problem with a reply you have not opened yet is marked New reply, and the button in the footer shows how many there are. You also get an e-mail when we reply. Once the problem is closed, report a new one if you still need help. Report a problem only appears for customer companies; use the Suggestions form for ideas and feature requests. If the team needs clarification, they will reply to the e-mail address of your account, so keep an eye on your inbox.

The Manual button

The Manual link (book icon) in the header, just left of the language switcher, opens this user manual as a PDF in a new browser tab. The manual is provided in English, Portuguese and French and the link automatically picks the edition matching the interface language you currently have selected; if a language has no edition of its own, the English manual opens instead.

> ℹ Because the PDF opens in a separate tab, you can keep it side by side with the app while you follow a procedure. Use your browser's search (Ctrl+F) inside the PDF to jump to a button or field name quickly.

42. Getting help

If you run into a problem this manual does not cover, open a support ticket in your own SLApp Helpdesk instance, check the in-app Knowledge Base for guides written by your own team, or ask the AI Assistant from the footer bar. To propose an improvement to the application itself, use the Suggestions button — every suggestion reaches the SL Apps product team.