



SPLASHTOP ON-PREM — SELF-HOSTED EDITION

Splashtop On-Prem Security Kit

How Splashtop On-Prem keeps your remote sessions secure

Contents

Select a section to jump straight to it. Each one covers what a control does, why it matters, and how it's applied.

Introduction

PAGE 3



At a glance

PAGE 3



01

Encryption and connection security

PAGE 4



02

Identity: authenticating your users

PAGE 5



03

Access origin and device posture

PAGE 6



04

Entitlement: what each user can reach

PAGE 8



05

Session control: what happens once connected

PAGE 9



06

Assurance: logging, recording, and audit

PAGE 11



07

Hosting, tenancy, and data control

PAGE 12



08

Compliance and certification

PAGE 13



Frequently asked questions

PAGE 14



Introduction

When you connect to a computer remotely, that connection has to be trustworthy in three ways: the channel itself has to be private, the people using it have to be who they say they are, and there has to be a record of what happened. Splashtop On-Prem is built around all three, and because you host it yourself, you control each one directly.

Splashtop On-Prem is a fully self-hosted deployment. The Splashtop Gateway — which handles authentication, user and device management, session relay, the web console, and the database — runs on a server you own, inside your own network. No Splashtop cloud service sits in the path of a session. Your user directory, access permissions, audit logs, and session recordings stay in your environment.

This document walks through the security capabilities available in Splashtop On-Prem — what each one does, why it matters, and how it fits together. It's organized as a practical reference you can use to answer the question “how does Splashtop keep our sessions secure?”, whether that question is coming from your own team, a customer, or an auditor.

For full technical and architectural detail beyond what's summarized here, see the Splashtop On-Prem Products Security Overview. For current certification status, see splashtop.com/compliance.

At a glance

Splashtop secures a session through independent layers of control. Because they're independent, no single weak point exposes a session — each layer asks a different question.

Layer	Question it answers	Key capabilities
Encryption & architecture	Is the connection itself private?	End-to-end TLS and AES-256, your own CA-signed certificate, a single controlled ingress port, signed software
Identity	Who is logging in?	Active Directory / LDAPS, SSO (SAML 2.0), SCIM provisioning, two-step verification, device and browser authentication, password policy and expiration, account lockout
Access origin & device posture	From where, from what device, and is that device healthy?	IP address allow list, MAC address allow list, Conditional Access
Entitlement	What are they allowed to reach, and for how long?	Role-based access, group permissions, scheduled access, authentication at the endpoint
Session control	What can they do once connected?	Granular feature controls, request permission to connect, privacy protections, enforced endpoint configuration
Assurance	Can we prove what happened?	Full activity logging, enforced centralized session recording, syslog/SIEM export
Hosting & tenancy	Where does our data live, and who can reach it?	Self-hosted on your infrastructure, your database, private team tenancy, air-gap capable

1. Encryption and connection security

A Splashtop On-Prem deployment has three parts, all of which you operate: the **Splashtop Gateway** (the central server — API, web console, relay, and database), the **Splashtop Streamer** (the agent on the computers you want to reach), and the **Splashtop On-Prem app** (the client used by technicians and end users, with a browser-based Web App as an alternative).

Every session is encrypted end-to-end. Web and API traffic uses HTTPS with HSTS; the session itself is carried over a relay-brokered tunnel using TLS and AES-256, where the two connecting devices negotiate their encryption key **directly with each other**. The relay passes the encrypted traffic through — it never holds the key and cannot decrypt the session, even though the relay runs on your own server.

Only **one TCP port** needs to be permitted through your firewall — 443 by default, and changeable from the Gateway console. That single port carries client authentication, session relay, web console access, and Streamer registration. The Streamers and client apps

themselves listen for nothing; they only make outbound connections to your Gateway, so if all of your users are internal, the Gateway never has to face the internet at all.

A few additional details:

- You install **your own CA-signed SSL certificate** (2048-bit RSA, SHA-2) on the Gateway, used consistently by the API, web, and relay components to verify identity and prevent man-in-the-middle attacks. The certificate is issued against your own FQDN and managed under your existing certificate lifecycle.
- Passwords are never stored in plain text. They're hashed with SHA-256 using a salt, and the Gateway's connection to its database is protected with SCRAM-SHA-256.
- Client and Streamer software is distributed **from your own Gateway** over HTTPS rather than from a public download site, and is code-signed — Windows executables with organization-validated certificates, Windows drivers with extended-validation certificates, macOS binaries per Apple's developer requirements. Administrators also control which software versions are enabled for their users.

2. Identity: authenticating your users

Active Directory integration

Splashtop On-Prem binds directly to your **Active Directory** over LDAP or LDAPS. You can add individual AD users or entire AD groups, and group members inherit the role and access permissions assigned to that group. Users authenticate with their existing AD credentials, which means your password policy and — most importantly — your offboarding process apply directly to Splashtop access. When someone loses their directory access, they lose their Splashtop access.

Multiple AD servers are supported, with UPN-suffix mapping so users can log in either with their UPN or with sAMAccountName plus domain. Best practice is a dedicated, read-only, non-interactive service account for the bind, scoped at a parent OU or domain level.

Single Sign-On (SSO)

Splashtop On-Prem supports SSO via **SAML 2.0**, with documented setup for Microsoft Entra ID, Okta, ADFS, JumpCloud, OneLogin, and generic SAML 2.0 identity providers. Authentication happens against your own identity provider rather than a separate Splashtop password, so your existing conditional access rules and session policies apply.

SCIM provisioning is supported for Okta and Microsoft Entra ID, so users and groups can be created, grouped, permissioned, and removed automatically from your identity provider, keeping access current without manual upkeep on the Splashtop side.

Two-step verification (2FA)

Splashtop On-Prem supports **TOTP-based** two-step verification using any standard authenticator app (Google Authenticator, Microsoft Authenticator, Cisco Duo Mobile, Okta

Verify, and others), as well as **email-based one-time passcodes**. Once enabled, logging in requires the second factor in addition to the password.

Enforcement is configurable rather than all-or-nothing: the Team Owner can require 2FA for Admins and Group Managers, for Members, or both, and can delegate that configuration to admins per user or group. Trusted devices are managed centrally — administrators can see every trusted device, revoke any of them, and set how long a device stays trusted (1, 7, or 30 days, or indefinitely) before it has to re-authenticate.

Device and browser authentication

Independent of 2FA, a login from an unrecognized device or browser triggers a mandatory verification step that confirms the person genuinely controls the email address behind their account. This is what stops a stolen password alone from producing a session — an attacker would also need access to the account holder’s mailbox. Verification requests can be routed to administrators instead of end users, giving IT explicit approval rights over every new device, and both the authentication duration and the request-expiry window are configurable.

Password requirements and expiration

The complex-password policy is enforced in the Gateway and is tunable: minimum length, complexity requirements, an optional expiration period, and an option to enforce the policy at next login. Password expiration closes the risk window on a credential that may have been quietly compromised without the user’s knowledge, and is a common requirement in regulated environments where standing credentials can’t be left unchanged indefinitely.

Account lockout policy

Accounts can be configured to lock automatically after a defined number of consecutive failed login attempts, for a defined lockout duration. Locked accounts unlock automatically when the duration expires, or manually by an Admin or the Team Owner from the Users page.

Session timeouts

Idle sessions can be logged out automatically — from the On-Prem client app after 15 minutes, 1 hour, 8 hours or 24 hours, and from the web console after 5 minutes to 8 hours — closing the window in which an unattended, already-authenticated session could be misused. Where policy requires credentials on every launch, the “remember app login” convenience option can be switched off.

3. Access origin and device posture

IP address allow list

You can restrict access by source IP address, using single addresses or CIDR ranges, with three scopes configured independently: the **Gateway web console**, the **On-Prem client app**,

and the **Open API**. Connections from outside the list are refused. This is commonly used to require that administration only happens from the corporate network or a specific VPN egress, while allowing client-app access more broadly — or the reverse.

Attempts blocked at this layer are recorded in the Gateway's debug logs rather than the web console. That's deliberate: these attempts occur before authentication, carry no account context, and surfacing them in the console would create a log-flooding target during a denial-of-service attempt.

MAC address allow list

For a tighter requirement — “only company-issued devices may connect” — client app logins can be restricted to a defined list of device MAC addresses, imported or exported by CSV, with the option to block mobile devices entirely. This is a strong policy-enforcement control and pairs well with device management (MDM) tooling for environments that don't allow a BYOD path. As with any MAC-based control, treat it as one layer among several rather than a standalone barrier, since MAC addresses can in principle be altered by someone with local administrative rights on a device.

Conditional Access

IP and MAC restrictions answer *where* a connection is coming from and *which* device it is.

Conditional Access answers a different question: is that device actually in a healthy state right now?

Before a session is allowed to start, the client device is checked against a policy you define. If it passes, the session proceeds normally. If it fails, the session is blocked and the user is shown exactly which check failed and what to do about it. Checks available include:

- **Antivirus** — that a real-time antivirus product is running (any product, or a specific one such as Microsoft Defender or CrowdStrike Falcon), and that its definitions are current.
- **Windows Security Center** — that the firewall, the antivirus protection it monitors, and the Security Center service itself are all reporting a healthy state.
- **Operating system** — that the device meets a minimum Windows version, and that a specific update (by KB number) is installed.

Policies are assigned the same way as other granular settings — a team-level default, overridden at group level, overridden again at user level — so you can apply strict posture requirements to contractors or high-privilege technicians without imposing the same on everyone. Devices that can't be evaluated (Windows Server, or older Windows versions) are blocked by default, with an explicit option to allow them.

For assurance, administrators get a device list showing the current pass/fail state of every managed device with per-check detail and last report time, plus a dedicated **Conditional Access failure log** recording every blocked session with the account, device, policy, and specific checks that failed. Both are exportable to CSV.

This turns endpoint hygiene from a policy you publish into a condition of connecting: a laptop with disabled antivirus or missing patches can't reach a production server until it's brought back into compliance. Conditional Access currently evaluates Windows client devices.

4. Entitlement: what each user can reach

Roles

Splashtop On-Prem separates administrative capability into distinct roles — **Team Owner, Super Admin, Admin, Group Manager, and Member** — so day-to-day administration can be delegated without handing over control of system-level or team-wide security settings. Group Managers, in particular, can be scoped to manage only the users and computers in specific groups rather than the whole team.

Super Admin sits between Owner and Admin, with two independently grantable privileges: full control over all Admins and Group Managers (role changes, removal, password resets), and the ability to manage team policies and settings including SSO and AD authentication. Granting one does not grant the other, so you can delegate user administration without also delegating authentication policy.

System management permissions

Separately from the role hierarchy, the Team Owner can grant individual users **System Management** — the ability to view and configure Gateway-level settings such as security policy, backup and restore, license management, system monitoring, network settings, SMTP, and syslog. Users without this permission don't just see these settings greyed out — the system-level pages don't appear in their web console at all.

Because this permission reaches into infrastructure configuration rather than just user and computer administration, it's granted and revoked independently of the Admin/Super Admin chain, and the Users page can be filtered to show at a glance who currently holds it. Given what it unlocks, treating two-step verification as mandatory for anyone with this permission — and revoking it immediately on a role change or offboarding — is worth making standing practice rather than an occasional review.

Group-based access permissions

Access to each computer is granted explicitly. A newly created user has access to no computers unless they're placed in a group that carries permissions — entitlement is something you grant, not something you have to remember to remove. Grouping lets permissions scale with your organization rather than being managed one user at a time, and where SCIM or AD groups are in use, membership can be driven entirely from your directory. Permissions can also be applied in bulk or updated via CSV, and user lists, access permissions, groups and computer lists can all be exported for review or audit evidence.

Scheduled access

Access can be granted only within defined time windows rather than standing access at all times — useful for contractors, third-party vendors, shift-based support, and lab or classroom environments. Scheduled access is additive to any standing permission a user already has, so for a fully time-bounded user, set their standing permission to “no computers” and grant the schedule alone. Sessions can be forced to disconnect when the timeslot ends, and **Exclusive mode** reserves a scheduled computer for one user at a time.

Authentication at the endpoint

Reaching a computer in the Splashtop console doesn't have to mean being able to use it immediately. A second authentication step can be required at the target machine itself: entering its Windows/macOS login credentials, or entering a security code specific to that computer. The Team Owner can also globally **disable the saving** of those credentials and security codes, so they must be re-entered for every session rather than cached on the technician's device.

Request permission to connect

One of the most frequently requested controls: the person physically at the target device can be required to explicitly approve a connection before it's allowed to proceed. When a technician attempts to connect, an on-screen prompt appears at the Streamer with a countdown, and the session only opens if the local user accepts it. This puts a human checkpoint in front of every connection to that machine, which matters most for personal or sensitive workstations where the person sitting at the computer should always have the final say over whether someone connects to it.

Behavior when the prompt times out — whether the connection is automatically rejected, automatically allowed, or something in between — is configurable, as is how the prompt behaves when the machine is sitting at a locked login screen with no one there to respond.

5. Session control: what happens once connected

Granular controls

Individual session capabilities — remote control, file transfer, copy and paste, remote printing, remote command execution, system tools, watermarking, session recording, and more — can be turned on or off per user, per group, or per role, rather than as a single team-wide switch. File transfer and clipboard can each be controlled by direction, so you can allow download but block upload, or vice versa.

This lets you apply different policies to different populations: for example, disabling file transfer and clipboard access broadly, while enabling it for a small, named group that genuinely needs it. Settings resolve in layers — team default, then group, then individual user override — so policy can be as broad or as specific as you need, and bulk actions let you apply a change across many accounts at once. An **Admin configurable** control determines

whether Admins and Group Managers may adjust a given feature for others, or whether it stays reserved to the Team Owner.

Attended (SOS) and unattended access are configured separately, so ad-hoc support sessions can carry a stricter policy than managed endpoints, or the other way around.

Session privacy protections

A set of controls protect the physical environment around an endpoint during a session:

- **Mandatory session notification** — a notification always appears on the target computer whenever it's being accessed, which cannot be suppressed. This is a built-in anti-surveillance guarantee.
- **Blank screen** — the target computer's physical display is blanked during the session.
- **Lock keyboard and mouse** — local input is disabled while a session is active.
- **Lock screen on connect / disconnect** — the machine can be locked automatically at the start of a session and returned to the OS lock screen when the session ends.
- **Automatic log-off and idle timeout** — sessions that disconnect unexpectedly or sit idle are handled automatically rather than left open.
- **Watermarking** — customizable text can be overlaid on the session, deterring and helping attribute out-of-band capture such as screen photography. Font size, colour, opacity and outline are all configurable. Note that watermarks are applied to the live session and are not burned into session recording video.

Enforced endpoint configuration

Settings that a local user could otherwise change are locked down centrally:

- **Deployment packages** apply your intended configuration — security settings, grouping, naming — at install time, and can be rolled out through your existing tools (Group Policy, Microsoft Intune, Jamf Pro, PDQ, etc.).
- **Streamer setting lock** ensures settings can only be changed by a Splashtop team administrator, not by anyone with local machine administrator rights.
- **Endpoint Policy** lets you configure and update Streamer settings across all deployed endpoints centrally from the Gateway web console, without touching the machine directly. Settings governed by policy are greyed out locally — and unavailable from the in-session toolbar — until removed from the policy. A policy can be attached to a deployment package so endpoints comply from the moment they're installed.
- **Auto-update Streamers** can be applied to all computers or scoped to specific computers and groups, keeping the fleet on a current, supported version.

6. Assurance: logging, recording, and audit

Activity logging

Every session is logged with start time, end time, duration, the target computer and its IP address, and the user, device, and device IP that connected. Logging covers the full range of session types — remote desktop, file transfer, chat, remote command, and support (attended) sessions — each with its own detailed log, and sessions currently in progress are flagged as such. File transfers are logged at the file level — name, size, direction, timestamp — which directly answers the question of what data moved and which way. Team management activity (permission changes, new devices, team setting updates, failed logins) is captured separately in its own operation log, and Conditional Access failures have their own log as described in Section 3.

Admins and the Team Owner see logs for the whole team; Members see only their own.

Log type	Retention
Session activity logs (remote desktop, file transfer, chat, remote command, support sessions)	1 year
Operation logs — team management history (web console)	1 year
CSV export	12 months
Syslog / SIEM export	Continuous — governed by your own retention policy

Session logs are viewable in the web console across any 90-day range within the past year. Administrators can also purge logs by month from System → Maintenance, with the two most recent months protected; removal is permanent, so agree an archival routine before using it.

Enforced centralized session recording

By default, session recordings save locally on the technician's own machine — useful, but not an audit control on its own, since the recording lives on the machine of the person being audited. Splashtop On-Prem includes **centralized session recording**, which enforces recording on all sessions, stores the recordings **on your own Gateway server**, and makes them available for playback or download from the web console for up to 30 days.

Technicians cannot start or stop recording when it's enforced. Playback, download, and delete rights are assigned by role, with the Team Owner holding all permissions by default, so the ability to review a recording and the ability to remove one can be held by different people.

One operational note: enabling centralized recording redirects what would otherwise be a direct LAN connection through the Gateway, so the encrypted stream can be captured intact. Size the Gateway and its storage accordingly.

Syslog / SIEM integration and open APIs

Session and team logs can be forwarded continuously to **up to two syslog servers** in standard RFC 5424 format, over UDP, TCP, or TLS, with configurable facility and severity filtering. This works with any syslog-capable SIEM — Splunk, IBM QRadar, ArcSight, LogRhythm, and Microsoft Sentinel are cited in Splashtop's documentation — so Splashtop activity is correlated with the rest of your security telemetry and retained under your SIEM's own policy rather than reviewed in isolation.

RESTful open APIs are available for custom reporting and monitoring integrations, and can be independently IP-restricted as described in Section 3.

7. Hosting, tenancy, and data control

Self-hosted by design

The defining property of Splashtop On-Prem is that **there is no Splashtop cloud in the path of a session**. All session traffic flows through your own Gateway. Your user accounts, access permissions, audit logs, session recordings, and configuration live in a database on your server. Questions about data residency, sub-processors, and cross-border transfer resolve to a single answer: your data is wherever you put your Gateway.

The Gateway can be deployed as a **single node** for smaller organizations and lighter session volumes, or **clustered** across multiple Windows nodes for scale and resilience. It can equally run in your own private cloud tenancy, or be hosted by a regional partner where local sovereignty rules require it.

Air-gapped operation

Splashtop On-Prem does not need internet access for remote sessions or for software updates — Streamer and client app updates are distributed from your Gateway. The only outbound requirement is licence validation, and an **offline licence** removes even that, for environments that must remain fully air-gapped.

Private tenancy

Every Splashtop Gateway hosts its own team — a private tenant, by default, for every self-hosted environment. Endpoints are enrolled and belong to one specific team. Visibility and reachability stop at that boundary: a user can only ever see and connect to the endpoints enrolled in their own team, and reaching an endpoint requires both team membership and an explicit access permission granted by that team's administrators.

This means access is never a property of an endpoint alone — it's always a property of team membership plus an explicit grant, both governed and revocable centrally. In a self-hosted deployment there is no shared infrastructure with any other organization at all, and no path by which a connection could originate outside your Gateway.

Backup and recovery

The Gateway console includes system backup and restore, plus a scheduled backup with configurable days and times, a choice of whether to include logs and computer inventory, a mandatory password on the resulting archive, and a retention rule for how many backups to keep. A system health view reports the current state of Gateway services.

Gateway updates

Splashtop releases regular Gateway updates throughout the year, and these aren't just maintenance — new security capabilities ship as part of the regular release cycle, alongside fixes for anything identified since the last release. Keeping the Gateway on a current version is the simplest way to make sure your deployment has the latest protections available, and it's worth building into the same patching cadence as the rest of your infrastructure.

8. Compliance and certification

Standard	Status
ISO/IEC 27001:2022	Annual audits
SOC 2 / SOC 3	Annual SOC 2 audits; maintains SOC 2 Type 2 and SOC 3 compliance
Penetration testing	Independent third-party penetration testing conducted annually
GDPR	Applies differently than in the cloud editions. Because an On-Prem deployment keeps all session data, logs, and recordings on your own infrastructure, Splashtop does not process your users' personal data as part of the product's normal operation — you remain both controller and processor for your own environment. Splashtop's own corporate practices align with GDPR principles.
PCI DSS	Splashtop does not store or process cardholder data as part of the remote session service, and partners exclusively with PCI DSS-compliant payment providers for its own billing and transaction processing.
NIST	Splashtop's access control, patch management, and policy enforcement capabilities are designed to support customer security programs aligned with NIST frameworks.
NIS2, DORA, ENS, TISA, Cyber Resilience Act	Splashtop's access control, audit logging, vulnerability management, and resilience capabilities support customer compliance efforts under these frameworks — an on-premises deployment additionally keeps all session data within your own infrastructure and jurisdiction.
SOX	The audit trails, encryption, authentication, and access controls described in this document support the remote-access controls organizations need for SOX compliance over financial systems.

Current certification status is always available at splashtop.com/compliance.

Frequently asked questions

Where is our data stored? On your own infrastructure. The Gateway holds user accounts, permissions, logs, and recordings in a database on your server, wherever you choose to run it. No session data traverses Splashtop's cloud, and Splashtop has no access to your Gateway.

Do we need to open firewall ports or expose anything to the internet? One port on the Gateway — 443 by default, and changeable — must be reachable by the clients and endpoints that use it. Streamers and client apps listen for nothing and only make outbound connections. If all of your users are internal, that port never has to face the internet at all.

Can it run fully air-gapped? Yes. Session traffic and software updates never require internet access, and an offline licence removes the only remaining outbound dependency.

What stops a stolen password from being used to access our systems? Several things, depending on what you've configured: device and browser authentication (an unrecognized device triggers a verification step), two-step verification, IP or MAC address restrictions, account lockout, and the requirement that the account belong to a team with an explicit permission to the target computer. Where AD or SSO is in use, the credential itself is governed and revocable centrally through your directory.

Can we stop data from leaving a session? Yes — file transfer, clipboard copy/paste, and remote printing can each be disabled team-wide or scoped to specific users, groups, or roles, and file transfer and clipboard can be controlled per direction.

Can the person at the target computer control whether someone connects to it? Yes — Request Permission to Connect requires the local user to explicitly approve each connection attempt from an on-screen prompt before the session is allowed to start.

Can we prove what happened during a session, after the fact? Yes. Sessions, file transfers, chat, and remote commands are logged in detail; logs can be forwarded continuously to your SIEM for longer retention; and centralized recording captures every session to your own Gateway for playback under role-based permissions.

Is our environment isolated from other Splashtop customers? Completely — it isn't shared infrastructure at all. Your Gateway is yours. Within it, teams act as private tenants, and users can only see and reach the endpoints enrolled in their own team, with an explicit permission.

This document summarizes security capabilities available in Splashtop On-Prem. Capabilities vary by Gateway version and licence, and evolve over time — please confirm current details with your Splashtop representative before relying on a specific control for a compliance requirement. For full architectural detail, see the Splashtop On-Prem Products Security Overview; for current certifications, see splashtop.com/compliance.