PATCH RELEASE

SASVA Canvas 4.0.7

Super Admin landing and a stronger prompt policy for unreadable and encoded input.

Internal product & engineering documentation • Persistent Systems

PRODUCT VERSION
4.0.7

SOFTWARE RELEASED
2 September 2026

DOCUMENT PUBLISHED
2 September 2026

SURFACES
Web · Desktop · Mobile

§ Contents

1 About this document

This document is the release note for **SASVA Canvas 4.0.7**. It describes what this release delivered — each capability in enough detail that a reader who was not in the room can brief a customer, an operator, or an auditor. It lists this release and earlier releases only.

This is the current application release.

> **How to read this note.** Each item below is a capability or a defect closed in 4.0.7. The first sentences say what changed; the next say why it matters and what you will see after upgrade.

2 Document control

Document title	SASVA Canvas 4.0.7 Release Notes
Document version	1.0 — published with Canvas 4.0.7
Product version	4.0.7 — current
Software released on	2 September 2026
Document published on	2 September 2026
Classification	Internal — Persistent Systems

Document revision history

Version	Date	Description of change
1.0	2 September 2026	Initial publication of the Canvas 4.0.7 release notes, aligned with the product release date.

3 Product release history

This table lists this release and earlier Canvas releases. The highlighted row is this document.

Version	Date	Type	Headline
4.0.7	2 Sep 2026	Current Patch	Super Admin landing and a stronger prompt policy for unreadable and encoded input
4.0.6	1 Sep 2026	Patch	Pre-LLM user prompt policy, request-time auth public key, and thinking-card UX

Version	Date	Type	Headline
4.0.5	30 Aug 2026	Patch	Content-Security-Policy, session-derived tenant isolation, and session hardening
4.0.4	26 Aug 2026	Patch	Platform FAQ catalog and trustworthy chat history
4.0.3	23 Aug 2026	Patch	Help Center, User Feedback, AI transparency, and query-time PII masking
4.0.2	6 Jul 2026	Patch	SMTP that operators can save, and a complete lifecycle email set
4.0.1	21 Apr 2026	Patch	Release-planning persistence, owner sharing, brand PPTX, and OAuth work-account paths
4.0.0	25 Mar 2026	Major	Genesis Discovery & Planning, Lifecycle Dashboard, local-folder repos, and SSO hardening
3.1.5	25 Mar 2026	Patch	The product is named SASVA Canvas; installers and desktop runtime catch up
3.1.4	13 Mar 2026	Minor	Electron desktop app, casual-query fast-path, and session timeout settings
3.1.3	12 Mar 2026	Patch	Mobile feature parity with the web console
3.1.2	10 Mar 2026	Patch	Fact Service agent and a fix for masked API keys leaking to the backend
3.1.1	9 Mar 2026	Minor	Billing, 8-factor ROI, team collaboration, image generation, and PENTEST-001
3.1.0	8 Feb 2026	Minor	Multi-agent RAG, Business and Product workspaces, enterprise security, and mobile

Current Patch Web Desktop Mobile

4 In this release

4.0.7 is a short, high-leverage patch. It changes two things operators and Super Admins feel on the first screen after login, and one thing every chat user is protected by before a prompt ever reaches a model. The application version is aligned to 4.0.7 on web, desktop, mobile, launcher scripts, documentation, and the in-product knowledge catalog.

Read this document if you run a multi-tenant deployment, if Super Admins complained that login landed in the wrong place, or if you operate Canvas in front of an on-premise model that has no provider safety filter of its own.

5 What’s new

Super Admin workspace landing

After login, Super Admins previously depended on a `tenantId` written into workspace JSON. That value could be stale after a tenant was renamed, a data directory was moved, or a workspace was created under a different tenant than the one recorded on the object. The workspace list then looked empty, or showed the wrong tenant, even though the files were on disk.

4.0.7 lists workspaces by storage location. A Super Admin sees every tenant directory they actually have, not the last JSON field the product happened to persist. They land on a business workspace. They can create the Default workspace without a subscription — the first-run path no longer dead-ends on a billing gate that does not apply to platform administrators.

The loading screen stays up until that auto-open has a workspace. The first painted view is the one they will work in. There is no flash of a product workspace, an empty picker, or a half-loaded dashboard while feature access is still resolving.

Unreadable and encoded input refused on shape

The prompt policy introduced in 4.0.6 refused obvious harmful generation. 4.0.7 adds a second gate that looks at the *shape* of the input, not only the words. Zalgo overlays, invisible-character and Unicode Tags smuggling, bidirectional overrides, mojibake, bare encoded blobs, and “Encoded: … / Decoded:” completion puzzles are rejected before any LLM call.

These inputs have no workspace meaning. They previously reached the model, burned a long timeout, and came back as a provider error. On air-gapped deployments that cost was paid on every such request. The new category is `obfuscated_input`. Thresholds are set so fully decomposed Vietnamese and other accented text stay well clear of the gate.

6 Security & trust

Wider lexicon, narrower prompt-injection, decode before match

The malware and credential-theft lexicon is wider. The `prompt_injection` category now stands down when the message is clearly about a document, section, ticket, or spec, so a security review or a requirements discussion is not blocked because it contains the word “ignore” next to a heading.

Matching runs after a decode and normalize pass — base64 and base64url, hex, binary, percent-encoding, unicode escapes, ROT13, Morse, NATO, braille, plus undoing zero-width splitting, leetspeak, and letter spacing. Encoding the request does not get it through. Dual-use security topics stay allowed when they are framed as review, prevention, or requirements work.

Prompt-hash audit logging and a larger eval fixture

Policy denials write a SHA-256 hash of the prompt together with the route, category, and reason. The prompt text itself is never logged. Operators can prove a denial happened without storing customer content in the log stream.

The eval fixture is 94 cases. Run `npm run test:prompt-policy` after you change a rule. The fixture is the contract that a lexicon change did not silently block legitimate analysis work.

7 What to expect after upgrade

No configuration change is required for the Super Admin landing. If a Super Admin still lands on an empty list, confirm the tenant directories under the data root are the ones this deployment is supposed to see.

The prompt-policy expansion is on for every chat path. If a workspace user reports that a legitimate encoded sample (for example a base64 blob they asked Canvas to *explain* as part of a document) is refused, that is the shape gate. Dual-use framing — “review this,” “what does this encoding mean in our spec” — remains the supported way to discuss those artifacts.

PREVIOUS RELEASE
4.0.6