

Nasuni AI Activate: Frequently Asked Questions

Nasuni AI Activate connects an organization's AI applications to its existing file data and enforces permissions before AI sees any content. This FAQ answers the questions IT and infrastructure leaders ask most often, with added detail for security and technology leadership.

/ WHAT IS NASUNI AI ACTIVATE

What is Nasuni AI Activate?

Nasuni AI Activate connects any MCP-compatible AI application to the file data stored in a Nasuni environment. It has two parts: an indexing service that continuously reads and indexes file content, and a hosted MCP server that AI applications connect to at query time. Together, they let an AI application search and reason over an organization's own file content while enforcing existing permissions automatically.

How does AI Activate respond to a query from an AI application?

When a user asks a question through an AI application, the application calls the MCP server. The server checks what that user is permitted to see, gets only the authorized file content from the index, and hands it back to the application to generate an answer. Users work exactly as they do today; the AI application gains a new, permission-aware source of context.

Does AI Activate require building or managing new infrastructure?

No. AI Activate is a fully hosted, managed service. There is no server to deploy, no indexing pipeline to build, and no infrastructure to maintain. Turning it on is a configuration step in the Nasuni Portal, not a development project.

Does AI Activate create new copies of data or move it elsewhere?

No. Files stay where they are, governed exactly as they are today. AI Activate builds its index by reading from existing Nasuni volumes; it does not copy files into a separate AI platform or pipeline.

/ WHAT DOES IT ALLOW TEAMS TO DO

The value shows up in these kinds of work:

- **Find:** get a direct answer to a specific question, such as which document contains a particular clause or where a decision was recorded, without knowing which file holds it or what it's named.
- **Analyze:** compare information across many files at once to surface inconsistencies, recurring issues, or gaps that would take hours to find one document at a time.
- **Build:** draft new documents, such as summaries, reports, or first-pass responses, working from existing files instead of a blank page.

Most teams start with Find and grow into Analyze and Build as they get comfortable.

What does a typical query look like?

For example:

- *"Find the most recent version of the vendor services agreement with [Company] and summarize the termination terms."*
- *"Identify which open items across project files reference the same unresolved issue, and how frequently it recurs."*
- *"Draft a first-pass summary of this quarter's incident reports, using the format applied last quarter."*

The AI application does the reasoning; AI Activate makes sure it reasons over real, current files, with permissions intact.

Does AI Activate replace existing AI applications or tools?

No. AI Activate does not replace an AI application, a file system, or a workflow. It gives the AI application a team already uses, or is evaluating, a secure path to file data. Everything a team does today keeps working the same way.

How much difference does AI Activate make?

The difference shows up wherever teams spend real time hunting for information by hand: finding similar past projects, confirming the latest version of a document, or pulling together material scattered across many files. Most teams feel it first when they realize how much time they save.

/ WHAT DATA CAN IT ACTUALLY USE

Which file types are indexed and searchable?

Full content indexing covers common business file types: Microsoft Office documents (Word, Excel, PowerPoint), OpenDocument and Apple iWork formats, searchable PDFs, HTML and XML, email, plain text, and other common formats. AI Activate also indexes name and metadata of every files in selected scope regardless of format, and this information can be passed to your AI App, even if its the file content isn't indexed.

What happens with scanned PDFs or image-based documents?

Nasuni stores scanned PDFs and keeps them findable by name, but their content is searchable only if OCR ran on them. Without OCR, a scan turns up by name while its content stays not indexed.

Does AI Activate work with CAD, engineering drawings, or similar specialized formats?

An AI application can find and reference these files by name and metadata. Full content indexing of specialized engineering formats is not currently available.

Can specific folders or files be excluded from AI access?

Yes. An organization can scope indexing and AI access to specific folders within a volume, so it controls which parts of its file estate are exposed and does not have to open everything at once.

Can the AI reason over older versions of a file?

AI Activate reasons over the current version of a file. The underlying platform keeps version history, but the AI does not reason across earlier revisions.

What happens if AI Activate cannot find a file known to exist?

Usually one of a few things is happening: the file sits in a folder that has not been scoped for indexing; it was just added or changed and the index has not caught up; it is in a format Nasuni does not content-index, such as an un-OCR'd scan or an image; or the question needs to be more specific. Rephrasing with clearer terms, names, or context usually solves it.

/ WHICH AI APPS DOES IT WORK WITH

Which AI applications can connect to AI Activate?

Any MCP-compatible AI application connects through the same standard endpoint, including Anthropic® Claude, OpenAI® ChatGPT, and Microsoft® Copilot (Microsoft 365, Windows, and Copilot Studio). As more MCP-compatible applications reach the market, AI Activate works with them too, with no custom integration work.

How does AI Activate differ from an existing connector, such as a Microsoft Graph connector?

Built-in connectors usually handle a narrow set of file types, often Office documents only, and offer basic search. An AI application working through AI Activate gets deeper retrieval and reasoning against a purpose-built index across a much wider range of formats. And because AI Activate is vendor-neutral, that same access extends to any MCP-compatible AI application, not just one vendor's connector.

Can multiple AI applications use AI Activate simultaneously?

Yes. AI Activate is not tied to one AI application. Different teams can point different applications at the same index, and each user still sees only what they are permitted to see.

/ HOW IS DATA PROTECTED

Can the AI access files a user is not already permitted to see?

No. AI Activate enforces permissions before it retrieves any content, not after. It checks a user's existing Active Directory and file-level permissions at query time, so an AI application sees only what the user could already open directly. Off-limits content is never retrieved, processed, or shown to the AI model.

Why does that distinction matter?

Some approaches let the AI model retrieve everything, then filter what it shows the user. By then, the model has already processed content the user was never authorized to see. From a compliance standpoint, that matters no matter what finally appears on screen. AI Activate enforces permissions before retrieval, so that exposure never happens.

Who can access customer data and queries, including Nasuni staff?

Nasuni staff do not hold standing access to customer data or queries. A customer grants access only when it chooses to, for example to help resolve a specific issue.

Who controls the encryption keys?

Volume encryption keys stay customer-owned; turning on AI Activate lets Nasuni-hosted infrastructure use them but does not transfer ownership. For the search index, Nasuni encrypts it at rest under its own managed keys.

Does AI Activate introduce a new security surface or vendor risk?

No. AI Activate runs inside the Nasuni security posture a customer already has: the same vendor and the same permission model, now extended to AI. It is not a second system bolted on alongside the first. You can find more about Nasuni security at Nasuni Trust Center.

Does AI Activate meet existing compliance and certification requirements?

Nasuni's SOC 2 Type II coverage applies to AI Activate. For a compliance mapping against a specific requirement, ask the account team for the relevant documentation or check out Nasuni Trust Center.

Has AI Activate been independently security-tested?

Yes. Third-party penetration testing is part of Nasuni's ongoing security validation, and results are available on request.

Does the platform provide a record of what was accessed, by whom, and when?

AI Activate enforces permissions before it retrieves any content, and supports logging access activity. For specific audit or reporting requirements, talk with the account team.

/ GETTING STARTED

What is required to use AI Activate?

AI Activate runs on a current Nasuni platform, with file data in AWS- or Azure-backed storage and managed under Active Directory with SSO. Your Nasuni team can confirm the specifics in a short qualification call.

Is AI Activate available today?

Nasuni is rolling out AI Activate in phases. Your account team can confirm current availability and timing for your organization and region.

For an existing Nasuni customer, what is required to add AI Activate?

AI Activate is an add-on to an environment you already run, not a new platform. Contact your account team to confirm eligibility and next steps. In most cases it is a configuration exercise on infrastructure you already have in place.

For an organization not currently using Nasuni, how does onboarding begin?

Onboarding starts by consolidating the file estate onto the Nasuni platform, the same foundation AI Activate is built on. The account team can walk through what that looks like for your environment and how AI Activate fits the plan from day one.

Does using AI Activate require migrating data elsewhere?

No. Once files are on the Nasuni platform, AI Activate reads directly from existing volumes. There is no separate, AI-specific migration.

/ WHAT RESULTS SHOULD BE EXPECTED

What changes with AI Activate compared to current capability?

Without AI Activate, an AI application works from its training data, whatever a user types, and maybe a connected mailbox or a narrow connector. With AI Activate, the same application leverages factual, deductive reasoning based over the organization's own files, current and specific to the business, and answers from what is actually there instead of guessing at plausible-sounding generalities.

How quickly can results be expected?

With no infrastructure to build, most organizations are running real queries against real files within weeks, not months or quarters.

How does AI Activate compare to building an equivalent solution internally?

Teams that price out a DIY build usually find the permission-enforcement layer is the hard part. They must build, test, and maintain it alongside the indexing pipeline, and keep both current as the file estate changes. AI Activate already does both as a managed service, so most organizations save real development time and cost against an equivalent in-house build.

/ TO LEARN MORE

Visit nasuni.com or reach out to the account team.

nasuni

Nasuni is a leading unstructured data platform for enterprises where file data is mission-critical for both people and AI. We power the operational file layer where work happens — helping organizations manage, protect, and activate data so teams can work smarter, reduce costs, and operate securely without limits.

Built on a patented architecture that fuses cloud object storage with enterprise file services — including permissions, versioning, and a global namespace — Nasuni delivers high-performance file access, global data availability, and a scalable, governed, AI-ready single source of truth across every major cloud.

Trusted by more than 1,300 enterprises globally, Nasuni helps organizations modernize file infrastructure, strengthen data security, and support AI-driven operations. Learn more at www.nasuni.com.