

SecurITPosture

Customer guide

Frequently asked questions

What leaves the appliance, the AI, updates, support access, backups, and sizing.

Version 0.1.0

Generated 2026-10-16

Draft: steps are specified but not final.

Vesta Cyber Solutions. Made from the same source as the guide on the web, so the two always match.

Contents

The appliance	3
What leaves the appliance?	3
Where does my data live?	3
Who holds the keys?	3
Does Vesta harden the operating system?	3
The AI	3
Is the AI on?	3
Which model runs, and is it trained on our data?	3
Does the AI need a graphics card?	4
Can the AI approve anything?	4
Updates	4
How do updates arrive?	4
How do I know an update file is genuine?	4
How is an update applied?	4
Support access	4
Can Vesta get into our appliance?	4
Who are Vesta's support engineers?	4
Backups	4
How should we back up?	4
Are backups encrypted?	5
Sizing	5
What size should the virtual machine be?	5
Does the CPU type matter?	5
How large can files be?	5

Draft, pending the appliance build. Answers marked as not built yet describe decisions that are not in the appliance today.

The appliance

What leaves the appliance?

By default, nothing. The appliance's own firewall blocks every outbound connection, and every feature that would fetch from the internet is off. The CISA known exploited vulnerabilities catalog is imported as a file in Settings. The appliance never checks for updates and never downloads anything. Nothing in it calls home, and no agent or tunnel connects it to Vesta.

Integrations (AWS, Tenable, Jira, Microsoft Entra ID) are off until an administrator connects them. Once connected, the appliance talks only to the endpoints you configure. In the image as built today, the application has no route out of the virtual machine at all, so opening that path when an integration is connected is not built yet.

Where does my data live?

Inside the virtual machine, in its PostgreSQL 17 database and its encrypted file store. Vesta's own AWS is used for trials only; production always runs in your environment, and Vesta holds no customer data.

Who holds the keys?

The appliance makes its own secrets on its first boot: the database passwords, the session secret, and the field encryption key. They are readable by the root user only. Vesta never holds them. Keep a recovery copy of the field encryption key in a safe place outside the appliance, because a backup cannot be read without it.

Does Vesta harden the operating system?

No. Your organization hardens the operating system to its own baseline. Vesta is writing a short hardening guide that lists the settings the product needs, and boots each release once on a CIS Level 1 hardened copy of the image before it ships. The guide is not written yet.

The AI

Is the AI on?

Yes, by default in the appliance, because the model runs inside the virtual machine and nothing leaves it. A system administrator can switch it off in Settings, AI assistance. Switching it off always works. See [Working with the AI](#).

Which model runs, and is it trained on our data?

The plan is IBM Granite 4.1, an open source model; the final pick is confirmed after more testing. The appliance is planned to use the 8B size when the virtual machine has 16 GiB of memory or more, and the 3B size otherwise, and an administrator can switch. The size rule is decided but not built yet. The

model is never trained on your data.

Does the AI need a graphics card?

No. It runs on the CPU. When the virtual machine has a supported GPU, the appliance uses it.

Can the AI approve anything?

No. The AI drafts; people decide. Gates and the posture score count only what a person has accepted, and the package sign off refuses while any AI drafted part in the package is unaccepted. The AI never writes anything the assessor records or signs.

Updates

How do updates arrive?

As one file, tested and signed by Vesta, that you download from the Vesta support site and apply when you choose. Updates are never automatic. A new AI model arrives only inside an update, and the release notes say so. See [Getting updates](#).

How do I know an update file is genuine?

Check its SHA-256 checksum against the release page, and check its signature with Vesta's public key, on a trusted machine before you apply it. Vesta's release signing key is not created yet, so the signature step is not available yet.

How is an update applied?

With the written instructions that come with the release, after a backup or a snapshot. The way the appliance imports and applies an update file is not built yet.

Support access

Can Vesta get into our appliance?

Not without you. Vesta has no standing access. Your system administrator can grant a named Vesta engineer read only access for 1, 2, 4, or 8 hours, revoke it at any time, and see everything the engineer did in the audit log. Or you share your screen and drive. See [Support](#).

Who are Vesta's support engineers?

US citizens. Support is by email and phone during business hours. The grant form records the engineer's name and your case reference.

Backups

How should we back up?

The simplest choice is your platform's snapshots of the virtual machine (AWS, Azure, VMware, or Hyper-V). The product also has its own backup, which copies the database and every stored file into one archive with a manifest of SHA-256 hashes. A restore checks every entry first and changes nothing if any entry is missing, extra, or changed. See the Administrator guide.

Are backups encrypted?

Stored files are encrypted in the archive with the field encryption key, and the archive never contains a setting, secret, or key. The database is in the clear inside the archive, so keep archives on encrypted storage. Encrypted and scheduled backups inside the appliance are not built yet, and the appliance's own backup command is not written yet.

Sizing

What size should the virtual machine be?

Provisional until Vesta measures it: 4 virtual CPUs, 8 GiB of memory, and 100 GB of disk without AI; 8 virtual CPUs and 16 GiB of memory with the AI model on, which is the default. Vesta will publish a minimum and a recommended size with each release, and you choose the size when you create the virtual machine.

Does the CPU type matter?

Yes. In Vesta's tests, a current AWS m7a.xlarge (4 cores) answered AI jobs in seconds, and an older t3a.xlarge took about a minute per job. Drafting statements for a whole Moderate baseline is a background job of about 75 minutes on the 3B model and about 3 hours on the 8B on 4 cores.

How large can files be?

Uploads up to 1 GB. The appliance allows 30 minutes per upload, enough for 1 GB at about 5 Mbit/s. Scan files are limited to 100 MB each.