

SecurITPosture

Customer guide

Install on AWS

Size, import, first start, and what the appliance sends out, on AWS.

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

Before you start	3
What is inside the appliance	3
Size	4
Install	4
FAQ	5
What does the appliance send out?	5
Why can I not use my EC2 key pair to sign in?	5
Which AWS connection should the appliance use?	5
Can I put a load balancer in front of the appliance?	6
How large can uploads be?	6
Does the instance console keep the first password?	6
Troubleshooting	6
An integration cannot connect	6
The product says the server has no AWS instance role	6
Users cannot reach the site	6
The image import fails	7
The instance starts but never answers	7
Every change in the product is refused	7

Draft, pending the appliance build. The appliance image, its import steps, and the first start setup are specified but not built yet. Steps marked draft will change once the image is built and tested.

SecurITPosture is delivered as one virtual machine image (an OVA file). The same image runs on AWS, Azure, VMware, and Microsoft Hyper-V. This page covers AWS. As of 2026-10-15, the AWS install is the first one being proven end to end; the steps below will be corrected from that proof.

A Vesta deployment engineer installs the appliance with you, at your site and in your account. This page tells you what to prepare and what happens during the install.

Before you start

Prepare these before the Vesta engineer arrives:

- ☐ An AWS account and VPC where the appliance will run. Vesta's own AWS is for trials only; production runs in your account.
- ☐ A subnet for the appliance, and a static private IP address for it.
- ☐ A security group that allows inbound HTTPS (port 443) from your users, and inbound SSH (port 22) from your administrators only.
- ☐ An S3 bucket in the same region, to hold the OVA while AWS imports it.
- ☐ The VM Import service role, named `vmimport`, which AWS requires to import an image. Your engineer brings its trust and read policy.
- ☐ Someone with rights to import an image and launch an instance, available during the install.
- ☐ The instance size you will use (see Size below).
- ☐ The host name your users will type in the browser, and a DNS record for it pointing at the appliance's private IP address.
- ☐ Your own TLS certificate and key for that host name, issued by your certificate authority.
- ☐ Who will be the first administrator inside the product. They will set a password and enrol MFA at first start.
- ☐ Who will hold the operating system account, and their SSH public key. First start installs the key; SSH accepts no password.
- ☐ If the appliance will read this AWS account, an IAM role for the instance. See [Connect AWS with the instance role](#).
- ☐ A safe place, outside the appliance, for a recovery copy of the field encryption key.

What is inside the appliance

- Ubuntu 24.04 LTS Server.
- The SecurITPosture application.
- A PostgreSQL 17 database. Your data stays inside the virtual machine.
- Apache in front of the application, serving HTTPS. The current build still uses Caddy; Apache replaces it before release.

- A small open source AI model that runs inside the virtual machine. Nothing it reads leaves the virtual machine. The model ships as its own signed file beside the OVA and is not in the image yet.

Size

The size below is provisional. Vesta measures the appliance under a realistic load and will publish a minimum and a recommended size with each release. Until then, use:

- Without AI: 4 virtual CPUs, 8 GiB of memory, 100 GB of disk.
- With the AI model on (the default): 8 virtual CPUs and 16 GiB of memory, with the same disk.

The appliance picks the larger AI model when the instance has 16 GiB of memory or more, and the smaller one otherwise; an administrator can switch later. The AI model runs without a graphics card. If the instance has a supported GPU, the appliance uses it. A GPU is never required. The CPU generation matters: in Vesta's tests an m7a.xlarge answered AI jobs in seconds, and an older t3a.xlarge took about a minute per job.

Install

Draft, pending the appliance build. The first start setup described in steps 7 to 10 is decided but not built yet.

1. Download the OVA and its checksum file from the Vesta support site, and verify it. See [Getting updates](#).
2. Your engineer uploads the OVA to your S3 bucket.
3. Your engineer imports it with AWS VM Import, which reads it through the `vmimport` role and makes a machine image (AMI) in your account. The import takes a while; your engineer watches it to the end.
4. Launch an instance from that image, at the size you chose, in the subnet and security group you prepared, with the static private IP address.
5. Attach the instance role, if you prepared one, and set the instance to require IMDSv2 with a metadata hop limit of 2.
6. Start the instance. The appliance makes its own secrets on the 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.
7. First start asks for the network settings: the IP address, netmask, gateway, and DNS servers, or DHCP.
8. First start asks for the host name.
9. First start creates the product's first administrator, who sets a password and enrolls MFA.
10. First start takes the operating system account's SSH public key from a small disk labelled `SPSETUP` holding it as `ssh_authorized_key.pub`, which your Vesta engineer attaches before the first start. The prompt for it is on the virtual machine's first screen, which this platform's serial console does not show, so without that disk first start waits and the product does not start; a prompt on the serial console is planned. The disk may also carry `os_password.sha512`, a SHA-512 crypt hash of the account's password; without it, the password is asked on the first screen. The account `spadmin`

accepts an ssh-rsa key of 3072 bits or more, or an ecdsa-sha2-nistp256 or nistp384 key; an ssh-ed25519 key is accepted but stops working if FIPS mode (Ubuntu Pro) is turned on later. A key that does not qualify is refused with the reason, and first start asks again. First start then asks twice for a password for `spadmin`, at least 15 characters, used only for `sudo` and for signing in on the console to recover the virtual machine. SSH accepts the key and never a password, and never root.

11. Save the recovery copy of the field encryption key in your safe place. Stored files and integration secrets are encrypted with it, so a backup cannot be read without it. Vesta does not hold a copy.
12. A self signed certificate is created so first start always finishes. Install your own certificate for the host name.
13. Sign in as the first administrator and create user accounts. See the Administrator guide.
14. Harden the operating system to your own baseline. Vesta does not apply STIG hardening to the appliance; your organization hardens the operating system.
15. Decide how you will back up the appliance. AWS snapshots of the instance's volume are the simplest choice. The product also has its own backup. See the Administrator guide.

Until first start is built, the image's first boot does steps 6 and 12 by itself, takes its address from DHCP, and prints a generated password for the product's first administrator once on the instance console. The first sign in must change it. Open the product by its IP address until then, because the site answers for that address only.

FAQ

What does the appliance send out?

By default, nothing. 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. The AI model runs inside the virtual machine. Integrations (AWS, Tenable, Jira, Microsoft Entra ID) are off until an administrator connects them; once connected, the appliance talks only to that integration's endpoints, on port 443, and only from the application: an AWS connection reaches the AWS endpoints of the regions you chose, Tenable and Jira reach the host you entered, and Microsoft Entra ID reaches its sign in and Graph hosts. While any integration is connected the appliance also asks its own DNS servers for those names, at a limited rate. Removing the integration closes its endpoints. Addresses on your private networks stay closed unless your administrator allows the range (see the troubleshooting entry below), and Tenable and Jira must answer on port 443. Vesta support has no standing access and no connection into the appliance. See Support.

Why can I not use my EC2 key pair to sign in?

The appliance turns cloud-init off after it is built, so AWS does not install the key pair you choose at launch. First start installs the operating system account's SSH key instead.

Which AWS connection should the appliance use?

Inside your AWS account, the instance role. It stores no keys. See Connect AWS with the instance role.

Can I put a load balancer in front of the appliance?

Tell your Vesta engineer first. The application trusts exactly one proxy in front of it (`TRUSTED_PROXY_HOPS=1`), and a load balancer adds a second hop.

How large can uploads be?

Files 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.

Does the instance console keep the first password?

It can. AWS keeps the instance's console output, and today's first boot prints the first administrator's password there. That is why the first sign in must change it. First start, when built, sets the password instead of printing one.

Troubleshooting

An integration cannot connect

The appliance opens an integration's endpoints only while the integration is connected, and logs every address it opens, closes, or refuses.

What to check: sign in over SSH as `spadmin` and run `journalctl -t sp-egress --since today`. A line naming a private address means the range is not allowed: add it, as narrow as you can (ideally one address as a `/32`), to `/etc/securitposture/egress-private.allow` with `sudo`, owned by `root` and writable only by `root`. A Tenable or Jira address on a port other than 443 is refused. For the instance role, the metadata address opens only after `sudo install -o root -g root -m 0644 /dev/null /etc/securitposture/egress-imds`; see [Connect AWS with the instance role](#).

What to send Vesta: the `sp-egress` lines from the journal, the integration's type, and the exact message the product showed.

The product says the server has no AWS instance role

The product offers the instance role only when the instance metadata service answers within one second, using IMDSv2.

What to check: an instance role is attached; the instance requires IMDSv2; the metadata hop limit is 2, because the product runs in a container (`aws ec2 modify-instance-metadata-options --http-tokens required --http-put-response-hop-limit 2`). On the appliance the metadata address also stays closed until `root` creates `/etc/securitposture/egress-imds` (see [Connect AWS with the instance role](#)).

What to send Vesta: the instance ID, the output of `aws ec2 describe-instances --instance-ids <id> --query 'Reservations[0].Instances[0].MetadataOptions'`, the role's name, and the exact message the product showed.

Users cannot reach the site

What to check: the security group allows inbound 443 from the users' network; the subnet's network ACL allows it and the replies; the DNS record points at the instance's private IP address; the users

open the address the site answers for (the IP address until first start sets the host name).

What to send Vesta: the security group's inbound rules, the instance's private IP address, and what the browser shows.

The image import fails

What to check: the `vmimport` role exists with that exact name, is trusted by `vmie.amazonaws.com` with the external ID `vmimport`, and may read the bucket; the bucket is in the region you import into; the OVA's checksum matched before upload. AWS checks the guest operating system against its supported list; Ubuntu 24.04 is expected to be on it, and this is confirmed at the first import.

What to send Vesta: the import task ID and the status message from `aws ec2 describe-import-image-tasks --import-task-ids <id>`.

The instance starts but never answers

What to check: the instance's status checks; the console output (EC2, Instances, Actions, Monitor and troubleshoot, Get system log) for lines starting [`firstboot`]; the security group.

What to send Vesta: the console output with any password line removed, the instance type, and the time you started it.

Every change in the product is refused

What to check: the browser's address is the one the site was set up for. The application compares the browser's origin with the site's own host name and refuses every change when they differ. If you changed the web server settings, keep `ProxyPreserveHost On`.

What to send Vesta: the address in the browser, the address you set at first start, and the version shown in Settings, then About.