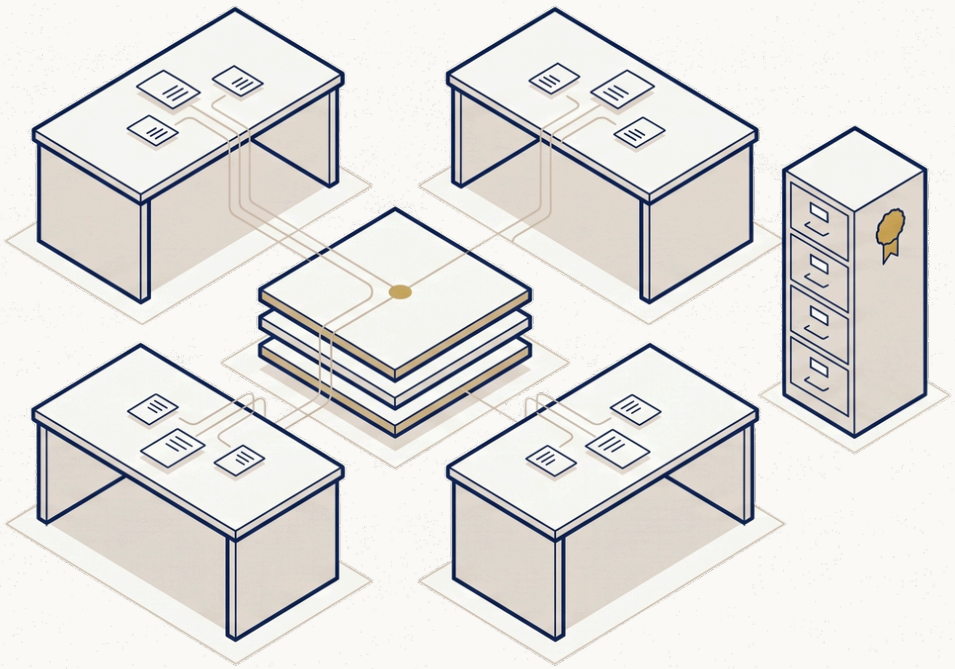
Adverant Nexus Alpha

For everyone



Adverant Limited

Release 1.0.0-alpha.1 · November 2026

CONTENTS

1.1 **What Nexus is, and who it's for** 5
What Adverant Nexus does for an organisation, who it is built for, and how its parts fit together, told in plain words with Harrow Port Authority as the example. It ends with the three editions in one paragraph.

1.2 **A working day with Nexus** 12
One working day at Harrow Port Authority, from the first look at Today to the last approval, showing what people actually do in Nexus and where each kind of work lives.

1.3 **Your organisation, its people, and what stays private** 20
How Nexus models an organisation and the roles its people hold, who may buy and claim on the Adverant account, and the line between what administrators can see and what stays private to each person.

1.4 **AI on your terms** 27
How Nexus uses artificial intelligence (AI) on the organisation's own terms. The organisation brings and owns its model keys, one switchboard routes every model call to providers it chose, local models run on its own servers, and nothing is sent to Adverant behind its back.

1.5 **Trust and compliance** 33
The parts of Nexus a compliance officer lives in. It covers the audit log and sealed records, approvals that need a second person and artificial intelligence (AI) decisions waiting for review. It ends with personal-data requests and the compliance pages that keep evidence ready before an auditor asks.

1.6 **Editions and licences in plain terms** 41
The three editions of Nexus and the 19 add-ons and domain packs on top, then how a licence works for a business. It covers the offline check and what each licence state means in practice, the promises every licence keeps, and moving from Community to Enterprise without reinstalling.

1.7 **Domain packs and verticals** 49
What a domain pack adds to Nexus for one industry, the ten packs Adverant offers, how a pack is switched on for one unit, and what a pack needs in your licence.

1.8 **Where Nexus runs** 59
The places Nexus can run, from Adverant's EU cloud to a site with no internet and devices in the field, what sovereignty means in practice, and how to choose.

1.9 **Getting started** 67
The two decisions before you start, what your IT team does, the first run as the first administrator sees it, and a checklist of what to have ready on day one.

1.10 **Away from the desk** 75
Using Nexus on a phone, from the console's phone layout in every edition to the native iPhone and Android apps. Covers what works without signal and how a company manages the phones.

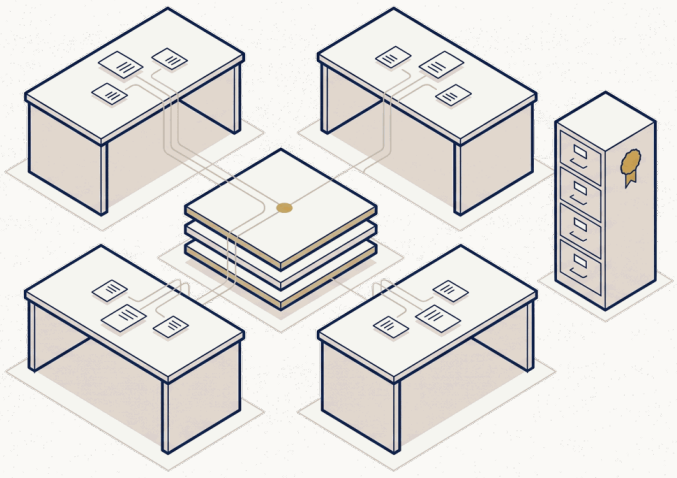
1.11 **Making it yours** 85
The ways an organisation shapes Nexus, from a single setting to a vertical of its own. Also how to brief developers or a partner, and what the platform guarantees whatever gets built.

1.12 **Questions people ask** 94
Plain answers to the questions buyers, managers and data protection officers ask most, including the uncomfortable ones about money and access to your data.

PART 1

For everyone

What Nexus does for an organisation, what it asks of the people who run it, and what they decide.



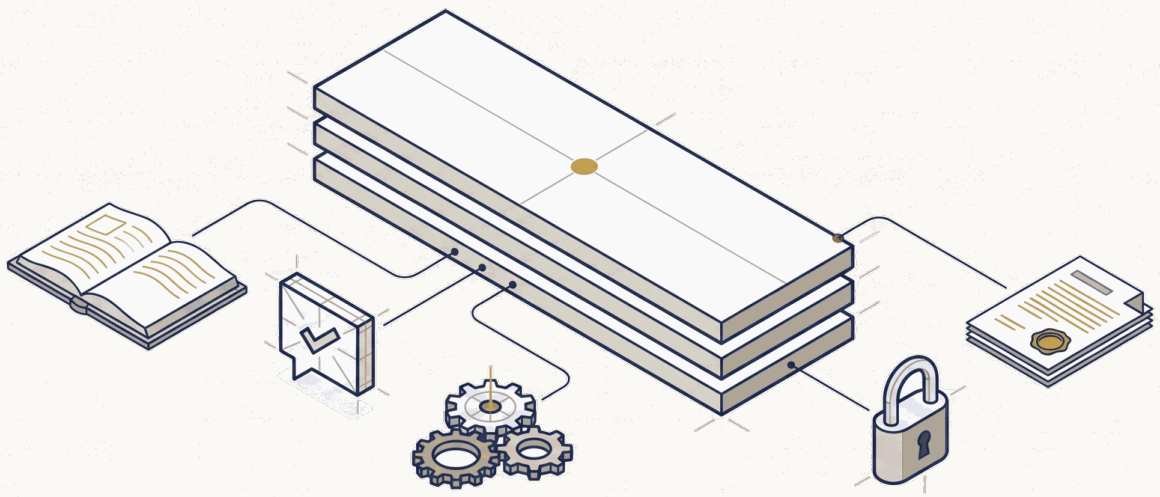
PART 1 · CHAPTER 1.1

What Nexus is, and who it's for

What Adverant Nexus does for an organisation, who it is built for, and how its parts fit together, told in plain words with Harrow Port Authority as the example. It ends with the three editions in one paragraph.

FOR EVERYONE

9 MIN READ



Ciara Holloway runs harbour operations at Harrow Port Authority. On a Tuesday morning she wants to know whether the dredging contractor's last survey changed the depth limits on the outer berths. A year ago that question meant three emails, a shared drive nobody trusted and a phone call to the one engineer who remembered. Now she types it into Nexus. The answer comes back in a few sentences, and each sentence points to the page of the survey and the clause of the contract it came from. If something was held back because she isn't allowed to see it, the answer says so.

None of it left Harrow's own servers. The port runs Nexus on its own hardware with no connection to the internet, and the model that wrote the answer runs there too.

That small exchange is most of what Nexus is for. Adverant's vision for it is short: every organisation thinking at its best. Its one-line summary is shorter still: bring your knowledge home.



A sovereign AI workspace

Nexus is a workspace for a whole organisation, built around artificial intelligence (AI) that the organisation controls. People ask questions and get answers drawn from what the organisation already knows. They talk in channels and threads, keep notebooks and hand routine jobs to workflows and agents. Every part shares the same people and the same access rules, and every part writes to the same record of what happened.

"Sovereign" is the word that sets it apart, and it means three practical things. Nexus runs where your rules say it may: in Adverant's cloud in the European Union, on your own Kubernetes cluster, on a site with no internet at all, or out in the field on a ship or a drone. It uses the models you choose, whether hosted ones you hold a key for or local ones on your own graphics cards. And the knowledge stays yours. It's read where it lives, it isn't copied off to someone else's cloud, and you can take all of it with you if you leave.



Who it's for

Nexus is built for regulated European organisations that must keep control of what they know. Insurers and banks are the obvious fit. So are public bodies and law firms, and ports like Harrow. They share a problem: sooner or later somebody will ask who allowed an AI to act, and on what evidence. It might be a supervisor or an auditor, or in a bad year a court. Most AI tools can't answer that question well, because the logs belong to the vendor and the data sits in the vendor's cloud.

Inside such an organisation, different people care about different parts. The executive who signs the budget wants to know AI costs won't run away and that the vendor can't hold the data hostage. The compliance officer wants every consequential decision to wait for a named person, and wants the evidence ready before the auditor asks. The head of IT wants something that runs on their own infrastructure without another fragile platform to nurse at three in the morning. And the people who use it every day want answers that show their sources, and an assurance that nobody is reading over their shoulder.

At Harrow, Aoife Brennan, head of IT, was the one who had to be convinced first. Her question was where the data would live. Her platform engineer, Tomás Reilly, wanted to know what happens when a model goes down. Both questions have plain answers, and you'll find them in chapter 1.4, AI on your terms, and chapter 1.8, Where Nexus runs.



The problem it solves

Most organisations already hold the answers their people are looking for. The trouble is where those answers sit: in mail, in shared drives, in a dozen line-of-business systems and in the heads of

people who are on leave. AI pilots promise to fix this, and many never get past the demo. The ones that do usually take one of two shapes.

The first is to rent a workspace from a large US vendor, bundled with office licences. It works, but the organisation's knowledge moves onto someone else's cloud and is tied to one vendor's models. For a regulated buyer whose procurement rules now ask where data lives and who can reach it, that's a hard sell.

The second is to build it yourself. A team of engineers joins search and chat to models, then adds access rules, then an audit trail, then budgets, project by project, for years. Every industry ends up rebuilding the same parts.

Nexus takes a third route. The parts every AI project needs around the model come ready: the knowledge, the access rules, the record, the routing of model calls and the budgets. It installs beside the systems an organisation already runs and reads them without changing them. Nothing has to be replaced to start, and a single team can go first.

How it's put together

Underneath, Nexus has five parts that work together. You don't need to know how they're built to use them, but it helps to know what each one promises.

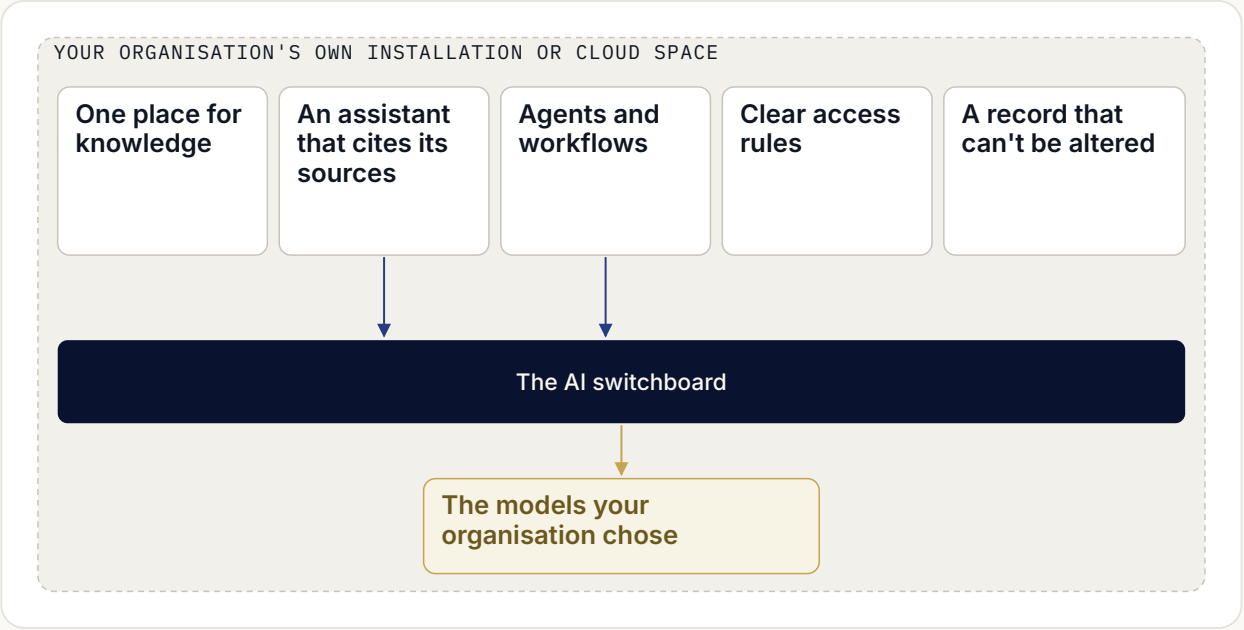


FIGURE The five parts of Nexus and the switchboard that every model call passes through.

One place for knowledge

Your documents and mail feed one knowledge graph, along with records from the systems you connect. Think of it as your own search engine for the business, except that it follows connections. A question about a contractor reaches its contracts and invoices, and the mail about them, even when none of them use the same words. Every file that comes in is opened in its own sealed machine first, so a dangerous attachment never touches anything else. And search shows each person only what they're allowed to open. Things they can't open are left out of both the results and the counts.

An assistant that cites its sources

Every answer lists what it was read with, like a bibliography, and each claim links to the page or record behind it. If the rules held something back, the answer says so in a plate beside it. The assistant says when it doesn't know instead of guessing.

Agents and workflows

Workflows start from things that happen in your business: a mail arrives, a hold is placed, a month ends. They run where you can watch them and stop for a named person at any step that needs a decision. Agents go a step further and work on their own, but every agent has an owner, a budget and an off switch, and none starts above "ask first".

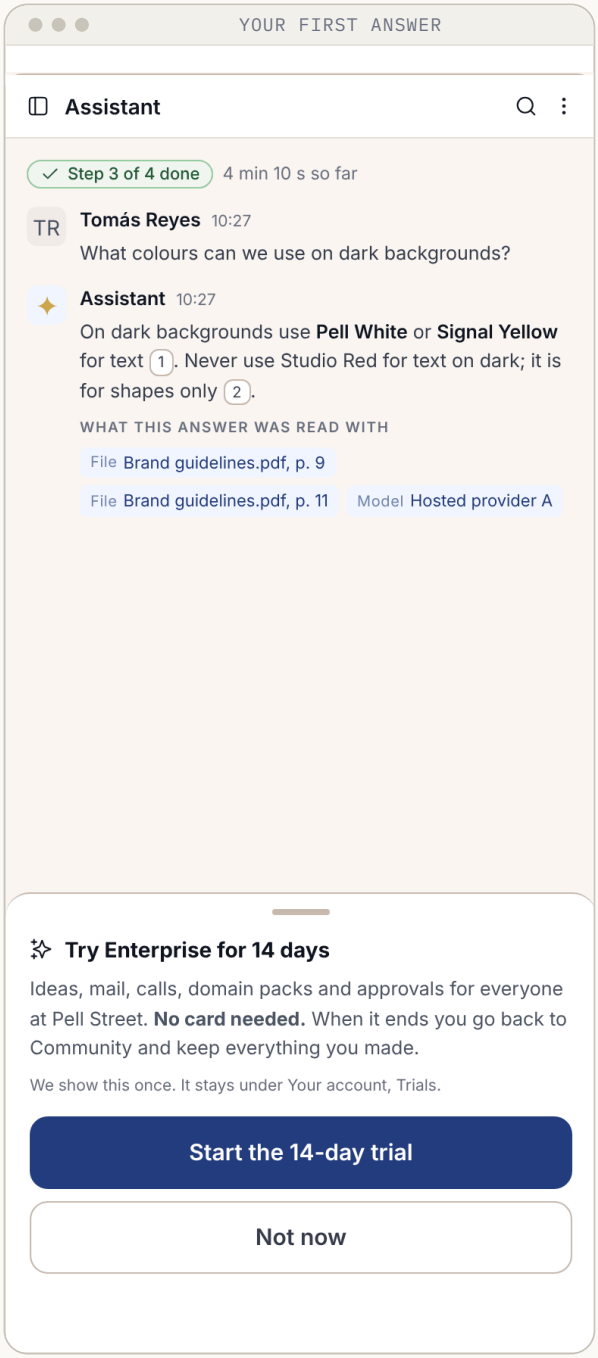
Clear access rules

You set who may see and do what once, for the organisation and its units, and every part of Nexus follows those rules, the AI included. Administrators govern without reading anyone's work. They see names, counts and states, never content. Chapter 1.3, Your organisation, its people, and what stays private, walks through this in detail.

A record that can't be altered

Every question and answer is written down as it happens, and so is every decision. Important decisions go into a sealed record with a time stamp from outside the installation. Change one character later and the check fails, and an auditor can run that check with a free verifier, no account needed.

Beside these sits the AI switchboard. Apps and agents never call a model directly. They ask the switchboard, which sends each call to a model and a place your rules allow, counts it against a budget someone owns, and never reads what was asked. Chapter 1.4, AI on your terms, is about it.



SCREEN A first answer from the assistant. Notice the citation to the exact page it came from, and the line at the end saying which files and which model it was read with.

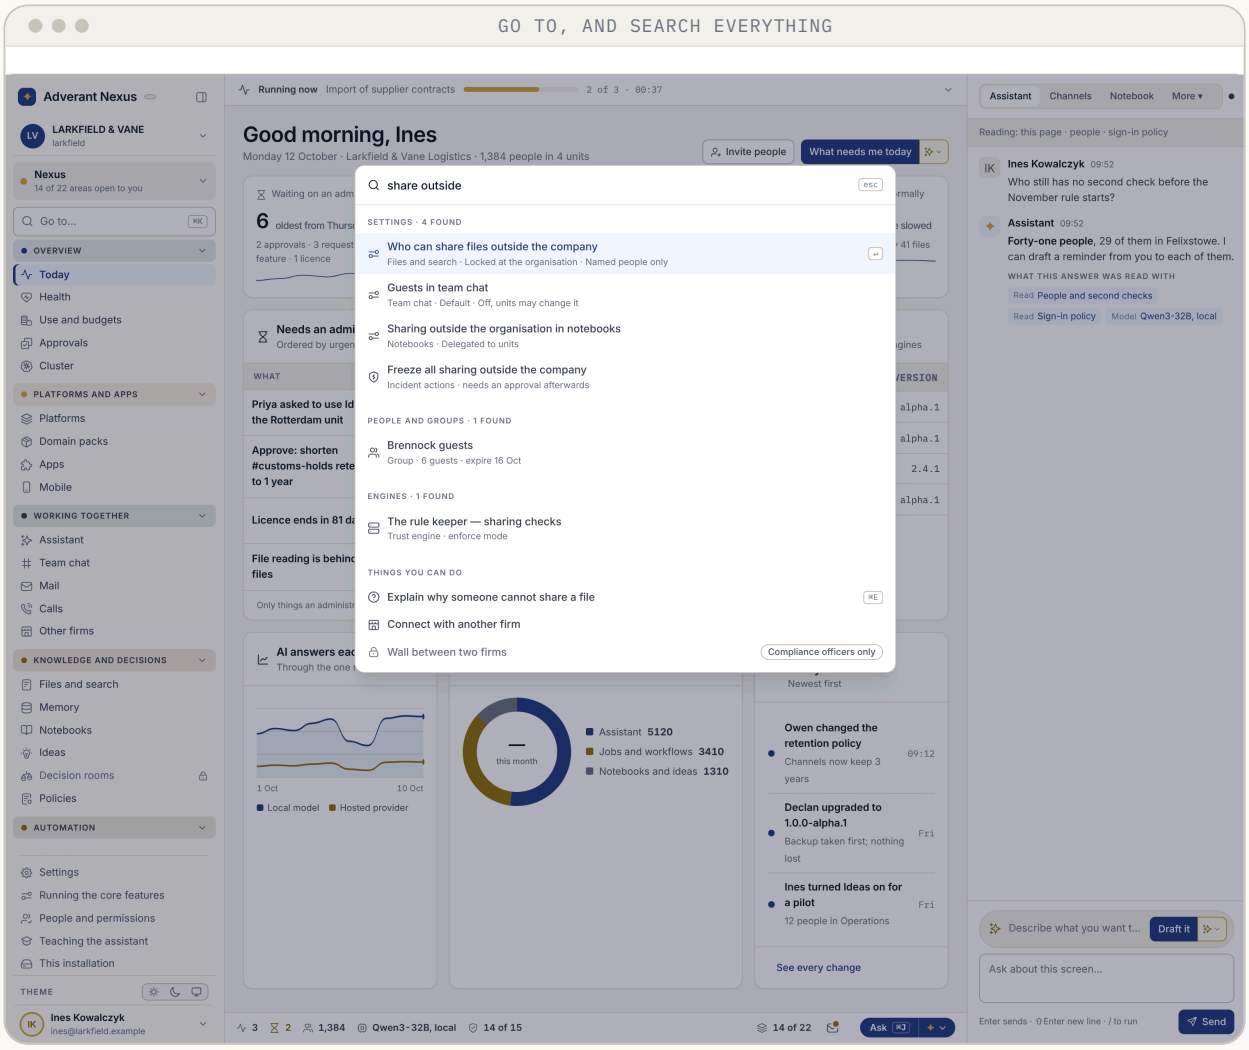
What it takes from you

A few things are worth knowing before you start.

It needs a model. An organisation that hasn't connected a model key or a local model gets no AI answers, and the screen tells you so plainly, with the way to connect one. There's no shared Adverant key quietly filling the gap.

It needs owners. Someone has to decide the access rules and budgets, and which steps should wait for a person. Nexus makes those decisions easy to set and easy to read back, but it can't make them for you.

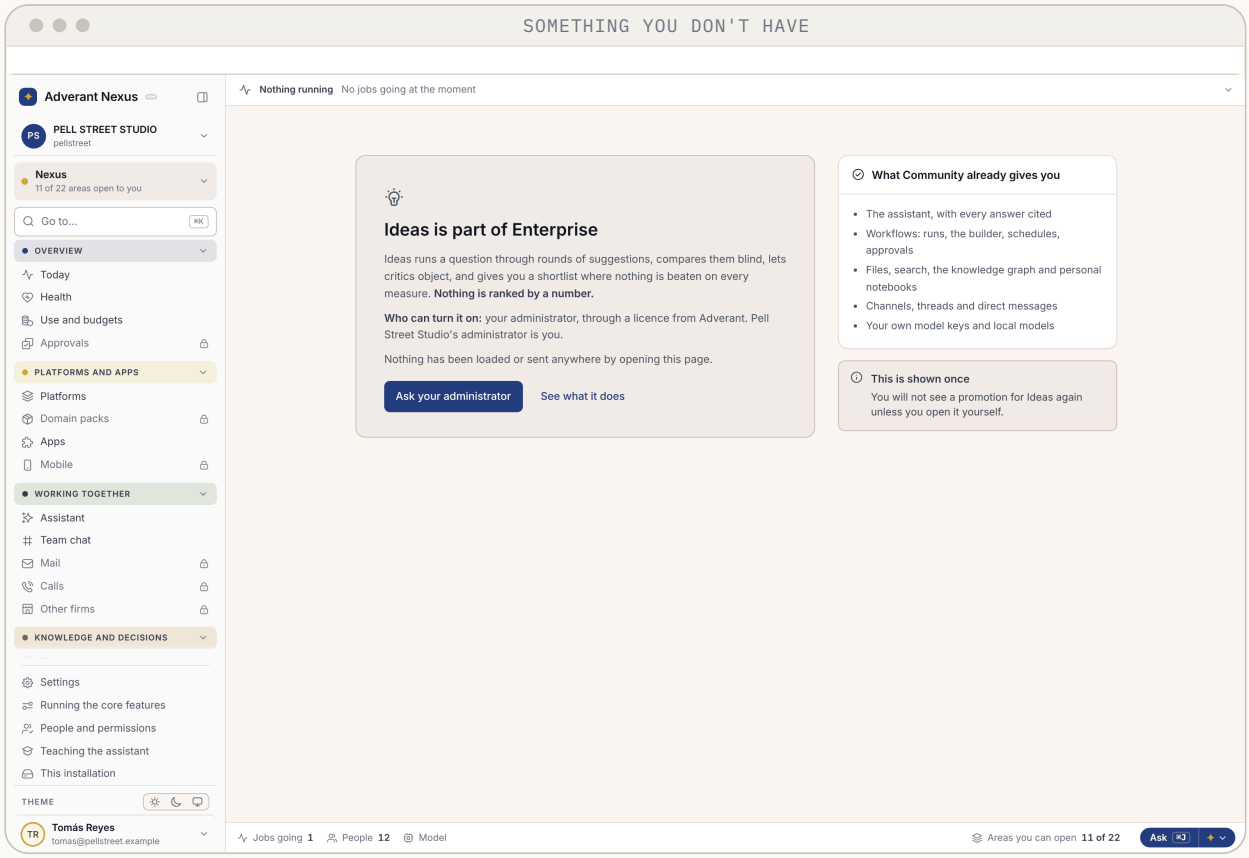
And some early choices are permanent. When Nexus is installed on your own infrastructure, a handful of settings are marked as fixed during set-up, the installation's name among them, because changing them later would break the records that depend on them. Your IT team will see them flagged before anyone signs in.



SCREEN Go to, and search everything. Results are grouped by kind, and anything the person may not open is missing from both the list and the counts, with no hint that it exists.

The editions in one paragraph

Nexus comes in three editions that share the same screens. **Community** is free and open source under the Apache-2.0 licence. It gives one organisation the assistant and search over its files, plus personal notebooks, basic channels and workflows, with a 30-day audit log. **Team** runs in Adverant's cloud and adds what a team shares: the work panel docked beside every page, for example, and shared notebooks. **Enterprise** is licensed. It brings business units and four-eyes approvals, enforced company sign-in, long audit retention and the compliance pages, along with the ways of running Nexus on your own infrastructure. Add-ons and industry domain packs sit on top of it. A signed licence decides what is switched on, never a different build. Chapter 1.6, Editions and licences in plain terms, explains the editions and how a licence works for a business.



SCREEN Something you don't have. An area outside your edition stays in the menu with a lock beside it. Opening it tells you what it does and who can turn it on, and it never shows a price.

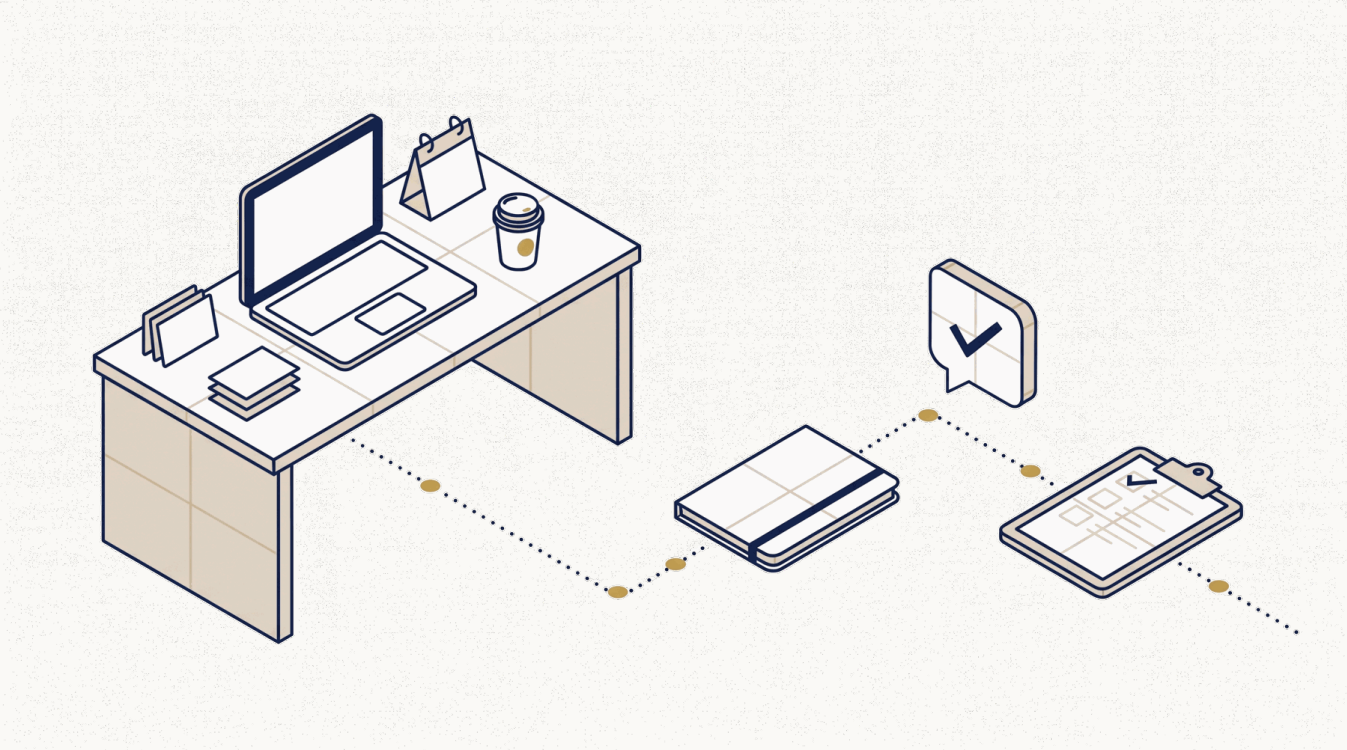
PART 1 · CHAPTER 1.2

A working day with Nexus

One working day at Harrow Port Authority, from the first look at Today to the last approval, showing what people actually do in Nexus and where each kind of work lives.

FOR EVERYONE

11 MIN READ



Ciara Holloway signs in at 07:40 with her passkey. She's harbour operations lead at Harrow Port Authority, which runs Nexus Enterprise on its own servers with no internet connection. The first page she sees is Today, and it's short: a mention in the berth planning channel, a workflow run waiting for her approval, and a question a colleague left for her in a thread. Three things. She can deal with all of them before the eight o'clock call.

This chapter follows Ciara through one ordinary day. It's the quickest way to see where things live in Nexus and how the pieces connect. Nothing in it needs technical knowledge, and almost all of it is the same in every edition. Where something belongs only to Enterprise or to an add-on, we say so.



The frame around everything

Every page in Nexus sits in the same frame, so it's worth knowing its parts before the day starts.

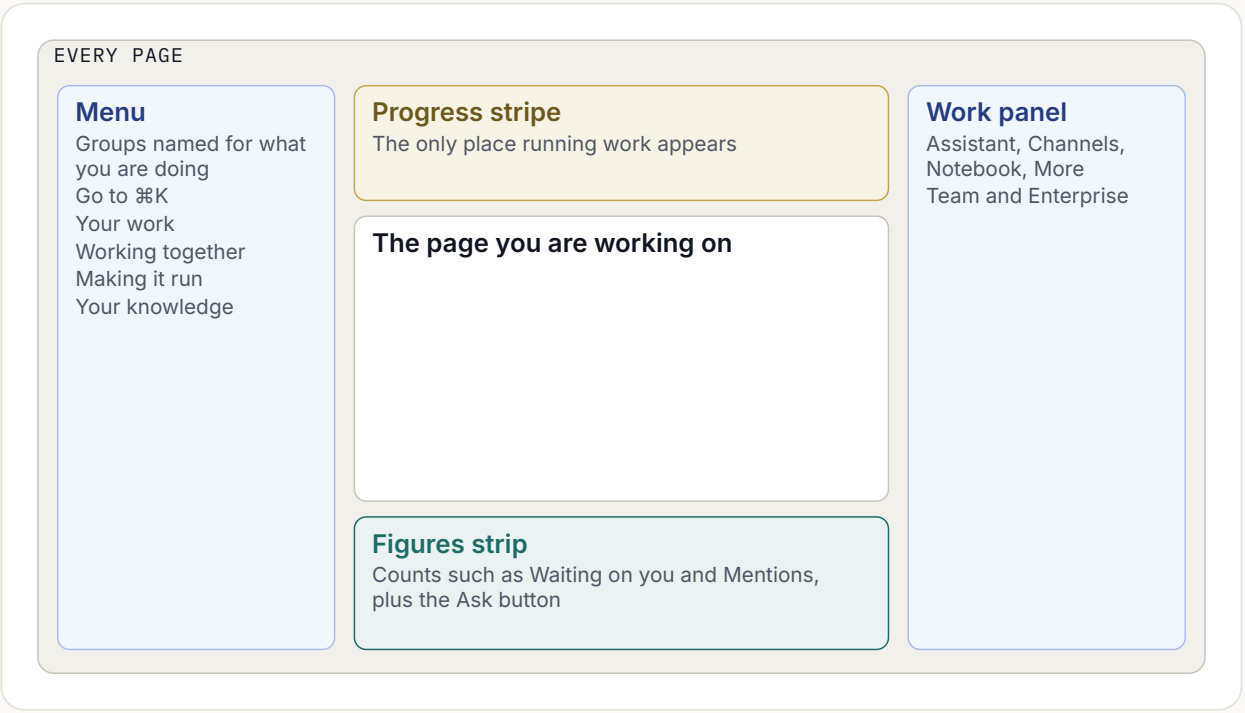


FIGURE The frame every page shares, with its five parts labelled.

The menu down the left is grouped by what you're doing rather than by product names, under headings such as Your work and Your knowledge. At the top of the page runs a thin progress stripe. It's the only place running work ever appears: an import, a workflow, a long answer being prepared. At the bottom is the figures strip, which counts things like runs going and items waiting on you, and names the model in use. On Team and Enterprise a work panel sits on the right of every page with the assistant and your channels in it. You can ask about the record in front of you without opening anything else. On Community the assistant and channels are pages of their own.

Two keyboard shortcuts are worth learning on day one. Go to (⌘K) jumps to any page or searches everything. Ask (⌘J) asks the assistant about the screen in front of you.

There are no pop-ups that appear and vanish. Nexus doesn't use them anywhere. Outcomes show up in place, as a changed row or a plate on the page, and running work stays on the stripe until it's done. It takes a day or two to stop looking for a toast in the corner of the screen.

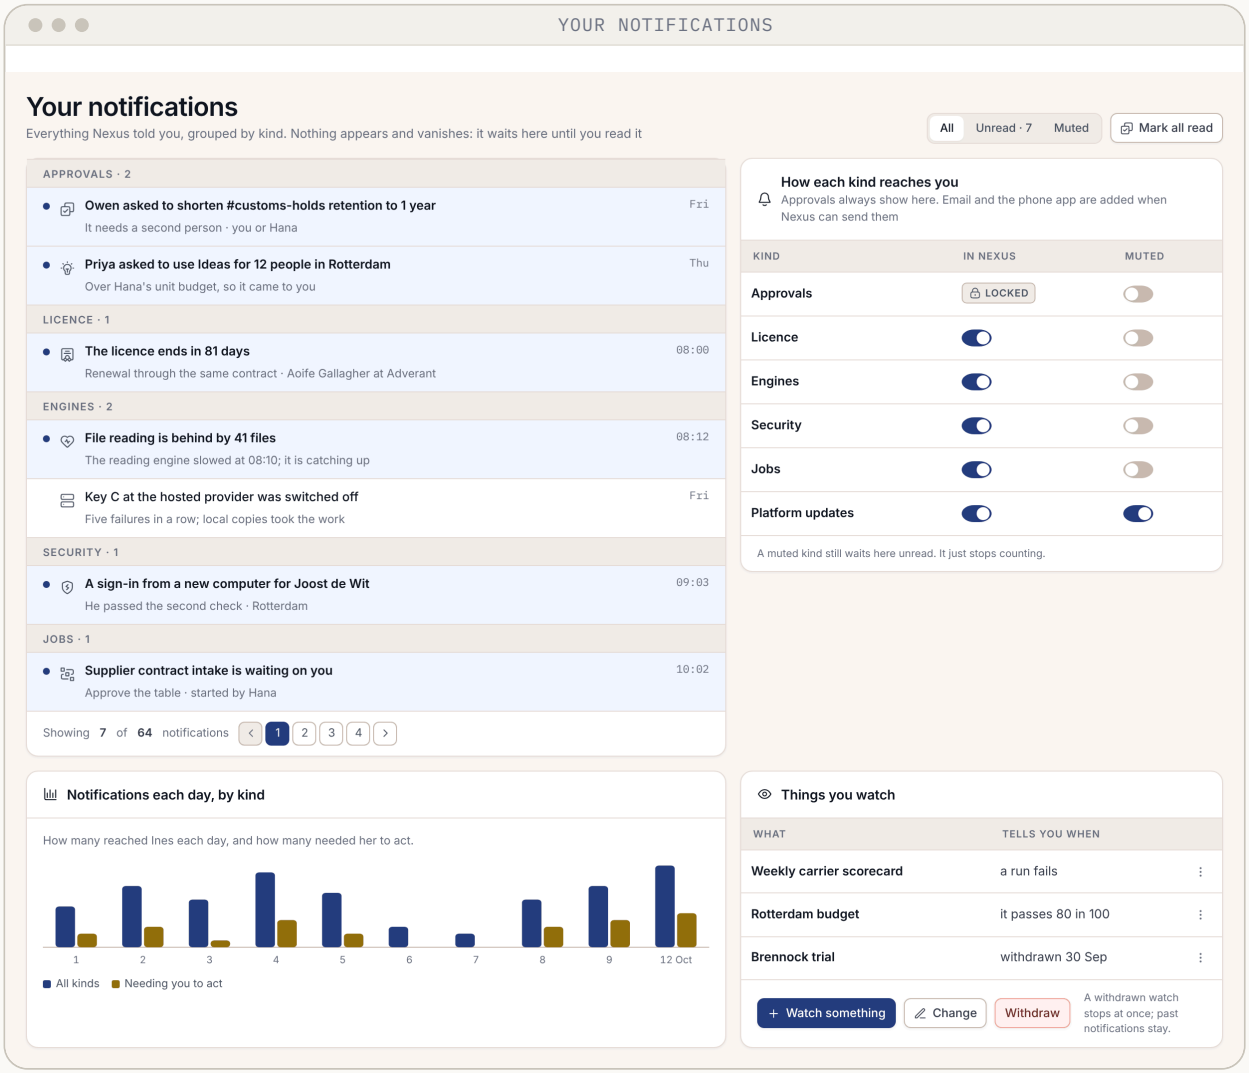


Today and notifications

Today is a list of what needs you, ordered by urgency with the reason in words. On Ciara's first day it was nearly empty, and each empty area said what would appear there and how to start ("Nothing

is waiting on you yet"). Now it holds her mentions, the runs waiting on her approval and anything someone has asked her to look at.

Behind the bell is a fuller inbox of notifications, grouped by kind, so approvals sit together and job alerts sit together. A notification waits there until she reads it. She can mute any kind except approvals, and a muted kind still waits unread; it just stops adding to the count. She can also have a kind sent by email or to her phone as well as shown in Nexus.



SCREEN The notifications inbox, grouped by kind. Each item says what happened and why it reached this person, and the panel on the right sets how each kind arrives. Approvals can't be muted.



What changed since you were last here

Ciara was off yesterday. Rather than scroll back through every channel, she opens What changed, which lists everything that moved since she last signed in, newest first and grouped by kind. A survey file was replaced with a corrected version. Two workflow runs finished. A colleague answered a question she'd left in a thread. It's the catch-up a good colleague would give you over coffee, minus the coffee.



Asking the assistant

At 08:20 Ciara opens the record for the outer berths and asks the assistant, in the panel beside it, whether last month's survey changed the depth limits. The answer arrives in four sentences. Each one carries a small citation, and clicking it opens the exact page of the survey beside the answer.

Underneath sits a line headed "What this answer was read with". It lists the files and the model the assistant used. If anything was taken out before the answer reached her, because the rules say she may not see it, a plate says so. Two people who ask the same question can get different answers this way, and each of them can see why without seeing what the other saw.

The assistant runs on the model Harrow chose, on Harrow's own hardware. If no model were connected, the assistant wouldn't pretend. It would show a plate saying it can't answer and why, with a button to connect one.

Typing `/` in the composer opens the skills palette: things the assistant can run, grouped by what they do, such as drafting a reply or summarising a long document. Skills that can't run here are shown greyed with the reason beside them. Many buttons around Nexus carry a small gold mark for artificial intelligence (AI). Right-clicking one shows "About this button": the skill it runs and what that skill reads, and who changed it last.



Team chat and channels

The berth planning channel is where Ciara's team works through the day. A channel holds messages and threads around a pinned record, and the assistant can be asked inside a conversation. A thread opens beside its channel, so one argument can be followed from start to finish without losing the rest. Direct messages are for one-to-one conversations, and the channel search finds a past answer with the messages around it.

Every AI answer in a channel shows what it was read with, just as it does in the panel. Private details stay private even in a shared channel.

Enterprise adds the controls an organisation needs around chat: retention and legal holds on channels, exports, and guests from other firms who see only the channels they were invited to and are marked clearly. If Harrow licensed the Calls and voice add-on, a huddle could start from the channel too, with a recording notice for everyone on it.



Files and search

Mid-morning, a contractor sends 40 files for the quay wall inspection. Ciara drops them in. The upload shows on the progress stripe, and each file joins the list as it's finished. Before anything reads a file, Nexus opens it in a throwaway sealed machine with no network, so a dangerous attachment never reaches anything else. Then the right reader handles it: text, scans and tables each go to a reader built for them.

The files page shows two halves of the picture: what Harrow has, and what it's still owed. Missing papers being chased sit next to the ones that arrived, with who was asked and when. A shared viewer opens any document with a column of the facts drawn from it alongside.

Search covers messages and files as well as records and runs, with results grouped by kind. Items Ciara can't open are left out of the results and out of the counts, with no line hinting they exist. For bigger questions the knowledge graph shows how things connect around one contractor or berth, and a timeline shows what happened when.



Notebooks

After lunch Ciara opens her notebook, "Quay wall inspection, autumn". A notebook holds the papers behind a piece of work and the questions she has asked of them. She ticks which papers it may read. The notebook points to those papers and never copies them.

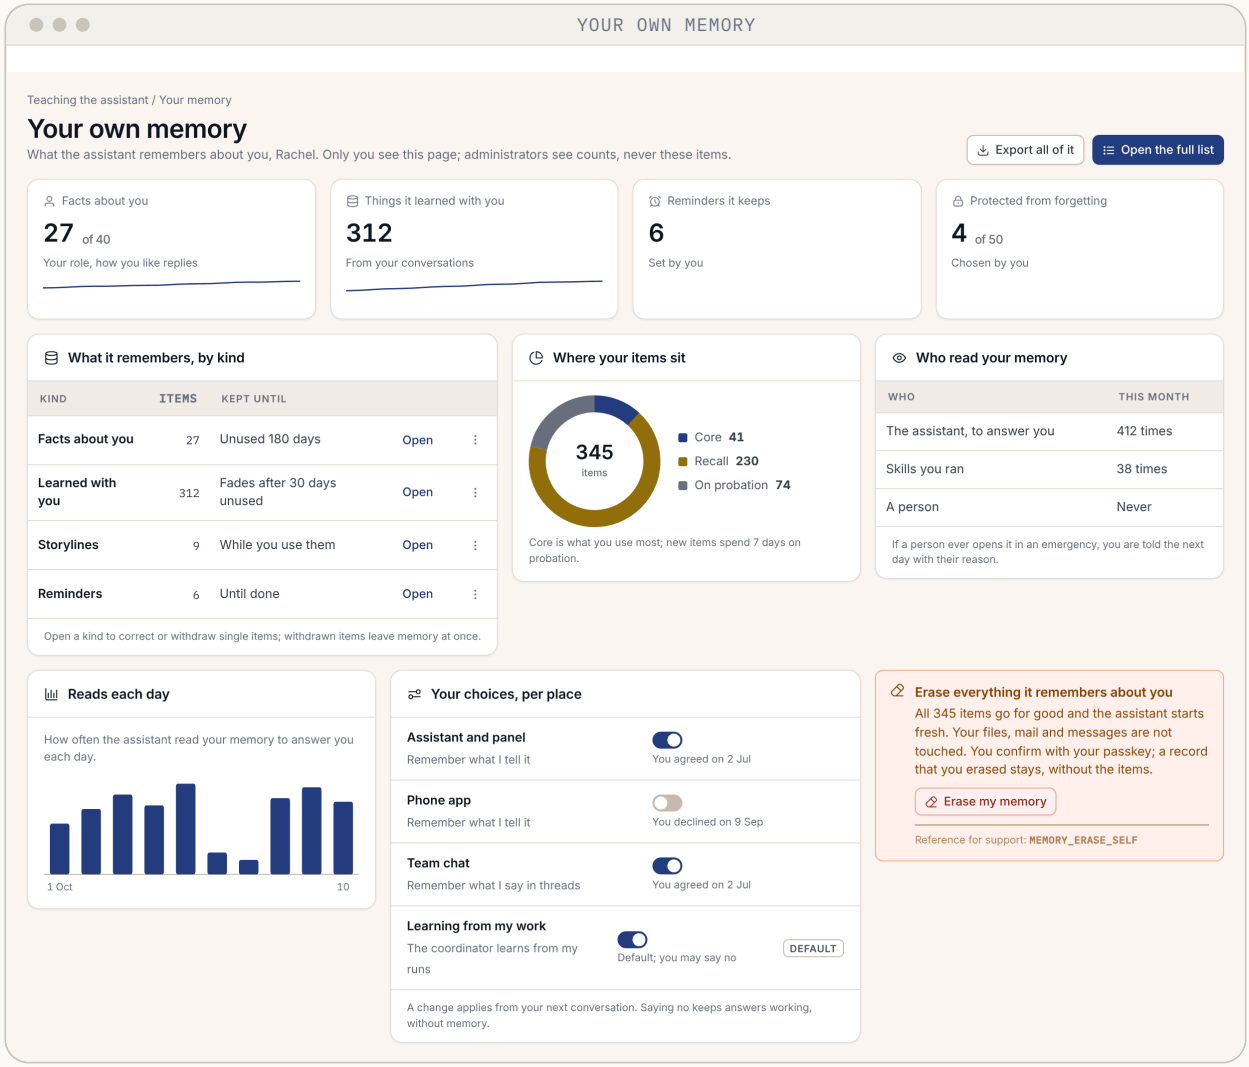
She asks whether the two inspection reports agree on the crack at the north corner. The answer cites a page on every sentence. When two papers disagree, both are quoted in their own words, and a compare view puts them side by side with the differences marked.

Personal notebooks are part of every edition. Sharing a notebook with read, ask or change rights comes with Team and Enterprise, and the notebook studio, which turns a notebook into slides, a one-page sheet or an audio summary, is an add-on.



Memory, and who controls it

The assistant remembers things Ciara has taught it: that she prefers short answers, which berths are hers, a reminder she set for Friday. She can see all of it on her own memory page under Teaching the assistant. It shows counts by kind, who read her memory this month (the assistant, to answer her, and skills she ran) and a choice for each place, such as the panel or team chat.



SCREEN A member's own memory page. Only she sees the items themselves; administrators see counts. Notice "Who read your memory", the choice for each place, and Erase, which she confirms with her passkey.

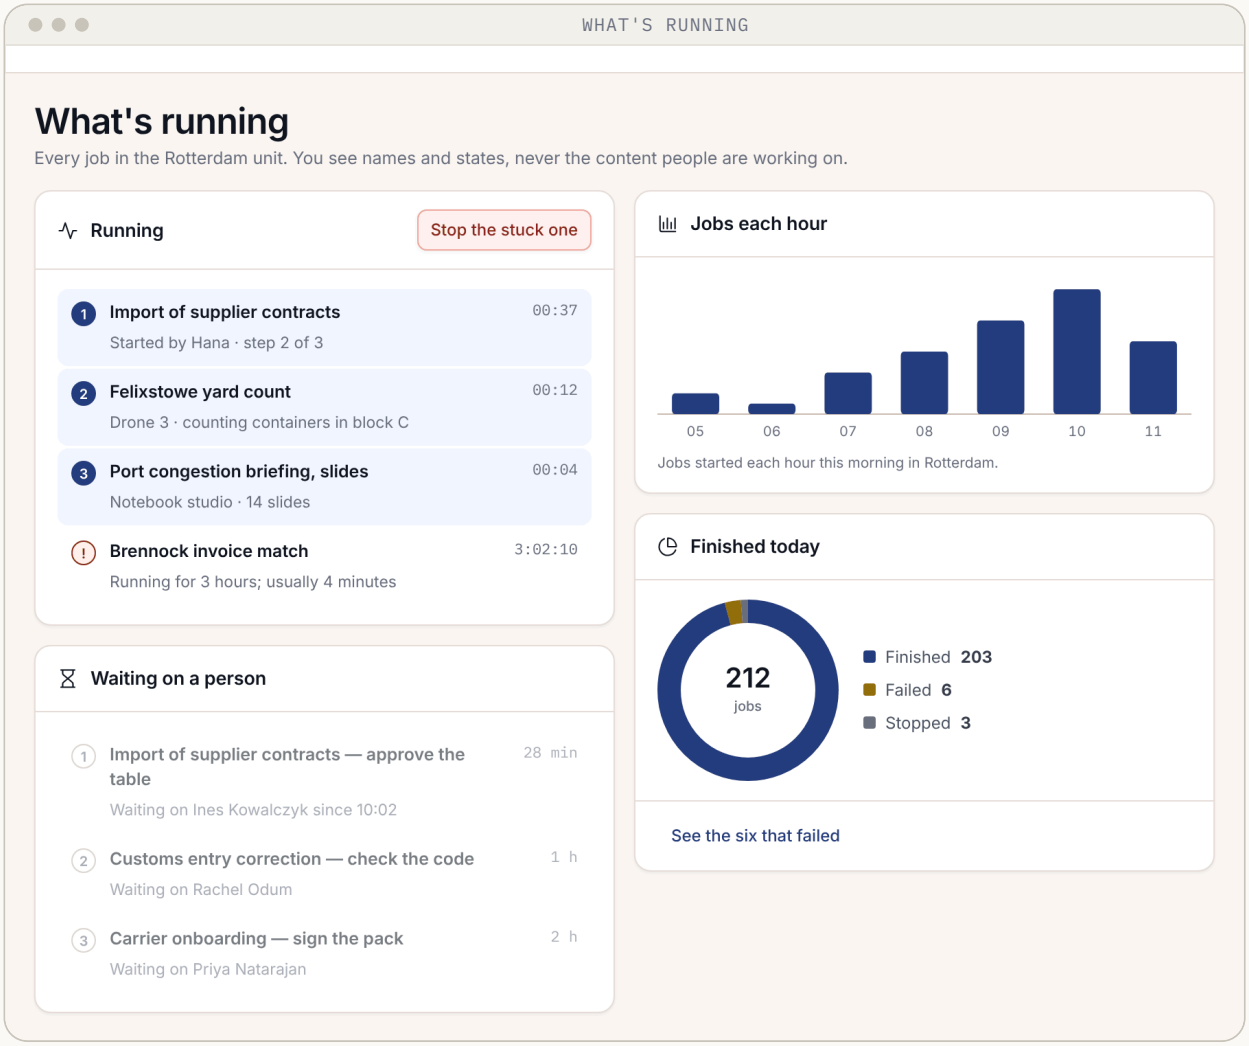
She can export everything it keeps or erase it all with her passkey, and a record that she erased it stays behind without the items. Saying no to memory in a place keeps answers working there, just without memory. If anyone ever opened her memory in an emergency, she would be told the next day, with their reason.

Agents and workflows

Harrow's operations team runs a handful of workflows. One builds the weekly berth schedule check from the port's booking records every Monday at 07:00, and the schedule page says exactly that, in words. The runs list shows what ran and who started it, with its state and how long it took. Opening one run shows each step in turn, and a failed run explains itself on a plate: what failed and why, and what to do next.

Workflows are put together on screen as steps on a canvas, each with its settings in a card beside it. Some steps are marked to wait for a person. Agents work in a similar way, but they keep going on their own within limits their owner set. Every agent has a named owner and a budget, and none starts above "ask first". Ciara has two agents of her own, and the My agents page shows what each did today, step by step.

The What's running page, opened from the progress stripe, lists everything running now, what's waiting on a person and what finished today. Ciara can approve or stop her own jobs from there.

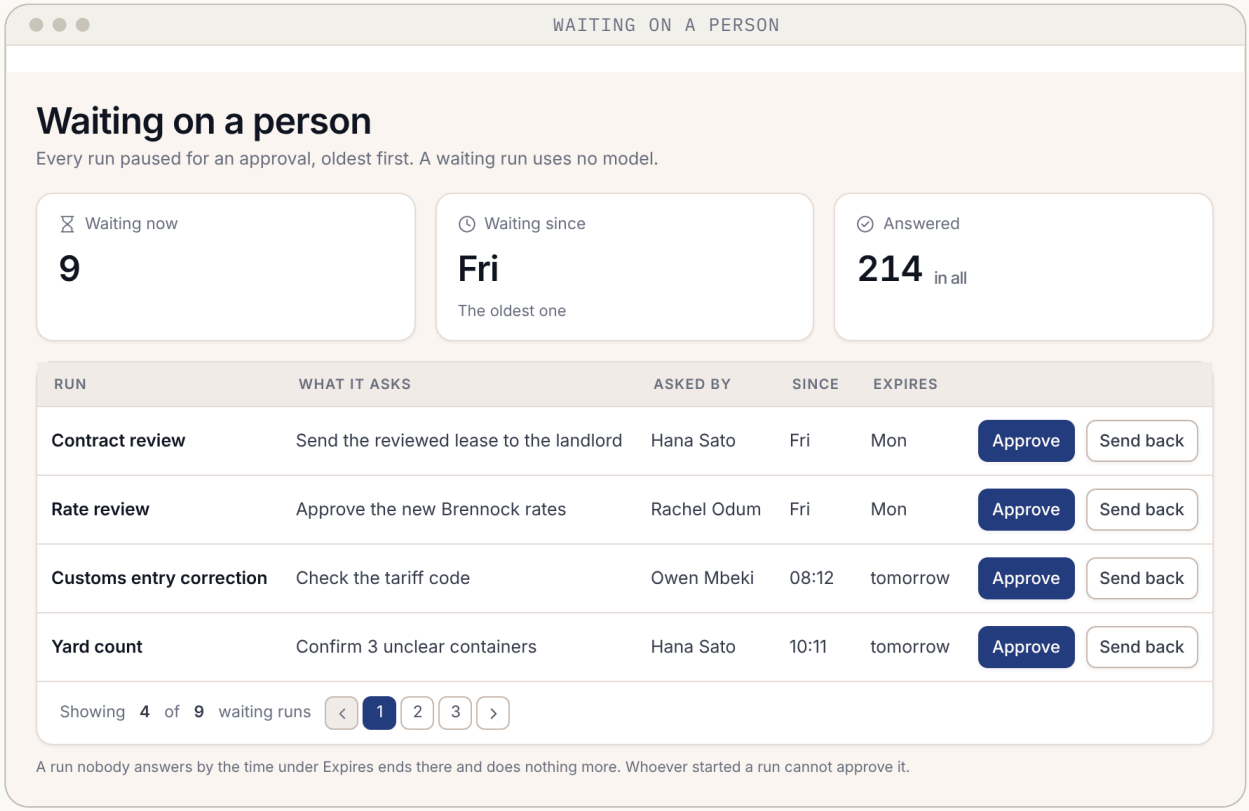


SCREEN What's running, opened from the progress stripe. Running and waiting work share one list with what finished today, and a stuck job can be stopped in place. An administrator sees what is running here but never its content.

Waiting on a person

At 15:30 the berth schedule run pauses at the step marked for her. Waiting on a person lists every task waiting for Ciara, oldest first. Each one says what it asks and who asked, and how long it has waited. She reads the proposed schedule and approves it. The run carries on. Had she sent it back, she'd have needed to write a reason.

Nothing waits silently. Rules can alert the right person when a task has waited too long and hand it to another approver if nobody answers. And nobody approves a run they started themselves.



SCREEN Waiting on a person. Every waitpoint is listed oldest first with who asked and since when. Approving or sending back is done by someone other than the person who started the run.

Away from the desk

Ciara leaves at five. On the bus she checks Today on her phone, where the web console fits a phone screen and the panel slides up from the bottom. If the connection drops, a plain plate says when it was lost and what still works. The native phone apps, which add approvals with a face or fingerprint check, are covered in chapter 1.10, Away from the desk.

The day had no dashboards to tend and no separate AI tool to copy things into. It was her work, in one frame, with the sources showing.

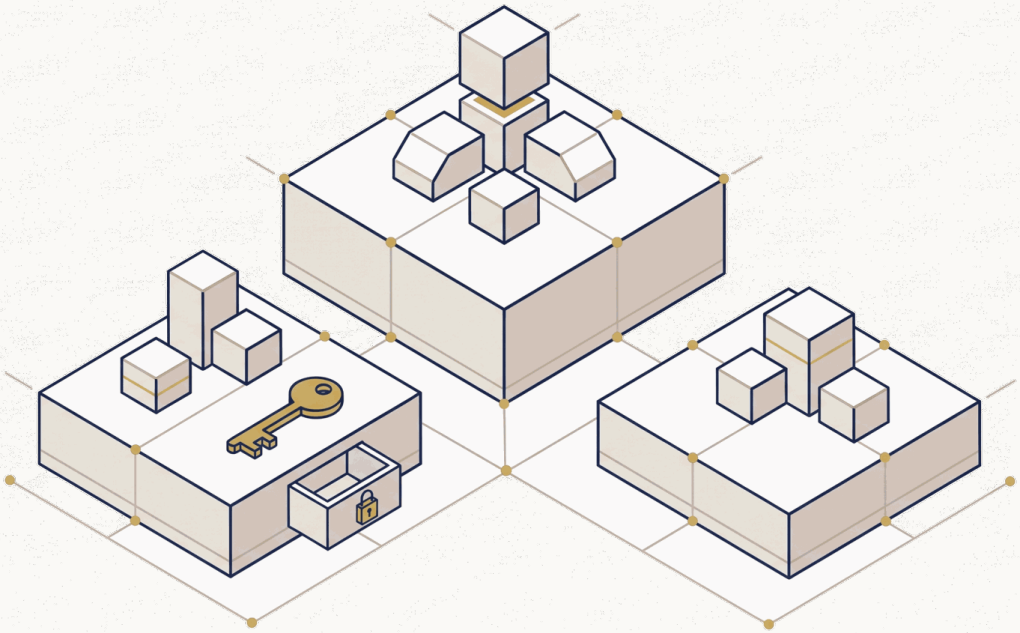
PART 1 · CHAPTER 1.3

Your organisation, its people, and what stays private

How Nexus models an organisation and the roles its people hold, who may buy and claim on the Adverant account, and the line between what administrators can see and what stays private to each person.

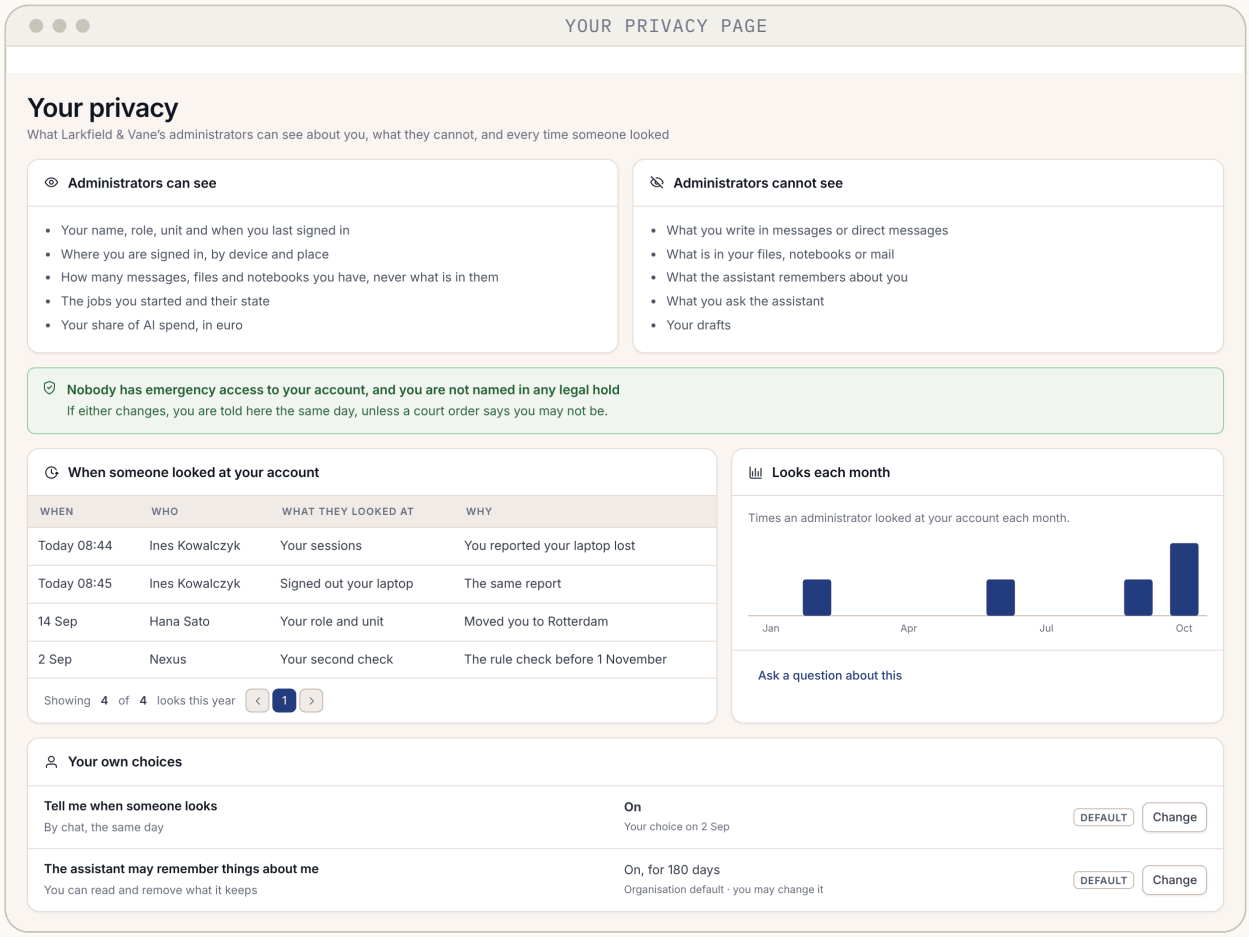
FOR EVERYONE

11 MIN READ



Aoife Brennan is head of IT at Harrow Port Authority and one of its two owners in Nexus. When Harrow first talked to Adverant, the question from the staff side wasn't about features at all. It was simpler: will IT be able to read our messages? Aoife opened her own privacy page in Nexus and turned the screen round so they could read it. Two columns. What administrators can see, and what they cannot.

That page is a good place to start this chapter, because it sets the tone for everything else in it. Nexus gives administrators real control over who may do what. It gives them almost no view of what people actually write.



SCREEN A member's privacy page. The left card lists what administrators can see, the right card what they cannot, and below them every time someone looked at this person's account and why.



Organisations and business units

Everything in Nexus belongs to an organisation. Its people and their files live inside it, along with its workflows and model keys, and nothing crosses to another organisation unless someone sets up a connection on purpose. Community holds one organisation per installation. Enterprise can hold several, and a person who belongs to more than one switches between them from the top of the menu.

Inside an organisation, Enterprise adds business units. A unit can be a terminal or a region. It can also be a function such as finance. Harrow has a unit for each of its two terminals and one for head office. Units matter because settings flow down through them. The organisation sets a value; a unit can narrow it for its own people but can never make it looser. A setting the organisation has locked shows who locked it, right beside the value, so nobody has to guess why they can't change it.

The same rule holds one level up. Your licence says what the organisation may use, and an administrator can narrow that, never widen it.

Groups cut across units. A group such as "Pilotage desk" might take people from both terminals and head office, and it can come from the company directory or be made by hand in Nexus. Groups are used wherever a list of people is needed: in settings and sharing rules, and in deciding who approves what.

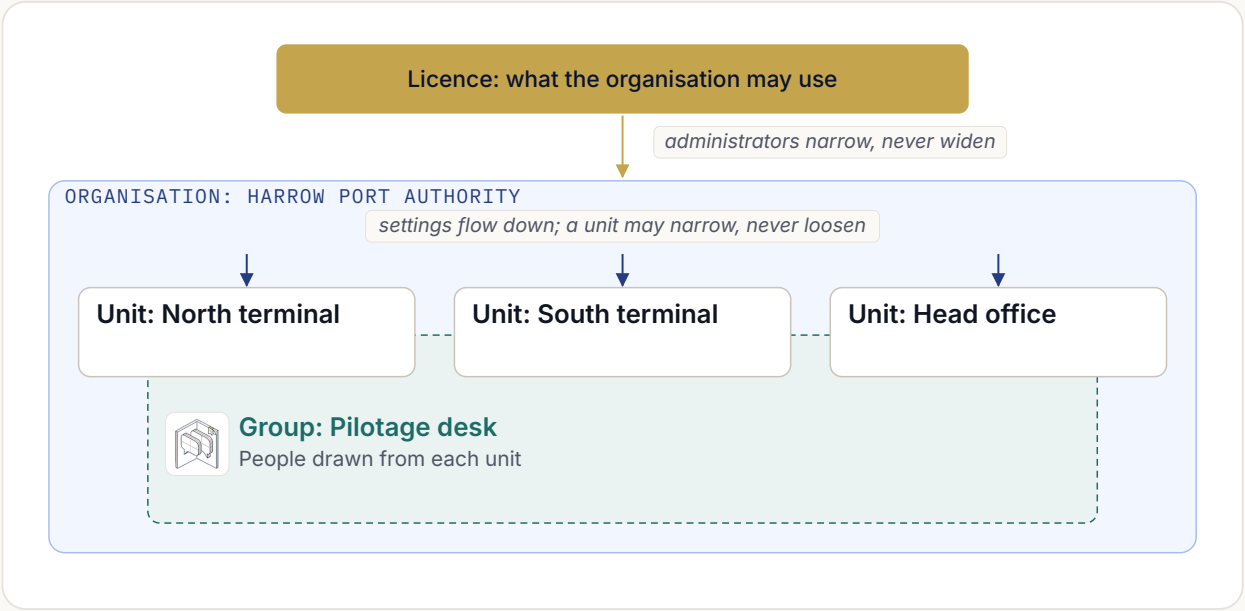


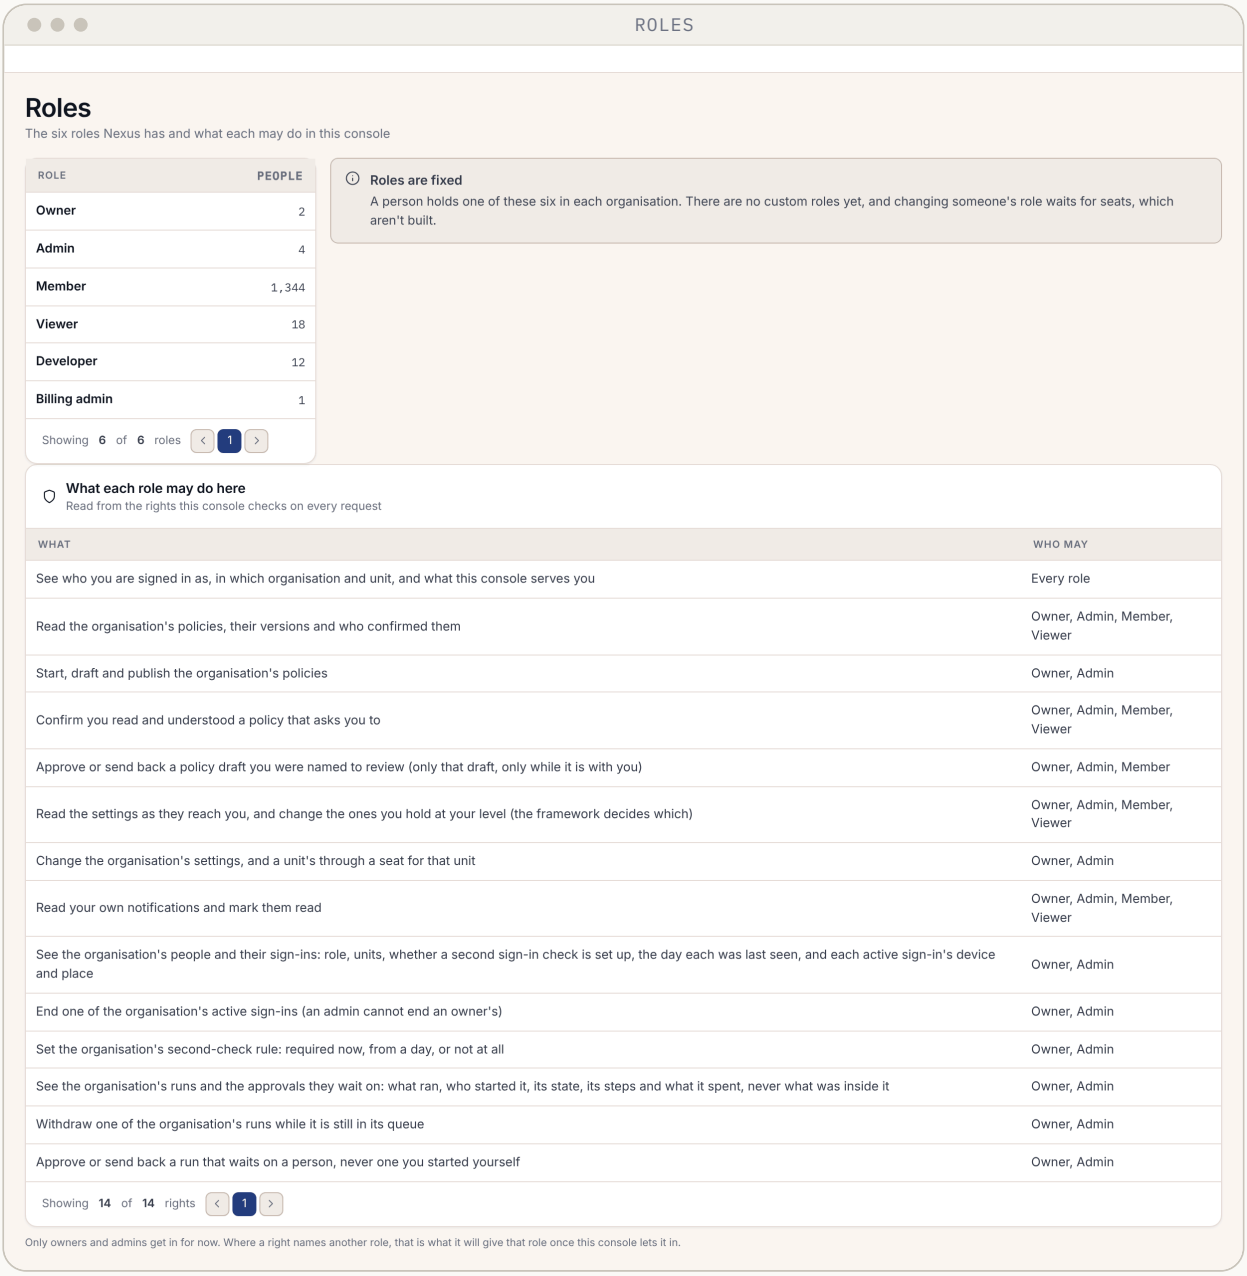
FIGURE How an organisation relates to its units and groups, and which way settings flow.

The six roles

Every person in an organisation holds one of six roles: Owner, Admin, Member, Viewer, Developer or Billing admin. Most people at Harrow are members, who do their daily work in Nexus. Viewers can read what reaches them and change nothing. Developers build on the platform. Owners and admins run the organisation: they see its people and sign-ins, set the rules for second checks and change its settings. They also see the runs and the approvals those runs wait on.

An owner holds every right by being an owner, and an admin can't end an owner's sign-in. Harrow keeps two owners, which is a sensible minimum: one person on holiday shouldn't leave the organisation without anyone at the top.

The Roles page lists every right in plain words, with the roles that hold it, read from the same register the system checks on every request. So the page can't drift from what the software actually enforces. In Community the six roles are fixed. Enterprise adds custom roles, built from the same list of rights, for organisations whose job titles don't fit six boxes.



SCREEN The six roles, how many people hold each, and every right written as a sentence with the roles that hold it.

A member who wanders into an administrator's page isn't shown a blank error. They see it read-only, every row greyed with "Only administrators can change this", and the page names their administrators and how to ask them.



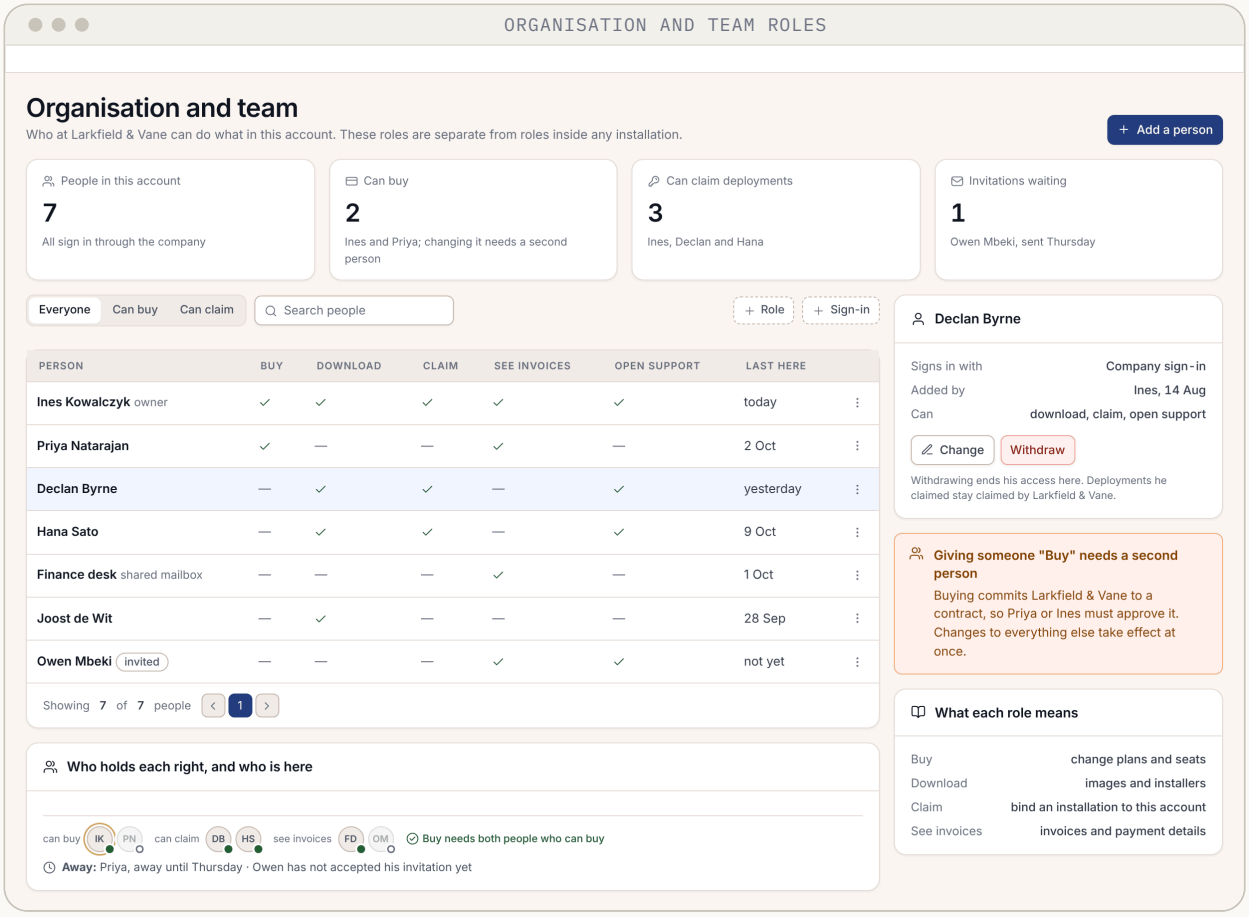
Rights on the Adverant account

There's a second set of permissions that people sometimes mix up with roles. Harrow also has an account with Adverant. That's where licences are requested and downloaded, and where installations are claimed. Who may act on that account is decided per person, apart from anyone's role inside the installation. There are five rights:

Right	What it allows
Buy	Change plans and seats
Download	Download images and installers

Claim	Bind an installation to this account
See invoices	See invoices and payment details
Open support	Open and follow support cases

Four of them take effect the moment they're given. Buy is different. Buying commits the organisation to a contract, so giving someone the right to buy waits for a second person who already holds it, and that person can't be the one who asked or the one it's for. Withdrawing a person ends their rights on the account, but any installation they claimed stays claimed by the organisation.



SCREEN The five rights on the Adverant account, decided per person. The plate on the right explains why giving someone Buy needs a second person.



Signing in, and second checks

People sign in with a passkey or through their company's own sign-in service, and Google and GitHub accounts work too. Community connects one company sign-in. Enterprise can require it for everyone and keep the people list in step with the company directory, so someone who leaves loses access at the next sync while their work stays with their unit. An administrator must test company sign-in with their own account before it can be required, and an emergency administrator sign-in stays available for the day it goes down. Every use of that emergency sign-in is recorded.

A second check is the extra step after the first: a passkey, or a six-digit code from an authenticator app, with backup codes shown once when it's set up. Owners and admins can require a second check for everyone from today or from a chosen day. When the rule starts, sign-in asks people to set one up and gives them a week of grace; after that, skipping it is refused. The Second checks

page lists who isn't ready, so nobody has to find out at the door. Administrators themselves never sign in with a password alone.

The Sessions page shows each person's sign-ins with the device and place of each. When a member reports a laptop lost, an administrator signs that one device out and leaves the phone alone. Signing everyone out at once needs a second person.



What an administrator can and cannot see

Here is the line, as it appears on every member's privacy page.

Administrators can see	Administrators cannot see
Your name, role and unit, and when you last signed in	What you write in messages or direct messages
Where you are signed in, by device and place	What is in your files or notebooks
How many messages, files and notebooks you have	What is in your mail
The jobs you started and their state	What the assistant remembers about you
Your share of spend on artificial intelligence (AI)	What you ask the assistant, and your drafts

The short version is names, counts and states, never content. The same principle runs through every administrator page. A run shows what ran and who started it. It shows the state and the spend too, but never what went in or came out. The notebooks directory lists each notebook by its owner and size. The memory pages show how many facts the assistant keeps about each person (at most 40), never the facts.

When an administrator genuinely needs to see what a member sees, to help with a problem, there's a way that leaves a trace. "See it as they do" opens a read-only view of the console as that person sees it. It lasts 30 minutes and needs a second check and a written reason first. The member is told, a banner runs across the top the whole time, and their private items stay hidden. For the commoner question, why can't Tomás open this, the Explain access page walks the chain step by step and names the step that blocks it.



Personal things stay private, even inside the organisation

Some things belong to a person, not to the organisation that employs them. Their mailbox and notebooks are theirs, and so are their direct messages and whatever the assistant remembers about them. Nexus keeps those private even from the organisation's own administrators.

This isn't a matter of a hidden menu item. If an administrator tries to open a colleague's mailbox, Nexus refuses with the code `PRIVATE_TO_ITS_OWNER` and points to the proper route for a legal case instead. Private items are encrypted under a personal key that belongs to that person, and erasing that key makes those private items unreadable everywhere, backups included.

Emergency access exists, because emergencies do. An administrator can open someone's memory in a crisis, giving a reason, but each administrator has a soft limit of 10 such uses a quarter and a hard limit of 25, past which Nexus refuses. Every use is recorded, and the person is told the next day with the reason. On their privacy page, members can see what administrators can and can't see about them, along with every time someone looked.



WATCH OUT

The privacy line is a design limit, and it applies to you as an administrator. There is no switch that lets an administrator read a member's private messages or mailbox for convenience. A legal case goes through legal holds and the personal-data pages, which record who asked and why.

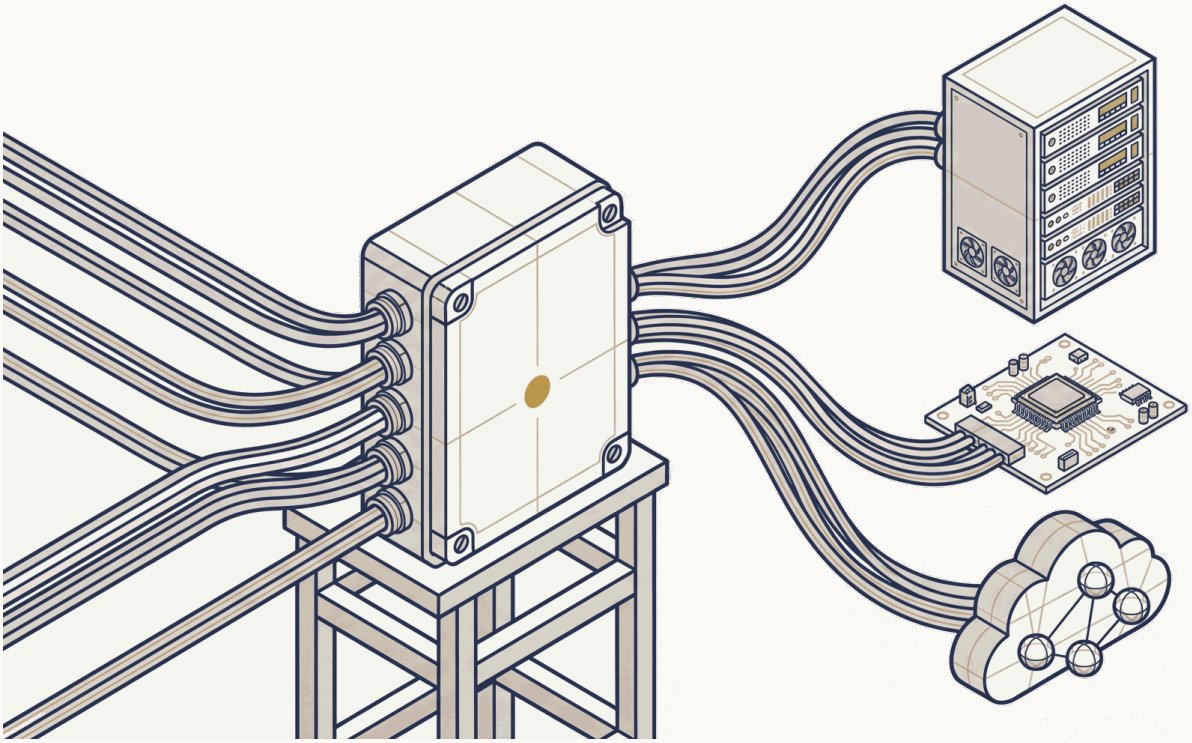
For Harrow's staff, the effect was what Aoife hoped for. People use the assistant for real work because they know nobody is reading over their shoulder. Chapter 1.5, Trust and compliance, covers the other side of the same coin: the record that proves who did what, for the auditor who needs it.

AI on your terms

How Nexus uses artificial intelligence (AI) on the organisation's own terms. The organisation brings and owns its model keys, one switchboard routes every model call to providers it chose, local models run on its own servers, and nothing is sent to Adverant behind its back.

FOR EVERYONE

10 MIN READ



Harrow Port Authority's operations installation has no hosted AI provider at all. It has no route to the internet, so every model it uses runs on graphics cards in its own machine room: a Qwen3-32B for most work, with a Llama 3.3 70B and a Gemma 3 27B beside it. When Aoife Brennan, Harrow's head of IT, opens the AI switchboard page, it lists exactly those models, which job each one does, and the place every call is allowed to go. That place is the building.

Most organisations will mix hosted and local models. The principle doesn't change: the organisation decides which artificial intelligence (AI) it uses, where its data may travel and how much it spends, and Nexus enforces those decisions on every call.



The key belongs to the organisation

When an organisation uses a hosted model, it brings its own key. Nexus keeps keys in the organisation's AI pool. They belong to the organisation rather than to any person in it. Owners and admins add and withdraw them. The pool takes keys for Anthropic, Google Gemini, OpenRouter and Scaleway, and it can point at any server the organisation runs that speaks the common OpenAI-compatible protocol.

A key is typed in once. Nexus encrypts it on arrival and never shows it again, anywhere. From then on the key appears only as a short hint and a line such as "set by Aoife on 3 Sept". When a key is withdrawn, the switchboard stops sending calls to it at once, and the key stays listed as withdrawn, so the history of who used what holds together. Nothing is deleted.

Each organisation has its own provider account and its own bill. There's no shared Adverant key sitting behind a tenant's AI, so one organisation's spending never mixes with another's, and an organisation that changes provider simply withdraws one key and adds another.

Keys for outside data sources are kept apart from model keys. Harrow holds a key for the ship positions feed it reads, for example, on a separate page of external data keys. That page refuses a model key outright, because model keys belong in the pool and nowhere else.



The switchboard routes every call

Apps and agents in Nexus never call a model directly, and neither do workflows. They ask the AI switchboard. It owns the keys and makes the final choice of provider and model. An app can say what it would like, perhaps a fast model for short replies, and the switchboard honours that wish when the organisation has that provider set up. When it hasn't, the switchboard picks another the organisation did allow. An app can never force a model the organisation hasn't chosen.

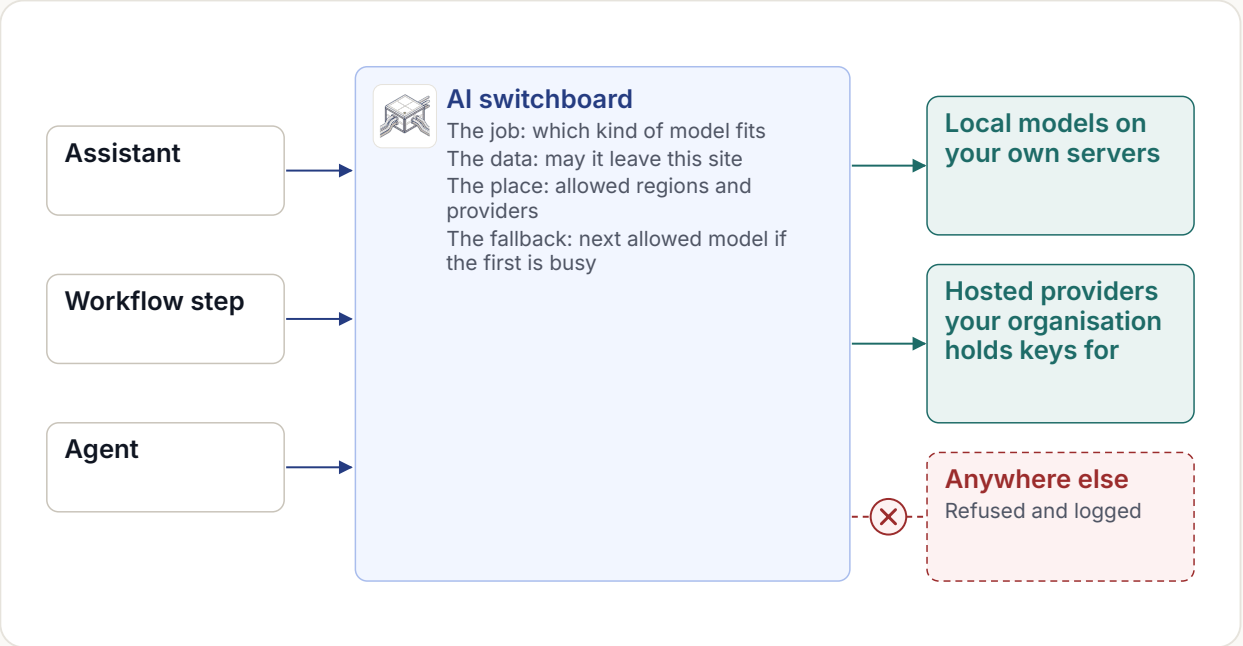
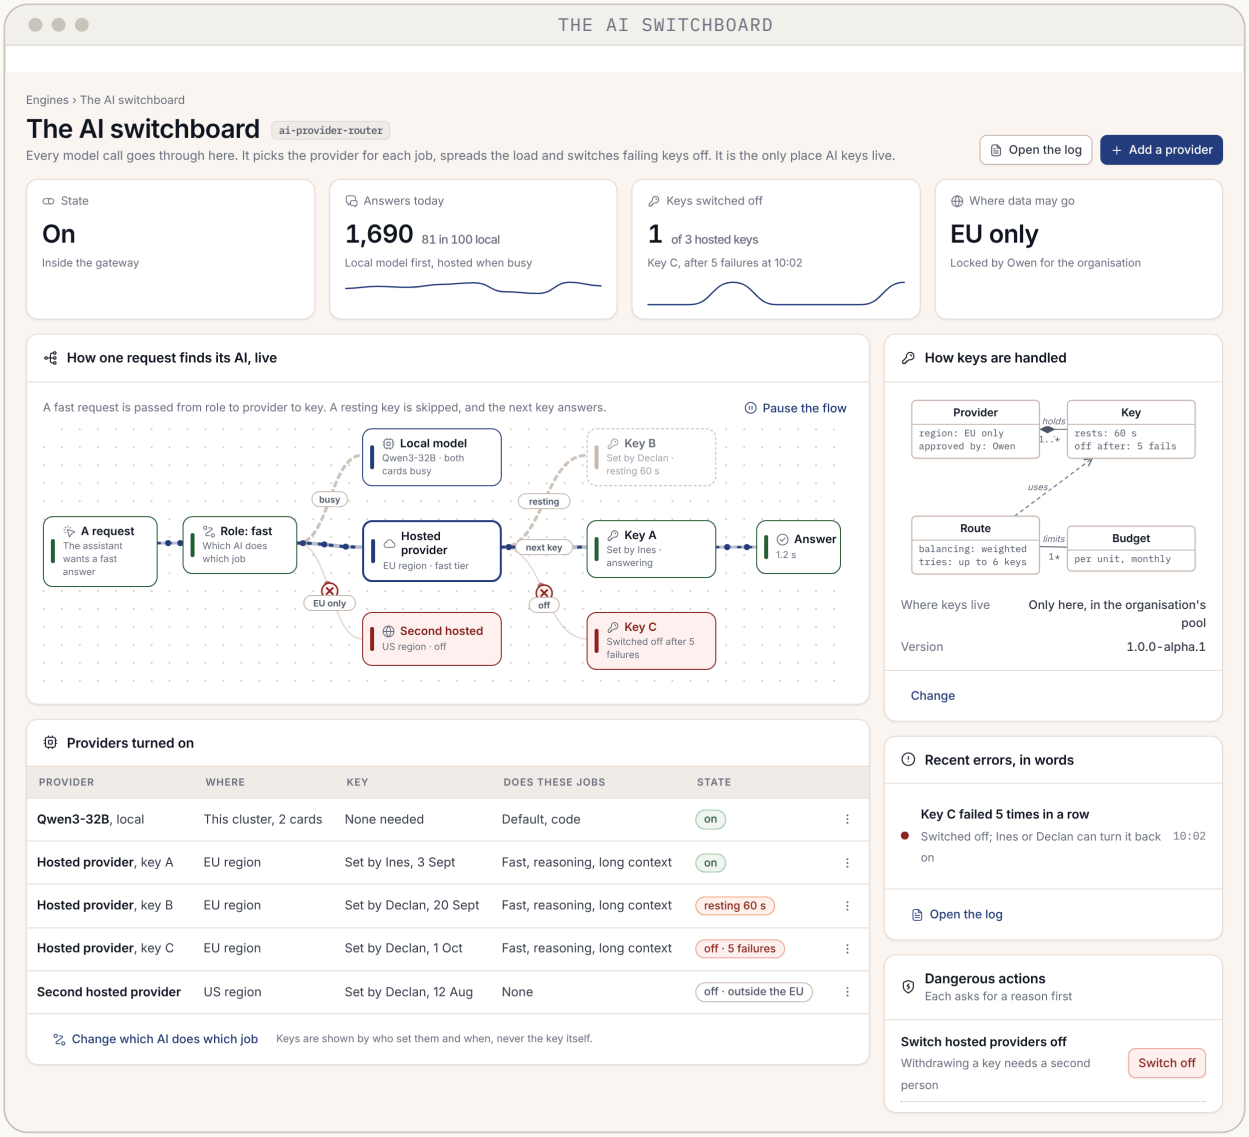


FIGURE What happens to one model call on its way through the switchboard.

Before a call leaves, the app labels it with the unit and role asking, and with how sensitive the data is. The switchboard routes on that label and never opens the call itself. It makes four decisions.

The first is the job. Reading a scanned form suits a different model from drafting a letter, so the organisation maps each kind of job (default, fast, reasoning, code and long context) to the models it prefers. The second is the data. Personal or secret data can be held to models on the organisation's own hardware, while public material may use more. The third is the place. Calls go only to the regions and providers the unit allows, and a call that would leave for anywhere else is refused, with the refusal written to the record. The fourth is the fallback. If a model is slow or down, the next allowed model takes the call. A key that fails rests for 60 seconds and is skipped; after five failures in a row it's switched off. The call reaches an allowed model or waits, and if neither works it fails with a clear reason. It never goes somewhere the rules forbid.

Rules start at the organisation and can be narrowed for each unit and app, and for each kind of data. A unit can be held to local models on its own hardware while another uses a hosted model in the European Union. A unit can never be looser than the organisation.



SCREEN The AI switchboard page. Notice which model does which job and where data may go. Each key appears only as who set it and when.

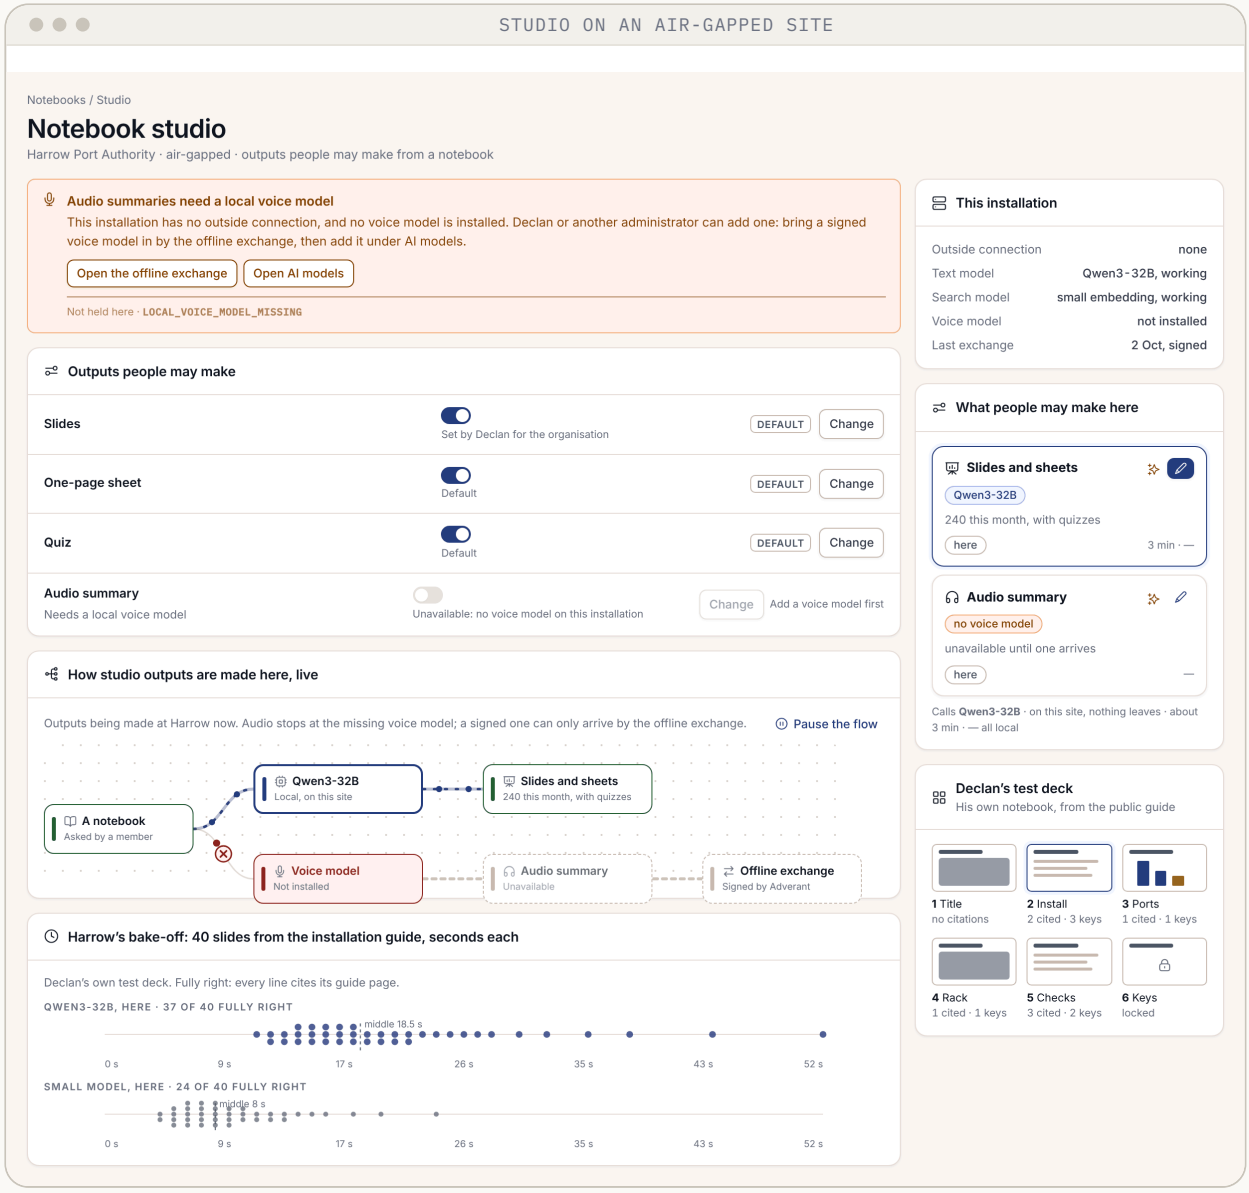
Every call is counted, never read

The switchboard counts every call by person and team, and by app and day. It measures each call's time and cost too. It never stores or reads what was asked or what came back. Each budget has a named owner and a limit. Budgets can be set at four levels, from the organisation down to one person, with units and skills in between. When a budget is reached, the rule its owner chose applies. The work can carry on or move to a cheaper model, or it can stop until an administrator agrees. Finance gets one statement that covers hosted and local models alike.

Models on your own servers

Any model the organisation runs on its own hardware can join the switchboard, as long as it speaks the OpenAI-compatible protocol. Your IT team gives the switchboard the server's address, sees the models it found, and chooses which job each one does. The connection page confirms "Nothing leaves this site" once it's done. A small model that turns text into searchable form also runs inside the installation by default, on ordinary processors, so search itself needs no outside service. Using a hosted service for that is a choice the organisation can make.

On an air-gapped site like Harrow, local models are the only models. That works well for most of what people do, and it brings one honest limit: a capability needs its model installed on site. Harrow's notebook studio can make slides and one-page summaries from the local models, but an audio summary needs a voice model, and until one is installed the studio says so plainly and names who can add it.



SCREEN The notebook studio on an air-gapped site. Audio summaries stop at the missing voice model, and the plate says who can add one and what the studio can make in the meantime.

Local models can also check each other. At Harrow, work done by one model family is reviewed by another (Qwen3-32B by Llama 3.3 70B, for instance), so a blind spot in one family isn't checked by the same family.

For organisations that want to go further, the Model factory add-on trains and tunes models on the organisation's own graphics cards, using only data the organisation is allowed to use.



No key means no AI, and Nexus says so

An organisation with no model connected gets no AI answers. That's a deliberate limit of the design, not a fault. Nexus doesn't quietly borrow a shared key to fill the gap, because then the organisation's data would go somewhere it never chose.

What people see instead is a plain plate. The assistant explains that it can't answer because no model is connected, gives the code `MODEL_NOT_CONNECTED` in small type, and offers "Connect a model" to whoever can. Jobs waiting for a model run as soon as one is back.



Nothing goes to Adverant behind your back

A self-hosted Nexus doesn't report on your people to Adverant. When something does leave, you read it before you send it.

An installation on your own infrastructure prepares a yearly usage statement holding counters only: how many seats and devices are in use, and which feature groups. Your administrator reads it line by line before uploading it or carrying it out on a disk. An air-gapped site asking for a new feature sends a request file that names the installation and the feature, along with the installation's fingerprint. Nothing about people or their work goes with it. When you need help, a support bundle shows what it will collect and a preview of what's redacted before anything is sent.

Adverant staff can't look inside a customer installation on their own initiative. If support needs to, the request comes through a support case, and your administrator says yes or no. A yes is for a set time of eight hours at most and a named reach, such as engine logs. It never covers anyone's messages or files, or their mail and memory.



A person's own subscription is never the organisation's

Some people pay for an AI subscription themselves, such as a Claude Max plan. Nexus never treats that as the organisation's. A personal subscription can't be added to the organisation's pool (the pool refuses it by name), and no page that shows the organisation's AI set-up ever lists one.

The reasons are practical. A personal plan is bought under consumer terms for one person, so routing colleagues' work through it would break those terms. The organisation can't see its use or set a budget on it, and it walks out of the door when that person leaves. The organisation's AI should rest on keys the organisation owns, and in Nexus it does.



TIP

If a team has been relying on someone's personal subscription, ask an administrator to add an organisation key for the same provider. The switchboard then routes the team's work through it, counted against a budget with a named owner.

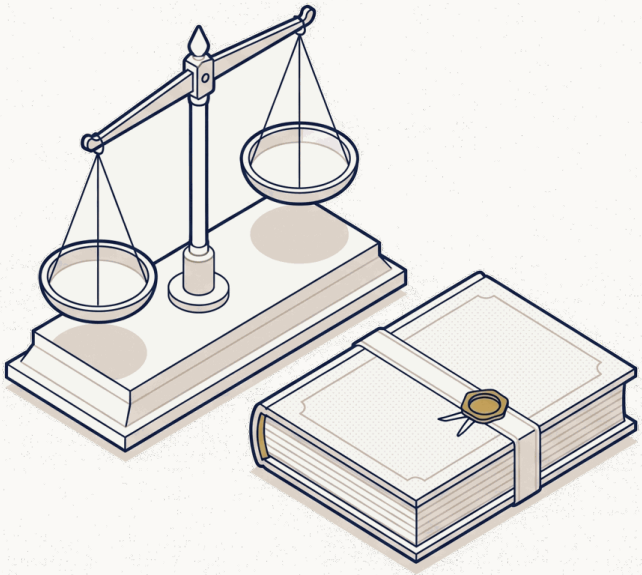
PART 1 · CHAPTER 1.5

Trust and compliance

The parts of Nexus a compliance officer lives in. It covers the audit log and sealed records, approvals that need a second person and artificial intelligence (AI) decisions waiting for review. It ends with personal-data requests and the compliance pages that keep evidence ready before an auditor asks.

FOR EVERYONE

12 MIN READ



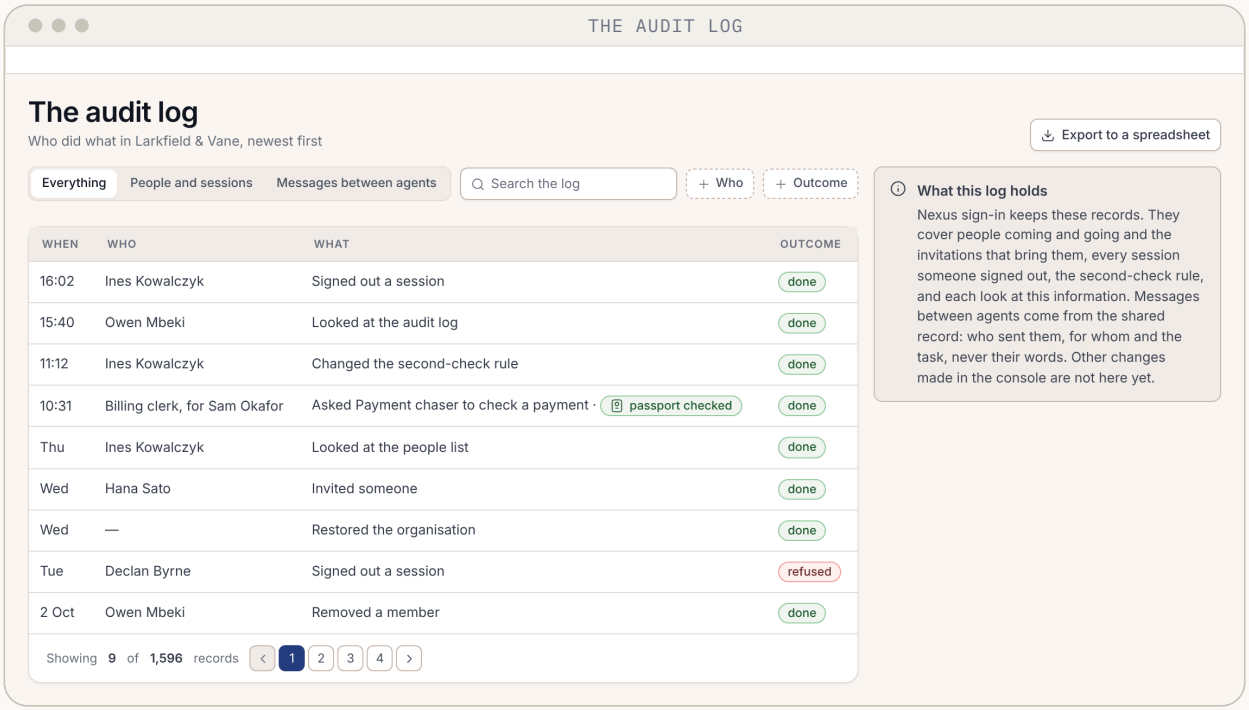
Ruairí Feltham is the data protection officer at Harrow Port Authority. His Monday starts with a queue of decisions an AI made overnight that can't take effect until a person looks at them. By ten he's answered a former employee's request for her personal data, with eleven days still in hand. After lunch an external auditor asks for evidence that last quarter's restore test happened. Ruairí builds the pack in Nexus, and before he sends it the pack tells him, in a sentence, the one thing it's missing.

None of that is a separate compliance product bolted on the side. It's the same record Nexus keeps of everything, read through pages built for the people who answer to regulators. This chapter walks through them.

The audit log

Every act in Nexus leaves a line in the audit log. The line records who acted and what they did. It also records when, and whether the act was done or refused. Each record says whether it came from a screen or a script. Looking at the audit log is itself an act, so that gets recorded too.

The log can be narrowed by area or outcome or by the person who acted, and searched by words. In Community it keeps 30 days and exports to a spreadsheet file at any time, and that stays true even though the rest of the compliance pages belong to Enterprise. Enterprise keeps the log for years and can stream it to your own security-monitoring system as it's written.



SCREEN The audit log. Each record says who acted and what they did, and when. Notice that it also says whether the act was done or refused.

Sealed records anyone can check

An audit log answers "what happened". For decisions that matter, a regulator will also ask whether the record could have been changed afterwards. Nexus answers that with sealed records.

A sealed record holds who decided and on what evidence, under which rules, and it's sealed with a time stamp from outside the installation. Change one character later and the check fails. Anyone can run that check with the open verifier, which comes with every edition and needs no account. It

sends nothing anywhere. On an air-gapped site like Harrow, time stamps are witnessed inside the site, and the records still check out later.

The Decision rooms and sealed records add-on takes this further for the big calls. A decision room gathers the question and the evidence, with seats for the people who must be heard and AI seats that argue each side. A named person chairs it, who is never the one who proposed the change. The chair rules with a written reason and signs with a passkey, and the ruling is sealed with the rest of the record.



Second checks and four-eyes approvals

Two different safeguards share the idea of "a second", and it helps to keep them apart. A second check is the extra step at sign-in, such as a passkey after a password, and chapter 1.3, Your organisation, its people, and what stays private, explains how an organisation requires it. Four-eyes approval is about changes: some changes wait for a second person before they happen.

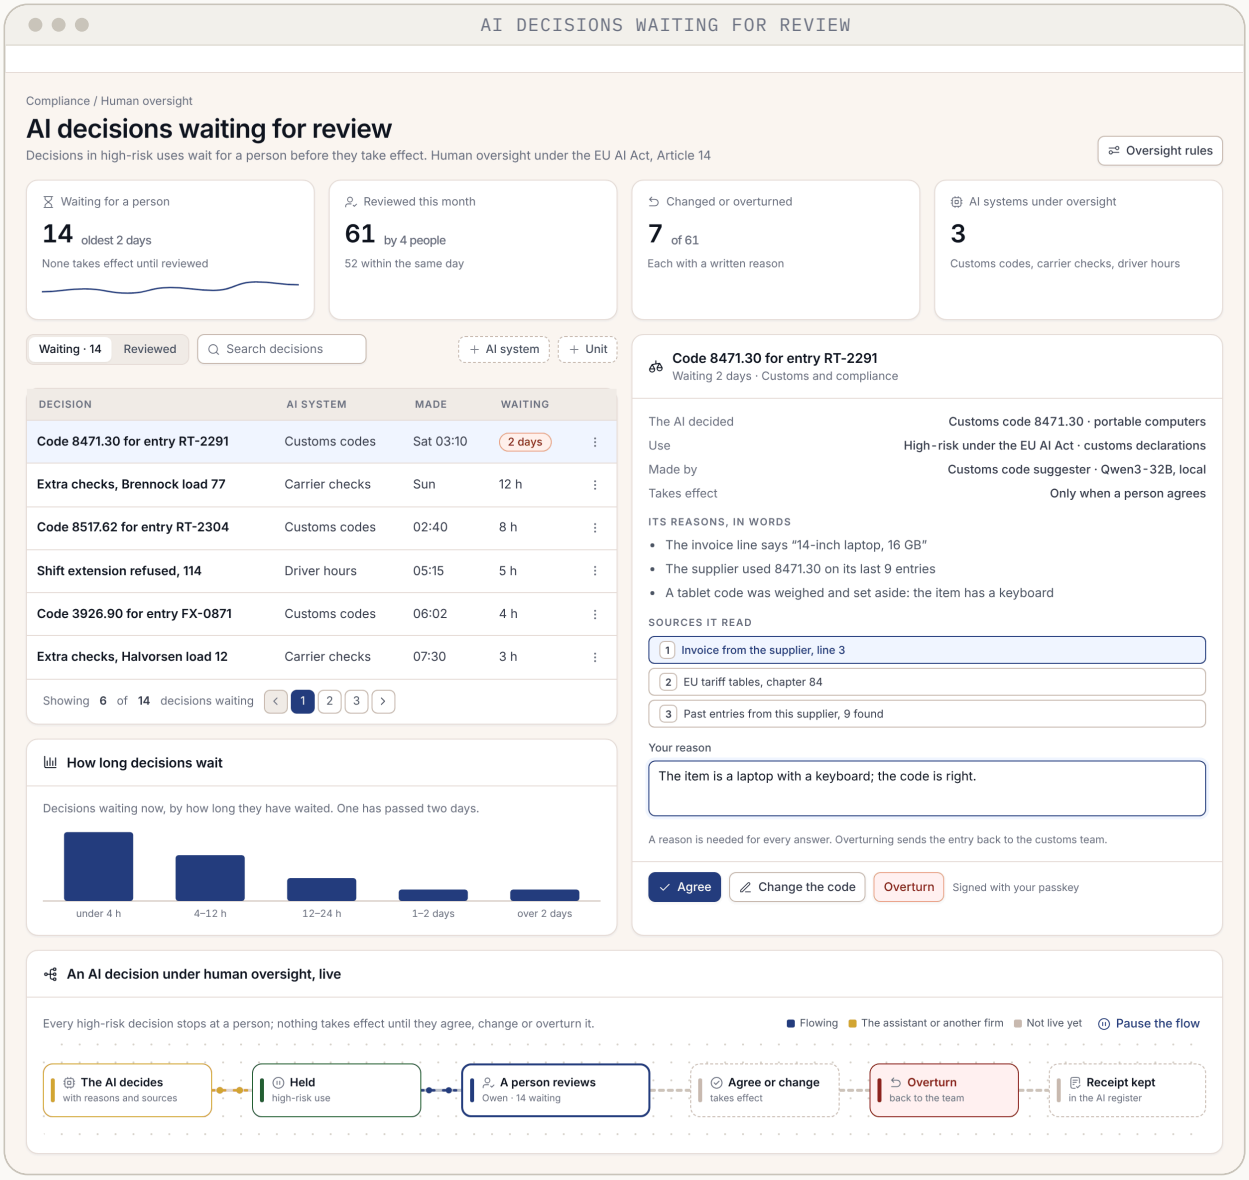
In Enterprise, the organisation's policies decide which changes need a second person. Typical ones are shortening how long records are kept, lowering a budget that's in use, withdrawing a platform from a unit and exporting a team chat channel. A request waits on the Four-eyes approvals page until it has every approval it needs. One person turning it down ends it, with their reason. A request nobody decides in time lapses, and the change isn't made.

Nobody approves their own request. When an emergency needs full access, break-glass gives an administrator a short window with a second person watching, and every step is recorded.



AI decisions waiting for review

Some uses of AI are high-risk under the EU AI Act, and the law expects human oversight of them. In Nexus, a decision made in a high-risk use is held. It waits on the AI decisions page until a person reviews it. The reviewer sees what the AI decided and why, with the sources it read. They can agree with it or change it, and they can overturn it. Each answer needs a reason. All three answers are signed with the reviewer's passkey. Overturning sends the work back to the team, and every answer leaves a receipt in the register of AI uses.



SCREEN AI decisions waiting for review. The list shows how long each decision has waited; the card on the right shows the decision with its reasons and sources, above the answers a reviewer can give.

The queue shows how long each decision has waited, so a backlog is visible before it becomes a problem.

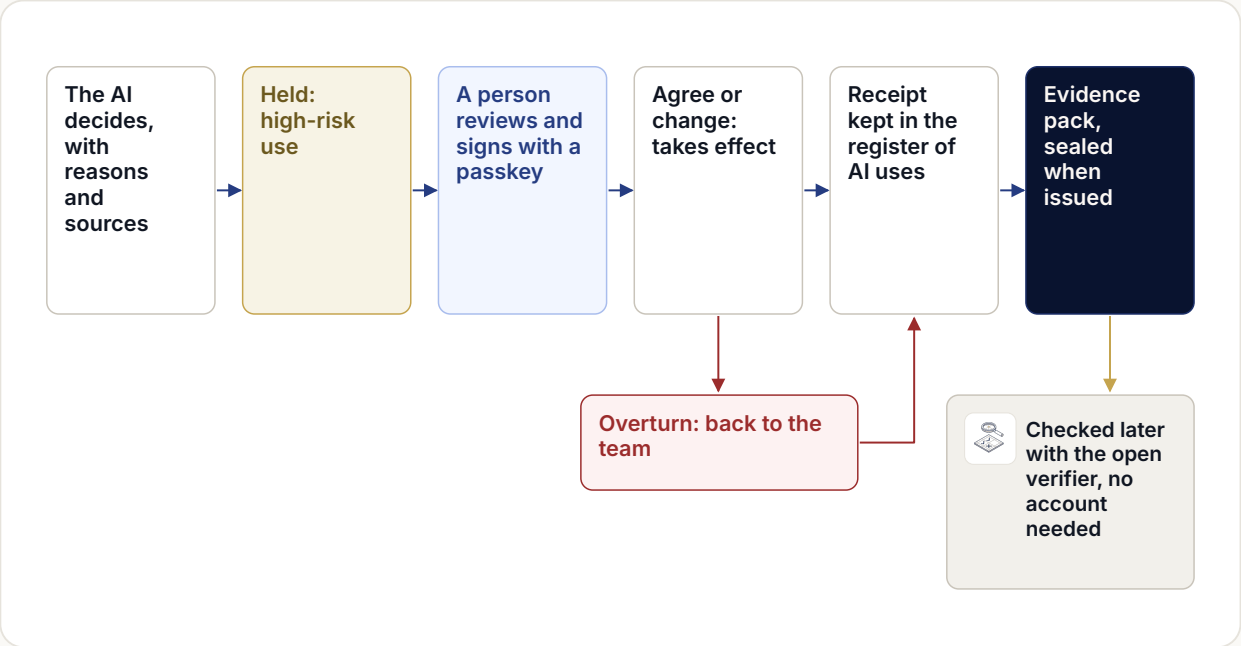


FIGURE How one high-risk AI decision becomes evidence an auditor can check.

Personal-data requests

Under the General Data Protection Regulation (GDPR), a person can ask what an organisation holds about them and ask for a copy, or ask for it to be erased, and the organisation normally has a month to answer. Every member of Nexus can export their own data and erase what's theirs, in any edition. Requests from customers or former staff are logged in Enterprise with their due dates.

For each request, Nexus finds the person's records across files and mail, in chat and in the assistant's memory, and shows them in one place. Legal exceptions are applied rather than ignored. Payroll records the law says must be kept for seven years stay, and Nexus shows why. Erasure itself is done with Erase, the one destructive action in Nexus. It needs a written reason and a second person, and it leaves a permanent record that the erasure happened (what was erased and when, and who did it) without keeping any of the content.

WATCH OUT

Erase can't be undone, by design. Check the records Nexus found before you confirm, and make sure the second person reads them too, not just the request.

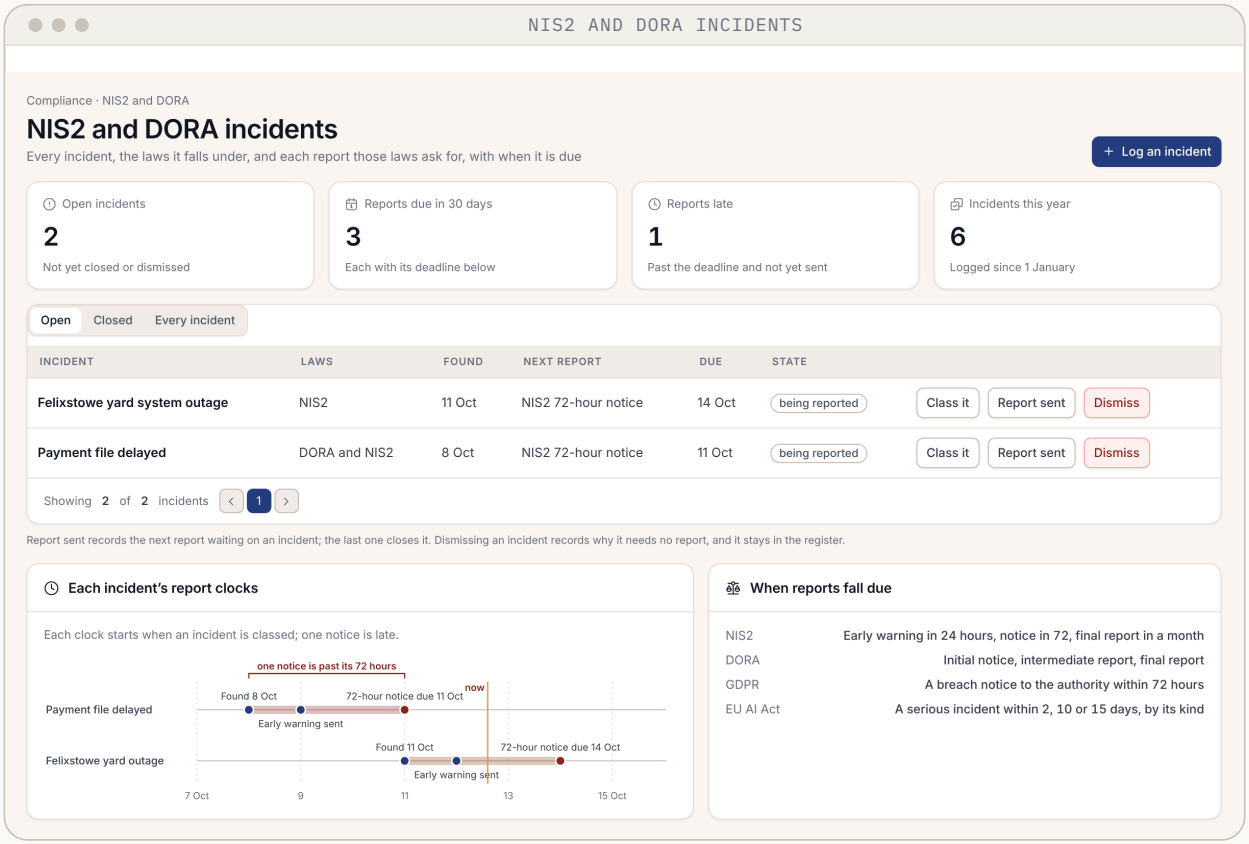
The compliance pages

Enterprise has a Compliance area with one page for each body of rules an organisation answers to. The overview lists every pack, with evidence counts and the open gaps written in words with due dates. The packs cover the EU AI Act, the GDPR and ISO 27001, along with the Network and Information Security Directive 2 (NIS2) and the Digital Operational Resilience Act (DORA). There are no scores or grades. "Exit plan missing for the hosted model provider, due 15 November" tells you more than a percentage would.

The **EU AI Act** page keeps the register of every AI use in the organisation. Each entry records its risk class and the person who oversees it. It also says how affected people are told, and holds the impact assessment and any incidents. The high-risk rules apply from 2 December 2027 for uses such as credit scoring and life and health insurance pricing, as amended by the Digital Omnibus.

The **GDPR and data protection** page holds the records of processing: purpose, lawful basis, the data involved, retention and transfers. Requests sit beside them with their deadlines. Consent records and breaches have their own lists.

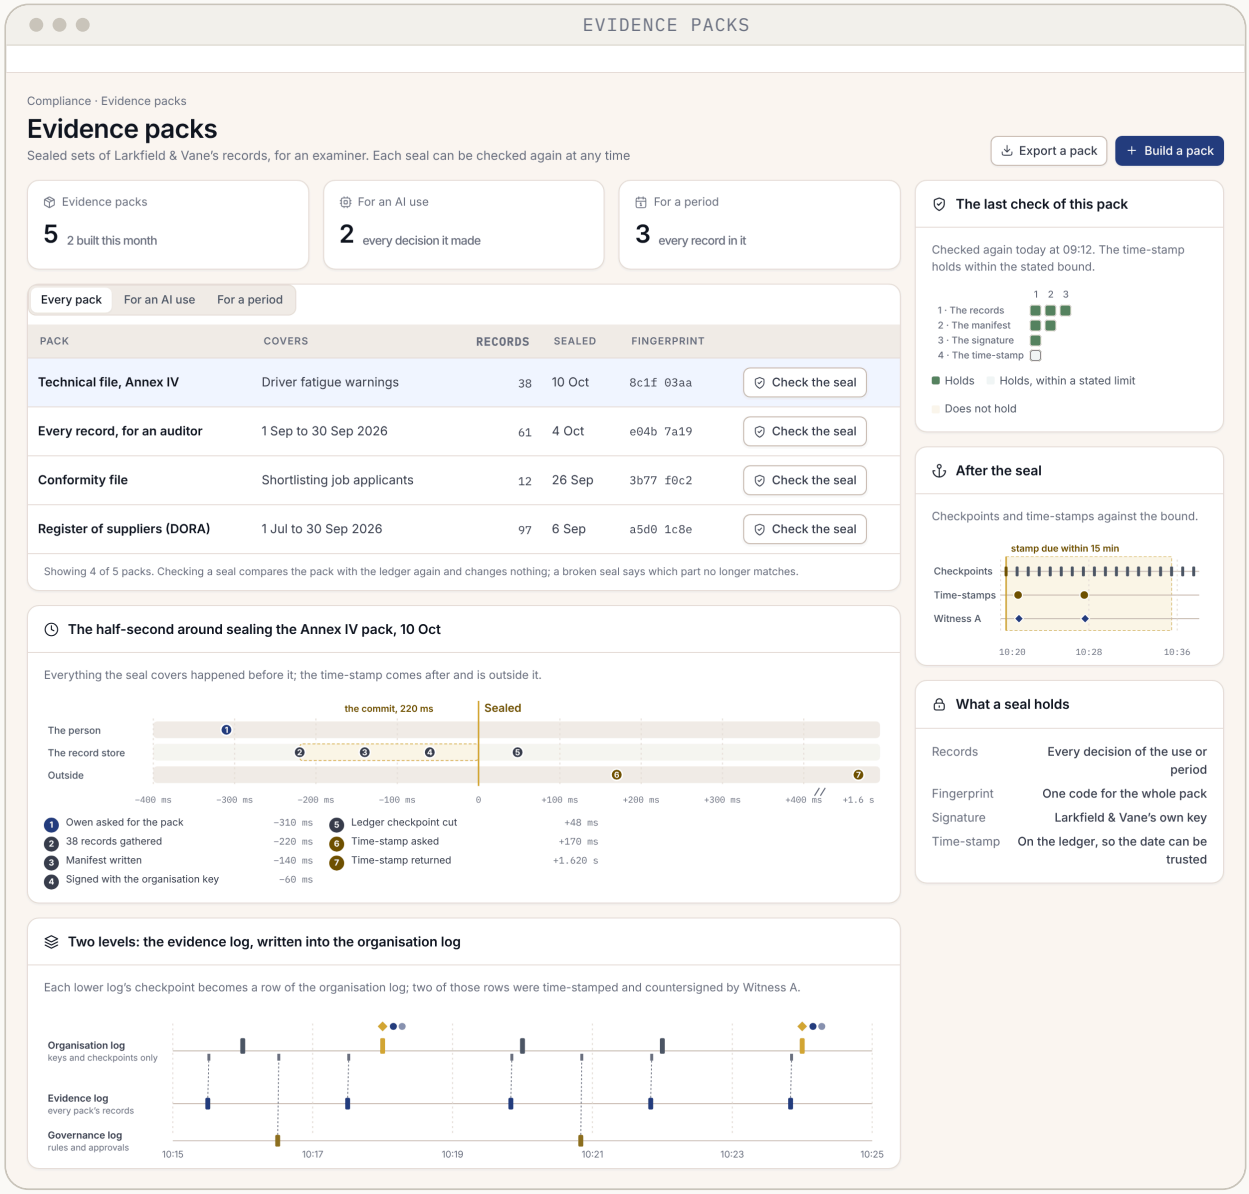
The **NIS2 and DORA incidents** page keeps the incident register with each reporting clock drawn as a timeline: the early warning within 24 hours, the notice within 72 hours and the final report within a month. A clock starts when an incident is classed, so a late notice shows as late.



SCREEN NIS2 and DORA incidents. Each incident's reporting deadlines appear as dates on a timeline, so the next one due is obvious at a glance.

Evidence packs and controls

An auditor rarely wants everything. They want exactly what they asked for. An evidence pack is built from live records for a named examiner and mapped to each control. It's sealed when it's issued. Before you send it, the pack lists what's missing in plain words, such as a restore test that was run but whose result was never attached. Months later, you can check that an older pack still holds. Legal holds work alongside packs: a hold stops deletion of what it covers, an export takes exactly what's held, and releasing it leaves a record.



SCREEN Evidence packs. Notice the line explaining what the seal covers: everything inside happened before it, and the time stamp comes after.

One control often answers several rules. "Human oversight of high-risk AI" satisfies an article of the GDPR and one of the AI Act at the same time. The controls page maps each control once to every framework it meets: the GDPR, NIS2, DORA, ISO/IEC 27001 and the EU AI Act. Each control has an owner and the date it was last tested. One piece of evidence then serves several frameworks.

Assessments, such as a quarterly ISO 27001 review, are worked through control by control with answers and evidence. Gaps are found along the way, and a second person signs the result off. The officer who ran it can't sign their own. For each AI system, a technical file gathers what the AI Act

asks for (its purpose, data, testing, oversight, risks and logs) and shows how complete it is. Export one before it's complete and every page is marked as a draft.



Suppliers and suspicious sign-ins

The supplier register lists every outside firm that handles your data or runs a service for you. For each one it records what the firm processes, when the contract ends, when it was last reviewed and how you'd get out. A critical function with no way out turns amber, which is exactly what a DORA examiner looks for. A withdrawn supplier keeps its records for seven years.

Suspicious sign-ins are flagged as they happen. Impossible travel is one sign, such as a sign-in from Harrow and then from another continent forty minutes later. New devices and runs of failures are others. Each appears on its own page, with what Nexus already did about it. An officer confirms or clears each one. Allowed and blocked address ranges sit beside it, and Nexus refuses a rule that would lock out the administrator writing it.



What Nexus doesn't do

Nexus keeps the records, maps them to the rules and builds the evidence. It doesn't certify you, and it doesn't give you a compliance score. Certification stays between you and your auditor. The regulatory dates and mappings in Nexus aren't legal advice, and your counsel should read the rules that apply to you.

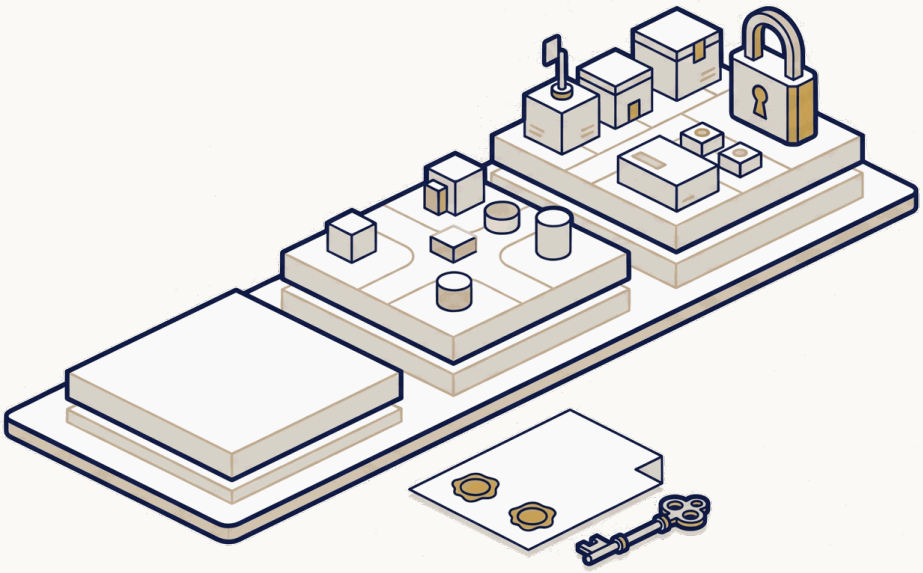
PART 1 · CHAPTER 1.6

Editions and licences in plain terms

The three editions of Nexus and the 19 add-ons and domain packs on top, then how a licence works for a business. It covers the offline check and what each licence state means in practice, the promises every licence keeps, and moving from Community to Enterprise without reinstalling.

FOR EVERYONE

11 MIN READ



Harrow Port Authority's licence arrived on an encrypted drive, carried by one of its own staff. That's how an air-gapped site works: a signed request file goes out, and a signed licence file comes back. Aoife Brennan imported it on the Licence page, and the installation checked Adverant's signatures there and then, with no call to anyone. A minute later the page listed every feature group in the licence. Beside each it showed the seats and the term.

This chapter explains what Harrow bought, what it could have bought, and what the licence does and never does once it's in place. There are no prices here. Adverant quotes them, and they depend on the way you run Nexus and what you include.



One product, three editions

Every edition of Nexus has the same screens. What a person may use is decided by a licence, never by a different build. A feature outside your licence doesn't disappear. It stays in the menu with a lock beside it. Opening it shows a plain page that says what it does and who can turn it on. There's no price on that page, and it doesn't nag. A member can send their administrator a request from it, with a reason.

Community is free and open source under the Apache-2.0 licence. It's what one person needs to finish a job in one organisation. It has the assistant with cited answers, and search over files and the knowledge graph. People keep personal notebooks and talk in basic channels and direct messages. Workflows and skills are included, along with apps from the catalogue. Administrators get members with fixed roles and one company sign-in with second checks, and an audit log kept for 30 days. Agents can watch and suggest, or ask first before they act. The organisation brings its own model keys or its own local models. Community runs on Adverant's cloud, where a free tier has usage limits, or on your own server with no licence at all.

Team runs on Adverant's cloud. It's Community plus what a team shares. The work panel docks beside every page, and notebooks and assistant contexts can be shared. Sharing a notebook sets who may read it and ask it questions, and who may change it. Administrators get a skills library that shows spend per skill and a memory and consent policy for the organisation. They also get workflow operations, such as promoting a workflow from test to live.

Enterprise is licensed per feature group and runs wherever your rules allow, including your own Kubernetes cluster and sites with no internet. It includes everything in Team. On top it adds several organisations and business units, plus enforced company sign-in with directory sync. Custom roles and four-eyes approvals come with it. The audit log is kept for years and can be streamed. Retention and legal holds are there, with the data protection and compliance pages and the pages for running your own installation. Add-ons and domain packs are licensed on top of Enterprise.

Some things are never paywalled in any edition. They form the floor. Administration and security basics are on it, and so is the separation between organisations. Signing out is on it, along with every refusal screen and the export of your data.



The add-ons

Add-ons are capabilities any kind of organisation might want, each licensed separately on top of Enterprise. An add-on you don't have shows as a single greyed row in the menu rather than a group of locked rows.

Add-on	What it adds
AI Gateway at scale	Spreads calls to artificial intelligence (AI) models across copies and providers. It scales them automatically and adds budgets per team and a cache

Model factory	Trains and tunes your own models on your own graphics cards and manages their lifecycle
Ideas	Runs idea rounds with blind comparison and critics. It searches prior art too, and a team approves the shortlist within a budget per run
Calls and voice	Huddles and live calls with a transcript, and phone numbers with local voices. Recording consent follows each country's rules
Mail	Personal and team inboxes with rules, and checks before mail leaves. Mail can be filed into the record
Other firms and walls	Shared conversations with another firm through cleaned copies, and walls between firms that must not see each other
Notebook studio	Makes slides and one-page sheets from a notebook, as well as audio summaries and quizzes. It also finds papers from outside
Map and GeoAI	Adds imagery and drone layers to the map, with models that read them
Data connectors	Reads outside drives and wikis, mail stores and older systems as they change. It also bridges channels to Slack and Microsoft Teams, or to Discord
Agent-to-agent	A registry of agents and a live map of who asks whom. Permissions are set per pair, and any request can be traced step by step
Decision rooms and sealed records	Rooms where a named person rules and signs with a passkey. Each ruling becomes a sealed record with a receipt and a decision passport
Revenue operations	Pipeline and accounts, forecast and deal desk, quotes and contracts, with territories on the map
Storefront / white-label	Your own brand and domain on Nexus, plus payments through the gateways you choose
Nexus Edge	Runs Nexus on board satellites and drones, vessels and vehicles. Mission packs and contact windows for each fleet are managed from the ground
Mobile apps	Native iPhone and Android apps for approvals and offline work, with company device management
Workforce	Always-on agents at higher autonomy levels. Each runs under a budget, and kill switches and fleet watch keep them in check
Shield	Protects every domain with certificates and a firewall, rate limits and bot checks. A waiting room and flood defence keep sites up on busy days
Sandboxes	Administers the sealed machines that risky steps run in, with profiles and quotas. New limits run in shadow before they are enforced
Quality (NexusQA)	A testing platform with outside testers and recordings. AI repairs are signed off by a person

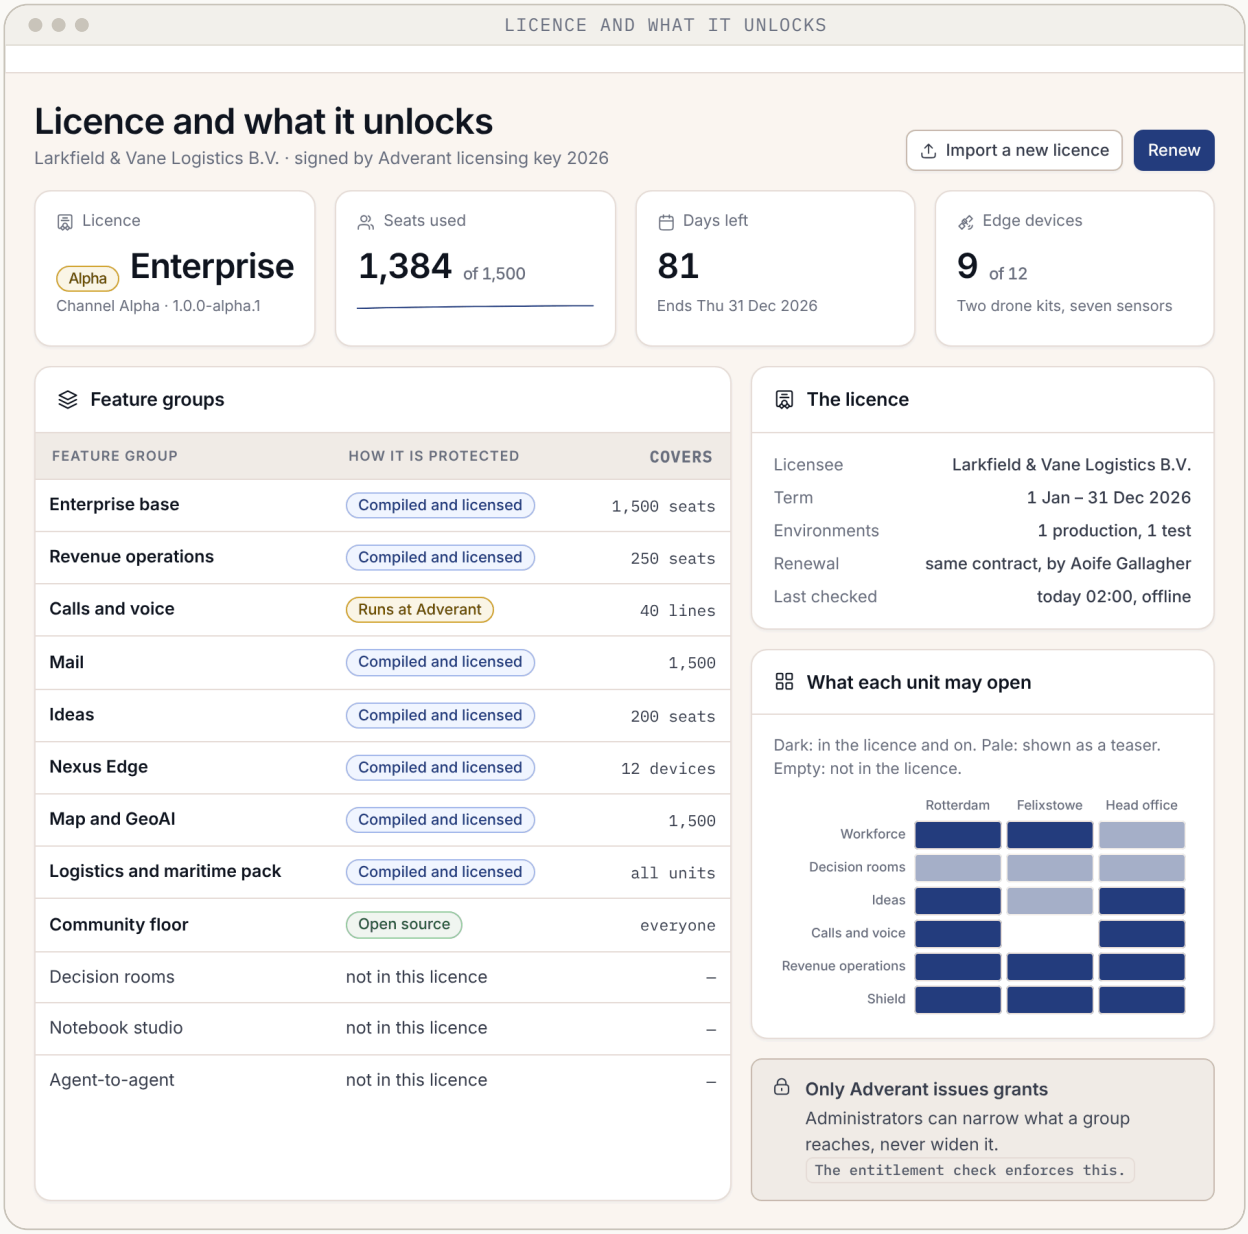
Some add-ons need others for parts of their work. Workforce agents on edge devices also need Nexus Edge, and approving agent work on a phone also needs Mobile apps.

Domain packs bring an industry's own words and rules to the same core, along with its screens and checks. There are ten: Insurance; Banking and payments; Capital markets and asset management; Government and public sector; Healthcare; Life sciences; Critical infrastructure and energy; Defence and security; Legal and professional services; and Logistics and maritime. A pack may need an add-on (Insurance needs Decision rooms and sealed records), and a pack's phone screens also need Mobile apps. Chapter 1.7, Domain packs and verticals, covers them.

How a licence works

A licence is a signed document. It names the licensee and the term. It also lists every feature group it covers, with what each group counts: seats for most, and devices or phone lines for a few. Adverant signs every licence with two independent keys.

The important part for a business is where it's checked. Your installation checks the licence itself, offline, against keys built into its own signed software. There's no call to Adverant, so the same licence works on a site with no internet. A licence file that has been altered, or wasn't signed by Adverant, counts as absent. The Licence page says so, and nothing the file claims is switched on. Each licensed module is also checked against Adverant's signature before it runs, so a changed copy is never loaded.



SCREEN The licence page. Each feature group is listed with its seats or devices and its term, and with how it's protected. Groups the licence doesn't cover stay listed, greyed.

How the licence arrives depends on how you run Nexus. On Adverant's cloud there's nothing to import. A connected installation fetches its licence from your Adverant account. An air-gapped one

exchanges signed files carried by a person, as Harrow does. A licence file made that way works only on the installation that asked for it.

You can try Enterprise first. On Adverant's cloud, a 14-day trial needs no card. On your own infrastructure, an evaluation licence runs for 30 days for five people named before it starts. Neither turns into a paid plan by itself. When an evaluation ends, its features return to teasers and everything made during it is kept, ready to open again the moment a licence is imported.

If more people join than the licence covers, perhaps because a directory sync brought in a batch of contractors, nobody is locked out. The page explains where they came from and gives 30 days to remove people or add seats.



What the licence states mean

A licence moves through a short ladder of states. The ladder is written out in words, with your dates, on the "If the licence runs out" page, so nobody has to guess.

State	What you see	What still happens
Valid	Everything the licence lists is on	Everything
Warning	From 30 days before the end, unless the contract sets another number, administrators are told in the product as well as by email	Everything; nothing changes
Grace	After the end date, a plate on Today counts down 30 days	Everything still works, and renewing ends grace at once
Restricted	Starting new Enterprise work says it needs a renewed licence, and who to ask. A licence that simply runs out stays here until it's renewed	Running work finishes, reads and exports keep working, and security fixes still arrive
Ended	Only if Adverant withdraws the licence: Enterprise parts stop being offered	The Community floor and every export keep working, and sealed records stay readable

An evaluation licence has no grace period, which is the one exception. Renewing at any rung goes straight back to Valid.

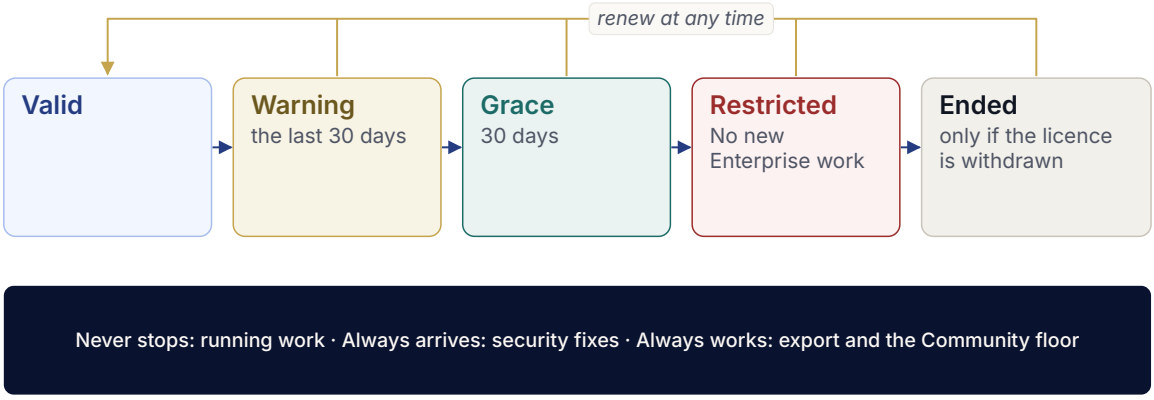


FIGURE The licence ladder, from valid to restricted, with what never changes underneath it. A licence that runs out stops at Restricted; only a withdrawn licence ends.



The promises every licence keeps

Whatever state a licence is in, four things hold. Nothing running is stopped by a licence: a job started before grace ended runs to its end. Export always works in every edition and every state, because your data is yours. Security fixes always arrive for every version, whether or not you've renewed. Sealed records stay readable, so the evidence you've built is never held back.

Buyers sometimes ask, at the end of a meeting, what happens if Adverant disappears. The answer is on the same page. Running work never stops. The Community floor and every export keep working, and the source of the compiled parts is held in escrow under the contract.



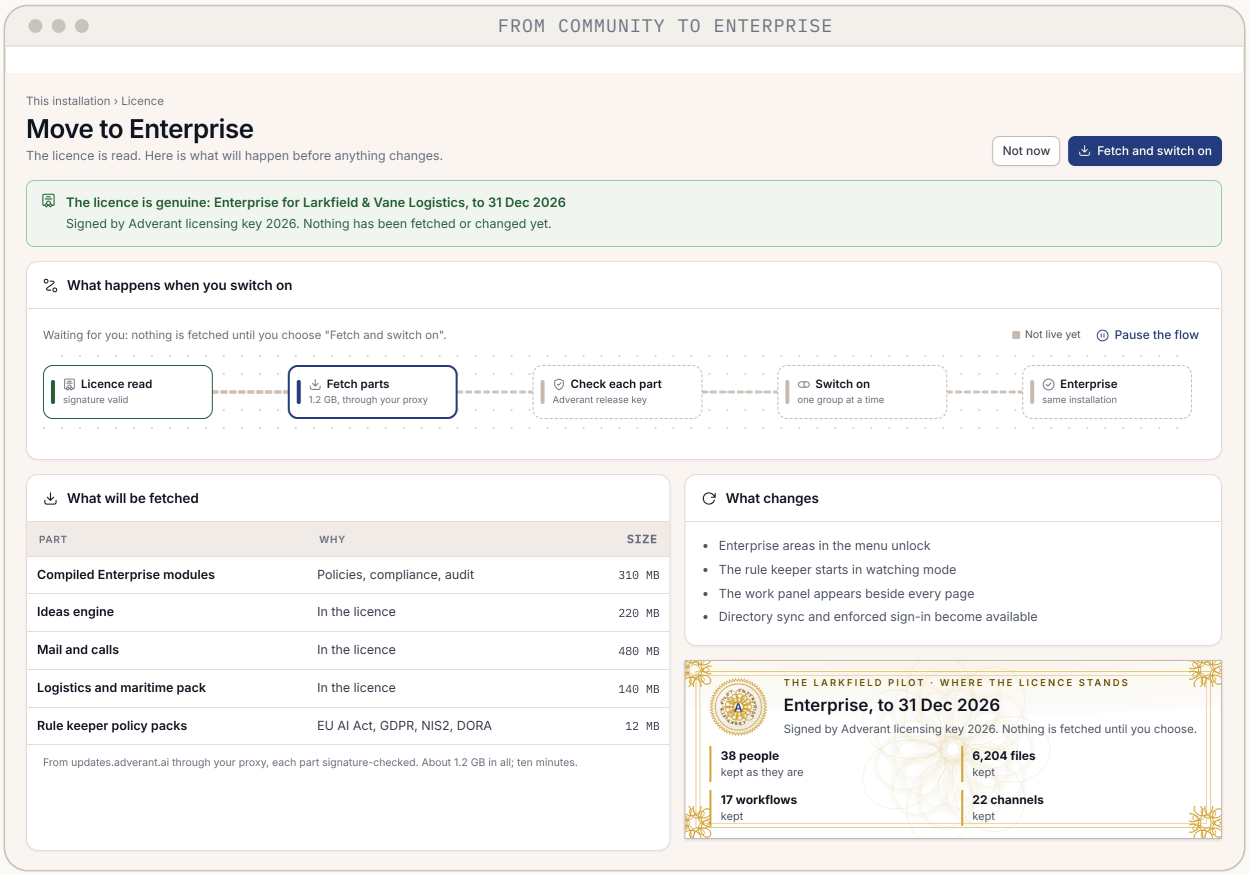
SCREEN If the licence runs out. The ladder is written in words with this licence's own dates. Beside it are cards for what never stops and what always arrives, and a third for what always works.

From Community to Enterprise, in place

Many organisations start on Community and move up once the work is proven. The move happens on the same installation. Nothing is reinstalled and nothing is migrated.

An administrator imports the Enterprise licence on the Licence page. Nexus first checks the licence and shows what it would fetch, before anything changes: the compiled Enterprise parts and any add-ons or packs the licence includes, each checked against Adverant's signature. It also lists what will change, such as the work panel appearing and enforced company sign-in becoming available.

Nothing is fetched until the administrator chooses "Fetch and switch on". Then the feature groups switch on one at a time. People and their files stay exactly as they were, and so do workflows and channels. On an air-gapped site the parts arrive on signed media instead of over the network.



SCREEN Moving from Community to Enterprise. The licence has been read and is genuine, and the page shows what will be fetched and what will change. Nothing happens until "Fetch and switch on".



NOTE

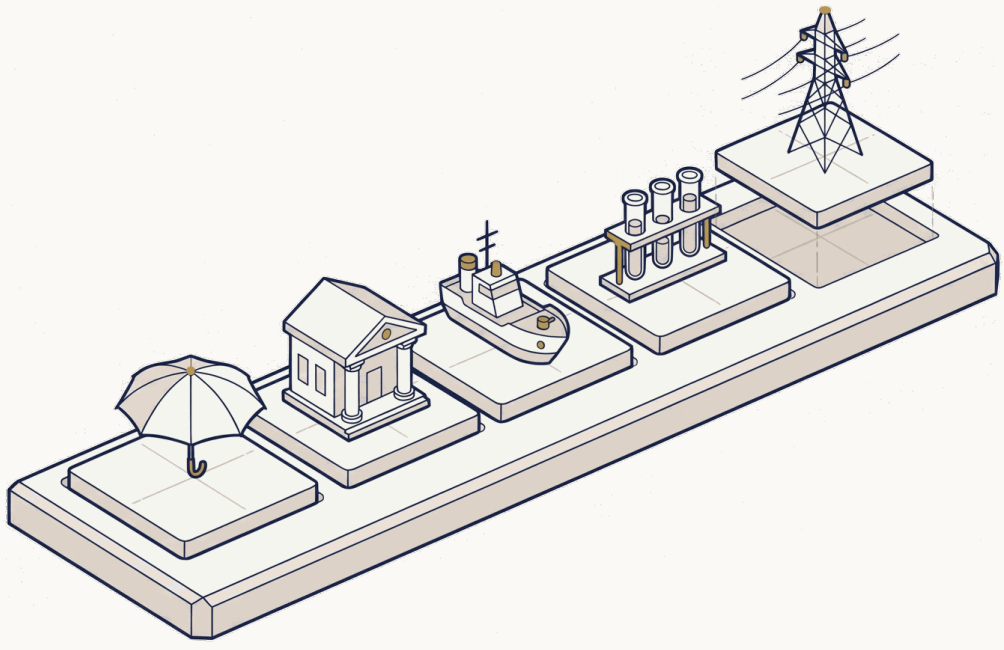
An administrator can narrow what the licence allows for a unit or a role but can never widen it. A feature the licence doesn't cover can't be switched on from inside the installation, however senior the person asking.

Domain packs and verticals

What a domain pack adds to Nexus for one industry, the ten packs Adverant offers, how a pack is switched on for one unit, and what a pack needs in your licence.

FOR EVERYONE

12 MIN READ



Sinéad Kavanagh runs claims at Larkfield & Vane. When she opens Nexus in the morning she sees the same menu and the same assistant as the finance team down the corridor. What's different is one extra group in her menu, a claims lens that puts the claims waiting on a decision at the top, and an assistant that knows what a reserve is and which policy wording applies. That extra layer is the Insurance pack.

A general chatbot doesn't know what a customs hold is. A domain pack teaches Nexus one industry's vocabulary and rules, and the shape of its work, without changing anything else about the product. This chapter explains what's inside a pack, walks through the ten packs Adverant offers, and shows how one gets switched on.



What a pack is

Think of Nexus as the building and a pack as the fit-out for one trade. Sign-in, permissions, files, chat and the artificial intelligence (AI) switchboard stay exactly as they are. The pack brings five things on top.

What the pack brings	What it is
Lenses for each role	A lens is what one role sees. A claims handler and an auditor can look at the same claim, and each sees the part their job needs.
Templates	Agents, workflows, skills and seat layouts for the industry's routine work, such as claims triage or a credit paper draft. Each is switched on or off per unit, and switching one off keeps its past runs.
A knowledge model and its connectors	The things the industry deals in and how they link, plus its vocabulary. Connectors read your existing systems and never write to them.
An evidence pack	The rules the industry answers to, each split into what it asks for and matched to the records Nexus keeps. Gaps are written in words, each with an owner.
A mobile overlay	The cards and capture steps the pack adds to the phone app, chosen per role.

A pack never changes the shell. It adds its own menu group and its own words, and it uses the components and writing rules every other part of Nexus uses. It also can't widen what anyone sees. A lens only narrows what the person's role already allows, and permissions still come from roles and policies.

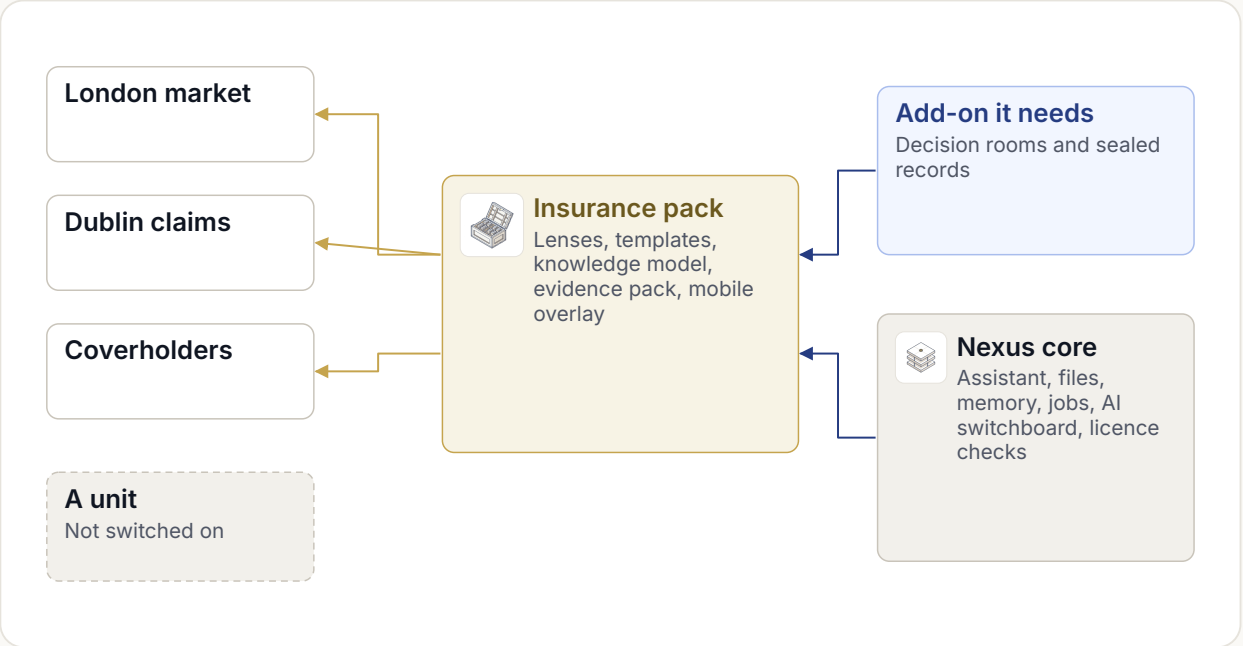


FIGURE A domain pack sits on top of the shared Nexus core and reaches only the units it is switched on for.

NOTE
The evidence pack is a starting point that your own counsel confirms. It maps rules to records and names the gaps. It isn't legal advice, and the page says so.

The ten packs

Every pack is built from the same parts. The sample customers in the screens are illustrative, and your own administrators decide which lenses and templates you use.

Insurance. Claims from the first call to a sealed decision, with underwriting, reinsurance, delegated authority and the Lloyd's market alongside. The assistant cites each fact to its page and raises fraud signals with their reasons. No claim is refused by a model: a named person decides, and the decision is sealed with its evidence. The worked example below goes through this pack in detail.

Banking and payments. It keeps the register of outsourced technology the rules ask for, with incident clocks and exit plans for critical providers, and then covers lending and card disputes. At Larchmont Bank, Omar Haddad's credit lens shows which applications wait and which rule held each one. The model drafts the credit paper. It never approves or declines one, because the pack locks that rule for the AI Act.

Capital markets and asset management. Calls on trading desks are recorded and kept for five years under the EU investment-services rules (MiFID II). Research stays behind a wall until it's released, and a wall-crossing log records who was brought over. Phones get the morning research briefing and release approvals. Orders never reach a phone; they stay in the order system.

Government and public sector. Benefit claims and planning applications, plus information requests. A caseworker records a home visit on a phone with no signal, and the decision is left for the office. Each decision carries a plain-language reason the citizen can challenge. Data stays in the national cloud, and a connector that would send it elsewhere is refused before it runs.

Healthcare. Access to patient records and sharing them with other providers. Every time a person or an agent opens a record, the purpose is checked against what they may do and then logged. The clinician summaries lens writes handovers with each line cited. It never suggests a diagnosis or a treatment, and a template that would is greyed with the reason.

Life sciences. Deviations and corrective actions, batch record review and validated change control. The drug safety lens takes in adverse-event reports and runs the reporting clock, and a person confirms each case. Each signature states what it means. Study documents are never edited on a phone, since editing stays validated on the desktop.

Critical infrastructure and energy. Incident clocks, inspections and work orders for grid and network operators. The Network and Information Security Directive 2 (NIS2) wants an early warning within 24 hours and a notification within 72, then a final report within a month. The incident manager lens carries all three clocks, locked by the pack. Switching instructions never go to a phone.

Defence and security. Field reports sit beside alerts from Nexus Edge devices, each item marked with who it may be shared with. The pack runs with no outside connection on accredited hardware. Its mission planner lens drafts plans, but a model never commits an action. A person signs every plan, and nobody can change that rule on the installation.

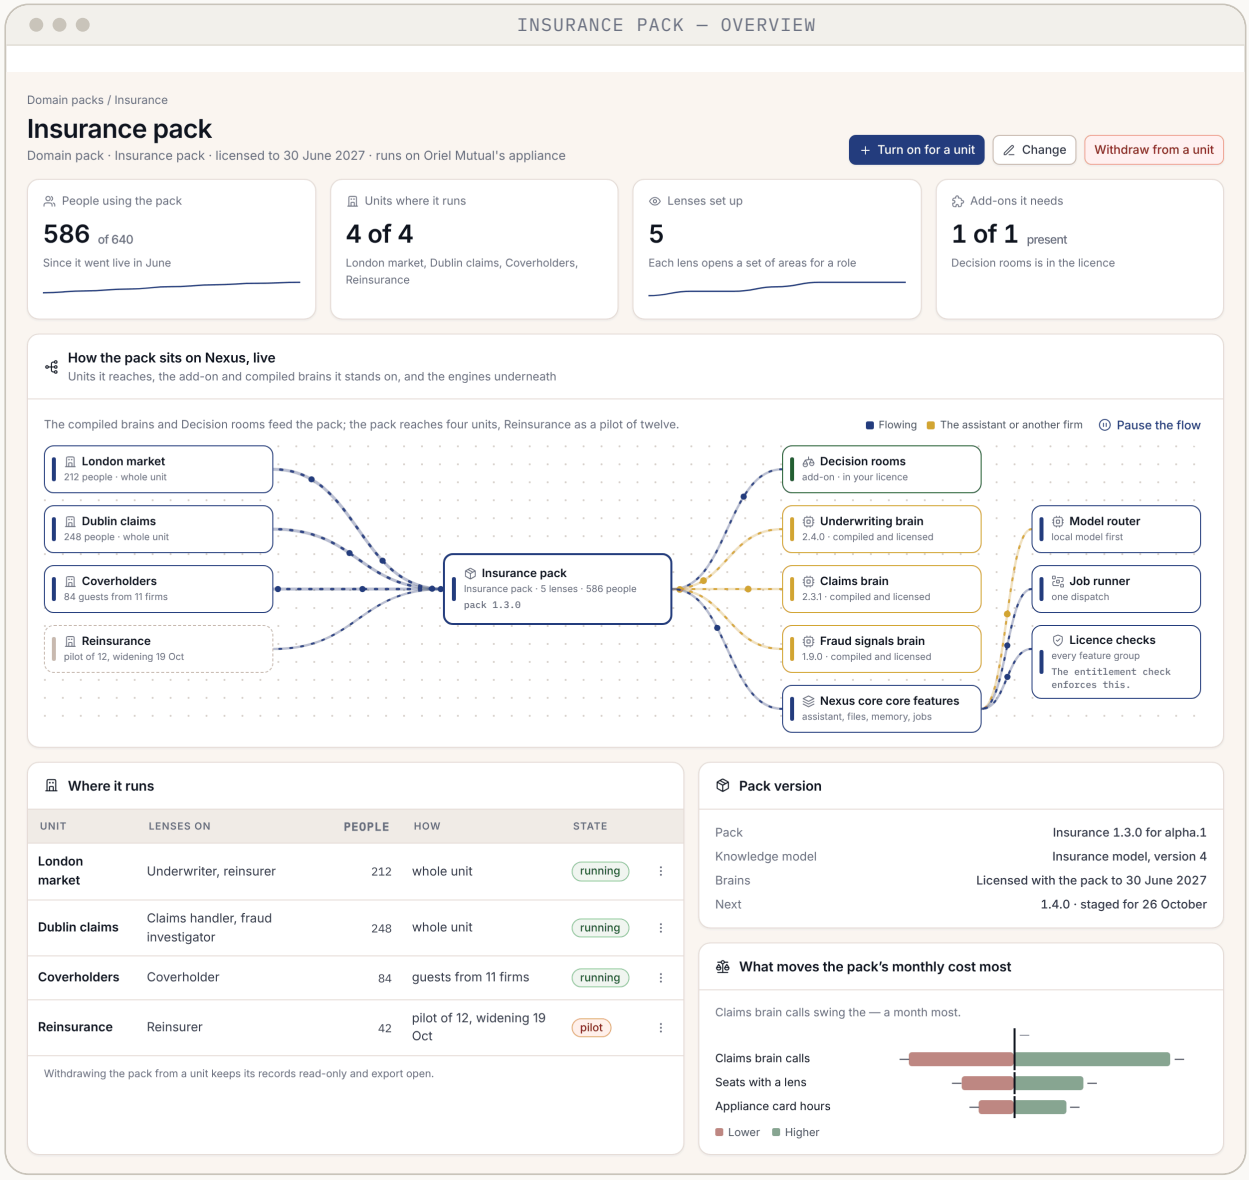
Legal and professional services. Every new matter gets a conflict check against parties and contacts. When a conflict is cleared, an ethical wall goes up, and it binds the assistant as well as people. Lifting a wall takes two partners. Lawyers capture time on their phones, but privileged files never reach one.

Logistics and maritime. The knowledge model knows ports and vessels, carriers and customs holds. Every new carrier is checked against sanctions lists before its first booking, and a person releases every hold. Port inspectors record yard and berth inspections on their phones, and releasing a hold there needs a passkey.



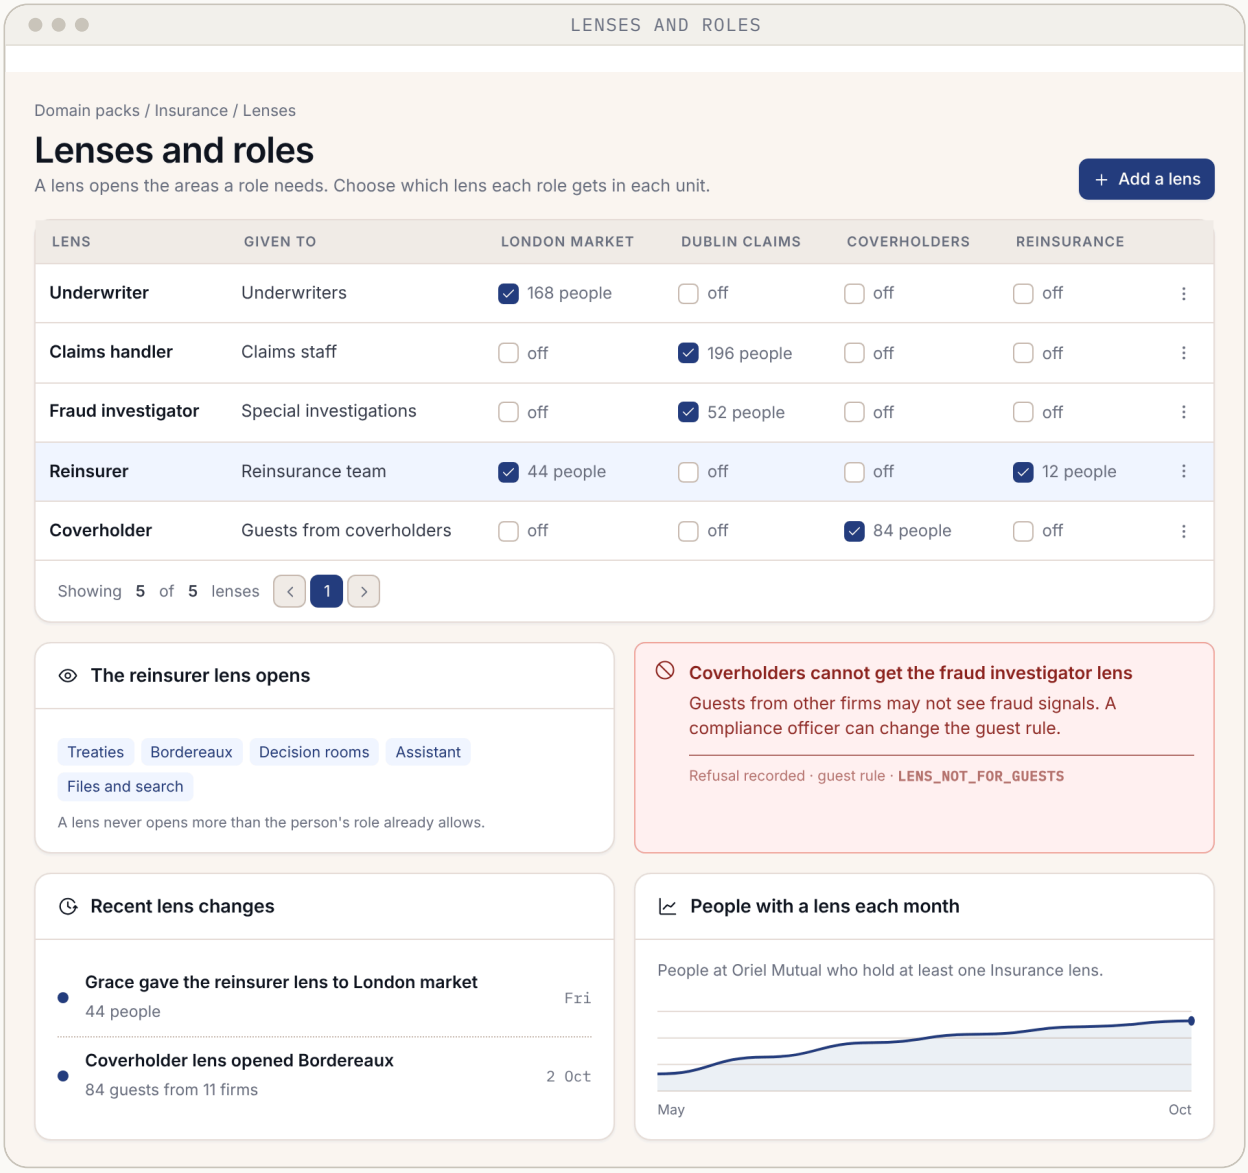
Worked example: Insurance at Larkfield & Vane

Larkfield & Vane runs the Insurance pack on Enterprise, with the Decision rooms and sealed records add-on the pack needs. Here's what its administrators set up, and what Sinéad and Priya Natarajan, the compliance officer, get from it.



SCREEN The pack overview shows where Insurance runs unit by unit, the add-on it stands on and the compiled decision models licensed with it. Notice the pilot unit, which widens on a date the administrator chose.

The **lenses** are set in one matrix of role by unit. Underwriters in the London market unit get the underwriter lens; claims staff in Dublin get the claims handler lens. Guests from coverholder firms get a coverholder lens, and that one matters. Coverholders see the bordereaux they send and claims on their own business, but never a fraud signal. If an administrator tried to give guests more than their role allows, Nexus would refuse the change with the reason.



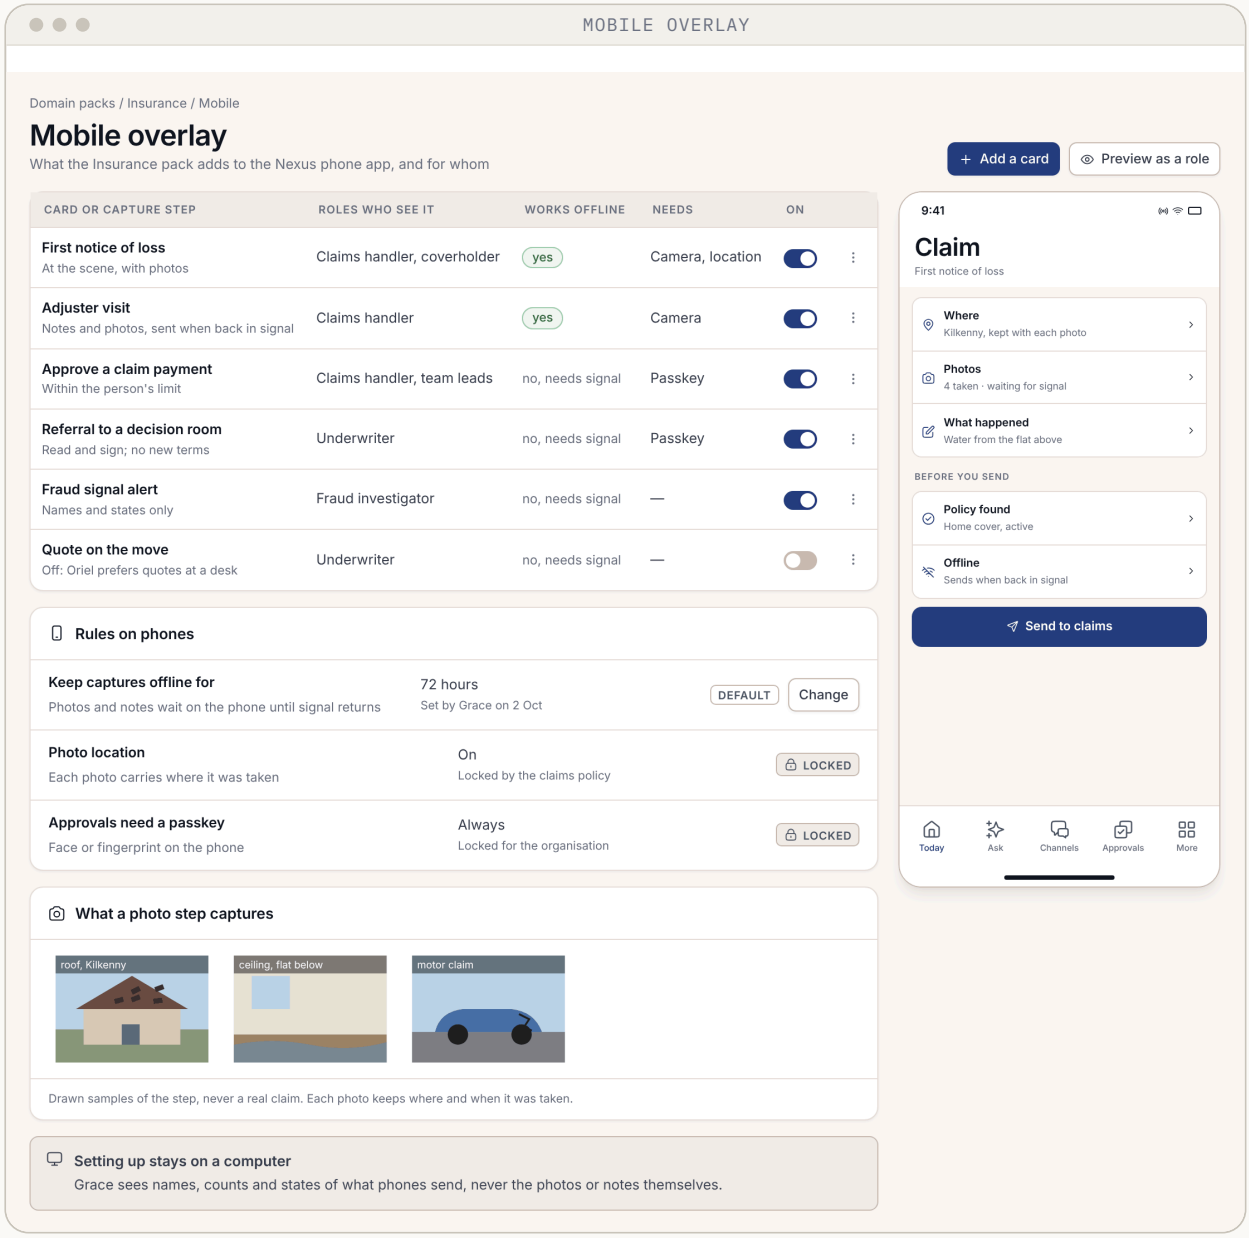
SCREEN Lenses are given per role and per unit in one place, with the number of people each choice reaches. A lens never opens more than the person's role already allows.

The **templates** arrive with the pack. Claims triage and a bordereaux check are on in Sinéad's unit. A template that needs an add-on Larkfield & Vane doesn't hold is greyed with that reason, never hidden.

The **knowledge model** knows policies, claims and treaties, and its vocabulary includes Larkfield & Vane's own terms. Connectors read the policy system nightly and the claims system every 15 minutes. Nothing in either system is changed from Nexus, and each connector has a test you can run from the page.

The **evidence pack** is where Priya spends her time. It lists what Solvency II and the Insurance Distribution Directive ask for, with the Digital Operational Resilience Act (DORA) and the AI Act beside them, and against each item the records that answer it. A gap reads "two coverholder contracts not loaded" and has an owner. When the supervisor asks, she exports the covered items. Each claim decision adds sealed records as it goes, so the evidence exists before anyone asks.

The **mobile overlay** decides what the pack adds to phones. A claims handler at a storm-damaged house takes photographs and notes, which wait on the phone for up to 72 hours until there's signal. A repair estimate goes back as a suggestion the office decides. Signing off a claim can't be done out there at all. Chapter 1.10, Away from the desk, explains why.



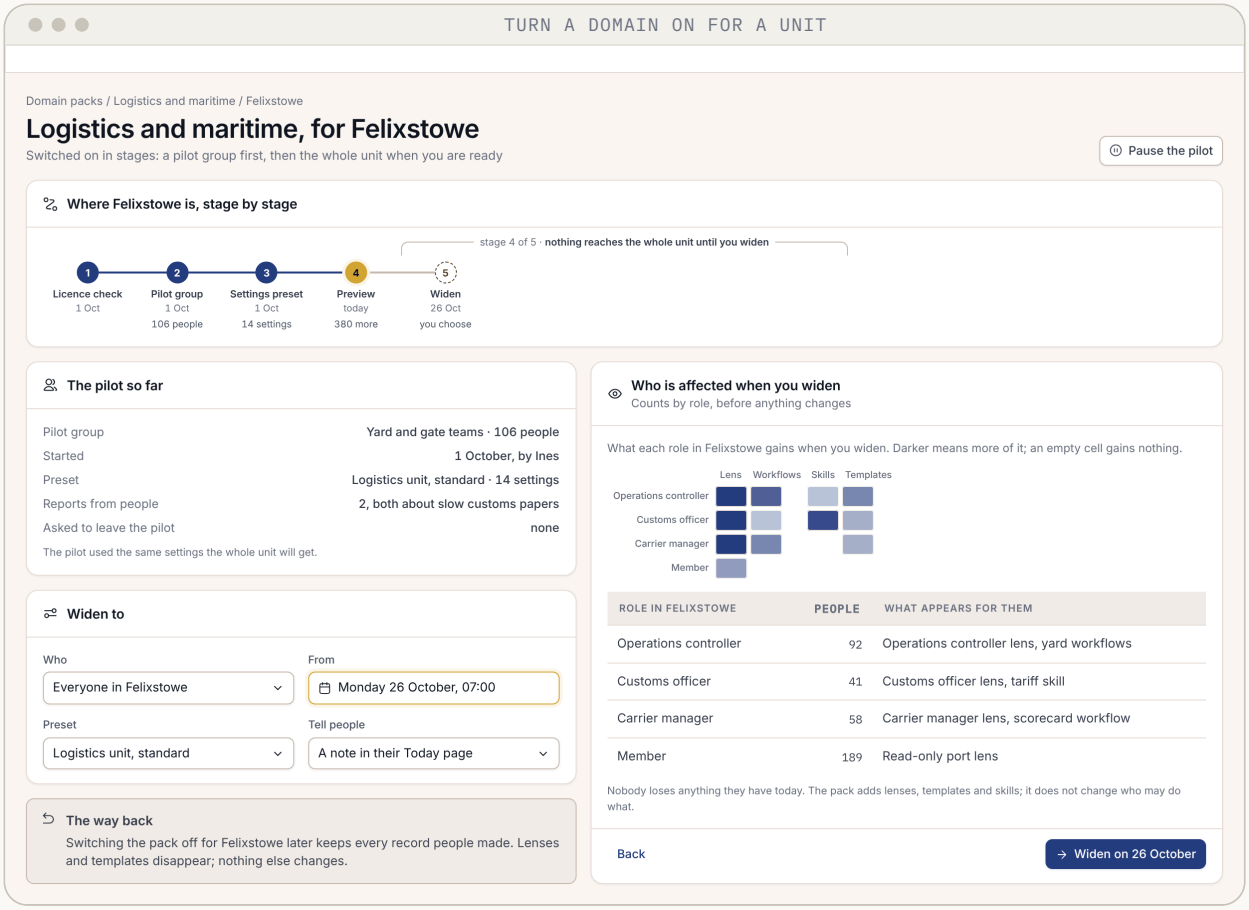
SCREEN The mobile overlay chooses which pack cards and capture steps appear on phones, and for which roles. The administrator sees names, counts and states of what phones send, never the photos or notes themselves.

Switching a pack on for one unit

Adverant grants a pack in your licence. Where it runs is your decision, made unit by unit. Larkfield & Vane's reinsurance team started as a pilot of twelve people, and that's the pattern to copy.

The switch happens in stages, on the page for that unit:

1. The licence check confirms the pack and the add-ons it depends on are in your licence.
2. You choose a pilot group, for example one team of about a hundred people.
3. You pick a settings preset for the unit. You can change it later.
4. You preview who is affected: which roles gain which lens and templates, and how many people each.
5. You set a date to widen to the whole unit, and choose whether people get a note on their Today page.



SCREEN A pack is switched on for a unit in stages: licence check, pilot group, settings preset, a preview of who gains what, then widening on a date you choose.

Nobody loses anything they already have when a pack arrives. The way back is written on the same page. Switching the pack off for a unit keeps every record people made: lenses and templates disappear, the records turn read-only, and export stays open.

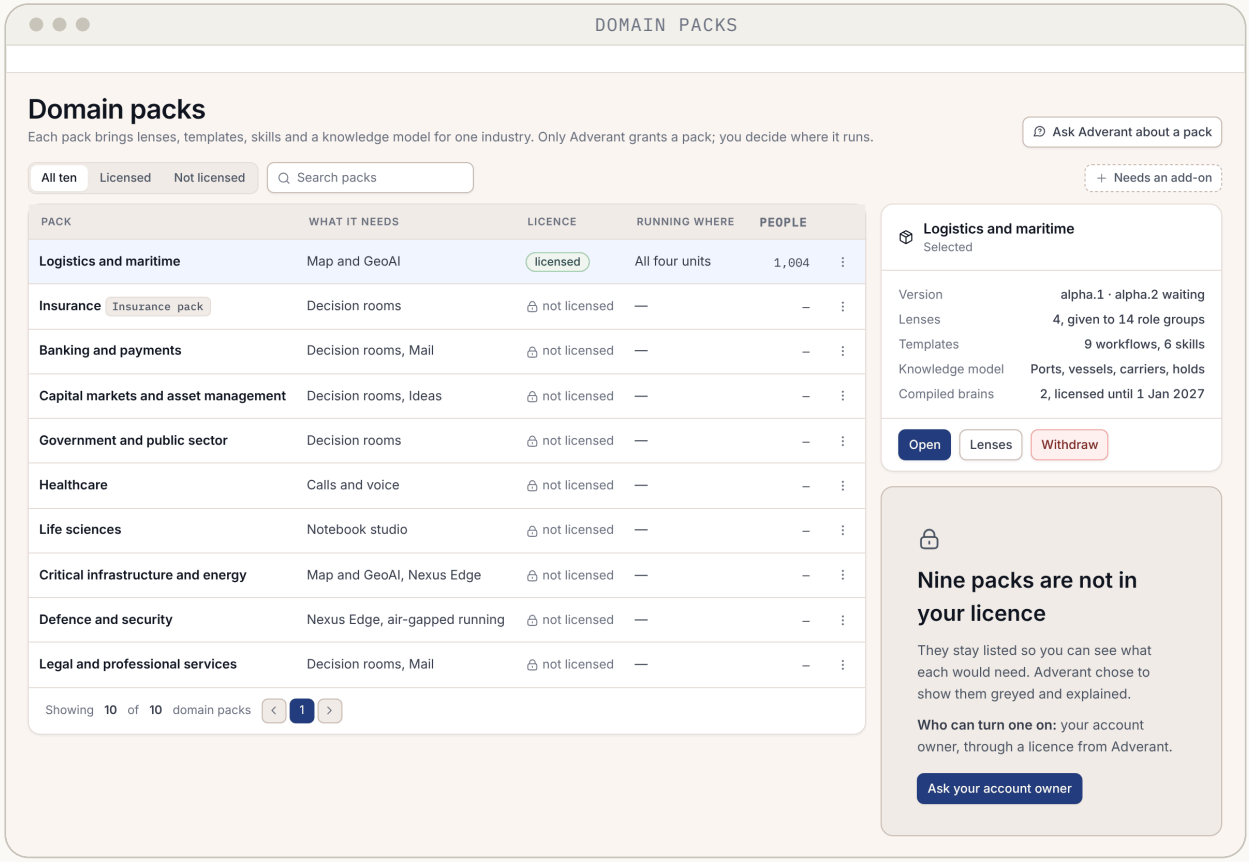
What a pack needs

Every pack needs Enterprise. Some need particular add-ons as well, and the licence carries those dependencies.

Pack	Also needs
Insurance	Decision rooms and sealed records
Capital markets and asset management	Decision rooms and sealed records
Legal and professional services	Decision rooms and sealed records, and Mail
Logistics and maritime	Map and GeoAI
Defence and security	Nexus Edge, for alerts from devices in the field
Any pack's phone screens	Mobile apps

Adverant won't issue a pack without the add-ons it depends on, which means a pack never arrives half working. Many customers add more than the table lists. A bank often wants Mail, for instance, but that's a choice. Adverant quotes the price of each pack and add-on.

The packs you don't hold stay listed on the Domain packs page, greyed, with what each would need and who can turn it on: your account owner, through a licence from Adverant. Opening one loads nothing and sends nothing anywhere.



SCREEN All ten packs, with what each needs, whether it's licensed and where it runs. Packs outside the licence are greyed with the upgrade path.



WATCH OUT

A pack's phone screens need both the pack and the Mobile apps add-on. If you plan field work around a pack, put both in the same licence request.



Building your own vertical

If your industry isn't on the list, a pack is built from the same parts as these ten. You describe the work, model the domain starting from the nearest existing pack, add lenses and rules, then test on sample data and release it unit by unit. Adverant can build it with your experts, or your own developers can use the developer kit, and a pack you build stays yours. Chapter 1.11, Making it yours, explains how to brief that work. Chapter 2.24, Building a custom vertical end to end, is the technical walk-through.

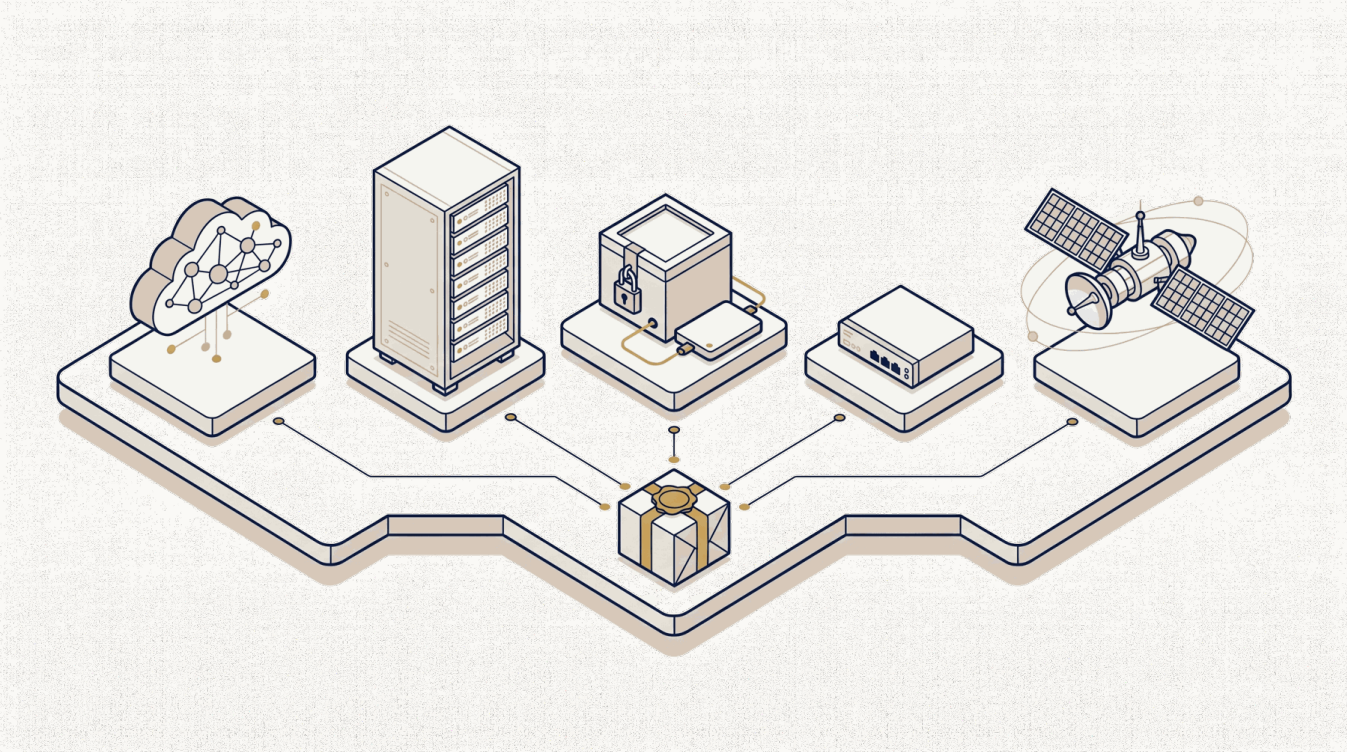
PART 1 · CHAPTER 1.8

Where Nexus runs

The places Nexus can run, from Adverant's EU cloud to a site with no internet and devices in the field, what sovereignty means in practice, and how to choose.

FOR EVERYONE

11 MIN READ



Harrow Port Authority has no internet on the network where Nexus runs. When a new release comes out, someone carries it in on a disk. Aoife Brennan, the head of IT, watches the import page check the disk's signature against the release key Harrow installed months ago, and only then does anything get copied. People keep working on the version they have. The upgrade waits until her team chooses to run it, with a backup first.

A few kilometres away, an insurer might run the same product in Adverant's cloud and never think about disks at all. It's the same software in both places. What changes is who runs the hardware and how it connects. This chapter goes through the options and what each costs you in effort, then gives a short guide to choosing.



One product, several places

Every delivery form gets the same platform, the same domain packs and the same kind of sealed records. A screen in this documentation looks the same wherever Nexus runs. The differences are practical: who patches the servers, where the data sits, which models are in reach, and how updates arrive.

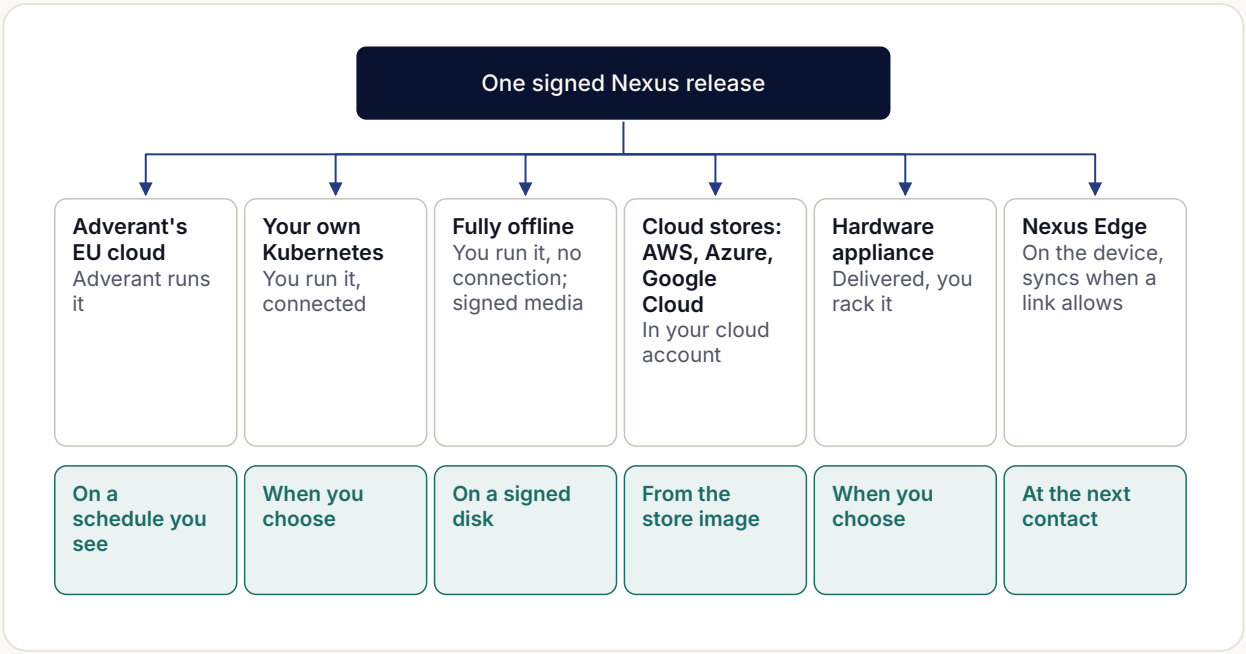


FIGURE The same Nexus release reaches six places, and only the hardware owner and the connection change.



Adverant's EU cloud

Adverant runs Nexus on sovereign compute in the European Union: graphics processing units (GPUs), open models and Kubernetes, under European law. Each organisation gets its own private space with its own keys and budgets. Data is stored and processed in the EU, and an Enterprise administrator can set where each kind of data lives on the Where data lives page.

It's the quickest way to start, because nobody on your side installs anything, and Adverant applies updates on a schedule you can see. Teams that need to keep data on their own premises often begin here and move later, once the work has proved itself.



Your own Kubernetes

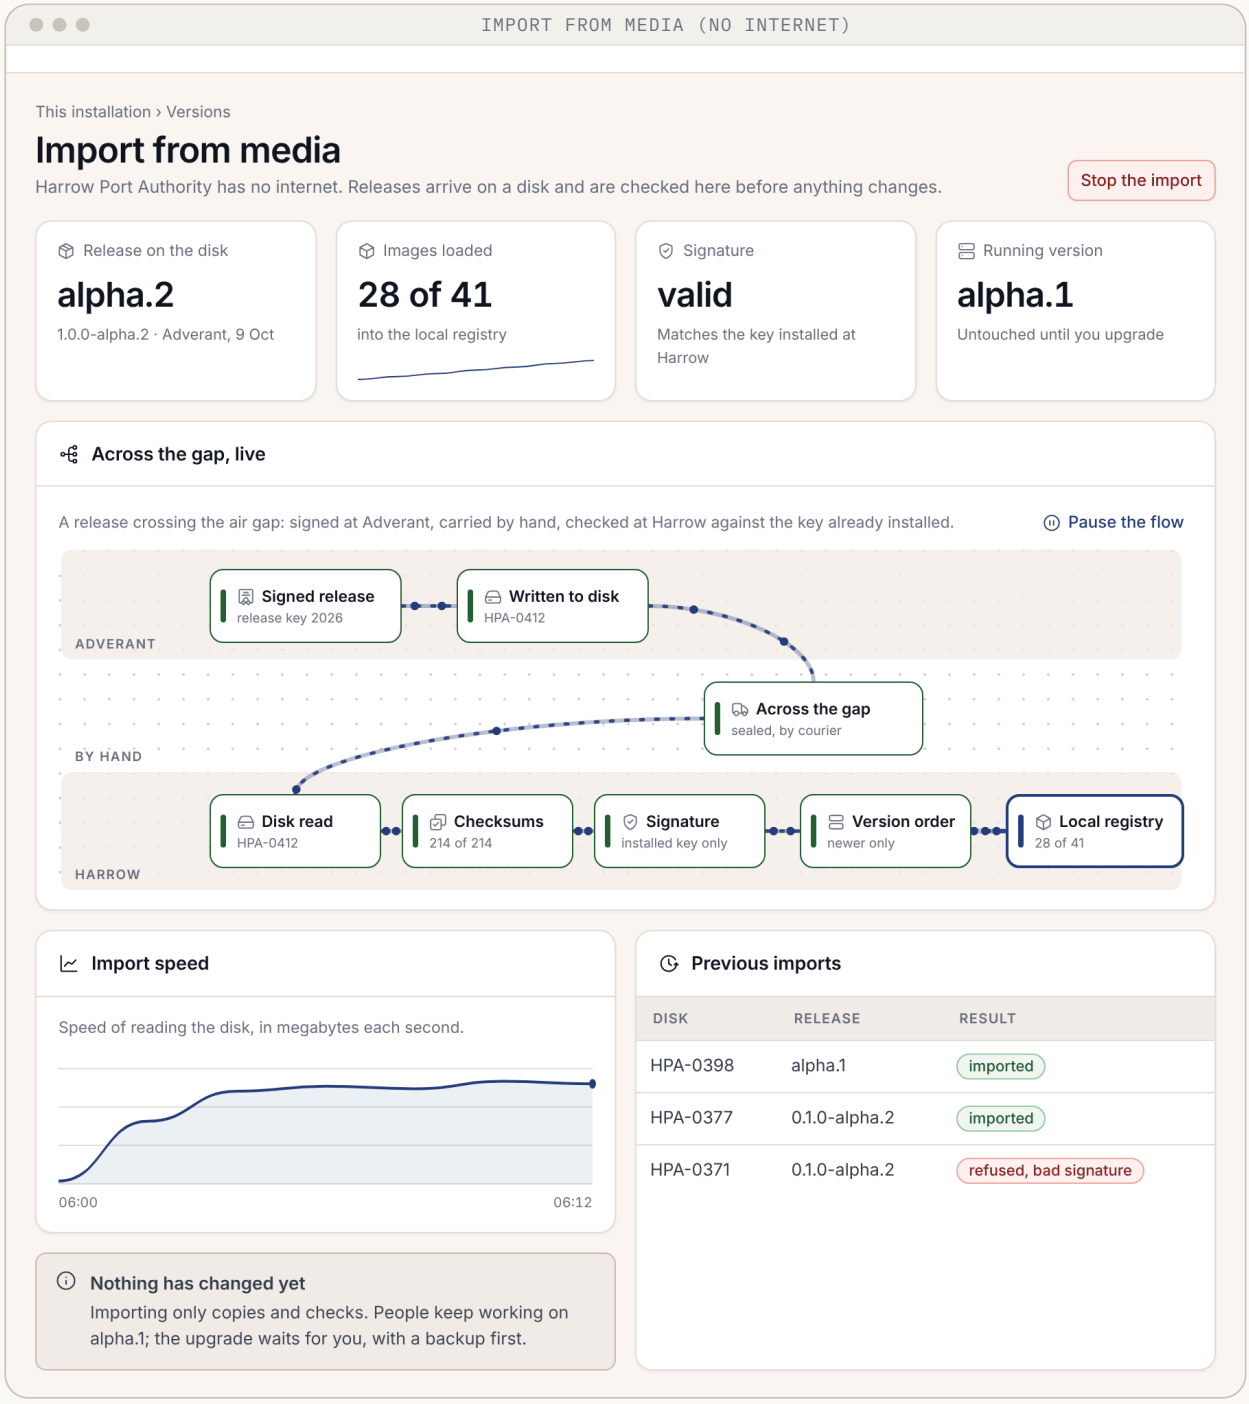
If you have a platform team and a Kubernetes cluster, Nexus installs there with a standard chart. Your IT team runs a preflight check first, which lists every problem with its fix before anything is installed. The install itself takes about twelve minutes on a cluster that passes.

You connect your own company sign-in, your own models on your own GPUs and your own key store. You choose when to upgrade. Upgrades are staged, take a backup first, and keep a way back with a date on it. While the installation is connected, a renewed licence can arrive by itself and wait for an administrator to accept it. Chapter 2.3, Planning an installation, is where your IT team starts.



Fully offline

For classified sites and sealed networks, Nexus runs with no connection at all. Updates and new packs arrive on signed media. The installation checks each disk's signature against the release key you installed, never against a key the disk carries, and anything that fails the check is rejected before it is opened. Nothing changes during an import: people keep working on the version they have.



SCREEN Importing a release from a disk at Harrow Port Authority. Notice that the signature is checked against the key installed at Harrow, and that the running version stays untouched until someone upgrades.

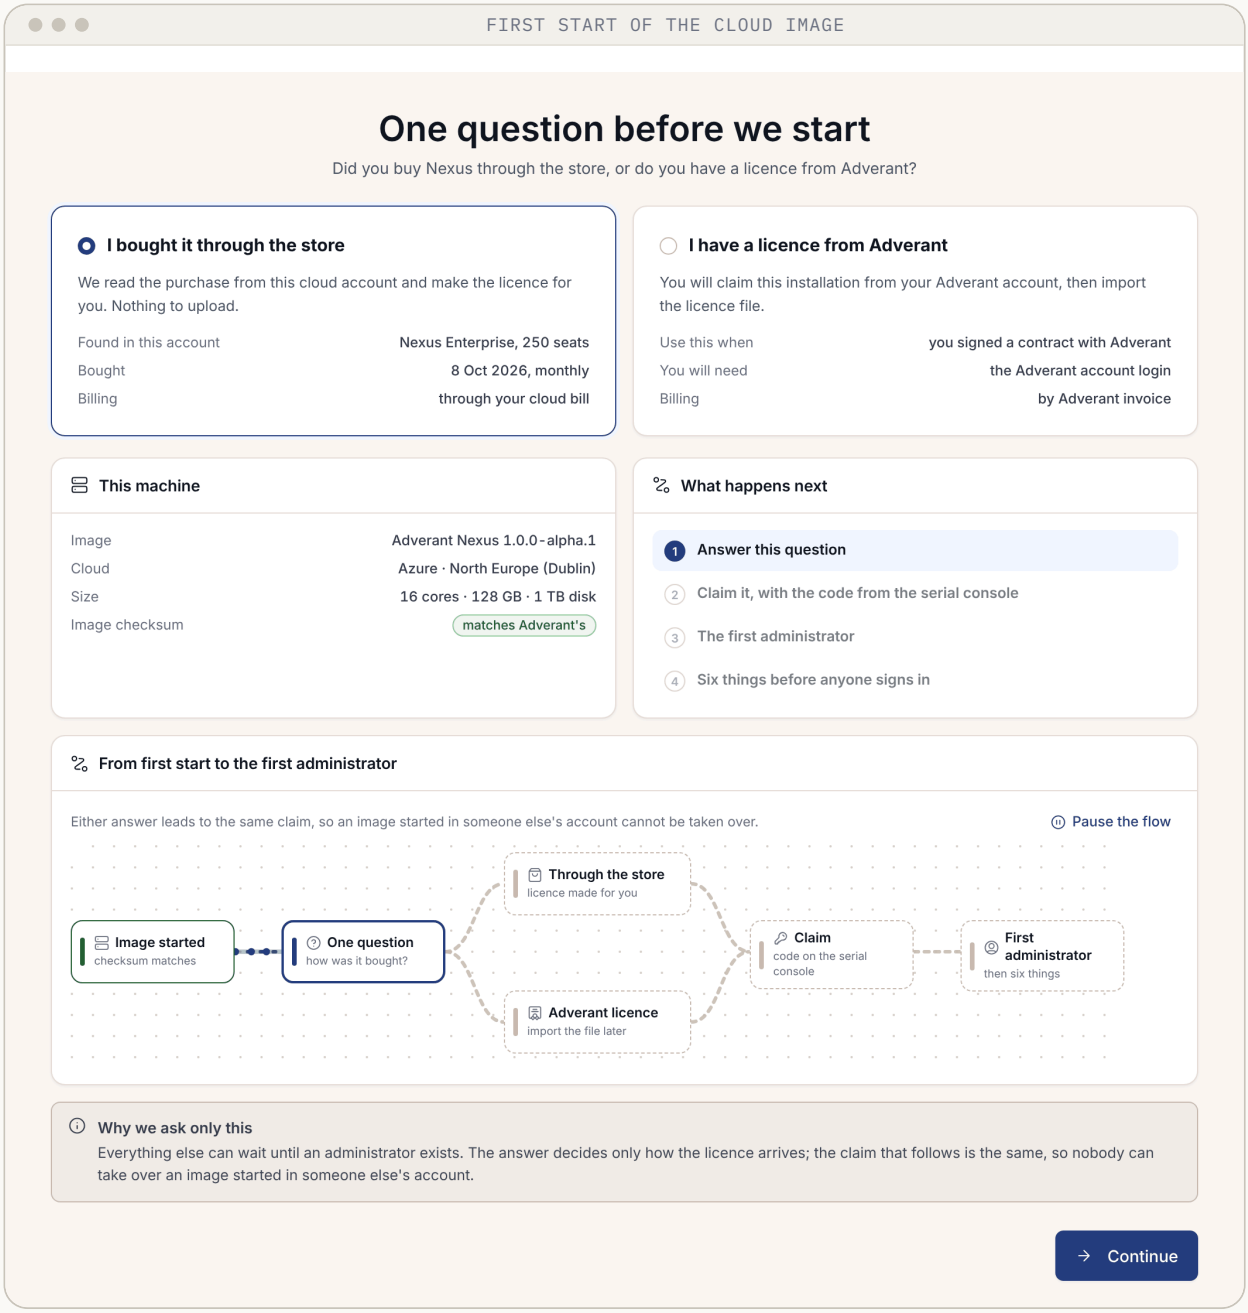
Licences travel the same way. Your administrator exports a small signed request file, a person carries it out, and a signed licence file comes back. The request names only the installation and the feature, with a fingerprint. The models are the ones on your own hardware, and the artificial intelligence (AI) switchboard uses nothing else. Time stamps on sealed records are witnessed inside the site and still check out later.

The cost is effort. Someone has to fetch, carry and import each release, and nothing arrives by itself. If a licence ever had to be withdrawn from an offline site, that would be a matter for the contract, because no signal could reach it.

In your cloud account, from the store

If your organisation buys through Amazon Web Services (AWS), Microsoft Azure or Google Cloud, you can buy Nexus in that cloud's store and pay through your cloud bill. The store sends you to an Adverant page where you sign in with a work account, because the subscription belongs to an organisation, never to a person. Then you start Nexus as a machine image in your own cloud account.

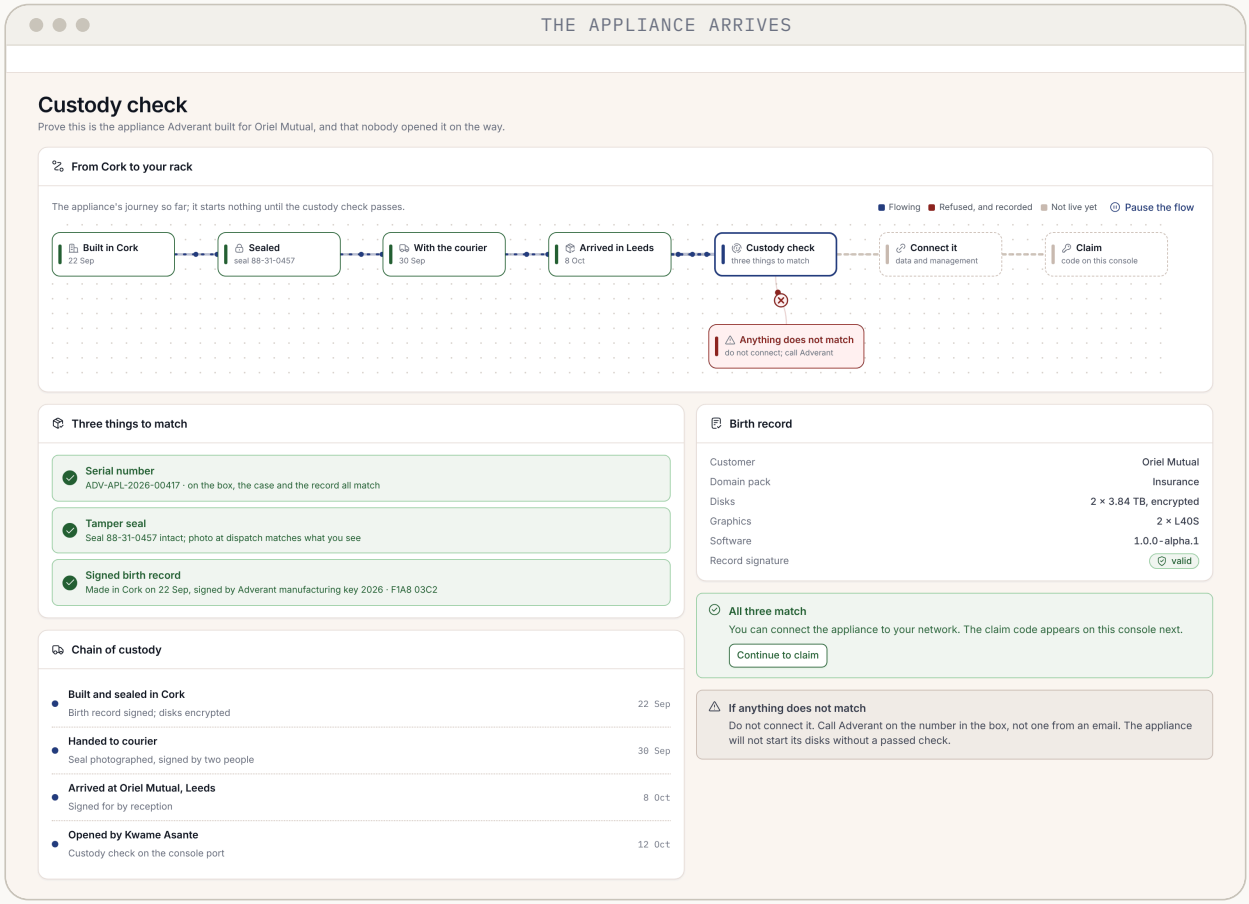
The first page of a new image asks one question: did you buy it through the store, or do you have a licence from Adverant? The answer decides only how the licence arrives. The claim that follows is the same either way, with a code printed on the machine's serial console, so nobody can take over an image someone else started. Nothing is billed until the licence exists, and if a store deadline lapses, the page says so plainly and nothing is lost.



SCREEN The first start of a cloud image asks only how Nexus was bought. Everything else waits until an administrator exists.

A hardware appliance

Some customers want servers delivered ready to run, with GPUs for their own models. The appliance comes from one of Adverant's hardware partners, on request; Adverant doesn't sell the box itself. When the box arrives, a laptop plugged into its console port runs a custody check, and the disks don't start until it passes. Three things must match: the serial number on the box, the case and the page; an unbroken tamper seal; and a birth record signed by Adverant.



SCREEN The appliance checks its own custody before anything starts. If one of the three checks fails, the page says not to connect it and how to reach Adverant safely.

WATCH OUT

If anything on the custody check doesn't match, don't connect the appliance. Call Adverant on the number printed in the box, not a number from an email.

Nexus Edge, in the field

Nexus Edge runs on the computers aboard a satellite, a drone, a vessel or a vehicle. Those computers are separate from the craft's own flight or control computer. On board, Nexus runs the knowledge graph, reads files and runs the agents for one mission, and it keeps working when the link drops. Results come home at the next contact with a ground Nexus.

Corrán Orbital, for example, flies an Earth-observation satellite whose images are read in orbit, so only what matters comes down in each contact window. Three rules keep a device in check:

- Mission packs are signed before loading. A pack holds what one mission needs (maps, models, rules). A new pack goes into a second slot, so the previous one stays ready if the new one fails

its checks, and when the mission ends, or a tamper switch trips, the keys on board are wiped and a signed receipt says so.
- Limits are set on the ground. A person decides what the device may do alone, what waits for approval however long the link is down, and what it may never do. With no link, anything above the limit waits.
- The device can ask, never command. Requests to the flight computer go to that computer's own arbiter, which accepts, refuses or defers each one.

What comes down first is set too. The outbox on board sorts everything into four kinds: alerts first, then changes to the knowledge graph, then finished products, then raw data. Raw data comes down only when an analyst on the ground asks for it and gives a reason, and it travels on the next fast link.

When one of the on-board computers stops, the device drops to a lower step that keeps detection running, writes its reports from a template, holds questions until the link returns, and asks the flight computer's arbiter to restart the board. Each device's page shows those steps in advance, so the team knows what the mission will still do before anything fails.

Each device is enrolled before it leaves, and a missed contact window raises a flag for a person. Agents on a device run on a lease with an end date, so a lost device doesn't keep working for anyone. Nexus Edge is an add-on to Enterprise.



Sovereignty in plain terms

"Sovereign" gets used loosely. In Nexus it comes down to four promises, and each holds in every delivery form above.

Your data stays where you chose. At the first run, the first administrator sets where data is kept. On Enterprise you can set a place for each kind of data, and moving any of it later is a tracked job, previewed first, that a second person approves.

AI goes only to the models you configured. Every model call passes through the AI switchboard, which sends it only to a model your organisation has connected and allowed. Calls to any other outside address are refused with the code `EGRESS_DENIED` and counted. An organisation with no model connected gets no AI answers at all; Nexus says so and shows how to connect one, rather than quietly using someone else's.

Licences are checked offline. Adverant signs every licence with two independent keys. Your installation checks the signatures itself, without calling anyone, and a licence that fails the check is treated as absent.

Nothing phones home. On your own installations, health reporting to Adverant is off unless you turn it on. Once a year your administrator sends Adverant a usage statement of counters, which can be read in full before it leaves; an offline site carries it out on media. A support bundle goes only when you ask for help, after a preview of what it holds with the private parts removed.



NOTE

One thing does leave a connected installation if you use the native phone apps: notifications. They pass through Adverant's relay as a device address and sealed content that only the phone can open. The Where calls can leave page declares this. Customers with their own push setup send nothing to Adverant, and an offline site has no push at all.



WATCH OUT

The place where data is kept is fixed at the first run. Changing it later means a new installation and a move, so decide it before the first administrator is created.



Moving later

Many teams start in Adverant's EU cloud and move to their own data centre once the work is proven. Records, settings and packs move with you. Sealed records keep their seals, so an auditor can check them before and after the move. People keep their access once you connect your company sign-in on the new site.



Which is right for us

There's no wrong first answer, because you can move. Still, a few questions settle most cases.

If this sounds like you	Start with
You want one team using it this month, and EU hosting meets your rules	Adverant's EU cloud
You have a platform team and your rules say data stays on your premises	Your own Kubernetes
The network has no internet, by policy or by design	Fully offline
Your organisation has committed spend with AWS, Azure or Google Cloud	The cloud store
You want your own GPUs without building a cluster	A hardware appliance
The work happens on ships, drones, satellites or field kits	Nexus Edge, with a ground Nexus in one of the other forms

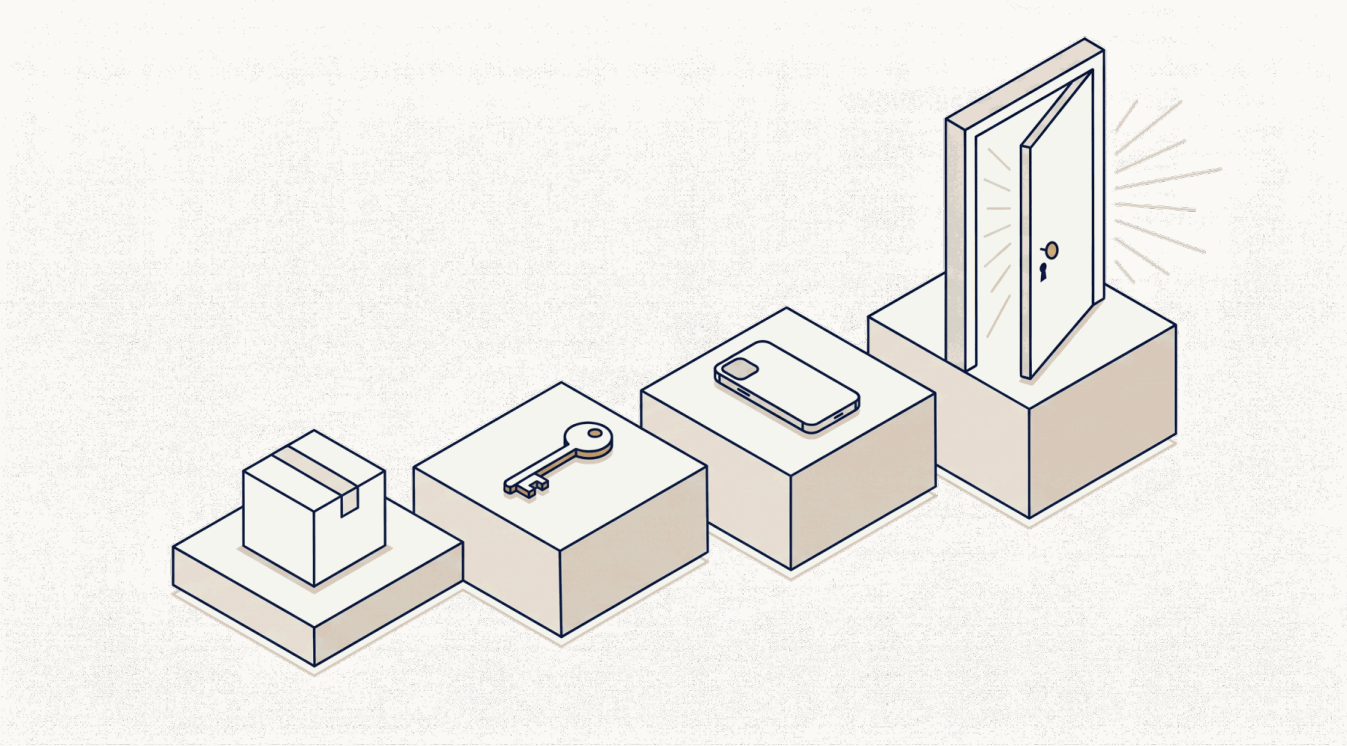
If two rows fit, pick the one your compliance officer can sign off now. Chapter 1.9, Getting started, covers what happens next, and chapter 2.3, Planning an installation, is the IT team's version of this decision.

Getting started

The two decisions before you start, what your IT team does, the first run as the first administrator sees it, and a checklist of what to have ready on day one.

FOR EVERYONE

13 MIN READ



Aoife Brennan's first hour with Harrow Port Authority's new installation went like this. Tomás Reilly, her platform engineer, finished the install and sent her a link with a code in three short groups. She typed the code and compared eight groups of letters on her screen with the eight on his. Then she ticked a box and filled in two short forms. Her laptop asked for a passkey, she touched the fingerprint reader, and she was looking at a Today page that said, honestly, that nothing had happened yet.

There is very little more to a first run than that. This chapter covers the two decisions before it and what IT does, then each step of that hour and the first day after it.



Two decisions first

Which edition. Community is the free, open-source edition for one organisation. Team adds a set of features on Adverant's cloud. Enterprise is licensed feature group by feature group, and it's the edition that runs domain packs and the add-ons, the phone apps among them. Chapter 1.6, Editions and licences in plain terms, sets out what each holds. You can move from Community to Enterprise later and keep everything you already have.

Where it runs. The choices run from Adverant's EU cloud to your own Kubernetes to a site with no internet at all, with cloud stores, appliances and field devices besides. Chapter 1.8, Where Nexus runs, explains each. On Adverant's cloud you sign up and skip straight to the first five minutes, described at the end of this chapter.



What your IT team does

For any installation you run yourself, your IT team does the groundwork. They check the release is genuine and run a preflight check that lists every problem with its fix. Then they install, and make a one-time claim code for the email address of the person who will be the first administrator. Part 2 is written for them. They start at chapter 2.3, Planning an installation, and chapter 2.12, The first run, covers their side of the hand-over.

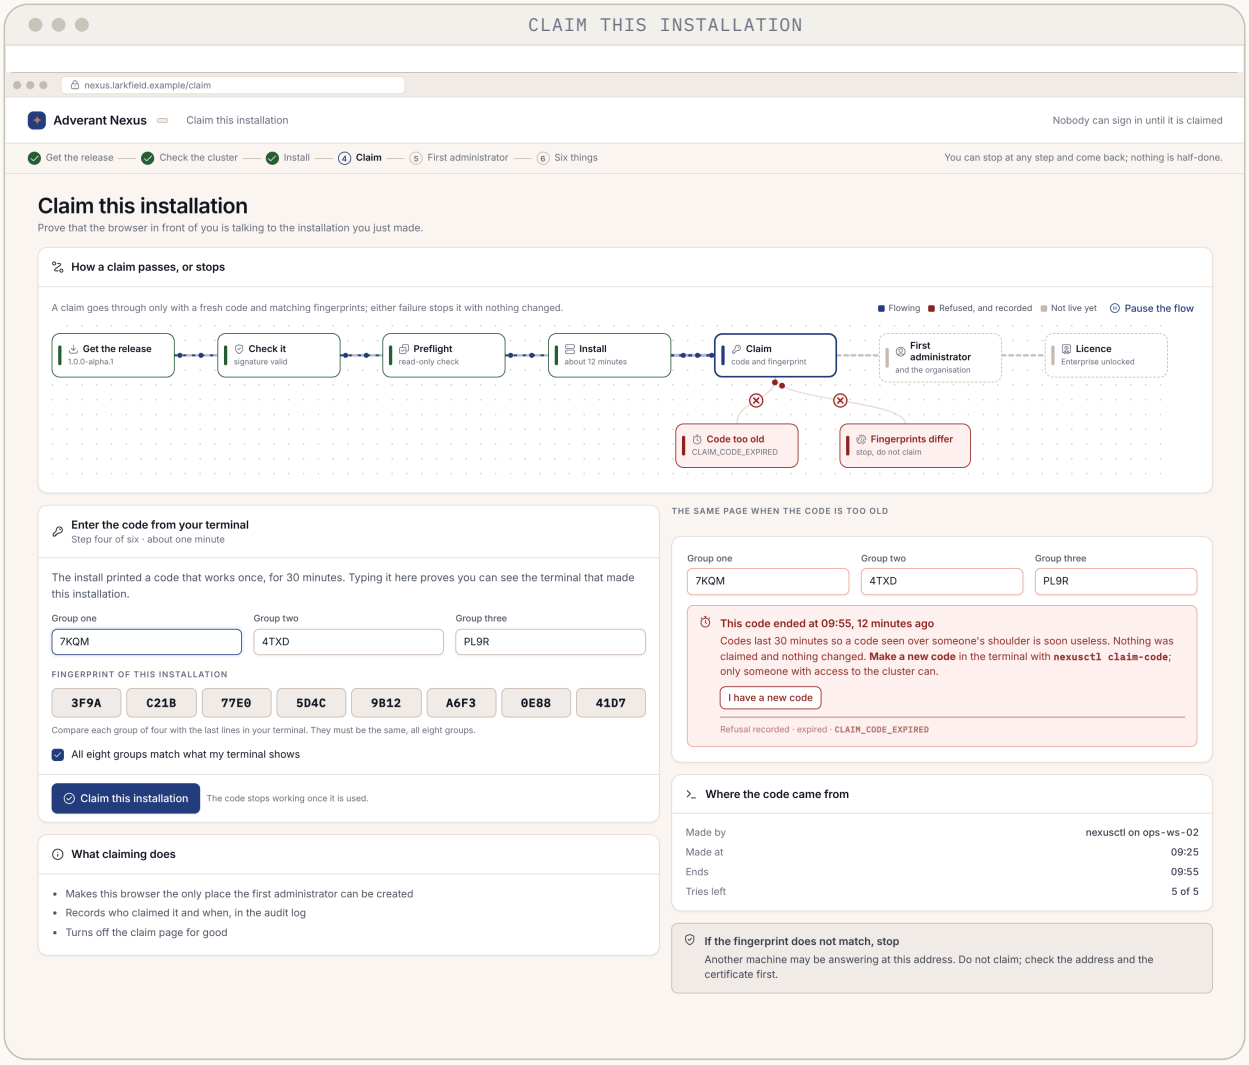
The claim code lasts 30 minutes and allows five tries. Agree a time with IT, so the person claiming is at a browser when the code arrives.



Claiming the installation

Claiming proves that the browser in front of you is talking to the installation your team just made. Nobody can sign in until it's done.

1. Open the claim address IT gives you.
2. Type the code into the three boxes, one group in each.
3. Compare the fingerprint on the page, eight groups of four characters, with the one IT reads you from their screen, group by group.
4. Type the work email address the code was made for.
5. Tick the box that says all eight groups match, then press **Claim this installation**.



SCREEN The claim page asks for the code from the terminal and shows the full fingerprint in groups of four, so a person can compare it line by line. An expired code is refused with what happened and how to get a new one.

Claiming does three things. It makes this browser the only place the first administrator can be created. It records who claimed and when. And it switches the claim page off for good.

If the code is more than 30 minutes old, the page says so and confirms nothing was changed. Press **I have a new code**. Ask IT for a fresh one.



WATCH OUT

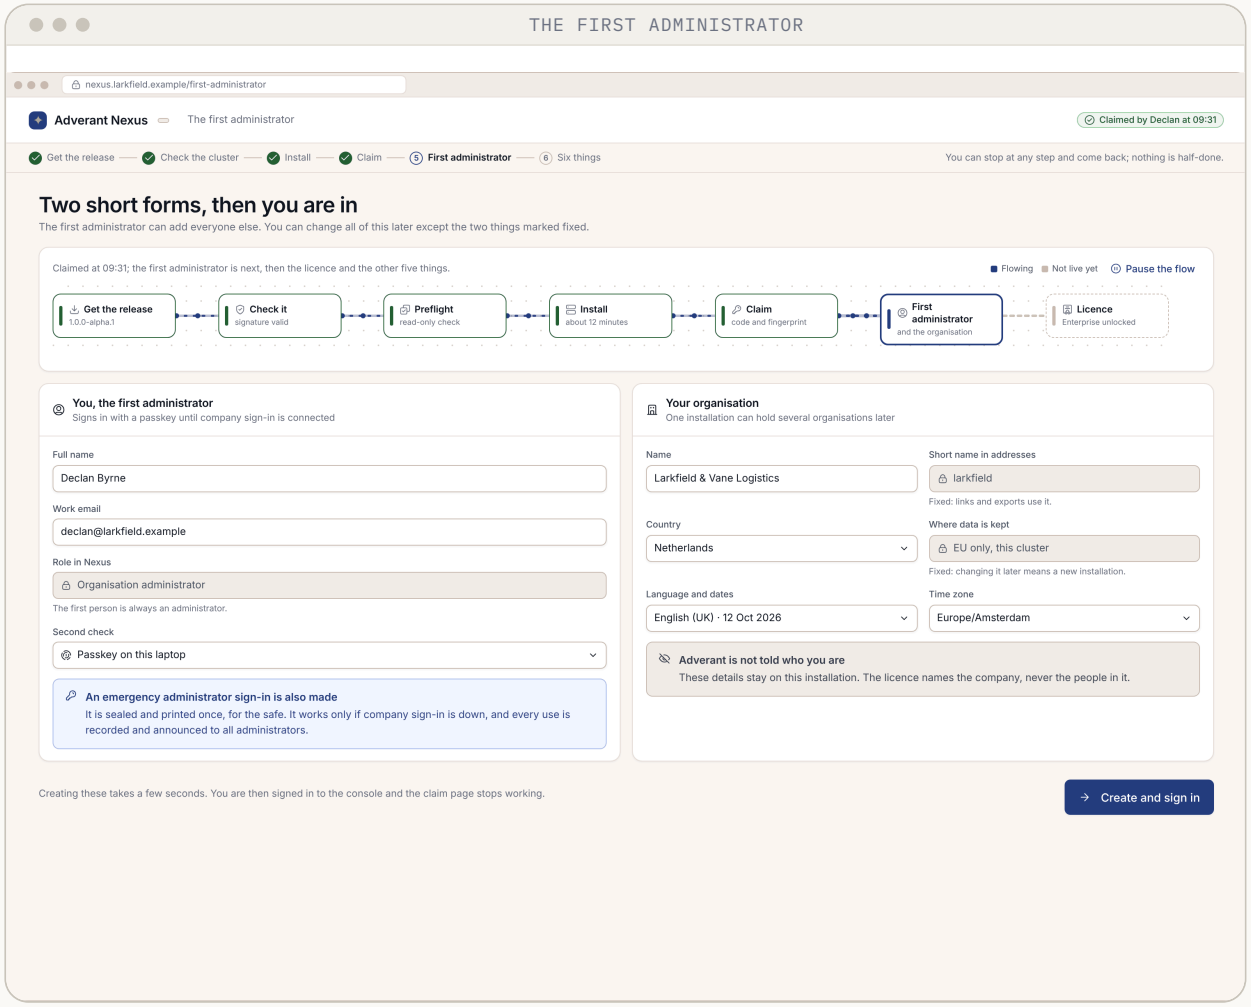
If any fingerprint group doesn't match, stop. Another machine may be answering at that address. Don't claim; ask IT to check the address and the certificate first.



The first administrator

Next come two short forms on one page. The first is about you: your full name and work email, plus a password. The first person is always an organisation administrator. The second is about the organisation. It asks for the organisation's name and a short name used in links, then where data is kept, then language and time zone.

Two of those boxes are marked as fixed. The short name appears in every link and export, and it can't change later. Where data is kept can't change either; moving data later means a new installation. Think before you fill them in.



SCREEN The first administrator and the organisation, on one page. Notice the two fields marked fixed, and the note that the licence names the company, never the people in it.

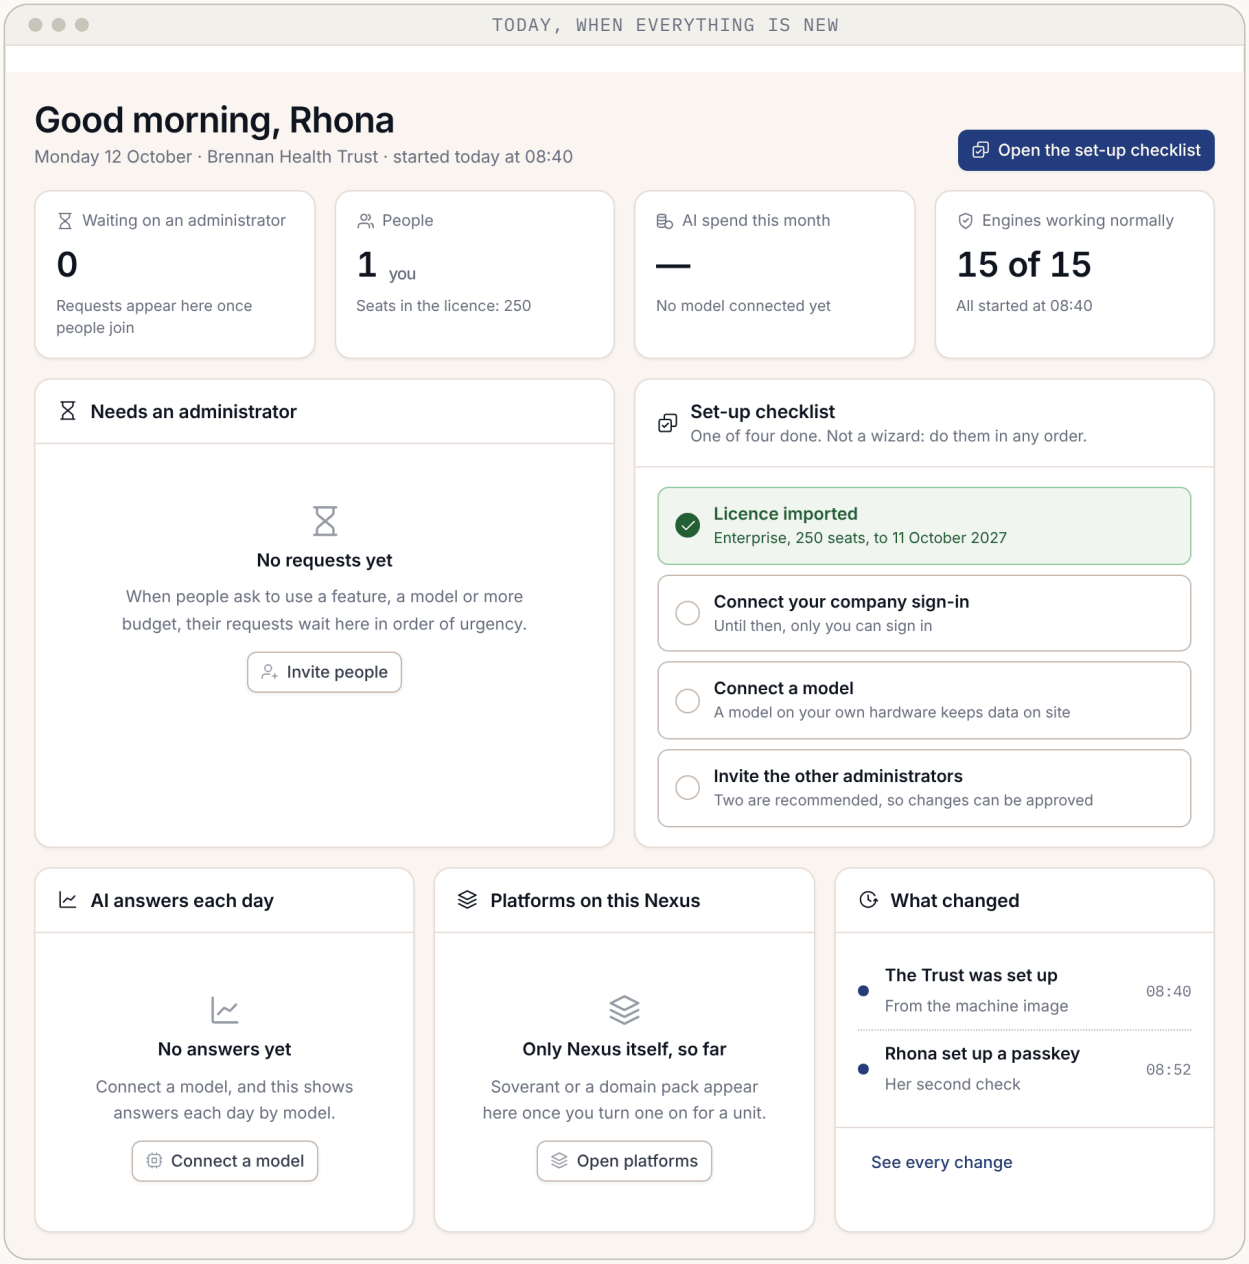
When you press **Create and sign in**, your computer asks you to make a passkey, a sign-in that uses your fingerprint or face (or your phone) instead of a typed secret. An administrator never signs in with a password alone, so the passkey is the second check every time. If your computer can't make one, the page shows a square code for your phone's camera instead. You have ten minutes in the same browser to finish this step.

Nexus also makes an emergency administrator sign-in at this point. It's sealed and printed once, for the safe. It works only when your company sign-in is down, and every use is recorded and announced to all administrators.



Your first day

Once the passkey is set, you land on Today. On a new installation it doesn't pretend to be busy. Each empty list says what kind of empty it is ("No requests yet") and how to start. The assistant in the side panel says the same thing in a sentence: nothing has happened here yet, and the usual first steps are connecting company sign-in and a model.



SCREEN Today on the first day of a new organisation. Every list is empty and says so, and the set-up checklist offers the next steps in any order.

The set-up checklist on Today has four items, and you can do them in any order. Import the licence if you're on Enterprise. Connect your company sign-in and a model. Invite the other administrators. Two other administrators are recommended, because some changes need a second person to approve them.

Behind it sits a page for the installation itself, **Six things before anyone signs in**. Its six items are the licence and backups, the address and its certificate, company sign-in and mail, and an artificial intelligence (AI) model. Each one shows "true" when it's done. The **Invite people** button opens when all six are true.

Connecting company sign-in

Most organisations want people to sign in with the accounts they already have. You choose your provider, and Nexus finds most of the settings from the provider's published details. You paste the application's identity and secret, which is never shown again, and press **Test sign-in**. The result says in plain words what the provider sent and how Nexus will use it. The name of the standard your provider uses sits in small type beside a short gloss.

Community connects one company sign-in through OpenID Connect (OIDC). Enterprise adds Security Assertion Markup Language (SAML) and two more things: you can require company sign-in from a date you choose, after which members' passkeys stop working on their own, and you can copy people and groups from your directory every 15 minutes. Until company sign-in is connected, only you can sign in.

Connecting a model

Every AI answer in Nexus goes through the organisation's AI switchboard, and the switchboard needs at least one model to send work to. You can connect a key for a hosted provider (Anthropic, Google Gemini, OpenRouter or Scaleway) or point Nexus at a model server your organisation runs on its own hardware. Nexus tests the connection with a short question before saving, and none of your files are sent in that test.

The key goes into the organisation's key pool. After you save it, only its last four characters are shown again, to anyone, including you. Because the key belongs to the organisation and not to you, it keeps working if you leave. A model on your own hardware is tried before any hosted one.

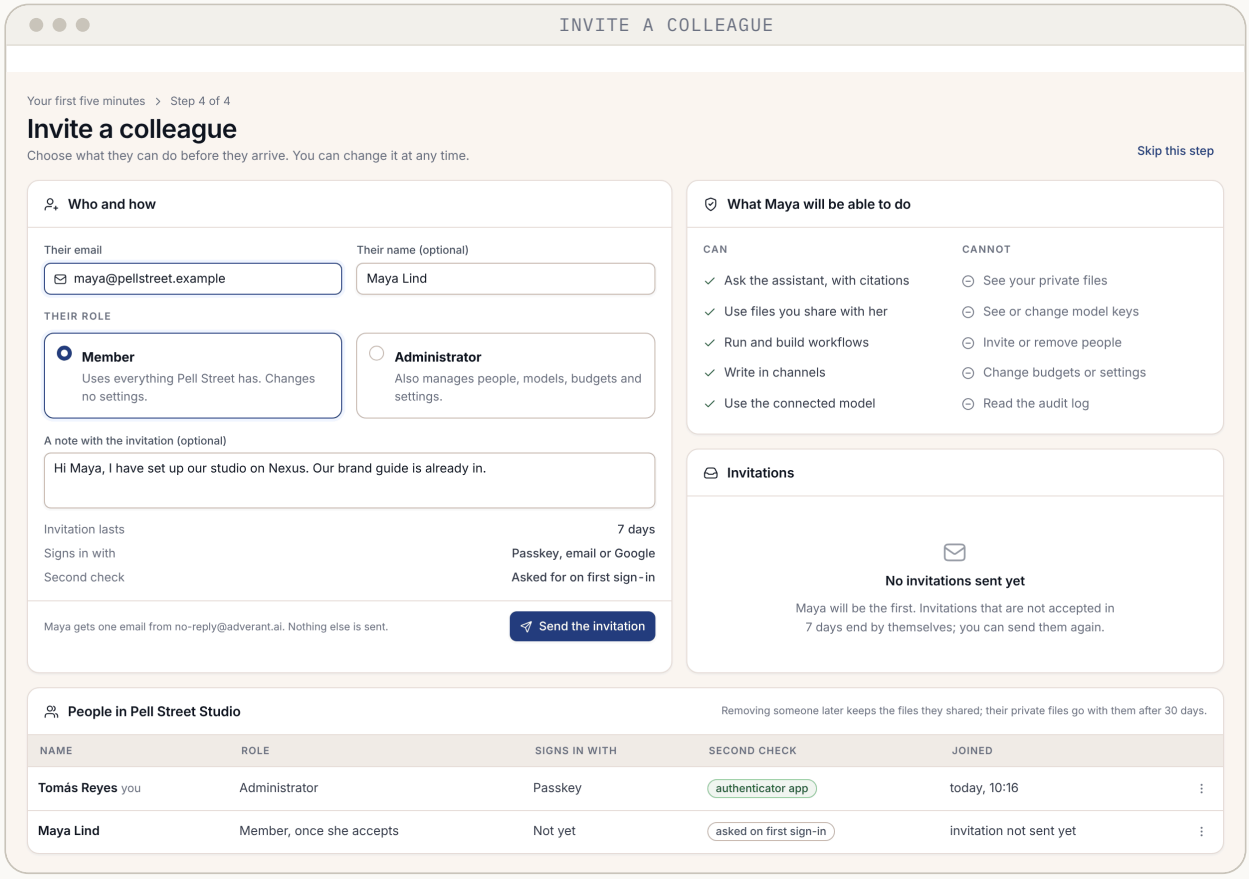


WATCH OUT

Until a model is connected, Nexus gives no AI answers. It says so plainly and points to the page where an administrator connects one. It never borrows a key from somewhere else to fill the gap.

Inviting people

An invitation is decided before it's sent. You type the email address and choose a role, and the page shows exactly what that role can and can't do. A member, for example, can't see model keys or read the audit log, and can't invite or remove people. The person gets one email, and the invitation ends by itself after seven days if nobody accepts it.



SCREEN Inviting a colleague. What the chosen role can and cannot do is shown before the email goes out.

Administrator roles are never handed out by email domain. Someone whose address ends in your domain can ask to join, but that makes them a member, and an administrator decides.



The first five minutes for everyone else

When an invited person follows the link, the join page tells them who invited them and the role they'll have. For an administrator of one unit, it also says plainly what they won't see, which is the content of anyone's messages or files. They sign in with a passkey, with Google or GitHub, or through the company sign-in, and set up their second check if the organisation asks for one.

Then they land on their own Today page in the console, which is where members do their daily work: asking the assistant, team chat, files and notebooks, agents and workflows. What each person sees follows their role. Administrators also see the administration pages, and a member who opens one gets the "Not an administrator" page, which says who their administrators are.

On Adverant's cloud, a new organisation's first administrator gets a short list instead of an installer: run a sample agent on made-up data, connect your own model, add one document, invite a colleague. Most people finish in under five minutes. The first answer cites the page it came from and lists what it was read with: which files, and which model. The list goes away by itself when all four are done, or after seven days.



What to have ready on day one

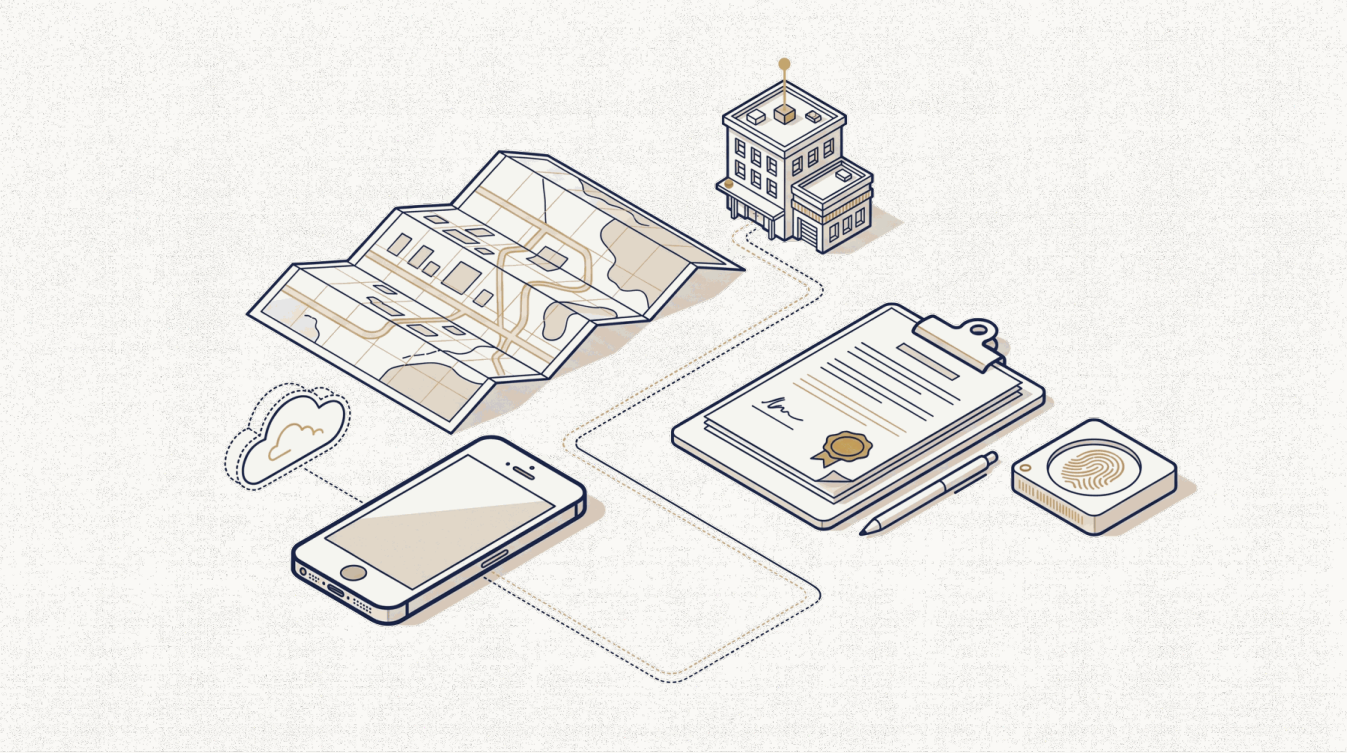
- The name and email address of the first administrator, and a time agreed with IT for the claim.
- The organisation's name, its short name for links, and where its data must be kept. The last two are permanent.
- At least two other people who will be administrators.
- Your company sign-in details, or the contact in IT who manages your identity provider.
- A model: a key for a hosted provider, or the address of a model server on your own hardware.
- The licence file, if you're on Enterprise.
- The mail server details Nexus will use to send invitations.
- One real document your team would like answers from, for the first test.
- A note of who keeps the printed emergency sign-in, and where.

Away from the desk

Using Nexus on a phone, from the console's phone layout in every edition to the native iPhone and Android apps. Covers what works without signal and how a company manages the phones.

FOR EVERYONE

11 MIN READ



A claims handler from Larkfield & Vane is standing in a storm-damaged house with no signal. On her phone, the claim visit screen sorts everything into three piles. Fourteen photographs and her notes are saved on the phone and will join the claim later. Her repair estimate and a change to the reserve are queued as suggestions for the office to decide. And three things are shown as rows she can't press at all: signing anything off, checking who gets paid, and starting a legal deadline. Each row says why.

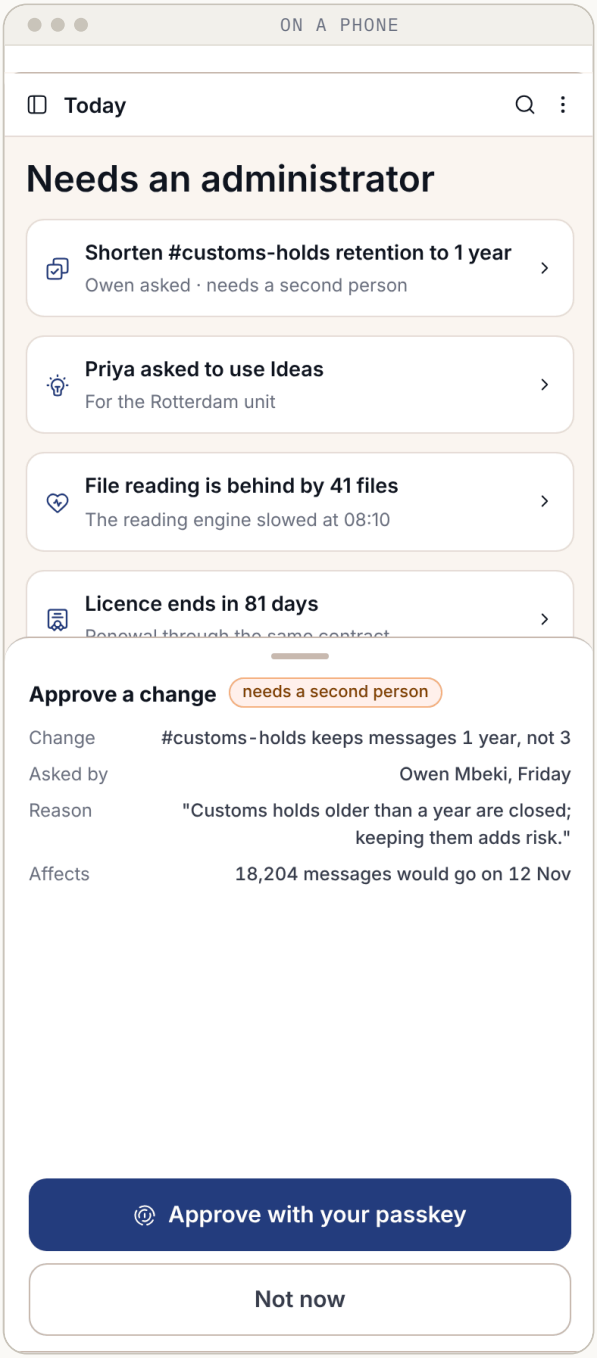
That screen sums up how Nexus treats phones. A great deal of work can happen away from the desk, and some of it waits for a connection on purpose. This chapter covers the two ways to use Nexus on a phone and what works offline. It ends with how your company keeps the phones in line.



The console on a phone, in every edition

Every edition, Community included, gives you the web console in a phone's browser. The menu folds away, the assistant opens as a sheet from the bottom of the screen, and the pages rearrange to fit.

For an administrator, the phone layout is for things that can't wait. From a phone you can approve a change with your passkey, stop an agent, freeze sharing outside the company and read warnings. Full administration stays on a computer, where a change can be previewed before it's saved.



SCREEN An administrator at an airport approving an urgent change from a phone. The page says plainly that full administration stays on a computer.

Community has no native app. If someone points the app at a Community installation, it refuses with `NOT_IN_YOUR_PLAN` and sends them to the web console instead.

The iPhone and Android apps

The native apps come with the Mobile apps add-on to Enterprise. There's one public app, called Adverant Nexus, in both stores. Companies that don't allow public stores can get it as a private app through Apple Business Manager or Managed Google Play. Air-gapped and defence sites install a build through their own device management only.

Signing in uses the phone's own passkey prompt. Company sign-in opens in the phone's browser, never inside the app. If your company manages the phone, the address of your installation is already

filled in, marked "Set by your company". The phone then makes its own key and joins your list of devices. After 15 minutes idle the app asks for your face or fingerprint, which works with no signal; after three failed tries it falls back to the passkey.

The bottom bar has four items: Today, Conversations, Ask and More. Here's what that opens up:

Area	On the phone
Today	Lists what needs you, with the reason for each in words. The home-screen widget shows counts only.
Ask	Answers with citations and the list of what it read. You can dictate a question, or send in a photo from the camera. A draft the assistant writes is only a draft until you send it yourself.
Conversations	Holds your channels and threads, with shared channels labelled. Retention and legal holds set on the desktop apply on the phone too.
Approvals and signing	Ask for your face or fingerprint every time. Signing works only in the open app, never from the lock screen, and turning something down asks for a reason.
Search	Finds only what you may open, and leaves the rest out of the counts as well as the results.
Capture	Uses the phone's own document scanner and saves straight to a folder.

Under More sit the rest. A notebook opens on the phone with its papers and an answer that cites a source for every sentence, and tapping a citation opens the page it came from. When it's your turn to judge in an idea run, two ideas arrive without their authors' names, and you pick one, choose neither, or say you can't tell. Mail that needs a home is listed with a suggested record to file it under, and a reply goes through the same checks before it leaves as it would on the desktop. What's running opens the progress stripe as a sheet: running, waiting on a person, and finished today. A long job can also be followed from the lock screen, as a Live Activity on iPhone or an ongoing notification on Android, showing only its name and step, and a stuck job of your own can be stopped there after a face or fingerprint check. Whoever is on call can trace one request between agents, hop by hop with timings, read-only.

Some of those need their own add-on as well. Signing a ruling in a decision room needs the Decision rooms add-on, for example.

Notifications that carry nothing

A notification from Nexus says that something is waiting and nothing more. Your company sets how much the lock screen may show. At the strictest setting it's "Something needs you in Nexus", and you open the app to see the rest.

Apple and Google carry only an encrypted envelope with a key reference, never the content. The phone decrypts the notification itself. How the envelope travels depends on how Nexus runs: through Adverant's relay for a connected installation, through your company's own push setup, or by site paging on an air-gapped site.

What works with no signal, and what never does

The app keeps an encrypted copy of what you need on the phone. The key is held by the phone's secure chip, and the copy is erased after an age limit your administrator sets, for example seven days. When you're offline, the app says since when, shows how old its copy is, and lists what's waiting to send.

Every field on an offline screen belongs to one of three kinds, and the screen shows which:

1. **Saved on this phone and joined up later.** Notes, photos and readings. They go into the record as soon as there's a connection.
2. **Sent as a suggestion the office decides.** An amount, a finding or a decision. A suggestion never goes straight into the record: when it arrives, the server runs the same authority checks as it would at a desk, and a person in the office decides.
3. **Not possible out here.** A signature, a seal, or a check against a list such as sanctions. These are drawn as rows with their reason, never as buttons.

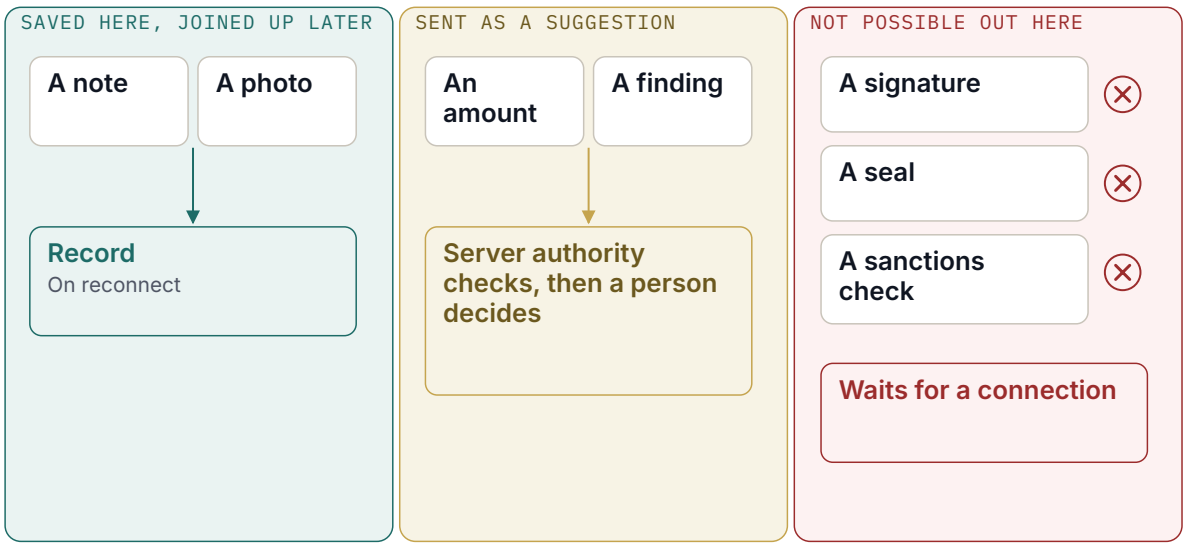
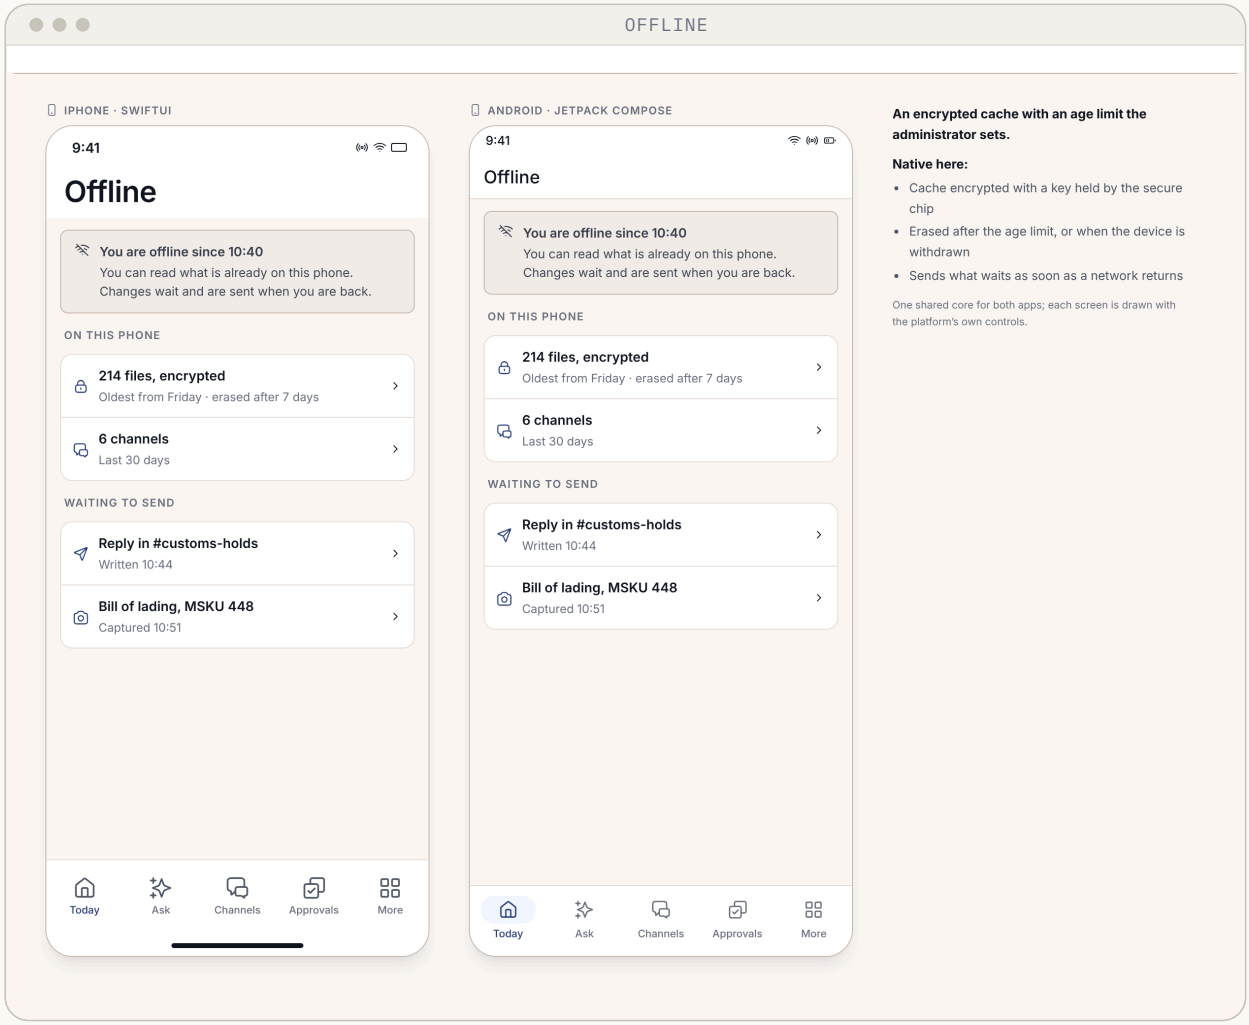


FIGURE What a phone with no signal does with each kind of field.



SCREEN Working on the quay with no signal. The phone shows what it holds and how old it is, and what waits to send. The copy is encrypted and erased after the age limit the administrator set.

This is deliberate, and it costs something. A field worker can't close a claim or release a hold from a basement, however urgent it feels. Approvals and signatures need a connection, so that they're recorded at once and checked against the latest state.

> TIP
> Before a day of field visits, download the day's work while you still have signal. A pack's phone screens can be kept for offline use, and the app shows how much space they take, until when they stay usable without signal, and the storage limit your company set.

Company device management

Most companies already manage their phones, and the app fits into that. It reads the managed settings your device management sends, such as the address of your installation. Microsoft Intune and Workspace ONE are both supported, and the company can remove the app's data remotely.

Administrators set the device rules in Nexus as well:

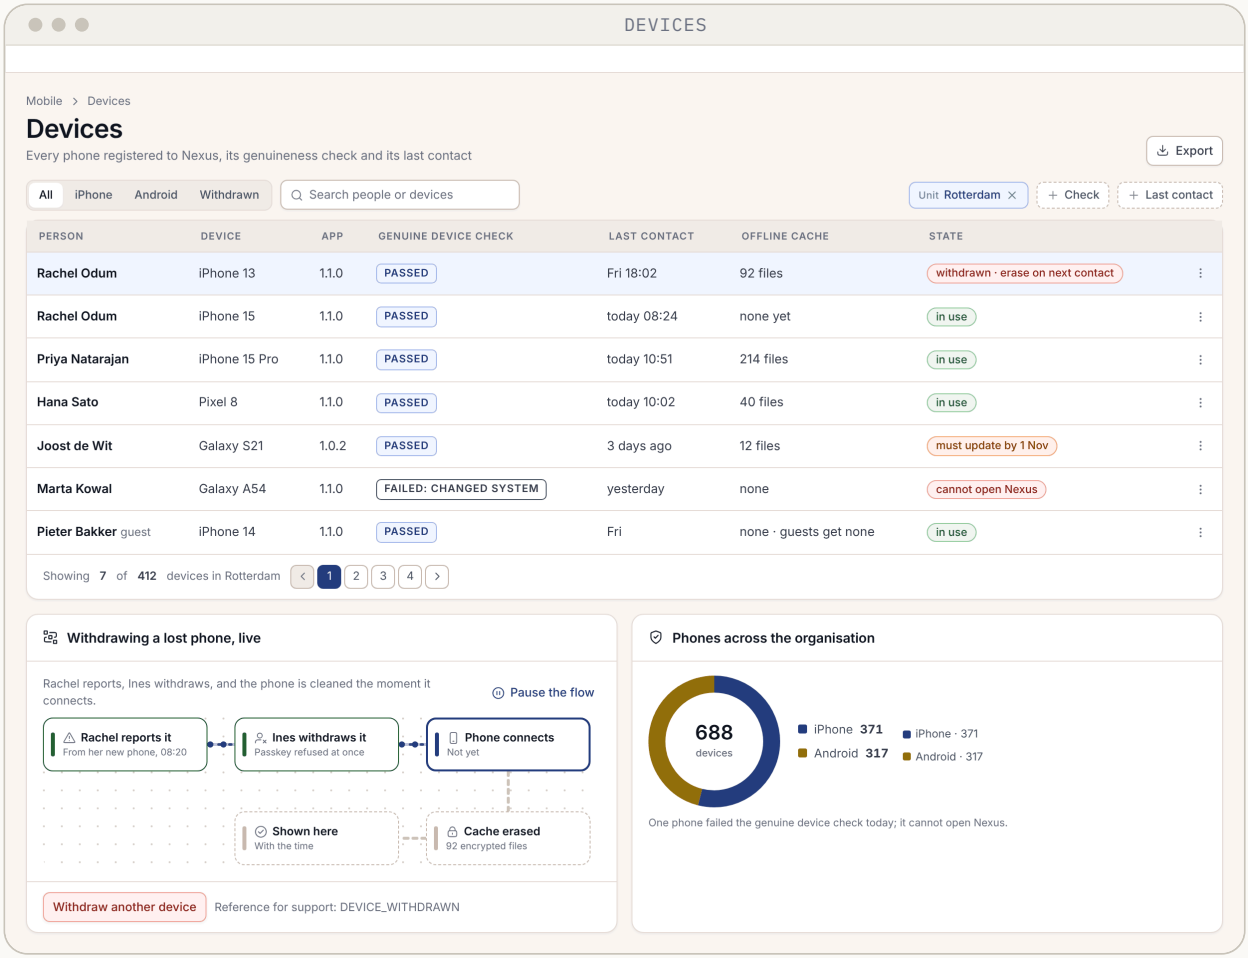
- what a phone needs before it can open Nexus, such as a screen lock, a minimum app version and a check that the phone and the app are genuine;
- whether screenshots are blocked on sensitive screens, and whether files may be shared outside the company;

- how long the offline copy may stay on the phone.

Each rule shows where it came from, and loosening one needs a second person. When a phone fails its genuineness check, the app says so with `DEVICE_CHECK_FAILED` and names who can help. When a rule blocks copying, the plate reads `COPY_NOT_ALLOWED_HERE`.

The company manages the app's data only. It can't see your photos or messages, and the registration screen says so. Administrators see app versions and devices as counts, never what anyone does in the app.

When a phone is lost, it's withdrawn in one step, by its owner or by an administrator. Its passkey is refused at once, and its encrypted copy is erased the next time it connects. If someone switches the lost phone on later, it shows that it was reported lost, with `DEVICE_WITHDRAWN`.



SCREEN Withdrawing a lost phone. Each phone shows its genuineness check and when it last made contact, and a withdrawn phone's encrypted copy is erased on its next contact.



WATCH OUT

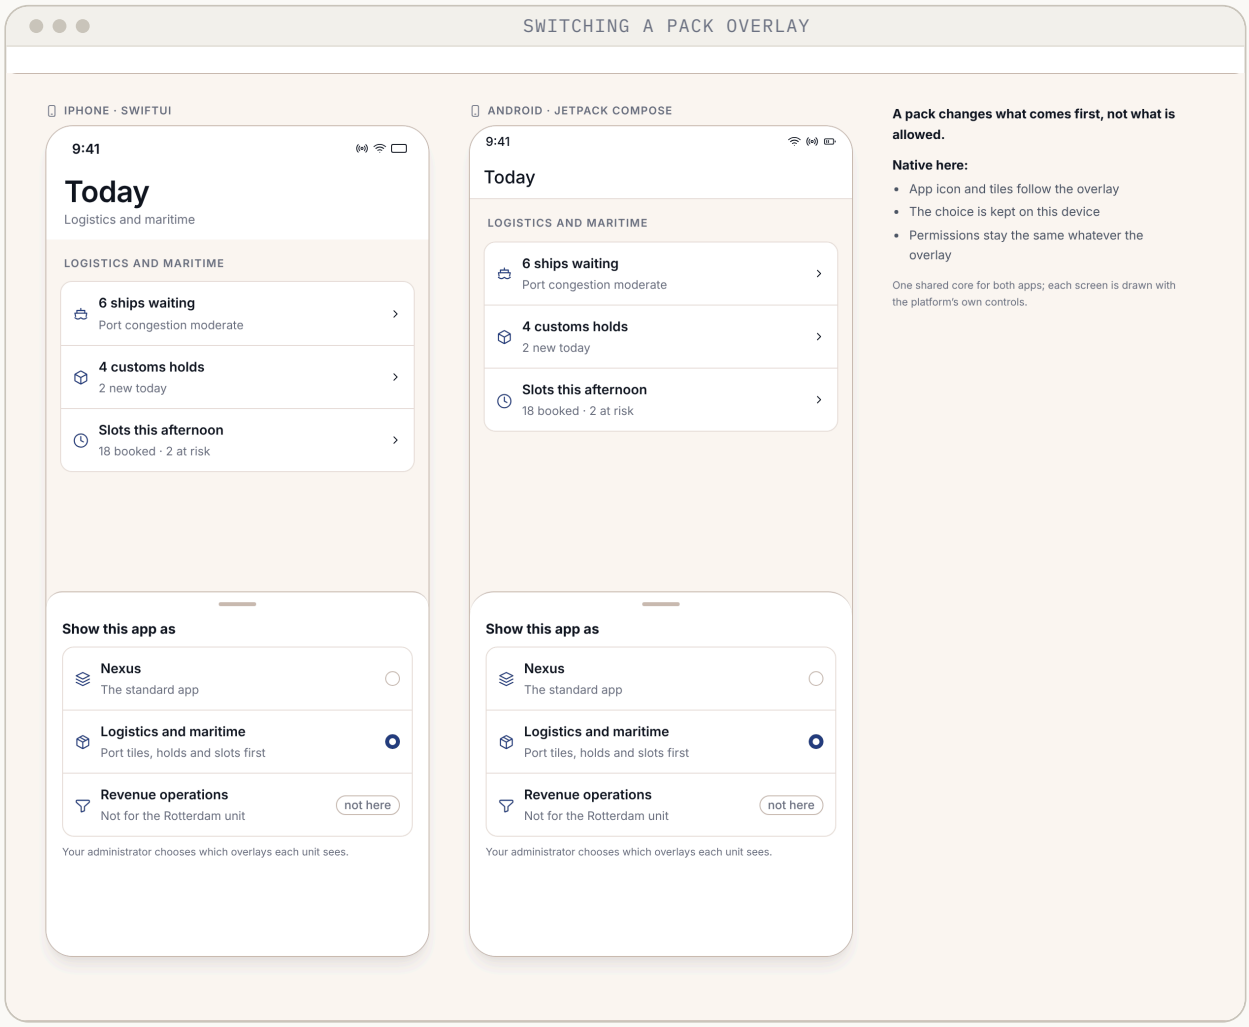
Until a lost phone connects again, its copy stays on it. The copy is encrypted, and it can't be read without the owner's passkey or biometrics. Keep the offline age limit short for roles that carry sensitive work.



Pack phone screens

A domain pack can add its own phone screens, called an overlay. An overlay reorders the app around the pack's work, putting port tiles and customs holds first for a logistics team, for example. It

never widens what a person may see or do. Administrators choose which overlays each unit gets, and anyone can switch back to the standard app.



SCREEN Switching the app to the Logistics and maritime overlay. A pack changes what comes first, never what a person is allowed to do.

Each pack decides what belongs on a phone and what doesn't:

Pack	On the phone
Insurance	Claim visits with photos that wait up to 72 hours for signal, and payment approvals within a person's limit.
Banking and payments	Credit approvals within an officer's limit. Screenshots are blocked on approval screens.
Healthcare	Ward-round summaries and handovers. Notifications show the ward and bed, never the patient's name, and shared ward phones sign out after ten idle minutes.
Life sciences	Site-monitoring visits and adverse-event intake. An adverse event saved offline raises a warning after four hours, because its reporting clock starts when it arrives.
Defence and security	Field reports on accredited rugged phones, which wait up to fourteen days until the phone docks.
Logistics and maritime	Yard and berth inspections; releasing a customs hold needs a person and a passkey.

Capital markets and asset management	The morning research briefing, downloaded at 06:30, research release approvals and surveillance alerts. Never orders, and calls made from the app are recorded like desk calls.
Government and public sector	Offline housing inspections and field visits with located photos, and signed briefings for leaders that are read and signed but not edited. Captures are stored only in the national cloud.
Critical infrastructure and energy	Offline pylon and substation inspections, incident first reports that start the reporting clocks, and drone findings a person reviews. It never switches anything, and lone workers check in hourly.
Legal and professional services	Time capture, matter updates and approvals, never privileged files. Notifications show a matter number only.

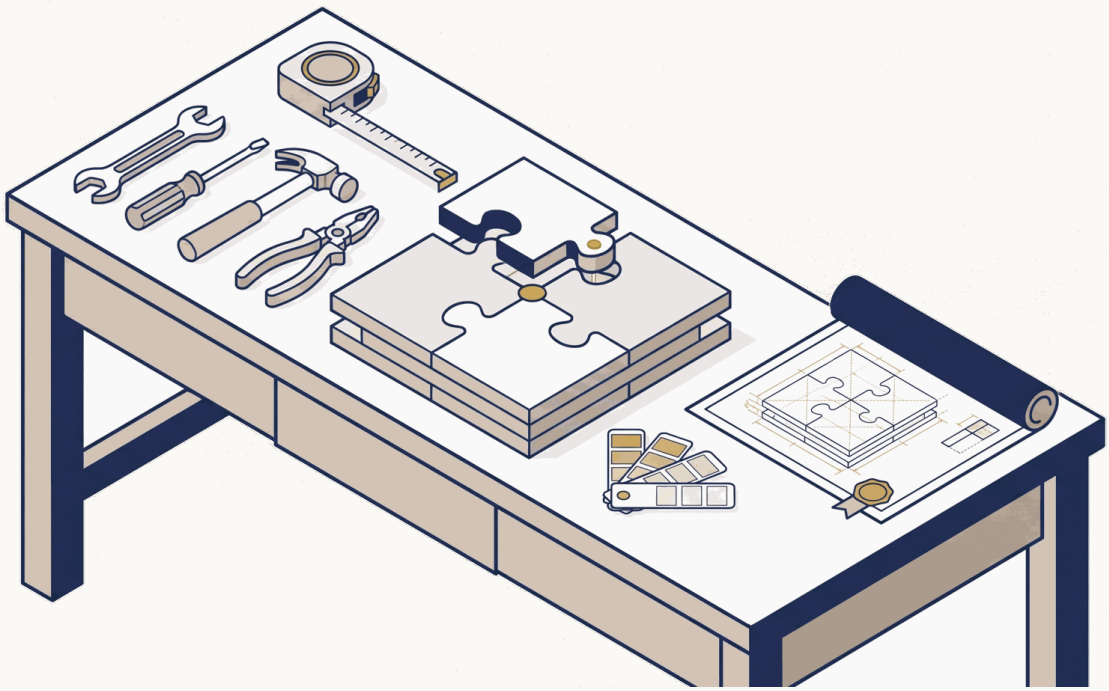
A pack's phone screens open only with both the pack and the Mobile apps add-on in your licence. Chapter 1.7, Domain packs and verticals, lists what each pack needs.

Making it yours

The ways an organisation shapes Nexus, from a single setting to a vertical of its own. Also how to brief developers or a partner, and what the platform guarantees whatever gets built.

FOR EVERYONE

13 MIN READ



Priya Natarajan wanted to know what the "Explain this delay" button actually did before she let her team rely on it. She right-clicked it. A small panel told her which skill and version it runs, on which model. It listed the three sources the skill reads and said it writes nothing. It gave a rough cost per press, and it named who wrote each version and who reviewed it.

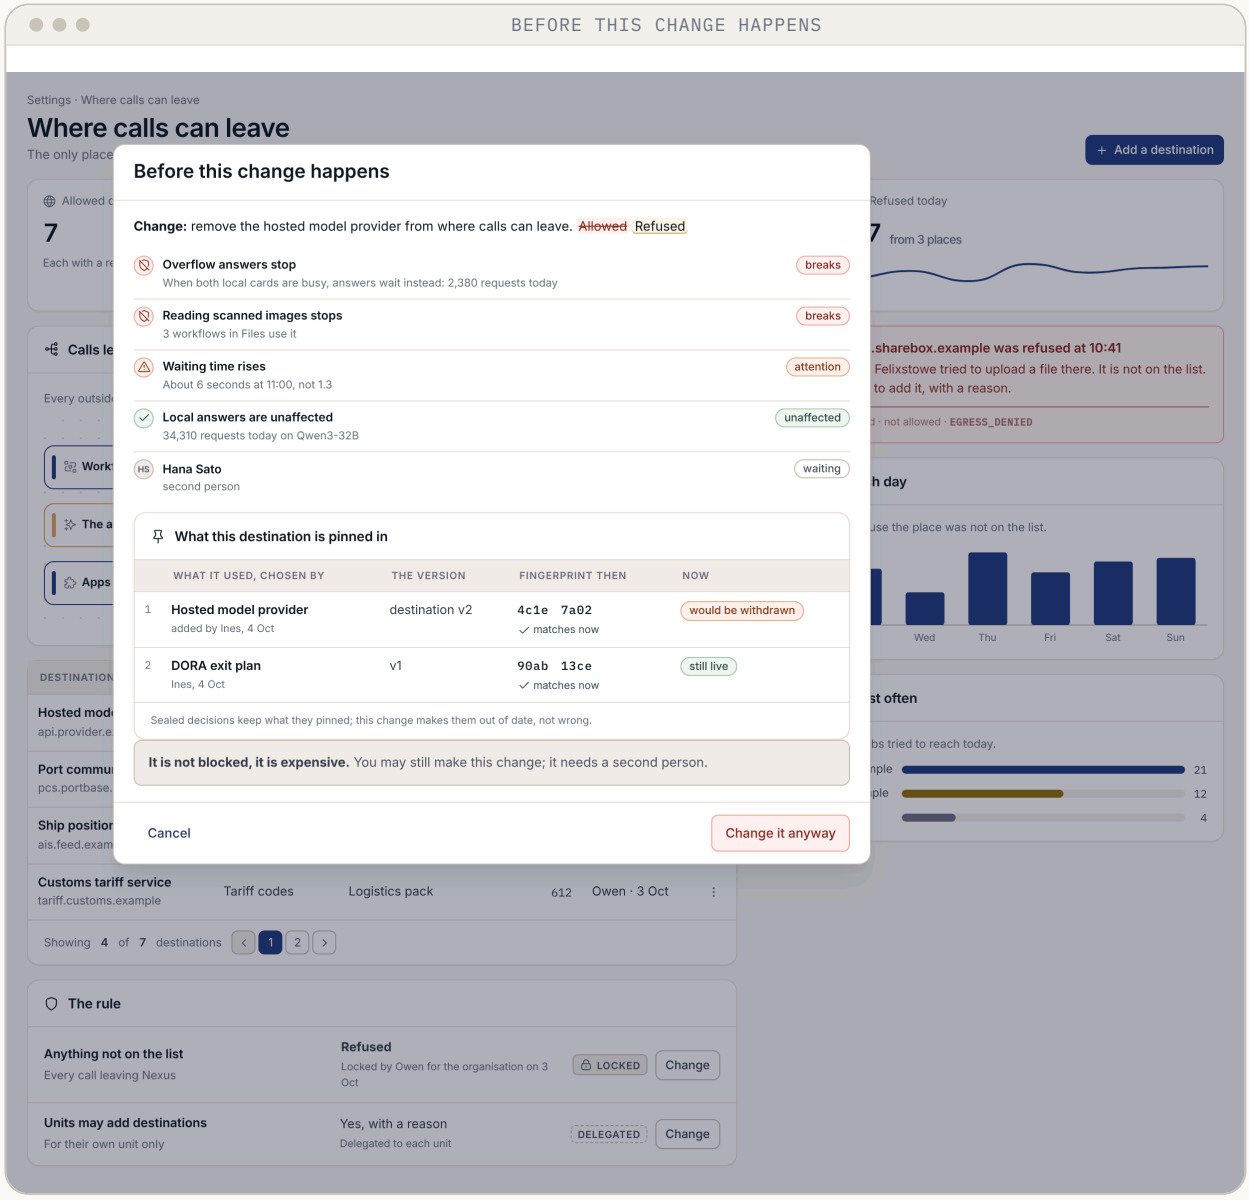
Every way of shaping Nexus works like that: you can see what you changed and who approved it, and you can go back. The options run from a setting an administrator changes in a minute to a vertical your developers build over several weeks. This chapter goes through them in that order.



Settings and policies

Settings live together on the Settings page. Beside its current value, each setting shows where that value came from and whether it's locked, and whether it has been handed down to a unit. Values flow from the licence to the organisation and then to units, so a unit administrator sees exactly which settings the organisation delegated to them.

Before a risky change saves, a preview called **Before this change happens** shows what would break, what needs attention and what is unaffected, in today's figures. A costly change needs a second person to approve it. Every setting keeps its versions, and putting an older one back creates a new version, so the history is never rewritten.



SCREEN The preview before a risky change. It says, in current figures, what would stop and what would carry on, and asks for a second person when the change is costly.

On Enterprise, company policies get the same care. Each policy has an owner and a review date. A policy owner can redraft a clause with the assistant's suggestions marked and cited, then run the draft against last quarter's records to see what it would have changed, without changing anything. Once published, people confirm they've read the new version, and Nexus sends reminders before telling the unit administrator.

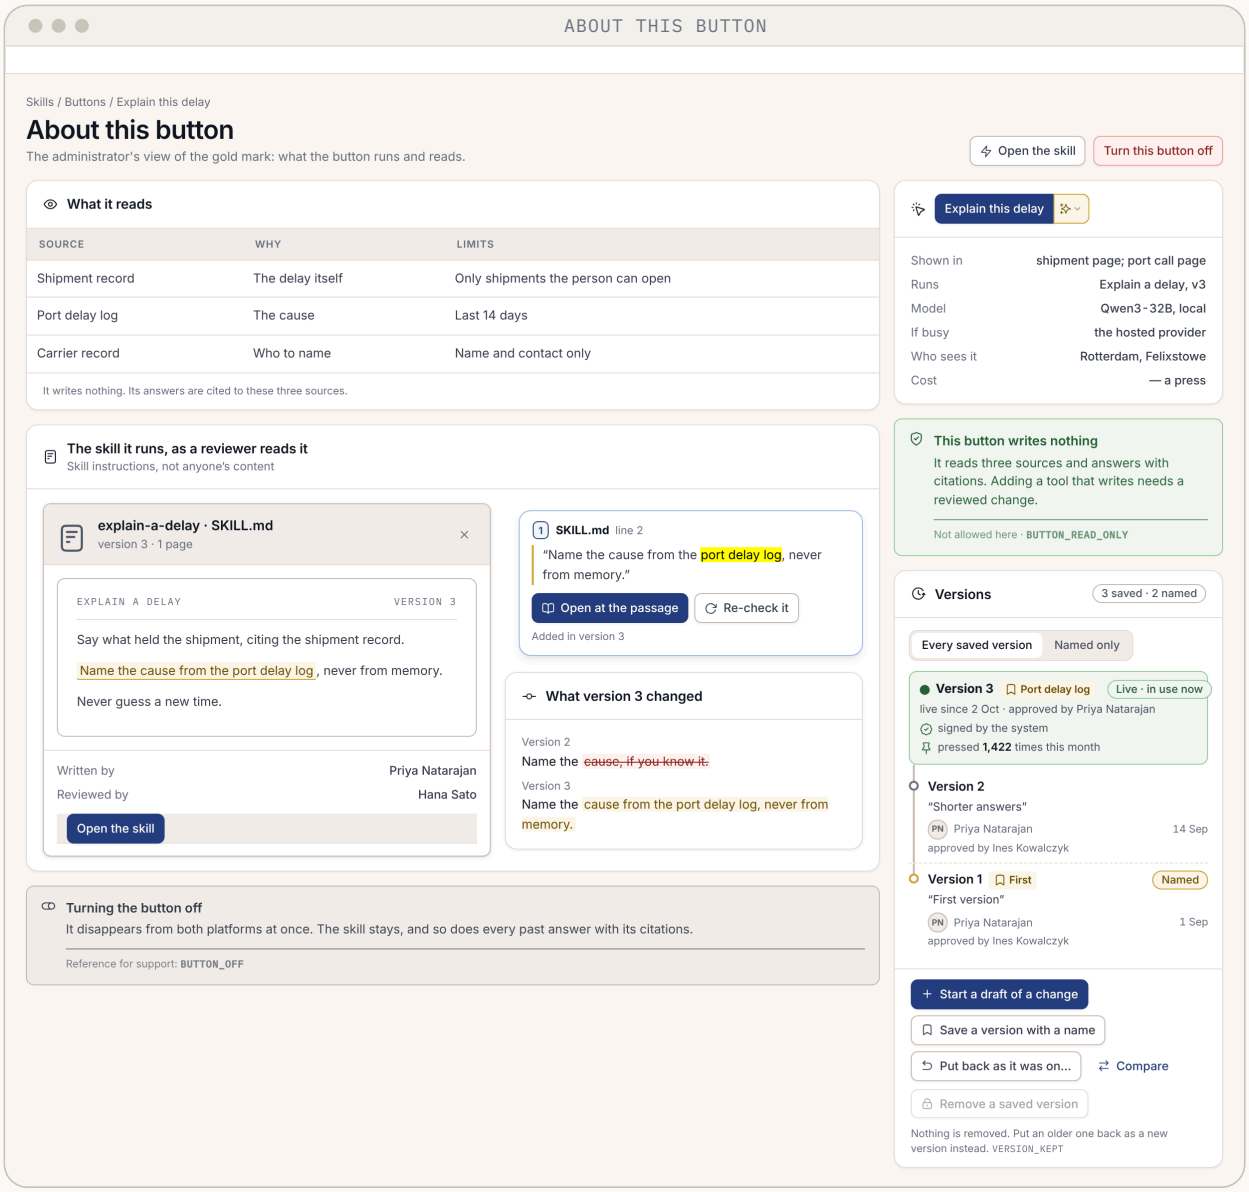
Skills: teaching the assistant a repeatable job

A skill is a repeatable job written down: the instructions, what it may read, the tools it may use and the model it prefers. Behind every artificial intelligence (AI) button in Nexus there's a skill, and every skill has an owner.

The quickest way to start a skill is to describe it in a sentence and let the engine draft it. You can also write one from scratch, import one, or combine skills you already have. Before anyone relies on it, it runs against saved test cases, then against past work in the background, and a model from a different company gives the results a blind review. Each run is sealed with the skill version and the model used.

Some limits are fixed. A skill can never use a tool the person running it doesn't have. On Enterprise, a change that lets a skill write rather than only suggest needs two named people to approve it, and the person who proposed it can't be one of them. New versions go to a small share of runs first and stop by themselves if quality drops.

There are two kinds of skill, and the difference matters when you brief developers. Skills Engine skills are the packages in the skills library, which your people create or upload and manage in the console. Manifest skills are the ones a plugin declares in its manifest; they arrive with the plugin and update when it does. Chapter 2.20, Skills, covers both.



SCREEN About this button shows the skill and model behind an AI button, what it reads and may write, and who wrote and reviewed each version.

Apps and plugins from the catalogue

The app catalogue lists ready-made apps. Before you add one, its page states what it reaches outside Nexus, what it may read and write, what it may never do, and how tightly it's sandboxed. Every app arrives switched off. Your compliance officer approves or refuses each outside address the app declares, and a call to any address it didn't declare is blocked and counted.

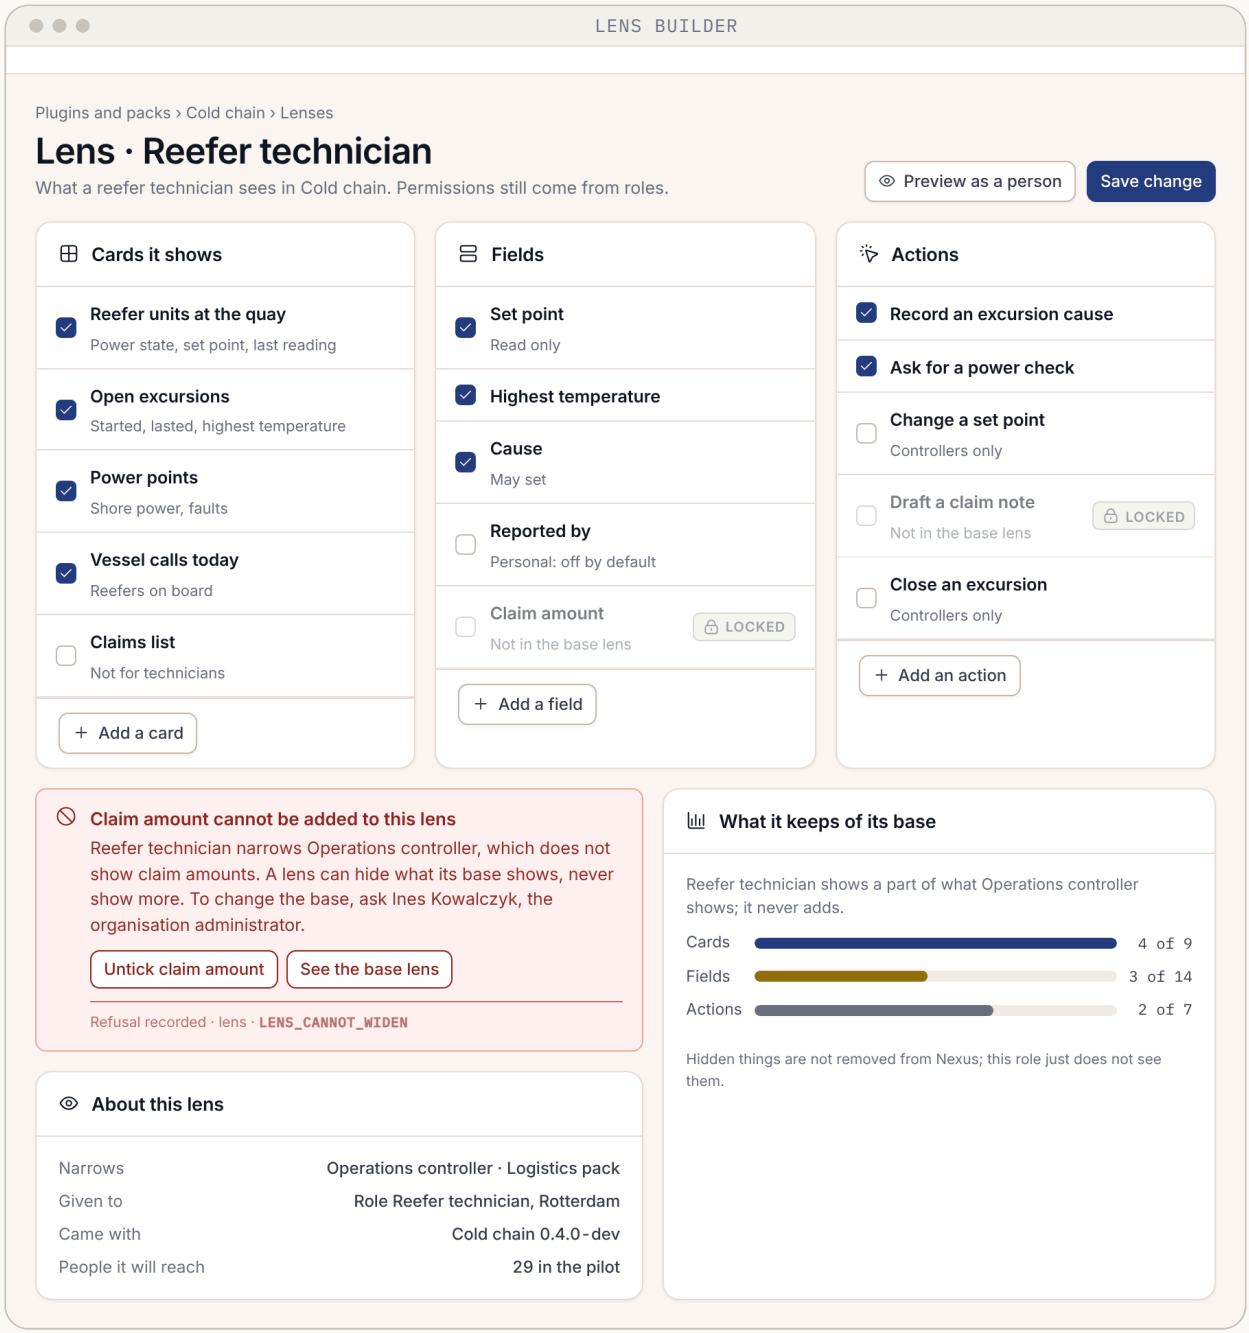
Withdrawing an app ends its access at once, and anything it made stays where it is. On Enterprise you can keep your own catalogue, synced from Adverant on a schedule, where new apps stay hidden until someone approves them. Third-party plugins get only the permissions they declared, and you can grant each one unit by unit, leaving some out.



A domain pack, or extending one

Chapter 1.7, Domain packs and verticals, covers switching a pack on. Extending one is the next step, and most organisations do some of it. You can add your own terms to the pack's vocabulary, add a connector to one of your systems, or add a template of your own beside the pack's.

You can also build a new lens for a role the pack doesn't have, with the lens builder. You start from an existing lens and tick the cards, fields and actions the new one shows. The lens builder refuses any attempt to show more than the base lens allows, with the code `LENS_CANNOT_WIDEN`.



SCREEN The lens builder. A new lens is ticked against its base lens, and anything the base doesn't allow is refused.



White-label

With the Storefront / white-label add-on, a platform you build on Nexus can carry your brand rather than Adverant's. Your customers pay through a checkout in your brand and get receipts in your name, and they manage their card or cancel in a customer portal you control. One customer app template can be built as several brands' own apps, which is how a broker or an insurer gives policyholders an app under its own name. If you publish your own phone app, its notifications run through your own setup and nothing about them goes to Adverant.



A vertical of your own

A platform on Nexus is the shared core plus a domain pack. If no pack fits your industry, you can build one. It has five parts: a knowledge model of the things your industry deals in, lenses for each role, a manifest that says what the pack may reach, your industry's rules mapped to the evidence that answers them, and small decision models trained on your own sealed records.

Every pack takes the same path. Its manifest is written and enforced, so a pack can't quietly do more than it said. It's previewed as each role will see it, on realistic sample data. Then a person reviews what changed before anything reaches real users, and it's released unit by unit. There are three ways to get one: use one of Adverant's ten packs, have Adverant build yours with your experts, or have your own developers build it and Adverant review the release. A pack you build stays yours. Chapter 2.24, Building a custom vertical end to end, is the technical guide.



Building with AI coding tools

Many developers now build by describing what they want to a coding assistant such as Claude Code and reviewing what it writes. Nexus is set up for that. A developer links their coding tool to Nexus through one address, and from then on the tool works inside the developer's own development space.

The platform's rules do the safekeeping, whether a person or a coding assistant wrote the code:

- The work runs on synthetic test data, or on a masked copy that the data's owner has approved. The shared live data is never written from a development space.
- The manifest says what the extension may reach. If a new version asks for more, such as sending mail or calling a new outside address, the release is blocked until two reviewers agree.
- Every release has a second reviewer, who sees the test results and the change in scope side by side. The author can't approve their own work.
- Releases move from development to a pilot unit to everyone, and no stage can be skipped. A unit can be rolled back with a reason and a second person.

An outside partner builds the same way, in a space on Adverant's cloud with sample data and no route to any customer's installation. Chapter 2.26, Building with AI coding tools, gives developers the detail.



How to brief your developers or a partner

A good brief fits on one page, and it saves weeks. Start with the work, in a paragraph your experts would recognise: what job the extension does, for whom, and what goes wrong today. Name the people who will use it by role, and say which unit tries it first. Then describe the things it deals in, in your own words, the way you'd explain them to a new starter, because that becomes the knowledge model. Say which existing systems it must read, and be explicit if it must never write to them. Finish

with the decisions it touches: which ones a person must always make, which need a second person, and which rules your compliance officer will want to see answered.

Attach a short list alongside the prose:

- the roles, and for each one what they must see and must never see;
- the systems to read, with an owner for each connection;
- the outside addresses it may call, if any;
- the rules and regulations it answers to, and who confirms the mapping;
- the decisions that stay with a person, and the ones that need two people;
- the pilot unit and the date you'd like it there;
- who reviews releases on your side.



What the platform guarantees, whatever gets built

Whatever your developers or a partner build, three things hold, because they're properties of the platform rather than promises an extension makes.

Each organisation's data is kept apart. Every read and write is scoped to an organisation, to the person who owns anything personal, and to a unit where you use units. Those come from the signed-in person's own sign-in and nowhere else. An extension can't pass in a different organisation, and a request that can't say whose data it means is refused.

AI goes only through the switchboard. An extension never holds a model key. It can say which kind of model it would like, and the AI switchboard decides, using only the providers your organisation has connected and allowed.

Every action is recorded. What people and agents do is written to the audit log, refusals included. Decisions that matter are sealed, so anyone entitled can check them later.



NOTE

These guarantees don't depend on anyone remembering them. Chapter 2.19, The rules every extension inherits, explains how the platform enforces each one.

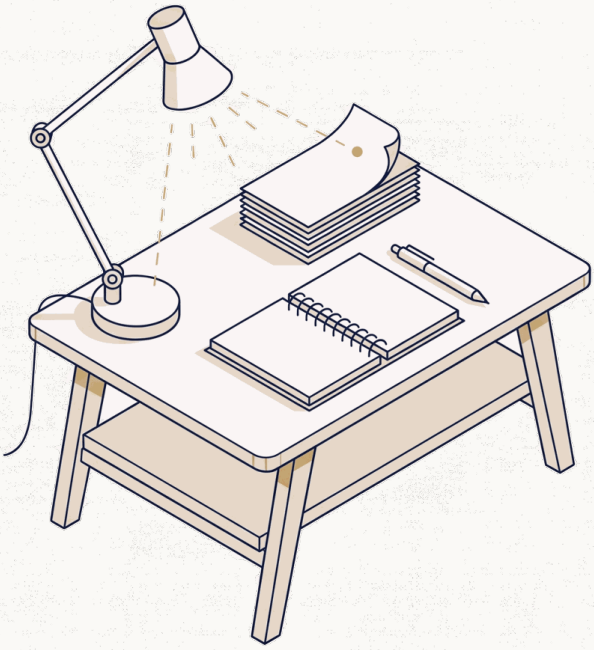


Questions people ask

Plain answers to the questions buyers, managers and data protection officers ask most, including the uncomfortable ones about money and access to your data.

FOR EVERYONE

12 MIN READ



These are the questions that come up in almost every first meeting, usually in this order: money first, then data, then artificial intelligence (AI), then whether it'll survive contact with the IT department. The answers are short. Where a fuller answer exists elsewhere, the chapter is named.



Licences and money

What happens if we stop paying?
Nothing that's running stops. When an Enterprise licence ends, the people who can act see a warning in the product, then a 30-day grace period in which everything carries on. After grace, new Enterprise work can't start. Running work continues, and security fixes still arrive. Your sealed records stay readable, and export keeps working alongside the free Community features. An evaluation licence has no grace period, though nothing is deleted when it ends.

Will a licence ever switch off something mid-task?
No. Running work is never stopped by a licence, in any state. Expiry stops new features and new versions, never security fixes.

How much does it cost?
Adverant quotes each edition and add-on for your situation, and this documentation gives no prices. Community is free and open source. If you buy through your cloud provider's store, the charge appears on your cloud bill, and nothing is billed until the licence exists.

Can we start small and add more later?
Yes, and most customers do. A Community installation moves to Enterprise by importing a licence and pressing **Fetch and switch on**, which brings in the Enterprise parts, checks each one is genuinely from Adverant, and switches them on in about ten minutes. Everything people already have stays where it is. Add-ons and packs arrive the same way.

What if Adverant goes out of business?
The core is open source under the Apache-2.0 licence, so the Community edition stays yours to run. Sealed records can be checked without Adverant, using a free verifier that runs on your own computer. Exports are tested, and the licensed parts' source is held in escrow, which your contract can rely on.



Your data, and who can see it

Can Adverant see our data?
On an installation you run, no. Health reporting to Adverant is off unless you turn it on. The licence names your company, never the people in it. Once a year your administrator sends a usage statement of counters, readable in full before it leaves. A support bundle goes only when you ask for help, after a preview with private details removed. On Adverant's cloud, each organisation has its own private space and keys, and model keys are sealed in your organisation's pool where Adverant staff can't read them.

Where is our data stored?
Where you chose at the first run, and that choice is fixed for the life of the installation. On Adverant's cloud, data is stored and processed in the EU. On Enterprise, the Where data lives page sets a place for each kind of data, and moving it is a tracked job a second person approves. Chapter 1.8, Where Nexus runs, explains the options.

Can our administrators read staff messages?
No. A person's mailbox and notebooks are private even inside the organisation, and so are their direct messages. Administrators see names and counts, never content. When an administrator helps someone by seeing the console as they see it, the view is read-only and ends after 30 minutes. Every page opened is recorded, and the person is told.

Can we get all our data out?
Yes, always. A full organisation export is never behind a paywall, and it keeps working whatever state your licence is in.

How do we remove data when the law requires it?
There's no everyday delete button. Things are withdrawn, and the record says so. Where the law requires removal, for example a person's right to erasure, an administrator uses **Erase**, which needs a written reason and a second person. A permanent record keeps what was erased and who erased it, with the date, never the content.

What does Nexus tell a data protection officer about processing?
On Enterprise, the data protection pages keep records of processing and consent, breaches, and requests from people about their own data with the deadline for each. The General Data Protection Regulation (GDPR) is one of the frameworks the compliance pages map controls to. Your counsel confirms the mapping; Nexus provides the evidence.



AI and models

Which AI models can we use?
Whichever your organisation connects. The organisation's key pool takes keys for Anthropic and Google Gemini, OpenRouter and Scaleway. Nexus can also use any model server your organisation runs that speaks the common OpenAI-compatible format. A model on your own hardware is tried before any hosted one. For search, an embedding model runs inside your installation by default.

What if the model provider changes its terms or withdraws a model?
You change it in one place. All AI in Nexus goes through the switchboard, and nothing else in Nexus holds a model key, so no workflow or app names a provider of its own. They say what kind of model they'd like, and the switchboard picks from the providers you've connected. Before you remove a provider, Nexus shows in today's figures what would carry on and what would wait.

Can staff use their own ChatGPT or Claude subscription?
Not inside Nexus. AI in Nexus runs on keys in the organisation's pool, which administrators manage, so every request is counted and recorded under your rules. A personal subscription is a credential that belongs to one person, and Nexus doesn't use personal credentials for shared work. If your team prefers a particular provider, connect an organisation key for it.

What happens if we don't connect any model?
Nexus gives no AI answers, and it tells you so, with the code `MODEL_NOT_CONNECTED` and a link to the page where an administrator connects one. It never falls back to someone else's key.

Can we run entirely on our own models?
Yes. An installation can use only models on your own hardware, and an offline site has no other choice. The switchboard then sends nothing outside the building.

Does the AI make decisions?
It drafts and suggests. Consequential decisions are made by a named person, and the domain packs lock that rule for the decisions their industry cares about most, such as credit decisions and claims. Where a decision needs more than one view, a decision room puts the case for and against in front of a person who chairs and signs.

How do we keep AI spend under control?
Each run's cost is estimated and set aside before it starts, and budgets cap spend by unit. When a budget is tight, work moves to a cheaper model rather than overspending. From Team up, spend is shown per skill, and Enterprise also shows it per unit.

Can agents act on their own?

Only as far as you allow. Agents start at the levels where they watch or suggest. Every level above asking first needs an administrator's approval, and the top levels need two. Every agent has a human sponsor. A kill switch stops a run, an agent, a team or the whole fleet within five seconds.



Security and compliance

What if we lose the custodian key?

Then the key store can't be unsealed, and the data it protects can't be read. The custodians are the people you name when the installation is set up, and every recovery share is encrypted to each of them, so any one of them can unseal after a restart. Adverant isn't a custodian and holds no copy. Name more than one custodian and keep their keys in different places.

How does Nexus help with the rules we answer to?

The compliance pages map controls to each framework, including the Digital Operational Resilience Act (DORA) and the Network and Information Security Directive 2 (NIS2). Each control shows the records that answer it, and gaps are written in words. Domain packs add their industry's rules the same way. The mapping is a starting point your counsel confirms; it isn't legal advice.

Can an auditor check a decision without an account?

Yes. A sealed record can be checked with the free verifier, on the auditor's own computer, with nothing uploaded. The answer comes back as plain yes-and-no lines.

What happens if our company sign-in goes down?

An emergency administrator sign-in, made at the first run and kept in your safe, still works. Every use is recorded and announced to all administrators.

What if someone loses a phone?

It's withdrawn in one step. Its passkey is refused at once and its encrypted data is erased the next time it connects. Chapter 1.10, Away from the desk, has the detail.



Running it

Do we need our own platform team?

Not on Adverant's cloud. For your own Kubernetes or an offline site you do, and Part 2 is written for them. An appliance or a cloud-store image needs less: someone who can rack a server or start a machine in your cloud account.

How long does installation take?

On a cluster that passes the preflight check, the install itself takes about twelve minutes. Claiming the installation and creating the first administrator take a few minutes more. Most of the first day goes on connecting company sign-in and a model, and inviting people.

How do upgrades work?

You choose when. An upgrade takes a backup first and rolls out in stages. It keeps a way back until a stated date. If a step fails, the previous version runs again and your data is untouched.

Can we move from Adverant's cloud to our own data centre later?

Yes. Your records and settings move with you, and sealed records keep their seals, so an auditor can check them before and after the move.



The name

What does "alpha" mean?

Alpha is the name of this release of the product: Nexus Alpha, version 1.0.0-alpha.1. The console shows a small Alpha chip beside the product name, and hovering it gives the version. Whatever the

release name, the same guarantees hold: every upgrade keeps all your data, export always works, security fixes always arrive, and a licence never stops work that's running.