NOVA IAM

NOVA IAM

Handbook

How Nova IAM manages identities, accounts and entitlements – openly documented.

As of: September 28, 2026

docs.nova-iam.com

Contents

CHAPTER 1

Introduction

1.1 About this handbook

Nova IAM is an Identity & Access Management platform. Nova manages **identities** – employees, external people and technical actors –, their **accounts** in the connected target systems and the **entitlements** they hold there. This covers the lifecycle of joiners, movers and leavers, access requests and approvals, recertifications and the logging of security-relevant changes.

This handbook describes how Nova goes about it – openly, for everyone to read.

Who it is written for

- **Users** – administrators, approvers and managers who work with Nova.
- **Auditors** – internal audit, information security and external auditors who want to trace how access is granted and removed again.
- **Evaluators** – anyone assessing Nova before adopting it.

What “planned” means

We only document what Nova does or what has been decided. Decided behaviour that no release ships yet is labelled:

- Whole pages carry a **Planned** notice at the top and the marker planned in the navigation.
- Individual statements on otherwise current pages are marked with planned.

As soon as a release ships the behaviour, we remove the label.

Languages

The handbook is published in German and English. Both versions have the same content; the language menu takes you to the same page in the other language.

As PDF

The whole handbook is also available as one PDF document ↗ – with a table of contents, for filing, printing or passing on. It is regenerated whenever this website is published, so it always matches it.

1.2 What is Nova

Nova IAM is an Identity & Access Management platform. Nova manages **identities** – people, technical identities and AI agents –, their **accounts** in the connected target systems and the **entitlements** they hold there. Nova grew out of decades of practice with SAP Identity Management.

What Nova does

- **Lifecycle:** Nova takes over personnel data from the HR system and handles joiners, movers, leavers and rehires through configurable routines – see HR integration and Joiners, movers, leavers.
- **Access model:** entitlements from the target systems can be bundled into business roles and assigned to identities, organisational units or positions – see Access model.
- **Access requests:** requests for entitlements and business roles go through defined approval workflows – see Requesting access.
- **Provisioning and reconciliation:** Nova creates accounts in the target systems, transfers entitlements, locks accounts and compares the state in Nova with the state in the target system – see How access is granted and removed.
- **Governance:** recertification, segregation of duties (SoD) and risk assessment.
- **Traceability:** audit log, provisioning log and change journal.

Architecture

Building block	Purpose
Web interface	used in the browser, in German and English; nothing to install on workstations
Application	Python application (Flask) holding the business logic: checks permissions, runs procedures, writes logs and provides the interface the web UI works with
Database	PostgreSQL for identities, roles, assignments, configuration and history; Nova updates the schema itself on start-up
Connectors	target systems SAP (via RFC), Active Directory and LDAP, Microsoft Entra ID, Keycloak and SCIM; source systems “Nova HR Stage” for personnel data and “SAP OM” for the organisational structure – see Target systems overview
Background jobs	scheduled tasks with a cron schedule in the Europe/Berlin time zone, such as the HR import or the leave-date scan; found under “Monitoring” → “Background Jobs”
Notifications	by e-mail; through an add-on also in Microsoft Teams
AI (optional)	local models via Ollama or a cloud provider; the core works without AI – see AI in Nova

Add-ons

Add-ons (plugins) extend Nova with further functions. After installation they are switched off; administrators enable them one by one under “Administration” → “Plugins”. Add-ons with their own interfaces require a restart afterwards. The add-ons include:

Add-on	Function
“Planstellen (OSP)”	positions between organisational units and people, with their own business roles
“SoD-Analyse (SAP ABAP)”	checks SAP roles against an SoD ruleset
“GRC-Berechtigungsanalyse (SAP Access Control)”	uses an existing SAP GRC Access Control as a source for risk analysis
“Native MFA (TOTP)”	second factor for password sign-in
“OIDC Identity Provider”	other applications sign people in through Nova (OpenID Connect)
“Microsoft Teams Notifications”	notifications in Microsoft Teams

How data flows

1. **Source system → Nova.** The HR system supplies personnel data. Nova creates identities from it, keeps them up to date and detects joiners and movers.
2. **Intended state in Nova.** Which entitlements an identity should have, and for which period, Nova records as assignments – from business roles, organisational units, positions, approved access requests and direct assignments.
3. **Nova → target systems.** Provisioning runs create accounts, transfer entitlements and lock accounts.
4. **Target systems → Nova.** Nova reads accounts and entitlements from the target systems. Where the intended state in Nova and the actual state in the target system differ, reconciliation shows the differences; administrators resolve them in either direction.

Operation

Today, Nova runs on the operator's own servers: as a container application with a PostgreSQL database, used through the browser. The editions available and those announced are described in Editions and deployment.

1.3 Editions and deployment

Nova is one product in three editions. They differ in where Nova runs and who operates it.

The editions

Edition	Operation	Status
Nova Native	on-premises: the complete stack on the organisation's own servers	available
Nova on BTP	on the SAP Business Technology Platform, embedded in the existing SAP landscape and its operating model	announced for Q4/2026 planned
Nova as a Service	in the cloud, operated as a service; no in-house infrastructure needed, ready to use in the browser	announced for Q2/2027 planned

Nova on BTP and Nova as a Service have been announced but are not yet available. Today, Nova runs on the operator's own servers.

Licence

Nova carries no licence cost per identity.

Technical deployment options

Containers on in-house servers

Nova Native runs in containers: one container for the application including the web interface, plus a PostgreSQL database. A Docker Compose configuration for the application and the database is available. What operation requires and how installation works is described in Requirements and Installation.

A single instance

Nova runs as a single instance with one application process. The background jobs run inside this process; several parallel instances would run them more than once. Nova handles more load through additional threads in the same process.

SAP via RFC

For SAP via RFC – as a target system and as the “SAP OM” source system – Nova needs the SAP NetWeaver RFC SDK, which is obtained from SAP. Without the SDK, the SAP functions stay switched off; the other connectors work independently of it.

SAP BTP with Kyma

For the Kyma runtime of the SAP Business Technology Platform (Kubernetes), there are templates that deploy the application and PostgreSQL – again with exactly one instance. They are currently used for demo installations. The Nova on BTP edition is announced for Q4/2026.

AI

AI functions are optional in every deployment option. Instead of a cloud provider, a local model can be connected via Ollama – see AI in Nova.

CHAPTER 2

Basics

2.1 Identities and their status

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

Every person and every technical actor whose access Nova manages is an **identity** in Nova. Its status describes where it stands in the employment relationship. Everything else follows from it: whether it may sign in, whether its accounts in the target systems are active and whether its entitlements take effect.

The principle

Only active identities may do anything. Every other status blocks sign-in to Nova, the accounts in the target systems and taking part in approvals and recertifications.

Nova implements this rule as an allow list: what counts is whether an identity is active – not whether its status is on a block list. A status added in the future is therefore always blocked at first.

The five statuses

Status	Sign in, approve, review	Accounts in target systems	Entitlements
Pre-hire contract signed, start date in the future	● no	● prepared, ● locked	● prepared, effective from the start date
Active employed	● yes	● active	● effective
Suspended temporarily paused, e.g. parental leave, sabbatical or after a security incident	● no – approvals go to the deputy	● locked	● kept
Terminated employment ended	● no – approvals go to the deputy	● locked immediately, ● deleted after the retention period	● removed after the grace period

Trash terminated and marked for permanent deletion	no	deleted	removed

THE CORE RULE

Suspended means: the entitlements stay. Someone returning from parental leave can continue working right away.

Terminated means: the entitlements go. Someone who is rehired later starts without old access.

Grace period after termination

On termination, Nova locks all accounts of the identity immediately. Its entitlements stay assigned for a configurable **grace period**, but have no effect because the accounts are locked.

The grace period serves two purposes:

- **Handovers:** while the entitlements are assigned, it remains visible which tasks and access need to be handed over.
- **Corrections:** a termination recorded by mistake or too early can be undone without requesting entitlements again.

When the grace period ends, Nova removes the entitlements – in Nova and in all target systems.

Transitions

A status only changes through defined procedures. Nova logs every change with the time, the triggering person or process, the reason and the state before and after.

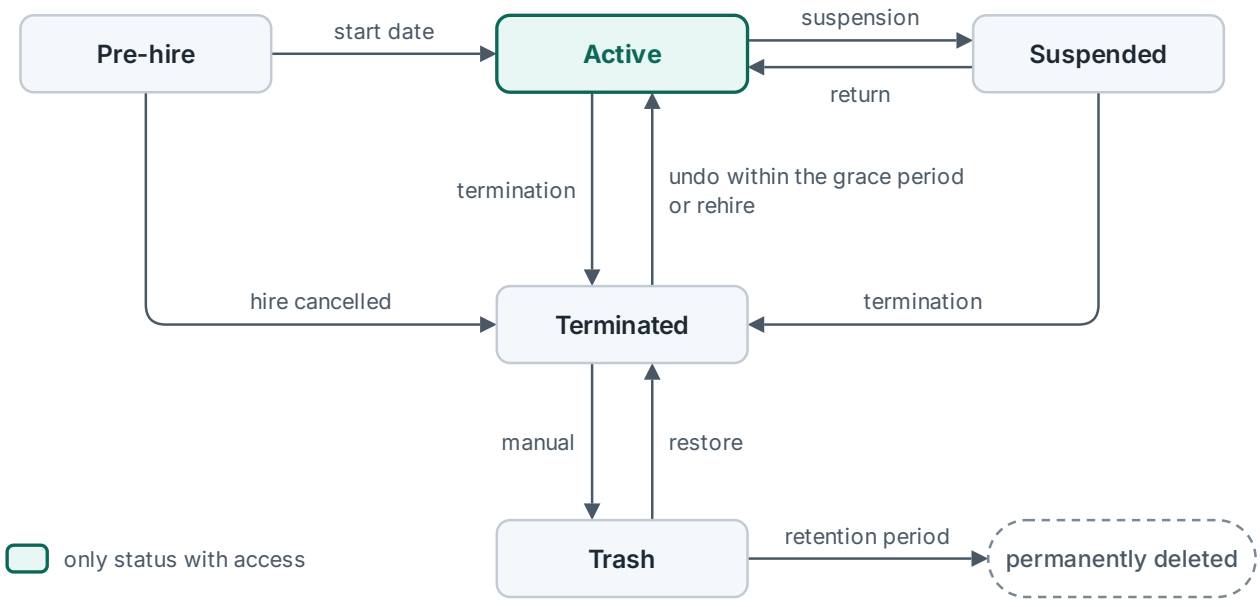

Transitions between the statuses. Only an active identity has access.

From → to	Trigger	What Nova does
Pre-hire → Active	start date reached	unlock accounts, make entitlements effective
Pre-hire → Terminated	hire cancelled	discard prepared accounts and entitlements
Active → Suspended	suspended employment reported by HR, or manual suspension	lock accounts, end active sessions, hand open approvals to the deputy
Suspended → Active	return	unlock accounts; entitlements unchanged
Active or Suspended → Terminated	end date reached, termination reported by HR, or manual termination	lock accounts, end active sessions, transfer ownerships, start the grace period
Terminated → Active, within the grace period	undo termination	unlock accounts; entitlements are kept
Terminated → Active, after the grace period	rehire	fresh start without old access; entitlements granted anew through roles and requests
Terminated → Trash	manual	hide the identity and mark it for permanent deletion
Trash → Terminated	restore	bring back the record only – without access
Trash → permanently deleted	retention period expired	remove personal data; audit log entries are kept

Mapping from SAP HCM

When Nova takes over personnel data from SAP HCM, the employment status (infotype 0000, field STAT2) determines the status of the identity:

Employment status (STAT2)	Status in Nova
3 – active	Active
1 – inactive	Suspended
0 – withdrawn	Terminated
2 – retiree	Terminated

A hire with a start date in the future results in **Pre-hire**. Other source systems are mapped onto the same statuses through the field mapping.

2.2 Access model

Nova organises access on three levels: **entitlements** in the target systems, **business roles** as bundles of entitlements, and **assignments** that link both to identities. Entitlements and business roles are listed under “Roles”, each in their own tab.

Entitlements

An entitlement is a single permission in exactly one target system – an SAP role (single or composite role) or a group in Active Directory and LDAP, Microsoft Entra ID, Keycloak or a SCIM-enabled application. Nova takes these entitlements over by importing them from the target system; they are listed in the “System Roles” tab.

Nova also knows entitlements without a target system. Nova does not transfer them to any target system:

- **“Nova Entitlements”** govern what someone may do in Nova itself. Pre-configured are “Admin” for administering Nova, “Security” and “SoD Reviewer” for decisions in certain approval steps, and “Manager”. Nova assigns this entitlement to people set as manager of an organisational unit.
- **“Nova AI”** are capabilities for AI agents. Only identities of the AI agent type receive them – see AI agents as identities.

Business roles

A business role bundles entitlements – often from several target systems – for a job function. In the “Assignments” tab of a business role, administrators define per target system which entitlements it contains. In addition, every business role carries:

- “Risk” – its risk score,
- “Owner” – the person responsible,
- “Area” – the business area, which also determines the colour,
- “Approval workflow” – the procedure for access requests for this business role.

Only business roles and entitlements with an approval workflow of at least one step can be requested – see Approval workflows.

When administrators change the content of a business role, Nova adjusts the assignments of all identities holding it as soon as the change is saved.

Assignments

An assignment links an identity to a business role or an entitlement. Nova distinguishes three types:

Type	Meaning

Business role	The identity holds a business role.
Indirect	An entitlement from an assigned business role. Nova creates it with the business role and removes it with the business role.
Direct	A single entitlement without a business role – for example from an approved access request for that entitlement, from taking over an account from the target system, or a Nova entitlement.

An identity holds each business role at most once. AI agents do not receive business roles.

Validity

Assignments have a validity period in Nova, “Valid From” and “Valid To”. It describes the intended state:

- A business role without a start date applies from the day it is assigned. The end date is optional and must not be before the start date; without an end date, the assignment has no time limit.
- Indirect assignments take over the period of their business role.
- The end date is inclusive: an assignment applies up to and including the day under “Valid To”.
- For SAP, Nova also records the period it last read from SAP – the actual state.

What happens at the start and end of a period is described in How access is granted and removed.

Organisational units and positions

Business roles can be attached to organisational units; the members of the unit receive them as assignments of their own. If “Inheritance” is switched on for a unit, its members also receive the business roles of all parent units. With the “Planstellen (OSP)” add-on, positions carry business roles too; their holders receive them together with those of the unit the position belongs to. Details are described in Organisation and positions.

Risk score

Every entitlement and every business role carries a risk score: “No Risk”, “Low”, “Medium”, “Critical” or “Not Scored”. The default is “Not Scored”. How Nova assesses risk beyond this is described in Risk assessment.

2.3 How access is granted and removed

In Nova, access is an assignment (see Access model). Nova brings assignments into the target systems through provisioning. This page describes how assignments come about and end, and when Nova transfers them.

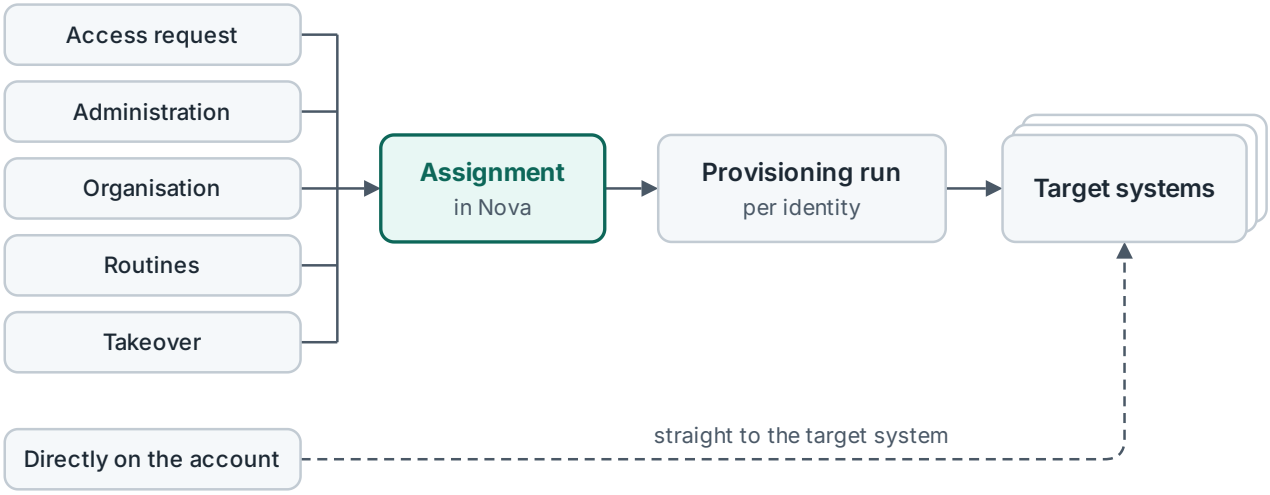

Assignments come about in five ways and reach the target systems through provisioning runs. Changes directly on the account are written straight away.

How access is granted

Route	What happens
Access request	Once an access request has been fully decided, Nova assigns the approved business roles and entitlements – see Requesting access.
Administration	Administrators assign business roles to an identity, optionally with a validity period.
Organisation	Members of an organisational unit and holders of a position receive its business roles – see Organisation and positions.
Routines	Steps in routines can assign business roles, for example via the position – see Joiners, movers, leavers.
Takeover	When Nova reads accounts from a target system, it can take over their existing entitlements as direct assignments – depending on the import settings.
Directly on the account	Administrators change the roles or groups of an account in the identity's “Systems” tab; Nova writes this change straight to the target system.

When Nova transfers to the target systems

Nova transfers assignments in **provisioning runs**. A run concerns one identity: for each affected target system, it checks the account, creates it if needed and transfers the entitlements. If a step fails in one target system, the steps for the other target systems continue. Runs are listed under “Monitoring” → “Provisioning Log”.

After installation, the setting “Auto-provision backend systems” is switched off (“Administration” → “Provisioning” → “Mapping”, tab “Provisioning settings”). When it is switched on, Nova starts a run when

- administrators change the business roles of an identity – assign, revoke or change the period,
- an access request has been fully decided; the approved items are transferred,
- administrators choose “Apply changes to users” after changing the content of a business role; a run then starts for each identity holding the business role. Until then, the business role shows “Pending changes”,
- a routine has changed assignments – including the mover routine that Nova runs after a person has been added to, removed from or moved to another organisational unit,
- the revocations are executed when a recertification is completed,
- a “Cleanup Expired Assignments” job removes expired assignments.

Routine steps that lock, unlock or delete accounts start runs of their own for the identity's linked accounts – even when automatic provisioning is switched off.

Nova only transfers entitlements that belong to a target system. If an identity holds the same entitlement through several assignments, Nova removes it from the target system only once none of these assignments applies any more.

Differences between Nova and a target system are shown in the identity's “Reconciliation & Provisioning” tab. There, administrators resolve them per target system in either direction.

Start and end of validity

Start. Nova transfers an SAP role together with its start and end date, so a future-dated assignment also reaches SAP with the run straight away. Nova transfers no validity periods to Active Directory and LDAP, Microsoft Entra ID, SCIM and Keycloak; for these, Nova holds a future-dated assignment back. The “Activate Pending Assignments” job transfers it for active identities once its start date has been reached; it is pre-configured to run daily at 02:30.

End. Expired assignments are removed by the “Cleanup Expired Assignments” job: it deletes assignments whose end date has passed and, when automatic provisioning is switched on, starts runs that remove the entitlements from the target systems. Nova does not set up this job by itself; administrators create it with “New job” under “Monitoring” → “Background Jobs” and give it a schedule – see Monitoring and jobs.

How access is removed

Route	What happens
Revocation	Administrators revoke a business role from an identity; Nova removes it together with its indirect assignments.
Expiry	The “Cleanup Expired Assignments” job removes assignments whose end date has passed.
Recertification	When completing a campaign, administrators can have the revocations executed; Nova then removes the revoked assignments.
Organisation	When a person leaves an organisational unit, they lose its business roles unless another unit or position grants them.
Leaver	By default, the leaver routine starts runs that lock the identity's linked accounts; it removes assignments only with an additional step – see Joiners, movers, leavers.

Whether and when one of these changes reaches the target systems is described in the section “When Nova transfers to the target systems” above.

2.4 Glossary

This handbook uses the following terms throughout. The German term is given in brackets.

Access request (Antrag)

A request for an entitlement or business role, decided by an approval workflow.

Account (Konto)

An identity's access to a target system, for example an SAP user or an Active Directory account. An account always belongs to exactly one identity.

Approver (Genehmiger)

A person who approves or rejects an access request. The approval workflow defines who approves.

Audit log (Protokoll)

A chronological record of security-relevant actions: who did what, and when.

Business role (Business-Rolle)

A bundle of entitlements – often spanning several target systems – that describes a job function. Whoever receives a business role receives all entitlements it contains.

Entitlement (Berechtigung)

A single permission in a target system, for example an SAP role or an Active Directory group.

Grace period (Karenzzeit) planned

The period after termination during which an identity's entitlements stay assigned before Nova removes them. See Grace period after termination.

Identity (Identität)

A person or technical actor whose access Nova manages. An identity can hold accounts in several target systems. See Identities and their status.

Joiner, mover, leaver (Eintritt, Wechsel, Austritt)

The three events in the lifecycle of an identity. For each, Nova runs a defined procedure, such as creating accounts, adjusting entitlements or locking accounts.

Provisioning (Provisionierung)

Transferring accounts and entitlements from Nova into the target systems.

Recertification (Rezertifizierung)

A regular review of whether existing entitlements are still needed. Nova removes entitlements whose need is denied.

Reconciliation (Abgleich)

A comparison between what Nova intends for an identity and what actually exists in the target system. Nova shows differences and lets you resolve them in either direction.

Rehire (Wiedereintritt)

The return of a terminated identity. Nova runs its own defined procedure for it.

Retention period (Aufbewahrungsfrist)

The period for which Nova keeps an identity in the trash before its data is permanently deleted.

Source system (Quellsystem)

A system from which Nova takes over data about identities – typically the HR system, for example SAP HCM.

Target system (Zielsystem)

A system in which Nova manages accounts and entitlements – for example SAP, Active Directory and LDAP, Microsoft Entra ID, Keycloak or SCIM-enabled applications.

Trash (Papierkorb)

Storage for deleted identities. From there, identities can be restored or permanently deleted.

CHAPTER 3

Lifecycle

3.1 Joiners, movers, leavers

For the events in the lifecycle of an identity – joiner, mover, leaver, rehire and archive – Nova has one **routine** each: a sequence of steps. Administrators define the steps under “Administration” → “Identity Management” → “User Lifecycle”; for each routine, steps can be added, reordered with the drag handle and switched on or off individually. The built-in steps carry German labels in both interface languages; the table under “The steps” below explains them.

The routines and their triggers

Joiner

- **Trigger:** on the dashboard, the “HR Integration” window lists the day's joiners from the HR import. “Onboard” starts the routine for the selected people, after a preview of the steps.
- **Default:** “Setze Identität aktiv - Entsperrung von Konten” and “Setze Identität auf Planstelle”.

Mover

- **Trigger:** the HR import, when a person's position has changed; adding a person to an organisational unit, removing them or moving them; “Run move” in the “HR Integration” window.
- **Default:** “Setze Identität auf Planstelle”.

Leaver

- **Trigger:** the “Leave-Date Scan” job (pre-configured to run daily at 02:00) for active identities whose leave date has passed; deleting an identity, which Nova then moves to the recycle bin.
- **Default:** “Ownership-Übergabe für KI-Agenten” and “Setze Identität inaktiv - Sperrung von Konten”. The default routine does not revoke assignments; a step such as “Entzug aller bestehenden Berechtigungen (Clean Start)” can be added for that.

A grace period after which Nova removes the entitlements of terminated identities is part of the status model in Identities and their status. planned

Rehire

- **Trigger:** the HR import finds an active person whose identity is in the recycle bin – recognised by personnel number or e-mail address – and brings the identity back.
- **Default:** “Entzug aller bestehenden Berechtigungen (Clean Start)”, switched off. Assignments the identity still has are therefore kept; with the step switched on, it starts without old access.

Archive

- **Trigger:** the permanent deletion of an identity from the recycle bin – with “Delete permanently” or by a job of the “Purge Deleted Users” type. This job deletes identities that have been in the recycle bin longer than the retention period; the default is 90 days, adjustable under “Retention (days)”.
- **Default:** no steps.

The steps

Step	Effect
“Setze Identität aktiv - Entsperrung von Konten”	sets the identity to active and starts a run that unlocks its linked accounts in the target systems
“Setze Identität inaktiv - Sperrung von Konten”	sets the identity to inactive and starts a run that locks its linked accounts
“Setze Identität auf Planstelle”	places the identity on the position from the personnel data (attribute P0001-PLANS); it receives the position's business roles. Requires the “Planstellen (OSP)” add-on
“Zuweisung Standardmitarbeiterrolle”	assigns the business role with the ID `standard-mitarbeiter`, if it exists
“Entzug aller bestehenden Berechtigungen (Clean Start)”	removes all assignments of the identity
“Irreversible Löschung von Backend Konten”	starts a run that permanently deletes all linked accounts in the target systems
“Ownership-Übergabe für KI-Agenten”	transfers the AI agents of the departing person to their deputy, otherwise to their manager derived from the organisation; without a successor, Nova marks the agent as orphaned

With “Manage methods”, administrators create methods of their own. These consist solely of building blocks from a fixed catalogue – such as setting a status, assigning or revoking a business role, creating, locking or deleting an account, or creating an access request – and can be tied to conditions on fields of the identity. When AI is switched on, “Generate with AI” proposes a method from a description in plain words. “Test” really runs a method against a selected identity; it is not a dry run.

How a routine runs

- Nova runs the switched-on steps in the defined order.
- If a step fails, Nova rolls back its changes and continues with the next step.
- If the steps have changed assignments, Nova starts a provisioning run per identity when automatic provisioning is switched on. Steps that lock, unlock or delete accounts start their runs even when automatic provisioning is switched off – see How access is granted and removed.

3.2 HR integration

Nova takes over personnel data from a source system, creates identities from it and keeps them up to date. Administrators set up source systems under “Systems” → “Source Systems” with “Add source system”.

Two kinds of source system

Type	Supplies	Connection
“Nova HR Stage”	personnel data: master data, employment status, position	staging area in the Nova database
“SAP OM”	organisational structure: organisational units and positions	SAP via RFC

The “SAP OM” source system is described in Organisation and positions. This page covers personnel data.

Personnel data in SAP HCM format

An export from the HR system provides the personnel data in the staging area of “Nova HR Stage”; how this export is set up is defined in the implementation project. “Nova HR Stage” expects a personnel number per person and attributes named after SAP HCM infotype and field – for example `P0002-VORNA` for the first name. Values can carry a validity period. By default, Nova takes over attributes from these infotypes:

Infotype	What Nova takes from it
0000 Actions	personnel number, employment status (STAT2), hire and leave date
0001 Organizational Assignment	department, position (PLANS), job description, location, building, floor, room
0002 Personal Data	first name, last name, salutation, academic title, name prefix
0105 Communication	e-mail address, phone, mobile phone, logon UID

In the field mapping, the field for the position is called “Plant”.

A person counts as active if their employment status (STAT2) has the value 3; any other value counts as not active. If the source supplies no employment status, the person counts as active. A finer mapping, for example for suspended employment, is described in Identities and their status. planned

Field mapping

Which attribute flows into which field of the identity is defined by administrators under “Administration” → “Provisioning” → “Mapping”, in the “Source system” tab. There are four options for each field: a source attribute, a “Fixed value”, an “Expression” that transforms or combines attributes, or “Disabled”. Without a mapping of its own, the default from the table above applies. The “Preview” shows the result for a personnel number; when AI is switched on, “Generate proposal” suggests an expression.

The status with which Nova creates new identities is set by the “Joiner status” in the source system's field mapping (button “Feld-Mapping”): “Onboarding” (default) or “Active”.

From source system to identity

The source system has buttons of its own for this. They carry German labels in both interface languages; the dialogs they open are translated.

1. **“Rohdaten einlesen”** (read raw data) takes the provided personnel data into the working table that simulation and import use. The source is authoritative here: empty values overwrite existing ones.
2. **“Simulieren”** (simulate) shows what an import would change and changes nothing itself.
3. **“Personen importieren”** (import persons) runs the import – for all people or for a selection. On a schedule, a job of the “Import HR users” type performs this step.

“Person debuggen” (debug person) shows the raw data for a personnel number, the mapping and the values calculated from them.

Nova finds a person's identity by personnel number, otherwise by e-mail address, and classifies each person:

Classification	Meaning	What the import does
“Joiner”	active, no identity in Nova yet	creates an identity with the “Joiner status”
“Mover”	position in the source differs from the one in Nova	writes the changes, then the mover routine runs
“Update”	other details changed	writes the changes
“Leaver”	no longer active, identity in Nova still active	nothing, see below
“Unchanged”	no change	nothing

Apart from rehires, the import does not change the status of existing identities. Administrators start the joiner routine on the dashboard with “Onboard” – see Joiners, movers, leavers.

Leavers. The import does not process people who are no longer active. While a person is active, however, it takes over their leave date if the source supplies one. Once that date has passed, the “Leave-Date Scan” job starts the leaver routine.

Rehires. If the import finds an active person whose identity is in the recycle bin, it brings the identity back and starts the rehire routine.

Dates and future changes

For each attribute, Nova uses the value that applies on the day the data is read. If no value applies on that day yet, Nova takes the next future one. A future change therefore takes effect when Nova reads and imports the data on or after its effective date. If a new position already applies, Nova takes its start date from the validity period in the source – even if the import runs later.

3.3 Organisation and positions

The organisation in Nova consists of organisational units and – with the “Planstellen (OSP)” add-on – positions. It determines which business roles people receive through their unit or position and who counts as their manager. Administrators maintain it under “Organization”.

Organisational units

Organisational units form a hierarchy: each unit has at most one parent unit. Nova assigns every unit an eight-digit object ID; units from SAP OM carry the object ID from SAP.

- **Members:** “Add person” adds a person to the unit. Nova sets their department to the unit's object ID.
- **Managers:** the “Manager” button on a member makes them manager of the unit; clicking it again reverses this. Nova assigns managers the Nova entitlement “Manager” and removes it again when the person is no longer a manager anywhere. Anyone moved to another unit loses their manager status.
- **Business roles:** a business role attached via “Assign business role…” is received by the members of the unit. If an administrator takes it off the unit again, Nova revokes it from the members; anyone who also receives it through another unit keeps it.
- **Inheritance:** when “Inheritance” is switched on (“Inherit business roles from parent units”), the unit's members also receive the business roles of all parent units. It is switched off by default.

Business roles a member receives through their unit cannot be removed individually on the identity – only by the person leaving the unit or by taking the business role off the unit.

When an administrator adds a person to a unit, removes them or moves them to another unit, Nova then runs the mover routine – see Joiners, movers, leavers.

Positions

Positions are available with the “Planstellen (OSP)” add-on. Administrators enable it under “Administration” → “Plugins”; after a restart, a unit shows the “Positions” tab.

- A position belongs to exactly one organisational unit and can have holders. A position marked “Chief” leads the unit: its holder counts as the unit's manager.
- Business roles can be attached directly to a position. The holder receives them together with the business roles of the unit the position belongs to – with “Inheritance” switched on, also with those of the parent units.
- An identity holds at most one position at a time. When it is placed on a new one, the previous holding ends; Nova revokes business roles that only the previous position granted.
- From the personnel data, the routine step “Setze Identität auf Planstelle” places an identity on the position named in the attribute P0001-PLANS. This requires a position in Nova with that object ID, for example from the sync with SAP OM.

Structure from SAP OM

The “SAP OM” source system reads the organisational structure from SAP via RFC (tables HRP1000 and HRP1001), limited to one “Plan variant” (default 01). “Root org unit (optional)” narrows the scope further.

- “Org-Struktur lesen” (dialog “Read org structure”) stores the result as a job run and changes nothing in Nova.
- “Org-Struktur abgleichen” (dialog “Sync org structure”) creates or updates organisational units and – with the add-on – positions in Nova, matched by object ID. Nova takes names, parent units and chief positions from SAP. The sync does not set holders; these come from the personnel data – see HR integration.

In the tree under “Organization”, a coloured dot shows whether a unit comes from SAP OM and whether its parent unit still matches SAP.

Managers

For access requests for an identity, approval steps of the “Manager” type go to its manager. That is the “Manager” entered on the identity. Without such an entry, Nova derives the manager from the organisation: what counts is the nearest manager in the identity's unit – a member set as manager or the holder of the chief position –, otherwise the nearest one in the parent units. In the identity's “Overview”, the derived manager is shown as “Org-derived”.

How access is granted and removed describes when Nova transfers changes to assignments to the target systems.

CHAPTER 4

Target systems

4.1 Overview

A **target system** is a system in which Nova manages accounts and entitlements. Nova talks to target systems through connectors; it ships five:

- SAP – ABAP systems via RFC and BAPIs
- Active Directory and LDAP
- Microsoft Entra ID – via Microsoft Graph
- Keycloak – via the Admin API
- SCIM – applications with a SCIM 2.0 interface

Target systems of type “Miscellaneous” have no connector; Nova exchanges no data with them.

What Nova manages in target systems

All five connectors provide the same basic functions:

- **Accounts:** create them, change their master data, lock, unlock and delete them
- **Passwords:** set them
- **Entitlements:** grant and revoke them – roles in SAP, groups in all other target systems
- **State:** read an account's state – whether it exists, whether it is locked and, where the target system provides it, when it was last used

Nova takes over the catalogue of roles or groups as entitlements through an import job. Administrators define which master data of an identity goes into which attribute of the target system under “Administration” → “Provisioning” → “Mapping”.

Target system	Entitlements	Locked via	Validity
SAP	single and composite roles	SAP user lock	in the target system
Active Directory	groups, including nested ones	`userAccountControl`	in Nova
LDAP	groups	password policy	in Nova
Entra ID	groups	`accountEnabled`	in Nova
Keycloak	groups	`enabled`	in Nova
SCIM	groups	`active`	in Nova

Validity: only SAP understands validity periods of role assignments itself. Nova passes the start and end dates on, even when an assignment only begins in the future. The other target systems only know "member or not". There, Nova holds back a future-dated assignment until its start date. A preconfigured job of type "Activate Pending Assignments" provisions such assignments of active identities daily at 02:30 (time zone Europe/Berlin).

Setting up a target system

Administrators create a target system under "Systems" → "Target Systems" → "Add target system". Only administrators can create, change or delete target systems. A target system with a connector has:

- **"Connection"** – address and credentials, with rarely needed settings under "Advanced". The pages of the individual connectors list the fields.
- **"Account setup"** – "Account ID template", "Password ruleset" and "Account ID ruleset". All three are mandatory; without them Nova does not save the target system. Administrators maintain the rulesets under "Administration" → "Provisioning" → "Password Rules" and "Account ID Rules".

Account IDs

The "Account ID template" defines how Nova forms the ID of a new account, e.g. `{first_initial_lower}{name_lower}` → `emustermann`. Placeholders exist for first and last name, initial, full name, email address, personnel number and department; name parts are also available in lower or upper case. Nova removes accents and umlauts (Müller → Muller). An unknown placeholder causes an error rather than a wrong ID.

When a provisioning run creates an account, Nova forms the ID from the template and shortens it to the maximum length of the "Account ID ruleset". If the ID is already taken in Nova or in the target system, Nova appends a number (`emustermann2`, `emustermann3` … up to 99 at most). In future, Nova appends the number only when the ID is taken in Nova; if the target system holds an account with this ID that Nova does not know, Nova reports a conflict instead – see Accounts and passwords. planned

The **"Account ID ruleset"** limits the length and the permitted special characters, or prescribes a regular expression. Nova checks the ID against it when an account is created manually or linked. Nova ships a default ruleset for each connector, for example at most 12 characters for SAP.

The **"Password ruleset"** applies to passwords entered manually when an account is created or a password is set. Initial passwords that Nova generates itself are random, follow the ruleset and are at least 12 characters long.

Checking the connection

"Check status" in the list of target systems or in a target system's detail view tests the connection and stores the result – "online", or "offline" with an error message. What is tested depends on the connector:

Target system	Test
SAP	call of `RFC_PING`

Active Directory, LDAP	sign-in (bind) with the service account
Entra ID	retrieval of an access token
Keycloak	token retrieval and counting the accounts in the realm
SCIM	retrieval of `ServiceProviderConfig`, otherwise of the base URL

A successful test proves reachability and sign-in, not every permission Nova will later need in the target system.

In addition, a preconfigured job of type “Target System Health Poll” checks the reachability of the target systems every ten minutes. When a system is newly found to be offline, Nova records this in the “Audit Log”.

Credentials

Nova stores passwords, client secrets and tokens of target systems encrypted in the database. The operator provides the key at start-up; it is not kept in the database. Nova returns stored secrets neither in the user interface nor through the programming interface; a placeholder is shown in their place. To change a secret, an administrator enters a new value. Secrets can neither be read nor set through the AI assistant.

Traceability

Nova records writes of the connectors in target systems in the change journal – with the values before and after.

4.2 SAP

Nova manages accounts – called users in SAP – and their roles in SAP ABAP systems. To do so, Nova calls standard BAPIs and the function module `RFC_READ_TABLE` via RFC.

> **PREREQUISITE: SAP NETWEAVER RFC SDK**
>
> For RFC, Nova uses the Python library pyrfc, which requires the SAP NetWeaver RFC SDK. Nova does not ship the SDK; the operator obtains it from SAP. Without the SDK, the SAP functions are not available.

Connection

One target system stands for one SAP system with one client. The fields: “SID”, “Application Server Host”, “System number”, “Client”, “User” and “Password” of the RFC user, and “Language”. Under “Advanced” are “Message server host”, “Message server port” and “Logon group” for load balancing, and the “SAP router string”. “Check status” calls `RFC_PING`.

Accounts

Action	Function module
Create	`BAPI_USER_CREATE1`
Change address data	`BAPI_USER_CHANGE`
Lock, unlock	`BAPI_USER_LOCK`, `BAPI_USER_UNLOCK`
Delete	`BAPI_USER_DELETE`

Nova commits every change with `BAPI_TRANSACTION_COMMIT`. For SAP, first and last name must be available separately; Nova does not split a full name itself. After creating an account and after name changes, Nova reads the name fields back and reports any deviation.

Nova passes **passwords** as initial passwords, not as productive passwords. SAP therefore requires a dialog user to choose a new password at the first logon. If Nova creates an account without a password, password logon is deactivated. Passwords that Nova generates for SAP have at most 40 characters and start neither with “!” or “?” nor with three identical characters.

Nova reads the **last logon** from table `USR02`.

Roles

Nova supports single and composite roles. A job of type “Import SAP Roles” reads the role catalogue, optionally filtered (for example `Z*`). Nova recognises composite roles by their flag in `AGR_FLAGS`, reads their single roles from `AGR_AGRS` and short texts in German and English from `AGR_TEXTS`. If SAP returns

the role structure incompletely, the import stops and leaves the previously imported definitions unchanged. Roles are scoped to their target system: roles with the same name in two SAP systems remain separate entitlements.

Validity

SAP manages validity periods of role assignments itself (`FROM_DAT`, `TO_DAT`). Nova passes on the start and end date of every assignment, even when it only begins in the future. Without a start date, Nova uses the current date; without an end date, 31 December 9999. If the same role has several periods with a gap in between for one account, Nova writes nothing and reports an error instead of bridging the gap.

Writing and reading back

SAP only accepts role assignments as a complete list (`BAPI_USER_ACTGROUPS_ASSIGN`). Nova therefore sends the account's complete target role set, with its periods, on every change.

> ⚠ **TAKING OVER EXISTING ACCOUNTS**
>
> Because SAP takes the list it receives as the complete role set, the existing roles of an account should be known to Nova before Nova assigns roles to it for the first time: import the role catalogue, then run a reconciliation with the target system as the leading side. In future, roles that Nova does not manage stay untouched when Nova assigns roles. planned

After the commit, Nova reads the account's roles back and compares names and periods with the target state, taking the single roles of a composite role into account. If anything differs, the change counts as failed and Nova names the roles concerned. Nova records the changes it detects per role in the change journal.

One role for many accounts

A background job of type “Bulk SAP account roles” grants one SAP role to many accounts or revokes it from them, with an optional validity period. Accounts are selected through filters such as identity type or department; without a status filter, Nova selects only active identities, and never AI agents. A run covers at most 500 accounts.

Nova processes the accounts one after another and writes and checks each of them as described above. The job does not revoke roles that a valid business role of the identity requires. Nova collects errors of individual accounts in the result without aborting the run. A cancellation by an administrator takes effect between two accounts.

Authorisations Nova needs in the SAP system

The SAP connector calls the following function modules. The SAP Basis team can derive a role for the RFC user from them.

Function module	Purpose
`RFC_PING`	connection test

`BAPI_USER_GETLIST`	list accounts (import)
`BAPI_USER_GET_DETAIL`	read address, lock status and roles
`BAPI_USER_CREATE1`	create accounts
`BAPI_USER_CHANGE`	change address data, set passwords
`BAPI_USER_LOCK`, `BAPI_USER_UNLOCK`	lock, unlock
`BAPI_USER_DELETE`	delete accounts
`BAPI_USER_ACTGROUPS_ASSIGN`	assign roles
`BAPI_TRANSACTION_COMMIT`	commit changes
`RFC_READ_TABLE`	read tables, see below

Table	Fields read	Purpose
`AGR_DEFINE`	`AGR_NAME`, `PARENT_AGR`	role catalogue
`AGR_FLAGS`	`AGR_NAME`, `FLAG_TYPE`, `FLAG_VALUE`	recognise composite roles
`AGR_AGRS`	`AGR_NAME`, `CHILD_AGR`	single roles of a composite role
`AGR_TEXTS`	`AGR_NAME`, `SPRAS`, `TEXT`	role texts
`USR02`	`BNAME`, `TRDAT`, `LTIME`	last logon

For orientation: SAP usually checks such calls through the authorisation objects `S_RFC` (function modules), `S_USER_GRP` (user maintenance), `S_USER_AGR` (role assignment) and `S_TABU_NAM` or `S_TABU_DIS` (reading tables). Which objects and values a system actually checks is decided by the SAP Basis team.

The add-on module “SoD-Analyse (SAP ABAP)” has to be switched on separately. It additionally reads table `AGR_1251` with the authorisation values of the roles; see Segregation of duties (SoD).

4.3 Active Directory and LDAP

One connector covers LDAP directories. Under “Variant”, administrators define how the directory is structured:

Property	Generic LDAP (ApacheDS)	OpenLDAP	Active Directory
Groups	`groupOfNames`, `member`	`posixGroup`, `memberUid`	`group`, `member`
Account ID	`uid`	`uid`	`sAMAccountName`
Unique ID	`entryUUID`	`entryUUID`	`objectGUID`
Lock	`pwdAccountLockedTime`	`pwdAccountLockedTime`	`userAccountControl`
Last logon	not preset	`authTimestamp`	`lastLogonTimestamp`, `lastLogon`

Connection

The fields: “Host”, “Port”, “Base DN”, and “Bind DN” and “Bind password” of the service account. Under “Advanced” are “User Search Base” and “Group Search Base” – if left empty, the base DN applies – and “Use SSL (LDAPS)”. For Active Directory, further fields are added, including “AD domain” and “Initial Account State”. “Check status” signs in to the directory with the service account.

Accounts

Nova creates new accounts directly below the “User Search Base”: for Generic LDAP and OpenLDAP as `uid=<account ID>`, for Active Directory as `cn=<name>`. In Active Directory, Nova sets `sAMAccountName` to the account ID and `userPrincipalName` according to the “UPN rule”, or to account ID@AD domain if there is no rule. The “Mail rule” fills the `mail` attribute for all variants. If “Initial Account State” is set to “Disabled”, a new AD account is created locked and stays locked until it is explicitly enabled. Nova writes further attributes according to the “Mapping”.

Nova finds accounts by `uid` or `mail`, in Active Directory by `sAMAccountName`, `userPrincipalName` or `mail`.

Groups

Nova takes over groups as entitlements with a job of type “Import LDAP Groups”. Nova recognises a group by its unique ID, provided the directory supplies one. If the group is renamed or moved within the “Group Search Base”, Nova finds it again through this ID and updates the stored DN. Once the ID is known, Nova accepts only exactly this object: a newly created group with the old name or DN does not count as the same group.

Nova changes memberships one at a time; the other members of a group remain untouched. Generic LDAP does not maintain `memberOf` itself. There, Nova writes it in addition, provided the directory schema knows the attribute.

Nested groups

Nova reads the group structure. The “Add group” dialog on an account shows the subgroups and parent groups of each group. Nova does not derive assignments from nesting. For Generic LDAP and Active Directory, administrators can nest groups in this dialog (“Add subgroup”) and remove nesting again. OpenLDAP groups (`posixGroup`) do not allow nesting.

Cycle protection: before Nova enters a group as a subgroup, it checks whether that group already contains the parent group, directly or across further levels. If it does, Nova does not make the change. The same applies if the check cannot be completed – for example because an object cannot be read or the limit of 64 levels, 512 reads or 5 seconds is reached. If the directory already contains a cycle, Nova shows a notice and does not run through the cycle repeatedly. The display reads at most 5,000 groups and 64 levels; if it is incomplete, Nova says so.

Passwords

- **Active Directory:** Nova writes `unicodePwd`. The directory only accepts passwords over an encrypted connection; without “Use SSL (LDAPS)”, Nova refuses to set the password.
- **Generic LDAP and OpenLDAP:** Nova uses the password modify operation according to RFC 3062 and, if the server does not support it, the `userPassword` attribute.

Locking

- **Active Directory:** Nova sets or clears the `ACCOUNTDISABLE` bit in `userAccountControl`; the other bits remain unchanged.
- **Generic LDAP and OpenLDAP:** Nova locks permanently through the password policy attribute `pwdAccountLockedTime` and removes it to unlock. The password is kept; Nova reads the result back. This requires an active password policy – for OpenLDAP the ppolicy overlay – and the right to write this attribute.

Last logon

In Active Directory, Nova reads `lastLogonTimestamp`, otherwise `lastLogon`; for OpenLDAP `authTimestamp`, provided the directory maintains it. For Generic LDAP no attribute is preset; one can be specified in the configuration key `last_logon_attrs`.

Directory permissions required

The service account (“Bind DN”) needs:

- **Read:** accounts below the “User Search Base”; groups with their member attribute and unique ID below the “Group Search Base”; for Generic LDAP also the schema, to check `memberOf`; for the last logon, the attributes named above.

- **Groups:** write the member attribute (`member` or `memberUid`) on the groups Nova manages; for Generic LDAP also `memberOf` on the accounts.
- **Accounts:** create, modify and delete below the “User Search Base”.
- **Passwords:** reset them – in Active Directory via `unicodePwd`.
- **Locking:** write `userAccountControl` or `pwdAccountLockedTime`.
- **Nesting:** write `member` on parent groups and read the subgroups involved.

Nova should connect via LDAPS (“Use SSL (LDAPS)”); Active Directory requires this for password operations.

4.4 Microsoft Entra ID

Nova manages accounts and their group memberships in Microsoft Entra ID. Nova talks to Entra ID through Microsoft Graph (version 1.0).

Connection

Administrators create an app registration with a client secret for Nova in Entra ID. Nova uses it to sign in without a user context (OAuth 2.0 client credentials) and acts with the app's application permissions.

The fields: “Tenant ID”, “Client (App) ID”, “Client Secret” and “UPN domain”. Under “Advanced” are “Authority (optional)”, “Graph endpoint” and “Scopes” (default: `https://graph.microsoft.com/.default`). “Check status” requests an access token. This confirms the app's credentials, not its Graph permissions.

Accounts

- **Create:** Nova creates the account enabled, with the account ID as `userPrincipalName` and the part before the “@” as `mailNickname`; further attributes follow the “Mapping”. If only the name part is entered when creating an account manually, Nova appends the “UPN domain”.
- **Change** and **delete** via the Graph endpoint `/users`.
- **Lock and unlock** via `accountEnabled`. Afterwards, Nova reads the state back and reports an error if it cannot be confirmed.

Groups

Nova takes over groups as entitlements with a job of type “Import Entra Groups”. Nova writes direct group memberships; the other members of a group remain untouched. Nova does not manage Entra directory roles.

Nova reads nested groups up to a limit of 500 groups and 32 levels. From this, Nova shows the group structure and, for an account, the section “Effective inherited memberships”. Nova reports cycles and incomplete reads as a notice. Nova does not write inherited memberships; they are not assignments in Nova.

Check after writing

During provisioning and reconciliation, Nova first reads all direct group memberships of the account. Nova then removes the groups to be revoked and adds the new ones. Afterwards, Nova reads the complete list again – across all pages of the Graph response – and compares it with the expected state.

Because Entra ID can make changes visible with a delay, Nova repeats the read for up to 30 seconds. If the state still does not match, the step counts as failed; Nova names missing and unexpected groups.

Last logon

Nova reads `signInActivity` and takes the most recent point in time from interactive, non-interactive and successful sign-ins. This requires the `AuditLog.Read.All` permission.

Graph permissions required

For each task, the table names an application permission that Microsoft documents for the calls concerned:

Task	Graph calls	Permission
Read accounts	`GET /users`	`User.Read.All`
Create, change, lock, delete accounts	`POST` , `PATCH` , `DELETE /users`	`User.ReadWrite.All`
Read groups and group structure	`GET /groups` , `…/members` , `…/memberOf`	`GroupMember.Read.All`
Read an account's groups	`GET /users/{id}/memberOf`	`Directory.Read.All`
Write memberships	`POST` , `DELETE /groups/{id}/members`	`GroupMember.ReadWrite.All`
Last logon	`GET /users` with `signInActivity`	`AuditLog.Read.All`

For setting passwords (`passwordProfile`), Microsoft requires rights beyond this table; they are described in the Microsoft Graph documentation for “Update user”. Microsoft adjusts permissions from time to time; the current documentation is authoritative.

4.5 Keycloak

Nova manages accounts – called users in Keycloak – and their group memberships in Keycloak. For this, Nova uses the Admin REST API; one target system corresponds to one realm.

Connection

Nova signs in through a confidential client with a service account (OAuth 2.0 client credentials). The fields: “Base URL”, “Realm”, “Client ID” and “Client Secret”.

The service account needs the roles `manage-users`, `view-users`, `query-users` and `query-groups` of the `realm-management` client. Depending on use, these are added:

- `view-events` – to read the last logon from login events
- `manage-clients` and `view-clients` – to create AI agents as clients

“Check status” requests a token and counts the accounts in the realm. The test therefore also checks whether the service account may read accounts.

Accounts

- **Create:** username, email address, first and last name; the account is enabled. Nova stores further attributes from the “Mapping” as Keycloak attributes of the account.
- **Change:** Nova reads the account, adds the changed values and writes it back. This keeps the account's other attributes intact.
- **Lock and unlock** via the `enabled` field.
- **Delete.**
- **Passwords:** Nova sets them through the Admin API.

Keycloak stores usernames in lower case; Nova therefore also writes and compares them in lower case. Nova finds accounts by username, email address or Keycloak ID.

Groups

Nova takes over groups and subgroups as entitlements with a job of type “Import Keycloak Groups”. Nova shows them with their path, for example “Engineering/Backend”, and recognises them by their Keycloak ID. Nova does not manage realm or client roles.

Keycloak lists an account as a member only of the groups it belongs to directly. Nova reads and writes only these direct memberships, one at a time. During provisioning, Nova first reads the current memberships so that the change journal contains only effective changes.

Last logon

Nova takes the most recent point in time from the realm's login events and from the account's active sessions. Login events exist only if the realm stores them and the service account has `view-events`. Without them, Nova only knows the last access from currently active sessions.

AI agents

Nova does not create identities of type “AI Agent” as Keycloak users, but as a confidential client with a service account. Their group memberships belong to the service account of this client. Nova does not store the client secret; “Rotate Secret” creates a new one, which Nova shows once. See AI agents as identities.

4.6 SCIM

SCIM 2.0 is an open standard for managing accounts and groups. Through SCIM, Nova manages accounts and group memberships in applications that offer a SCIM 2.0 interface.

Connection

- **“Base URL”:** Nova appends the SCIM paths itself, for example `/scim/v2/Users`. If the endpoints are located under `https://app.example.com/api/scim/v2/`, the base URL is `https://app.example.com/api`.
- **“Auth type”:** “Bearer” with a token, “Basic” with “User” and “Password”, or “None”.
- **“Last-logon attribute (SCIM)”** under “Advanced”, see Last logon.

“Check status” retrieves `/scim/v2/ServiceProviderConfig`, otherwise the base URL itself.

Accounts

Action	SCIM call
Search	`GET /Users` with filter `userName eq "…"`
Create	`POST /Users`
Change	`PATCH /Users/{id}`, operation `replace`
Set password	`PATCH` on the `password` attribute
Lock, unlock	`PATCH` on the `active` attribute
Delete	`DELETE /Users/{id}`

When creating an account, Nova sends `userName` (the account ID), `displayName`, `name` with first and last name taken from the identity's name, `active`, the email address as the primary address and, if present, the password. Further simple attributes follow the “Mapping”.

Groups

A job of type “Import SCIM groups” takes over the group catalogue; it reads the groups page by page.

Nova changes memberships with a `PATCH` on the group (`/Groups/{id}`, operation `add` or `remove` on `members`); the other members remain untouched. As the member value, Nova passes the account ID under which the account is linked in Nova. During provisioning, Nova looks up the group by its display name (`displayName eq "…"`). If a group is renamed in the SCIM server, its name has to be updated in Nova, for example by importing again with “Overwrite existing roles”. Before writing, Nova reads the group's members so that the change journal contains only effective changes.

Last logon

SCIM has no standard attribute for the last logon. Nova reads the attribute entered under “Last-logon attribute (SCIM)” – a path with dots is possible. If none is entered or it returns no value, Nova checks `lastLogin`, `lastLogon`, `lastLoginTime` and `loginTime`.

What differs between SCIM servers

SCIM leaves vendors some room. Before going live, it should be clear:

- whether the endpoints sit under a path ending in `/scim/v2/`
- which authentication the server requires – Nova supports bearer tokens and Basic
- whether the server accepts the account ID as a group member value or requires the account's internal SCIM ID
- whether the server supports the filters `eq` (for `userName` and `displayName`) and `co` (group search)
- whether `PATCH` is supported for accounts (`/Users`) and groups and whether `password` and `active` can be written with it
- whether, and under which name, the last logon is provided

A run-through with a test account – create, grant and revoke a group, lock, delete – shows whether a server fits these assumptions.

CHAPTER 5

Accounts and provisioning

5.1 Accounts and passwords

PLANNED

This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

An **account** is an identity's access to a target system, for example an SAP user or an Active Directory account. This page describes how Nova names and links accounts, where the account status it shows comes from and how Nova handles passwords. How Nova creates accounts during provisioning is described under Provisioning runs.

One account, one identity

Every account belongs to exactly one identity. If an account is already linked to an identity, Nova refuses to link the same account to another identity.

Account IDs

Nova builds the ID of a new account from the target system's naming rule – see Target systems – Overview.

- **ID already exists:** if an account with the computed ID already exists in the target system, Nova never takes it over silently. Nova stops and reports a conflict. An administrator decides: link the existing account after checking it, or assign a different ID.
- **Creation uncertain:** if it is unclear whether a create request reached the target system – for example after a timeout – Nova keeps the chosen ID. On the next attempt, Nova looks for exactly this account and links it instead of creating a second one.

Account status

Nova shows an account's status, such as active or locked, as it last read it in the target system – together with the time of that read. A newly linked account shows “unknown” until Nova has read it.

- **Locking and unlocking** count as done only once Nova has read the new state back from the target system.

- **Missing accounts:** Nova concludes that an account does not exist only when the search in the target system succeeded and was unambiguous. A failed or ambiguous search does not count as “not found”.

Passwords

- **Generated passwords** are random and follow the target system's password rules.
- **Passwords set by administrators** must be changed by the identity at the next sign-in – in every target system that supports this: SAP, Active Directory, Microsoft Entra ID and Keycloak.
- **Encrypted only:** to directories such as Active Directory, Nova transmits passwords only over LDAPS or StartTLS.
- **Not stored:** Nova does not store the passwords of accounts in target systems. They never appear in the logs or in the change journal.

Leavers

When an identity leaves, Nova locks all its accounts. If a lock fails, Nova retries it until it is confirmed.

Nova permanently deletes an identity only once all its accounts are confirmed as locked or deleted – see Identities and their status.

Orphaned accounts

For each target system, Nova shows the accounts that exist there but belong to no identity.

5.2 Provisioning runs

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

A **provisioning run** brings an identity's intended state from Nova into the target systems. A run concerns exactly one identity. It consists of steps per target system, such as checking the account, creating it and transferring entitlements.

When a run is created

When automatic provisioning is switched on – see How access is granted and removed – every change to the intended state creates a run for each affected identity, including:

- an approved access request,
- assigning or revoking a business role,
- attaching business roles to, or removing them from, organisational units and – with the "Planstellen (OSP)" add-on – positions; see Organisation and positions,
- a change to the definition of an entitlement or business role,
- deleting an entitlement or business role,
- the start and end of a validity period.

Nova stores every run before it executes. If a restart interrupts a run, Nova resumes it afterwards.

For any one identity, only one operation writes to the target systems at a time. This applies to runs as well as to every other way in which Nova writes to target systems.

Transferred and confirmed

Nova distinguishes two states:

State	Meaning
transferred	The target system accepted the write.
confirmed	Nova read the state back from the target system; it matches the intended state.

For each assignment, Nova shows whether it is transferred or confirmed and when Nova last confirmed it.

Removals: an entitlement counts as removed only once the removal is confirmed. If Nova cannot look up a group or role in the target system, the step fails and is retried. In that case, Nova never reports it as done.

Existing roles in SAP

Roles on an SAP account that Nova does not manage are preserved. Before its first write to an account, Nova reads the roles the account already has. In the reconciliation, they appear as present only in the target system until an administrator explicitly decides on them.

Retries

- Nova retries failed runs automatically, at increasing intervals.
- Nova re-evaluates every retry against the current state. Whether Nova locks or unlocks an account, for example, follows the status the identity has at the time of the retry.
- If a newer run has superseded an older one, Nova does not replay the older run.

Validity

Nova evaluates the start and end of validity periods in the configured business time zone.

Bulk changes and pausing

- **Preview:** when a change affects many identities – for example when a changed business role is applied to all its holders – Nova shows a preview with the number of affected identities and entitlements before it writes anything.
- **Pause:** each target system has a pause switch, and there is also a global one. While paused, Nova holds back all writes to the affected target systems; they stay queued.

What a run shows

Each run shows

- who or what triggered it: a person, an access request, a business role change, an organisational change or a job – with a link to the request if one triggered it,
- when it was triggered,
- its steps and their results.

Runs are listed under “Monitoring” → “Provisioning Log”. How runs relate to access requests and to changes in the target systems is described under Logs.

5.3 Reconciliation

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

In a **reconciliation**, Nova compares what is intended for an identity with what actually exists in the target system. Nova shows the differences; they are resolved deliberately, in one direction or the other. For a single identity, the “Reconciliation & Provisioning” tab shows the differences per target system.

Leading side

For each target system, it is configured and visible which side leads: Nova or the target system.

Only complete reads count

Reconciliation and imports act only on data that was read completely:

- If a target system returns its results page by page, Nova reads through to the last page.
- If a read remains incomplete – because of an error, a timeout or missing read permissions – Nova changes nothing and marks the result as incomplete.

What an import never removes

An import never removes assignments that are still pending, start in the future or come from an approved access request.

Validity counts

Nova never shows an expired or not-yet-valid assignment as “In Sync”.

Changes made outside Nova

If someone changes a target system directly, Nova detects the difference the next time it reads the system and shows it with the time of the last read. Administrators resolve it deliberately, in one of two directions:

- take the target system's state over into Nova, or
- restore Nova's intended state in the target system.

Bulk actions first show a preview of exactly what will change.

The change journal records only what Nova itself has written. Changes made outside Nova become visible through reconciliation.

CHAPTER 6

Access requests

6.1 Requesting access

Entitlements and business roles can be requested in Nova through an **access request**. Each requested role is decided by the approval workflow assigned to it. Only after that decision does Nova grant the access.

Who can request what for whom

- **Requester:** every signed-in identity can submit a request – for itself, for another identity or for a department. Whether access is granted is up to the approvers.
- **Requestable** are entitlements and business roles that have an "Approval workflow" assigned. Nova does not accept roles without a workflow in any request.
- **Several roles:** a request can contain several roles. Each runs through its own approval workflow.
- **AI agents:** access for an AI agent can only be requested by its accountable owner, the owner's registered deputy or administrators. Through requests, an AI agent receives AI capabilities (roles `nova-ai:*`) only; these cannot be requested for any other identity.

Submitting a request

Under "Access requests" → "My requests", "New request" opens the form:

1. **"Requested for"** – a person or a department. The signed-in identity is preselected.
2. **"Roles"** – the requestable entitlements and business roles, searchable and filterable by "Category". Roles that all target identities already hold are hidden.
3. **Pre-check** – Nova checks the selection against the segregation of duties rules and shows conflicts as "Risk findings". Depending on the risk policy, Nova only warns, adds a review step or blocks the request.
4. **"Submit request"**.

For a **department**, Nova creates one shared request: approvers decide once for the whole department, not per person. Which people receive the access is determined only at final approval, from the members at that time.

An access request carries no validity dates. An assignment resulting from it is valid from the approval onwards, without an end date.

Request status

Status	Meaning
“Pending”	At least one requested role has not been decided yet.
“Approved”	All requested roles are approved.
“Partially approved”	All roles are decided – some approved, some rejected.
“Rejected”	All requested roles are rejected.
“Cancelled”	The request was withdrawn or cancelled by an administrator.

The “Access requests” page has three tabs: “To approve” shows what the signed-in identity may decide – including as a deputy –, “My requests” its own open requests and “Completed” the finished ones. Administrators see all requests under “To approve” and “Completed”. Each request has a chat between requester and approvers; Nova records decisions there as system messages.

After the decision

As soon as every role in a request has been decided, Nova assigns the approved roles to the target identities. Existing assignments stay unchanged. Rejected and cancelled requests lead to no assignment.

If “Auto-provision backend systems” is switched on under “Administration” → “Provisioning” → “Mapping”, Nova also creates a **provisioning run** for each identity whose new roles affect a target system:

- Nova saves the assignment and the run in the same transaction. If saving fails, neither is created.
- Nova starts the run only afterwards. If the run cannot be started, it stays saved and Nova resumes it at its next start.

If the option is switched off, Nova assigns the roles in Nova only. In the request, the “Provisioning” section shows the state of the runs. “Approved” means: assigned in Nova. Nova tracks the execution in the target systems separately.

Withdrawing and cancelling

While a request is “Pending”, the requester can withdraw it with “Cancel request”. Administrators can cancel any open request. A cancelled request cannot be reopened. Nova notifies the approvers of open steps and logs the cancellation.

6.2 Approval workflows

An **approval workflow** – “Workflow” in the user interface – defines who decides on an access request for a role. Every entitlement and every business role points to at most one workflow through its “Approval workflow” field. Without a workflow it cannot be requested.

Maintaining workflows

Administrators maintain workflows under “Administration” → “Workflows”. A workflow consists of the fields “ID (slug)”, “Name” and “Description” plus an ordered list of steps (“Approval Steps”); each step has an approver type.

- Only administrators create, change or delete workflows.
- A workflow needs at least one step; without steps Nova accepts no request.
- Nova does not delete a workflow that roles still point to.
- On submission, Nova copies the steps into the request. Later changes to the workflow apply to new requests.

Who decides a step

Approver type	Who decides
“Manager”	the target identity's manager, recorded or derived from the organisation; for a department, its head
“Role Owner”	the “Owner” of the requested role
“Agent Owner”	the accountable owner of the AI agent
“Risk Owner”	the identity stored in the step
“Specific User”	the identity stored in the step
“Security”	every identity holding the “Security” entitlement
“SoD Review”	every identity holding the “SoD Reviewer” or “Security” entitlement

- For open steps, Nova keeps resolving the responsible person anew. If the manager changes, for example, the new manager may decide; the person resolved at submission remains authorised as well.
- If nobody can be resolved – for example because a role has no owner – only administrators can decide the step.
- An “SoD Review” step only becomes active when there are open risk findings; otherwise Nova skips it. Whoever approves it despite open findings must enter a comment. See Segregation of duties.

- Administrators can decide any open step and close a request as a whole with “Approve (admin)” or “Reject (admin)”. Nova logs this intervention as a separate event and notifies the requester and the approvers of open steps.

Order and outcome

- All steps of a workflow are open from submission and can be decided in any order. The step number determines the order of notifications.
- A requested role is approved when all its steps are approved or skipped. A single rejection rejects it.
- “Approve” and “Reject” in the request decide, in one go, all open steps for which the signed-in identity is the resolved approver. A comment is optional.
- If a workflow consists only of “SoD Review” steps and there are no findings, the role counts as approved on submission.

Deputies during absence

An identity can have a “Deputy” and, under “Absence”, the dates “Absent from” and “Absent until”. If both dates are set and today falls within this period, the deputy may additionally decide the absent person's open steps.

- The absent person remains authorised.
- Nova does not resolve chains: if the deputy is also absent, the authority does not pass on to the deputy's deputy.
- In the “To approve” tab, such requests carry the note “As deputy for …”. Nova logs the decision together with the person represented.
- During the absence, the deputy also receives the notifications addressed to approvers.

Reminders and escalation

The background job “Approval Escalation” makes overdue steps visible. Nova does not create it by itself; administrators set it up under “Monitoring” → “Background Jobs” with “New job”.

- From “Remind After (Days)” – default 3 – Nova reminds the responsible approvers, at most once a day.
- From “Escalate After (Days)” – default 7 – Nova alerts the administrators once. The reminders continue.
- A step's age counts from the latest decision on an earlier step of the same role, otherwise from submission.
- The job approves nothing and skips no step.

Notifications

Event	Recipients
“Access request submitted”	resolved approvers of the first active step
“Approval step completed”	resolved approvers of the next step
“Access request approved”	requester, target identity
“Access request rejected”	requester
“Access request cancelled”	approvers of open steps
“Admin override on request”	requester, approvers of open steps
“Pending approval reminder”	responsible approvers
“Overdue approval escalated”	administrators

Under “Administration” → “Notifications”, administrators define per event whether Nova notifies and through which channels; the texts can be adapted. The “Email” channel only sends once a mail server is configured. The “Send mode” (“Active”, “Paused”, “Off”) applies to all channels.

Microsoft Teams is an add-on module: after “Microsoft Teams Notifications” has been switched on under “Administration” → “Plugins”, Nova can post messages to a Teams channel through a webhook. If “Interactive approvals (approve in Teams)” is also set up, the notified approvers receive a card with “Approve” and “Reject” through a Power Automate flow. One click decides an open step of that person; Nova checks the authorisation just as for a decision made in Nova.

CHAPTER 7

Governance

7.1 Recertification

In a **recertification**, responsible people check whether existing access is still needed and confirm or revoke it. Nova groups this review into **campaigns**. Administrators manage campaigns under “Compliance” → “Recertification”.

Creating a campaign

“New campaign” creates a campaign:

- **“Scope”** – whose access is reviewed: “All users”, a “Department”, an “Org unit” (without sub-units), all identities with an account in a “System”, or “AI Agents”. Nova excludes disabled and deleted identities.
- **“Reviewer”** – who reviews: the identity's “Manager”, the role's “Role owner”, a “Specific user” or, for AI agents, their “Agent owner (AI identity)”.
- **“Deadline”** – optional.
- **“Campaign type”** and **“Standard”** (e.g. ISO 27001) serve for classification.

Nova immediately creates one item per identity and assignment. Reviewed are the business-role and direct entitlement assignments valid on the day of creation; entitlements an identity receives through a business role are covered by reviewing that business role. Each item records identity, role, system and risk rating as at creation. Nova rejects a scope without items. For the “AI Agents” scope, the capabilities of linked routine agents are added as separate items.

If no reviewer can be resolved for an item, the person who created the campaign reviews it.

Flags on items

Nova flags items that deserve particular attention; the “Flagged” filter lists them together:

- **“SoD conflict”** – the assignment is part of a conflict under an active SoD rule.
- **“External w/ high access”** – an external identity holds a role rated “Critical”.
- **“No peer holds this”** – within the scope, only this one identity holds the role.
- **“Reviewer defaulted to admin”** – no reviewer could be resolved for the item.

Starting and deciding

A campaign starts as “Draft”. “Start” sets it to “In review”; Nova notifies every reviewer of the number of their items.

- Each item is decided with “Confirm” or “Revoke”. A revocation requires a reason.
- Reviewers decide only the items assigned to them; administrators decide all.
- Several items can be selected and decided together. “Confirm all” confirms all open items – for reviewers only their own.
- A decision is final. Nova stores it on the item with person, time and reason.
- Items can only be decided while the campaign is “In review”.

The “Deadline” does not close a campaign automatically. The background job “Recertification Deadline Reminder” reminds reviewers with open items at most once a day as soon as the deadline lies within “Remind within (days)” – default 7 – or has passed. Nova does not create this job by itself; administrators set it up under “Monitoring” → “Background Jobs”.

Completion and revocation

Once all items are decided, administrators close the campaign with “Complete campaign”. Nova asks whether the revocations should be executed now:

- **Execute:** Nova removes the revoked assignments in Nova – for a business role including the entitlements received through it. If “Auto-provision backend systems” is switched on, Nova starts provisioning runs for the affected identities to remove the roles in the target systems as well. If the capabilities of a routine agent are revoked, Nova switches that routine agent off.
- **Do not execute:** the decisions remain documented and the access stays in place. The campaign is closed nevertheless; the revocations can no longer be executed through it later.

> ⚠ **REVOCATION ONLY ON COMPLETION**
>
> A “Revoke” only takes effect when the revocations are executed on completion. Until then, the access remains unchanged.

After completion, Nova notifies the person who created the campaign and the administrators. “Download CSV” gives administrators all items with reviewer, decision, reason, time, flags and revocation status as evidence. Completed campaigns show under “Risk reduction (revoked)” how many revoked assignments had which risk rating. Nova logs the creation, start, decisions and completion of a campaign.

7.2 Segregation of duties (SoD)

Segregation of duties (SoD) is meant to prevent an identity from combining entitlements that must stay separate – for example creating vendors and posting payments. For this, Nova checks **SoD rules**. Two add-on modules extend the check for SAP.

SoD rules

An SoD rule describes a **conflicting pair**: two entitlements or business roles (“Side A” and “Side B”) that the same identity should not hold. Each rule has a “Severity” – “Critical”, “High”, “Medium” or “Low” – and can be deactivated. Administrators maintain the rules under “Compliance” → “Risk Analysis” in the “SoD Rules” section. SoD rules are part of Nova and need no add-on module.

Nova takes into account all of an identity's assignments valid on the day of the check: direct entitlements, business roles and the entitlements received through business roles.

When Nova checks

Occasion	What is checked	Result
new access request	existing access of the target identity plus the requested roles; only conflicts the request contributes to	“Risk findings” on the request
scan	existing access of all identities except disabled and deleted ones	“SoD Violations” register
recertification	existing access of the identities in scope	“SoD conflict” flag on the item

Assignments that do not arise from an access request – for example those made directly by administrators or through org units – are not checked in advance. The scan picks up conflicts from such assignments.

Nova shows findings for a request already in the “New request” form, then in the list of requests and, within the request, next to the affected role.

Response to access requests

The “Risk policy” under “Risk Analysis” defines how Nova responds to findings. The default is “Warn only”.

“Risk policy”	Effect
“Warn only”	Nova shows the findings on the request.
“Require SoD review step”	If a requested role has findings and its workflow contains no “SoD Review” step, Nova appends such a step.

“Block request”	If there are findings, Nova does not create the request. Administrators can still submit it with an “Override justification”.

- An “SoD Review” step is decided by identities holding the “SoD Reviewer” or “Security” entitlement. Whoever approves it despite open findings must enter a comment; the findings then count as overridden. If such a step is already part of the approval workflow, it becomes active under “Warn only” as well as soon as there are findings; without findings Nova skips it.
- Administrators can override individual findings in the request with “Override” and a justification. “Approve (admin)” also requires a comment when findings are open.
- Nova logs overrides with person and justification.
- With the “SoD rules” and “GRC plugin” switches, administrators define which sources the check includes.
- If a check fails for technical reasons, Nova still creates the request and logs the error.

SoD violations register

“Run scan” under “Risk Analysis” or the background job “SoD Violation Scan” compares existing access with the active rules; Nova does not create the job by itself. The result is kept in the “SoD Violations” register:

- New conflicts appear as “Open”.
- Administrators set a violation to “Accepted” or “Mitigated” – both only with a justification – or reopen it.
- If a scan no longer finds an open or accepted violation, Nova sets it to “Resolved”. If a resolved conflict reappears, Nova reopens it.
- If administrators delete a rule, Nova also removes its register entries, including the justifications. When a rule is deactivated, the entries are kept.

Add-on modules for SAP

Both modules are switched off out of the box. Administrators switch them on under “Administration” → “Plugins”.

“SoD-Analyse (SAP ABAP)” checks SAP ABAP roles internally: for critical single permissions and for combinations that should not occur together in one role. It works on the role's authorisation values (table AGR_1251), which Nova reads from the SAP system via RFC. Administrators trigger reading and checking; neither runs automatically. If an SoD finding on an access request involves a role checked this way, Nova stores that role's results with the finding as technical evidence.

When switched on, the module takes over the supplied default rule set – a YAML file – provided no rules exist yet. A rule in it looks like this:

```yaml
rules:
  - key: vendor_create_and_payment   # stable key
    name: "Kreditor anlegen UND Zahlung buchen"
    risk: critical                   # critical | high | medium | low
    type: conflict_combo             # or critical_single
    requires_all:                    # all conditions must match
      - { auth_object: F_LFA1_APP, field: ACTVT, values: ["01", "02"] }
      - { auth_object: F_BKPF_BUK, field: ACTVT, values: ["01"] }
```

A rule of type `critical_single` names a single condition under `match` instead of `requires_all`. `values: ["*"]` matches any value of the field.

“GRC-Berechtigungsanalyse (SAP Access Control)” reads the rule set (risks and functions) as well as the assignments of roles to accounts from an existing SAP GRC Access Control. Nova then also reports, on access requests, the risks that the requested SAP roles would newly add for the target identity. At request time, Nova does not access SAP for this but uses the data read beforehand.

7.3 Risk assessment

Nova distinguishes two kinds of risk information:

- the **risk rating** of an entitlement or business role – a fixed classification maintained by administrators;
- **risk findings** – conflicts that Nova derives from rules for an access request. They are described under Segregation of duties.

The levels

A role is rated “Critical”, “Medium”, “Low” or “No Risk”. “Not Scored” stands for a missing classification and is the default for new roles. There is no “High” level for roles; it only occurs in risk findings.

Maintenance

- Administrators set the rating in the “Risk” field of an entitlement or business role.
- A business role's rating stands on its own: Nova does not derive it from the entitlements it contains. Exception: when role mining creates a business role, Nova takes over the highest rating of its entitlements.
- With the add-on module “GRC-Berechtigungsanalyse (SAP Access Control)”, administrators can transfer the risks determined with the GRC rule set as ratings onto the SAP roles in Nova. For each role the highest risk counts; “High” becomes “Critical”. Nova overwrites an existing rating.

Where the rating appears

- under “Roles”, in the list and in the detail view of every entitlement and business role;
- on an identity's assignments;
- in a request, next to the requested roles;
- in recertification, on every item – recorded when the campaign is created – and, after completion, under “Risk reduction (revoked)”;
- on the “Dashboard”:
  - “Risk Exposure” – number of assignments of critically rated entitlements and how many identities hold them;
  - “Risk Distribution” – all entitlement assignments, direct and through business roles, by level;
  - “Top Risk Users” – up to five identities with the most assignments of critically rated entitlements;
  - “Risk Scoring” in the “Compliance Posture” – the share of entitlements that carry a rating.

What the rating does

The risk rating on its own changes no approval workflow and adds no review step. Which steps a request runs through is defined by the role's approval workflow; additional steps only arise from risk findings. The rating feeds into the recertification flag “External w/ high access” and serves role mining and AI analyses as context.

AI analysis of an identity

If an AI mode is active, administrators can have an AI analysis created for an identity. It estimates a risk level and names anomalies and recommendations. The result can be edited and saved; it changes no ratings and no assignments. See Overview and modes.

7.4 Role mining

Role mining searches the **direct assignments** of an area for recurring patterns and proposes business roles based on them. Nova creates nothing on its own: administrators review every proposal, and only accepted ones are applied. The feature is available to administrators only, under “Roles” → “Role Mining”.

What data Nova analyses

- **Scope:** an org unit, optionally with its sub-units (“Include sub-units”, on by default). A run covers at most 10,000 identities.
- **Identities:** active ones only; technical identities and AI agents are left out by default.
- **Assignments:** direct ones only. What an identity receives through business roles is already bundled; a run leaves existing business roles and their assignments untouched. Nova does not analyse AI capabilities (`nova-ai:*`).
- **Structure:** Nova uses org units and positions to bind proposals to them.

At the start, Nova freezes the scope; review and preview work on this snapshot. Only one run can be active at a time.

How proposals are formed

Nova computes four layers one after the other. Whatever one layer covers no longer counts in the following ones. This way, every direct assignment ends up in exactly one proposal or stays direct.

Layer	Formed from (defaults)	Binding
“Global base role”	roles held by at least 90 % of the identities in scope	the scope's org unit
“Org-unit base roles”	per org unit, the roles held by at least 70 % of its members	that org unit
“Functional roles”	roles shared by all holders of a position, and frequent combinations of at least 2 roles held by at least 3 identities	position or identities
“Special functions”	per identity, at least 2 roles held by fewer than 3 identities in scope	that identity

The thresholds can be adjusted in the “New mining run” dialog, as can the “Maximum number of BRs” (default 12; special functions do not count). Candidates beyond this budget – ranked by the number of assignments they would convert – appear as “Over BR budget”. With the “Outlier mode” set to “Report only”, special functions only appear as “Outlier”. All figures, such as coverage and effects, are calculated by Nova itself.

AI review (optional)

If an AI mode is active, an AI assesses the candidates before they appear as proposals:

- It suggests a name and description, confirms a candidate or places a **veto**, and classifies special functions as a special function or an anomaly.
- A veto turns the candidate into the finding “No layer candidate” with the AI's reasoning; an anomaly appears as “Anomaly” with a recommendation.
- The AI cannot add roles or identities. Nova discards answers containing unknown identifiers and does not adopt any figures from the AI.
- The AI receives key figures, names, descriptions and ratings of the roles as well as names of org units, positions and existing business roles. For candidates bound to individuals, names of identities are added, for special functions also their department and position.
- Without AI, or if the AI fails, the run delivers the same statistical candidates with generic names.

Review and apply

- Proposals start as “Pending”. Administrators “Accept” or “Reject” them, can rename them and remove individual roles (“Remove role”). A proposal without roles counts as rejected.
- “Apply…” first shows a preview against the current data: how many direct assignments would be converted and which identities would newly receive roles (“New grants”).
- Nova applies accepted proposals only. For each proposal, Nova creates a business role – with the highest risk rating of its roles –, assigns it to the reviewed identities and replaces their covered direct assignments with assignments through the business role. Nova carries over validity dates in the process.
- “Bind business roles to org units and positions” (on by default) additionally attaches the business role to the org unit or position, so that people joining later inherit it. When switched off, only the reviewed identities receive the role.
- If the current data deviates from the reviewed state, Nova skips the proposal (“Skipped (drift)”). A new run then provides an up-to-date picture.
- “Findings” and “Remaining direct assignments” are a report only; Nova revokes nothing based on them. Revocations go through recertification or manual maintenance.

Nova logs the start, review decisions and application of a run, as well as every business role created, together with its origin in the run.

CHAPTER 8

Traceability

8.1 Change journal

The change journal records what Nova has written in target systems – with the values before and after. It answers the question “What has Nova changed in this target system?” and is the basis for reverting individual changes.

What an entry contains

The journal records writes of the connectors of all five target system types: accounts created, changed, locked, unlocked and deleted, passwords set, and group memberships and SAP roles added and removed. An entry contains:

- the time, the target system and the object – account, group or role
- the operation, for example “Group assigned” or “Account disabled”
- the affected values “Before” and “After”
- where known, the person who triggered it (“Actor”)
- whether the change can be reverted and whether it has already been reverted

For group memberships and SAP roles, Nova records only effective changes: if Nova finds, for example, that a membership already existed, no entry is created. For SAP, Nova determines the role changes by reading the account's role set before and after writing.

Passwords are not stored in the journal. Nova replaces password attributes such as `userPassword`, `unicodePwd`, `password` or `passwordProfile` with `<redacted>` – also in the snapshot of a deleted object.

Where to find the journal

Administrators find the journal under “Monitoring” → “Provisioning Log”. The selection “Backend change” shows journal entries only. Clicking an entry opens the details with system, object, actor and the values “Before” and “After”.

The view shows the most recent entries. Through the programming interface and the AI assistant, the journal can also be searched by target system, operation, account and period.

Reverting changes

For Active Directory, LDAP and Microsoft Entra ID, administrators can revert individual entries. The details then show “Revert”; after the confirmation “Revert change”, Nova performs the inverse operation, for example removing a membership that was added.

Beforehand, Nova checks whether the current state in the target system still matches the recorded after-state. If the object has been changed in some other way since, Nova refuses the revert in the user interface. The revert appears as an entry of its own, and the original entry is marked “Reverted”. Each entry can be reverted only once; Nova rejects a further attempt with “This change was already reverted”.

Target system	Reversible	Not reversible
Active Directory, LDAP	account created, changed, locked, unlocked or deleted; group memberships	passwords
Entra ID	group memberships	account changes, passwords
SAP, Keycloak, SCIM	–	all entries

Except for account creation, a revert requires a recorded before-value; if Nova could not read it, the entry is not reversible. Entries that cannot be reverted serve traceability only.

Deleted accounts in Active Directory and LDAP

Before deleting an account, Nova stores a snapshot of the entry. The revert recreates the account from it, provided no entry exists under the DN. The new object gets a new unique ID (`objectGUID`, `objectSid` or `entryUUID`); rights bound to the old ID do not pass to the new object. Nova does not restore passwords or attributes maintained by the directory itself. Nova restores group memberships where possible; it records failures as a warning in the “Audit Log”.

What the journal does not record

The journal describes writes that Nova itself performs in target systems. It does not record changes that others make directly in the target system, nor changes within Nova.

8.2 Logs

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

Nova keeps three records. Each answers a question of its own:

Record	Answers
“Audit Log”	Who did what in Nova?
“Provisioning Log”	What did Nova intend and execute for each identity?
Change journal	What did Nova actually change in the target system?

One chain, one correlation ID

A shared correlation ID links the entries that belong together:

Access request → approval → assignment → provisioning run → change in the target system → confirmation

This way, an access request can be followed into the target system – and, conversely, a change in the target system back to its request.

What Nova records

- **Identities and assignments:** every change to identities and to their entitlements and business roles.
- **Accounts:** every change to accounts, including password resets and locks. Nova files these entries under the affected identity.
- **Catalogue and organisation:** changes to entitlements, business roles, organisational units and approval workflows.
- **Configuration:** changes to target systems and settings, each with the values before and after; secrets masked.
- **Sign-ins:** successful and failed sign-ins to Nova – see Signing in to Nova.
- **AI assistants:** actions of an AI assistant, marked with the assistant – see AI guardrails.

Every entry names its **actor** – who or what initiated the action: a person, a job or an AI assistant.

Unalterable and permanent

- Log entries are append-only. No function in Nova changes or deletes them, and the database user Nova runs with has no right to do so.

- Nova stores timestamps with their time zone and shows them in local time.
- When an identity is deleted, its log entries are kept.
- Secrets such as passwords, client secrets, tokens and keys never appear in any log.

Who can see the logs

Only authorised roles can see the logs: administrators and auditors. The logs are found under “Monitoring” → “Audit Log” and “Provisioning Log”; the change journal is part of the “Provisioning Log”. Personal log views show each identity only its own entries.

How auditors use the logs is shown in the Auditor guide.

8.3 Auditor guide

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

This page maps typical audit questions to the places in Nova that answer them. They build on the three logs and the correlation ID that links their entries.

Starting point: an identity's history

An identity's "History" shows the whole chain as a timeline: access request → approval → assignment → provisioning run → change in the target system → confirmation. Questions about a single person start here. Questions about many identities are answered by reports under "Reporting" and their exports.

> **EXPORTING THE AUDIT LOG**
>
> The filtered "Audit Log" can be exported as CSV or JSON.

Typical questions

Who has access to X – and since when?

- **Where:** a report under "Reporting", as an export.
- **Evidence:** all identities with a valid assignment of the entitlement or business role, each with the start and end of validity and the time of the last confirmation in the target system.

Why does person Y hold entitlement Z?

- **Where:** the identity's "History".
- **Evidence:** who requested the entitlement, who approved it and when, when Nova provisioned it and when Nova confirmed it in the target system.

Which leavers still have an active account anywhere?

- **Where:** a report under "Reporting" on the accounts of terminated identities.
- **Evidence:** for each account, the status last read in the target system with the time of that read. How Nova locks accounts when an identity leaves is described under Accounts and passwords.

Which accounts in a target system belong to no one?

- **Where:** the target system's orphaned accounts – see Accounts and passwords.
- **Evidence:** all accounts that exist in the target system but belong to no identity.

What changed in SAP in a given period – and who triggered it?

- **Where:** the change journal, narrowed down to the SAP system and the period.
- **Evidence:** every change Nova wrote there, with the values before and after and the actor; through the correlation ID, also the related run and access request. Changes that others made directly in SAP are shown by reconciliation.

Which SoD conflicts exist – and who accepted which exception?

- **Where:** “Compliance” → “Risk Analysis”, “SoD Violations” register – see Segregation of duties (SoD).
- **Evidence:** open, accepted and mitigated violations; for accepted and mitigated ones, the person who decided and their justification. Nova also records overrides on access requests with the person and justification.

What is the evidence for a recertification?

- **Where:** “Compliance” → “Recertification”, then “Download CSV” – see Recertification.
- **Evidence:** for each item, the reviewer, decision, reason, time, flags and revocation status.

CHAPTER 9

AI in Nova

9.1 Overview and modes

Nova uses language models for analyses, suggestions and a chat that lets people operate Nova in plain language. AI complements Nova; by default it is switched off.

Principles

- **The core works without AI.** Identities, entitlements, access requests and approvals, provisioning, reconciliation and recertification do not call a language model. Without AI, only the AI features themselves are missing.
- **AI proposes, people decide.** Analyses, descriptions, role candidates, import parsers and steps for lifecycle routines are suggestions. Entitlements and master data change only when an authorised identity adopts or runs a suggestion.
- **The chat acts only on request.** It works with the permissions of the signed-in identity; actions marked as destructive run only after explicit confirmation. See AI guardrails.

Modes

Mode	Where requests go
Off	nowhere – the interface does not offer its AI features
Local	to an Ollama endpoint, for example in the organisation's own data centre; an administrator enters the endpoint URL and the model
Cloud	to Anthropic (Claude) or OpenAI, optionally through a custom base URL

Every identity chooses the mode with the "NOVA-AI" switch in the header: "Off", "Local" or "Cloud". Nova remembers the choice in the browser; the default is "Off". When switching, Nova checks whether the chosen mode is set up and otherwise switches to "Off". The healthcheck has an AI setting of its own.

AI is set up under "AI Configuration" in the menu that opens when clicking one's own name in the header; only administrators can save there. For the cloud, the following can be selected:

- **Claude (Anthropic):** Claude Opus 5.5, Opus 5, Sonnet 5, Sonnet 4.6, Opus 4.8 and Haiku 4.5
- **OpenAI:** GPT-4o, GPT-4o Mini, GPT-4.1 and GPT-4.1 Mini

Operations sets the API key as an environment variable on the server (`CLAUDE_API_KEY` or `OPENAI_API_KEY`). The interface only shows whether it is set.

Where Nova uses AI

- **The "Nova AI" chat** with assistants. Three are preset: "Access Requests" finds roles and files access requests, which go through the usual approval workflow; "Help & Navigation" answers questions and only reads; "Admin Chat" is open to administrators only and may use every tool.
- **"AI user analysis"** of an identity: risk assessment, anomalies and a comparison with other identities in the same organisational unit. Administrators can edit and save the result.
- **Reports:** from a description, AI generates the queries and layout of a report.
- **Role mining:** Nova computes the business role candidates without AI; AI names, describes and assesses them. See Role mining.
- **Migration Workbench:** AI writes the parser for an import file. See Migrating from SAP IdM.
- **Suggestions during setup:** the description of a business role, expressions for field mappings and steps for lifecycle routines, composed from fixed building blocks.
- **Further uses:** "Governance Analysis" (without AI, the questions remain open for manual assessment), the classification of findings in the "AI Healthcheck" and a pre-check under "Report a problem".

Which data is sent to the model

In "Local" mode, the following data goes to the configured Ollama endpoint; in "Cloud" mode, to Anthropic or OpenAI.

Feature	What the model receives
Chat	the messages; name, department, roles, business roles and organisational units of the signed-in identity; the results of the tools called
"AI user analysis"	name, email, department, location, status and type of the identity; its organisational units; names and risk of its business roles and entitlements; names and types of the target systems; the number of comparison identities
Reports	the description, a chosen template and the structure of the database tables; in the report dialogue, after each run, also the column names, row count and first result row of each query
Role mining	names, descriptions and risk of the entitlements, key figures, names of organisational units and positions; for person-bound candidates, names of identities
Migration Workbench	an excerpt of the file (the first rows, at most 20,000 characters), the chat messages and the result of the last dry run with up to 15 rejected records
Suggestions during setup	the description entered; names of the objects a suggestion may refer to, such as entitlements, business roles and organisational units – for open access requests including the requester's name
Governance Analysis	questions and target answers from the questionnaire, stored facts about Nova
AI Healthcheck	up to 100 open findings with subject and details

Report a problem	description, type of report, current page and an attached screenshot

If a model does not support tool calls, the chat falls back to a simpler mode. It then sends a data extract along, including the organisation chart with the names of its members and, in the “Access Requests” assistant, the requestable roles.

> **THE SERVER SET-UP DECIDES**
>
> Nova can only reach a model that has been set up. Without an API key on the server, Nova rejects every cloud request before any data leaves the system. Anyone who wants to rule out cloud providers sets no key and, if needed, runs a model through Ollama in their own infrastructure.

9.2 AI guardrails

The “Nova AI” chat can not only read data in Nova but also change it. Nova enforces the limits for this on the server – regardless of which model answers and what it suggests.

Only with the permissions of the signed-in identity

The chat does not access the database itself. It calls tools, and every tool runs one of Nova’s endpoints internally – in the session of the identity using the chat. Permission checks, validation, logging and follow-up actions are therefore the same as in the interface. The chat does not extend the identity’s permissions.

Some tools combine several endpoints into one flow. Writing tools of this kind are available to administrators only; self-service such as filing access requests is the exception.

One tool list per assistant

Under “Administration › AI Assistants”, administrators define for each assistant which tools it may use.

- If the model calls a tool outside this list, Nova rejects the call – including indirect routes, for example through an export.
- Releasing all tools is possible only for assistants restricted to administrators. The preset “Admin Chat” always has all tools and always remains restricted to administrators.
- Assistants open to administrators only are hidden from other identities, and Nova rejects their requests to them.
- The chat does not change the tool lists, the sign-in policy, the lifecycle routines or the AI routine assistants, nor settings that contain secrets such as the AI configuration. This is possible only in the administration interface.

Confirmation before destructive actions

The chat does not run tools that Nova marks as destructive – for example those that delete – straight away. It first shows the “Confirmation required” card with the action and its arguments; under “Current state” it shows the affected object where available. The action runs at the earliest with the identity’s next message: through “Confirm and execute” or, if the identity agrees in a typed reply, through a renewed model call with confirmation. The following applies:

- A confirmation counts only in the immediately following message from the same identity to the same assistant, and for 10 minutes at most.
- Tool and arguments must match the preview exactly; otherwise Nova asks for a new preview.
- The model cannot confirm on its own: Nova rejects a confirmation in the same message that produced the preview.
- “Confirm and execute” runs exactly the stored call. For the reply to it, the model has read-only tools only.
- The chat runs at most one destructive action per message.

Secrets stay out

No tool accepts passwords, client secrets, tokens or API keys as parameters. Nova checks this for every tool definition at start-up and on every call; the chat also refuses to clear a secret. Secrets are entered in the interface. In tool results and previews, Nova replaces such fields and values stored encrypted with a placeholder.

Limits per message

Limit	Value
Model calls in the tool loop	12
Destructive actions	1
Same tool call with the same arguments	3, then blocked
Size of one tool result	12,000 characters, truncated beyond
Tool results in total	60,000 characters, no further calls after that

Nova sets no token or cost budget per identity or period.

Traceability

- Nova marks log entries created during a tool call with `via: ai_chat:<assistant>` in their details. Under “Monitoring › Audit Log”, the “Origin” filter set to “AI chat” shows exactly these entries.
- What gets logged is what the called endpoint logs – whether it is called from the chat or from the interface. Nova writes no entry of its own per message or per tool call.
- Nova stores the chat history per identity and assistant so that it can be reopened under “Chat History”. In Nova, every identity sees only its own histories and can delete them – one by one or with “Clear all”. The history is therefore not evidence; the audit log is what counts.

Which data the chat sends to the model is described in Overview and modes.

9.3 AI agents as identities

Nova manages AI agents as an identity type of their own, “AI Agent”; in the list of identities they appear on the “AI agents” tab. The same means apply to them as to people – access requests, approvals, recertification and logging – supplemented by rules that apply to agents only.

Mandatory details

An administrator creates an AI agent. Nova requires:

- an **accountable person (owner)** – an existing identity of the natural person type,
- a **purpose** describing why the agent exists.

A date for the **next review** and an **expiry date** are optional.

Entitlements: capability roles only

AI agents receive entitlements as capability roles with the prefix `nova-ai:`. For every combination of object (such as identity, business role or account) and action (create, read, update, delete) for which a building block exists, Nova creates such a role – for example `nova-ai:identity.update`. By default, reading counts as no risk, creating and updating as medium risk and deleting as critical; administrators can change this. The roles take effect only in Nova; Nova does not transfer them to target systems.

Nova keeps these roles separate from all other entitlements:

- When assigning and requesting, Nova allows only capability roles for AI agents; it rejects business roles and other entitlements.
- Only AI agents receive capability roles. They cannot be members of a business role.
- Only the agent’s owner, the owner’s deputy or an administrator can file an access request for an AI agent. An approval workflow can name the “Agent Owner” as approver, see Approval workflows.

What the capability roles govern

Nova can compose steps for lifecycle routines from building blocks with AI. Administrators define which objects and actions AI may use for this under “Administration › AI Assistants” on the “AI Routine Assistants” tab:

- Without an active AI routine assistant with released capabilities, AI composes no steps.
- A routine assistant can be linked to an AI agent. The capability roles the agent validly holds then count as well.
- If the linked agent is inactive, deleted or past its expiry date, the routine assistant no longer counts.
- This only affects composing new steps: saved and built-in steps keep running.

The AI agent itself does not act in Nova. The link makes it the accountable point of reference, not an actor.

Reviews

- **Review date:** the “Agent Review Scan” job type finds active AI agents whose review date has passed and writes one audit log entry per agent and run – naming the owner and the linked routine assistants. An administrator creates such a job under “Monitoring › Background Jobs” with “New job”; none is preset. The job sends no notifications.
- **Recertification:** a campaign with the scope “AI Agents” reviews the entitlements of the AI agents; “Agent owner (AI identity)” can be chosen as reviewer. In addition, the matrix of released objects and actions of every linked routine assistant is up for review. If it is revoked, Nova switches the routine assistant off. See Recertification.

When the owner leaves

The preset leaver routine contains the step “Ownership-Übergabe für KI-Agenten” (ownership transfer for AI agents). It transfers the leaving identity’s AI agents to its deputy or, failing that, to its manager according to the organisation. Only an active natural person who is not leaving in the same run can become the successor. If Nova finds none, the agent is left without an owner and Nova writes a warning to the audit log; Nova logs every transfer as well. See Joiners, movers, leavers.

Accounts in Keycloak

In Keycloak, Nova creates a confidential client with a service account for an AI agent instead of a personal account; the client signs in machine-to-machine (client credentials). Its tokens are valid for five minutes by default. “Rotate secret” generates a new client secret; Nova shows it once and does not store it. See Keycloak.

CHAPTER 10

Security and privacy

10.1 Signing in to Nova

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

This page describes signing in to Nova itself. How Nova manages accounts and passwords in target systems is described under Accounts and passwords.

Who can sign in

Only active identities with an active Nova account – that is, an active account on the “Nova IAM” system – can sign in. The statuses an identity can have are described under Identities and their status.

Nova ends running sessions immediately when

- the identity's status changes to **Suspended** or **Terminated**,
- its Nova account is locked, or
- its password is changed.

In addition, every session has a maximum lifetime, after which a new sign-in is required.

Sign-in methods

Method	Note
Password	Nova stores only a salted hash of the password.
One-time code by email	“Email code” tab on the sign-in page
Single sign-on with Microsoft Entra ID	“Login with Microsoft Entra”, if set up
Second factor (TOTP)	with the “Native MFA (TOTP)” add-on

Every method passes the same checks – including the second factor where it is required. Administrators decide whether the second factor is mandatory for administrators or for all identities.

The “Native MFA (TOTP)” add-on is switched off by default; administrators enable it under “Administration” → “Plugins”.

Password rules

Passwords that identities use to sign in to Nova are subject to three kinds of rules:

- **Minimum length and composition** – set under “Administration” → “Provisioning” → “Password Rules”,
- **Password history** – recently used passwords cannot be set again,
- **Maximum age** – once it is exceeded, the identity must choose a new password when signing in.

Administrators set the history and the maximum age in the “Nova Login Policy” under “Administration” → “Identity Management” → “Authentication”.

Failed attempts

After repeated failed attempts, Nova locks the Nova account. Administrators set the threshold, the counting window and the duration of the lock in the “Nova Login Policy” as well. Wrong one-time codes and wrong second-factor codes count like wrong passwords.

Self-registration

Self-registration, where a person creates an identity for themselves when signing in, is switched off by default.

Nova as a sign-in service for other applications

With the “OIDC Identity Provider” add-on, Nova also signs identities in to other applications, acting as an OpenID Connect provider. The same rules apply as for signing in to Nova. The add-on is switched off by default.

Nova records sign-ins and failed attempts in the “Audit Log” – see Logs.

10.2 Data protection and encryption

> **PLANNED**
>
> This page describes decided behaviour that ships with one of the next releases. Until then, the current version of Nova may behave differently.

This page describes which data Nova stores and how Nova protects it.

Which data Nova stores

- **Identity master data** from the HR system and other source systems
- **Accounts** of the identities in the target systems
- **Entitlements and assignments**
- **Access requests and decisions**
- **Logs** – see Logs

Secrets of target systems

Nova stores the passwords, client secrets and tokens it uses to connect to target systems in encrypted form. Only the operator holds the key; it is not kept in the database.

- Nova checks the key at start-up.
- If the secret of a target system cannot be decrypted, Nova does not use that target system.
- The key can be rotated.

Nova does not store the passwords of accounts in target systems – see Accounts and passwords.

Encrypted connections

- **Web interface:** over TLS. Encryption is handled by a reverse proxy that the operator places in front of Nova – see Installation.
- **Target systems:** to directories, Nova transmits passwords only over LDAPS or StartTLS – see Accounts and passwords.

Least privilege

- The database user Nova runs with has no administrative rights in the database.
- “Free Reports” – self-written SQL queries – run read-only and with a restricted database role.

Retention and deletion

Nova permanently deletes identities in the trash once the retention period has expired – see Identities and their status. In doing so, Nova removes the personal data. Log entries are kept for traceability.

Demo and reset functions

Functions that load demo data or reset data cannot run against a production database.

AI

Which data Nova passes to a language model, and whether it leaves the system, depends on the AI mode – see Overview and modes.

Reporting vulnerabilities

A channel for reporting security vulnerabilities is published in a security.txt file.

CHAPTER 11

Operations

11.1 Requirements

Nova runs as a container with a PostgreSQL database and is used in the browser. Nothing needs to be installed on workstations. The installation steps are described in Installation.

Runtime environment

Component	Requirement
Platform	a server with Docker Engine and Docker Compose v2, or a cluster in SAP BTP, Kyma runtime
Database	PostgreSQL 16; the supplied templates run it as a container of its own
Application	Python 3.12 inside the image; the user interface is built when the image is built
User access	HTTPS through an upstream TLS proxy or the Kyma gateway

When the image is built, Docker pulls the base images for Node.js and Python as well as packages from the Debian, npm and PyPI repositories. The machine that builds the image needs access to these sources.

A single instance

Nova runs as exactly one instance with one Gunicorn worker (`GUNICORN_WORKERS=1`). Nova keeps the scheduler for background jobs and the state of running jobs in its process; a second instance or a second worker would run scheduled jobs twice. Nova handles more load through threads (`GUNICORN_THREADS`, default 8).

Sizing

There is no binding sizing recommendation. Points of reference:

- The Kyma template is sized for a demo installation. It reserves 0.2 CPU and 384 MiB of memory for Nova with a limit of 1 CPU and 1 GiB, 0.1 CPU and 256 MiB for PostgreSQL with a limit of 0.5 CPU and 512 MiB, plus 2 GiB of storage for the database.
- Nova opens up to 30 connections to the database.

Network

Nova opens the connections to the target systems itself. Only the routes to the systems actually connected need to be allowed:

Destination	Protocol and usual ports
SAP (ABAP)	RFC to the gateway of the application server, port 33‹system number›; when logging on through the message server, its port as well; optionally through a SAProuter
Active Directory, LDAP	LDAPS (636) or LDAP (389)
Microsoft Entra ID	HTTPS (443) to Microsoft sign-in and Microsoft Graph
Keycloak	HTTPS to the realm's sign-in and admin interfaces
SCIM applications	HTTPS to the application's SCIM endpoint
Email	SMTP, by default port 587 with STARTTLS
AI, optional	HTTPS to the cloud provider, or a connection to a self-hosted Ollama server

The same routes apply when a system serves as a source system – for example an LDAP directory with HR data, or SAP organisational data over RFC. Details on each connection are in Target systems – Overview.

SAP: the customer's RFC SDK

For SAP, Nova uses the SAP NetWeaver RFC SDK. It is SAP software and not part of the standard image: the image from `deploy/Dockerfile` is built without the SDK, and SAP functions stay switched off in it. To connect SAP, the operator provides the SDK for Linux and adds it to the image together with the Python package `pyrfc`.

Email

Nova needs an SMTP server for notifications. It is set up under “Administration” → “Notifications” → “Email (SMTP)”. If nothing is stored there, Nova uses the environment variables `NOVA_SMTP_*`; at least `NOVA_SMTP_HOST` and `NOVA_SMTP_FROM_ADDRESS` are then required.

Environment variables

Variable	Meaning
`SECRET_KEY`	Required. Key for sign-in sessions. If it is missing, Nova generates a new one on every start, and everyone has to sign in again after each restart. The Compose template does not start without it.
`NOVA_ENCRYPTION_KEY`	Required. Key Nova uses to encrypt stored credentials. Why it must be kept separately is explained in Updates and backup.
`DB_HOST`, `DB_PORT`, `DB_NAME`, `DB_USER`, `DB_PASSWORD`	Database connection

`TRUST_PROXY`, `SESSION_COOKIE_SECURE`	both set to `1` behind a TLS proxy

Optional variables include `NOVA_TIMEZONE` – the time zone Nova uses to determine the calendar day for validity dates, default Europe/Berlin –, `NOVA_SMTP_*`, and `CLAUDE_API_KEY` or `OPENAI_API_KEY` for a cloud provider's AI. `RUN_DUMMIES` starts demo target systems inside the container and stays at `0` in production.

11.2 Installation

There are two prepared ways to install Nova. Both build the image from `deploy/Dockerfile` and run Nova together with PostgreSQL 16. The templates are in the `deploy/` directory; the commands below run in the directory that contains `deploy/`. What has to be in place beforehand is described in Requirements.

Single server with Docker Compose

The template `deploy/docker-compose.yml` starts Nova and PostgreSQL as containers. The database lives in a Docker volume of its own and survives restarts and updates.

1. **Create the configuration.** Copy `deploy/.env.example` to `deploy/.env` and set at least `SECRET_KEY`, `NOVA_ENCRYPTION_KEY` and a dedicated `DB_PASSWORD`; `RUN_DUMMIES` stays at `0`. The two keys can be generated as follows, the second one with Python and the `cryptography` package:

   ```
   openssl rand -hex 32
   python -c "from cryptography.fernet import Fernet; print(Fernet.generate_key().decode())"
   ```

2. **Prepare for the TLS proxy.** In `deploy/.env`, set `APP_BIND=127.0.0.1` so that Nova can only be reached from the server itself, plus `TRUST_PROXY=1` and `SESSION_COOKIE_SECURE=1`.
3. **Build and start.**

   ```
   docker compose -f deploy/docker-compose.yml --env-file deploy/.env up -d --build app
   docker compose -f deploy/docker-compose.yml --env-file deploy/.env logs -f app
   ```

   The first run builds the image and takes a few minutes. Nova is ready as soon as the log shows Gunicorn listening on port 8000. The template also contains the `handbook` service, which serves this handbook; Nova does not need it.
4. **Set up the TLS proxy.** A web server on the host accepts HTTPS and forwards to `http://127.0.0.1:8000`. The nginx templates in `deploy/nginx/` show the settings: pass on the headers `Host`, `X-Forwarded-For`, `X-Forwarded-Proto` and `X-Forwarded-Host`, allow uploads of up to 25 MB and set timeouts of 120 seconds. The certificate can be obtained from Let's Encrypt with certbot, for example.

> ⚠ **DO NOT RUN WITHOUT TLS**
>
> Without an upstream proxy, Nova serves plain, unencrypted HTTP only. Put a TLS proxy in front of it before using it with real data.

SAP BTP, Kyma runtime

The templates in `deploy/kyma/` run Nova in a Kyma cluster: a Deployment for Nova, PostgreSQL as a StatefulSet with its own volume, and an APIRule that makes Nova reachable over HTTPS through the Kyma gateway – at `https://nova-iam.<cluster domain>`.

1. **Access and namespace.** Set up kubectl with the cluster's kubeconfig; sign-in runs through the OIDC plugin kubelogin. Create a namespace, for example `nova-iam`. If the image is in a private registry, also create an image pull secret.
2. **Provide the image.** Build the image from `deploy/Dockerfile`, push it to the registry and enter it in `deployment.yaml`. Give every new version a new tag: a node does not pull a tag it has already pulled.
3. **Create the secret.** Nova and PostgreSQL read their environment from the secret `nova-secrets`, template `secret.example.yaml`: `POSTGRES_DB`, `POSTGRES_USER`, `POSTGRES_PASSWORD`, `DB_NAME`, `DB_USER`, `DB_PASSWORD` (same as `POSTGRES_PASSWORD`), `SECRET_KEY` and `NOVA_ENCRYPTION_KEY`. Further variables such as `SESSION_COOKIE_SECURE=1` or `NOVA_SMTP_*` are added as extra keys. The real values do not belong in version-controlled files.
4. **Adapt for production.** The template is designed for a demo: in `deployment.yaml`, set `RUN_DUMMIES` to `"0"` and do not apply `dummies-expose.yaml`. Keep `replicas: 1` and the `Recreate` strategy – this way two instances never run at the same time.
5. **Apply.**

```
kubectl -n nova-iam apply -f deploy/kyma/postgres.yaml
kubectl -n nova-iam apply -f deploy/kyma/deployment.yaml
kubectl -n nova-iam apply -f deploy/kyma/apirule.yaml
kubectl -n nova-iam rollout status deploy/nova-iam
```

The first start

On start, Nova creates the database schema and brings it up to date on every subsequent start. On the first start, Nova also writes basic data, including:

- an administrator account with a preset password, which administrators should change right after the first sign-in,
- the entitlements “Admin”, “Manager”, “Security” and “SoD Reviewer”,
- the preset background jobs, see Monitoring and jobs,
- two example target systems pointing to local test services: “SCIM Dummy Target” and “LDAP Directory”.

Administrators then connect their own target systems under “Systems”, see Target systems – Overview. How updates and backups work is described in Updates and backup.

11.3 Updates and backup

Updates

An update is a new Nova image. On start, Nova brings the database schema up to date itself: it creates new tables and columns and runs data migrations. There is no separate migration step. Migrations only run forwards – the way back is the backup taken before the update.

Docker Compose: install the new release files, then run

```
docker compose -f deploy/docker-compose.yml --env-file deploy/.env up -d --build app
```

Compose rebuilds the image and replaces the Nova container; the database keeps running. Optionally, a systemd timer takes care of updates: `deploy/autodeploy.sh` checks a Git branch every five minutes and rebuilds when it changes. The directory is then used for operation only; the script resets any local changes in it.

Kyma: build the new image with a new tag and push it to the registry, enter the tag in `deployment.yaml`, apply the file and wait for `kubectl rollout status`. The `Recreate` strategy stops the old instance before the new one starts; Nova is briefly unavailable in the meantime.

Work in progress during a restart

- Jobs that were running during the restart are marked as failed by Nova on the next start.
- Provisioning runs interrupted by the restart are resumed by Nova on the next start, provided their execution plan was stored.
- Scheduled job times that fall into the downtime are not caught up; the job runs at its next regular time.

Backing up the database

Nova stores its data in the PostgreSQL database – identities, assignments, configuration, logs and job runs. Nova does not keep any data in the container itself. A backup is therefore an ordinary PostgreSQL backup of the Nova database. Example for Docker Compose; the command reads the user and database name from the database container's environment:

```
DC="docker compose -f deploy/docker-compose.yml --env-file deploy/.env"

# back up
$DC exec -T db sh -c 'pg_dump -U "$POSTGRES_USER" -Fc "$POSTGRES_DB"' > nova-$(date +%F).dump

# restore: stop Nova, load the backup, start Nova
$DC stop app
$DC exec -T db sh -c 'pg_restore -U "$POSTGRES_USER" -d "$POSTGRES_DB" --clean --if-exists' < nova-
2026-09-27.dump
$DC start app
```

In Kyma, the same commands run in the PostgreSQL pod, for example `kubectl -n nova-iam exec postgres-0 -- sh -c 'pg_dump -U "$POSTGRES_USER" -Fc "$POSTGRES_DB"' > nova.dump`. A fresh backup before every update is advisable.

Configuration export

In addition, Nova offers administrators a configuration export as a JSON file through the interface `GET /api/config/export`. It contains target systems with redacted secrets, roles, business roles with their members, SoD rules, approval workflows and settings. It does not contain identities, assignments or logs, so it does not replace the database backup. The import through `POST /api/config/import` – with `?dry_run=1` as a trial run – currently applies only the SoD rules.

Keeping the key safe

`NOVA_ENCRYPTION_KEY` is not stored in the database but in `deploy/.env` or in the Kubernetes secret. Nova uses this key to encrypt stored secrets, such as passwords and client secrets of target and source systems or the SMTP password. The database backup contains these values in encrypted form only.

- Without the key – or with a different one – the stored secrets cannot be decrypted; they then have to be entered again.
- Nova offers no key rotation for existing data. Do not replace the key of an existing installation.
- If the variable is missing, Nova does not fall back to a temporary key: operations that encrypt or decrypt a stored secret stop with an error message.

> **KEEP IT SEPARATE**
>
> Keep the key, together with `SECRET_KEY` and the database password, separate from the database backups, for example in a password vault. A backup can only be fully restored with the matching key.

11.4 Monitoring and jobs

Administrators follow Nova's operation in the “Monitoring” area. The background jobs come together there as well: tasks such as imports, deadlines and retries that Nova runs on a schedule or at the push of a button.

The “Monitoring” area

Tab	Content
“Background Jobs”	all jobs with their schedule and last run; the most recent runs with status, progress, duration and trigger; key figures such as the failed runs of the last 24 hours
“System Health”	CPU, memory and disk usage; state of the application and the database; reachability of the local AI server; result of the last check per target system
“Application Log”	warnings and errors of the application, filterable by level, period and text
“Audit Log”	logged actions in Nova
“Provisioning Log”	provisioning runs, individual provisioning events and changes in the target systems in one list, filterable by type and system
“Outbox”	outgoing notifications with their status; this is where administrators set sending to “Active”, “Paused” or “Off”

“System Health”, “Application Log”, “Audit Log” and “Outbox” are reserved for administrators. “Background Jobs”, “System Health” and “Provisioning Log” refresh every five seconds.

Managing jobs

Only administrators may create, edit, schedule, start, cancel and delete jobs; new jobs are created with “New job”. A schedule sets weekdays and a time, or a one-off date. “Run Now” runs a job immediately. “Cancel” marks a running run as cancelled; whatever it has written up to that point stays in place. A job's runs remain in the history even if the job is deleted.

Nova evaluates all schedules in the Europe/Berlin time zone. This time zone is fixed and does not depend on `NOVA_TIMEZONE`. On start, Nova writes the scheduled jobs and their next run time to the container log.

Preset jobs

After installation, four jobs are created, enabled and scheduled:

Job	Schedule	Task
“Retry Failed Provisioning”	every 15 minutes	retries failed provisioning runs of the last 24 hours at growing intervals (10, 20, 40 minutes), at most three times

“Leave-Date Scan”	daily at 02:00	runs the leaver routine for active identities whose leave date has passed
“Activate Pending Assignments”	daily at 02:30	pushes future-dated assignments on their start date to target systems without validity dates of their own
“Target System Health Poll”	every 10 minutes	checks that the target systems can be reached; when a system newly goes offline, Nova records it in the “Audit Log”

Nova recreates deleted preset jobs on the next start. To stop using one of these jobs, disable it.

In addition, internal tasks run that do not appear in the job list: Nova sends queued notifications about every 20 seconds and cleans up every day – completed outbox entries older than 90 days (default), plus expired exports.

Further job types

Administrators create further jobs as needed, among them:

- **Lifecycle and HR:** “Import HR users”, “Apply future changes”, “Purge Deleted Users” – permanently deletes identities from the trash once the retention period has passed
- **Entitlements:** “Cleanup Expired Assignments”, “Mass Role Assignment”, “Check System Roles”
- **Governance:** “SoD Violation Scan”, “Approval Escalation”, “Recertification Deadline Reminder”, “Agent Review Scan”
- **Target systems:** imports per connector, such as “Sync SAP Users”, “Import SAP Roles”, “Import LDAP Groups”, “Import Entra Groups”, “Import Keycloak Groups”, “Import SCIM groups” and “Sync SAP OM”

> **EXPIRED ASSIGNMENTS**
>
> “Cleanup Expired Assignments” is not created during installation. The job removes assignments whose validity has ended. If “Auto-provision backend systems” is switched on (“Administration” → “Provisioning” → “Mapping”), it also triggers the removal in the target systems. Anyone working with end dates should create the job and schedule it, for example daily.

Notification of failures

If a job run fails, Nova notifies the administrators by email and – if the Microsoft Teams add-on is enabled and set up – in Teams as well. Each recipient receives at most one message per job, error and hour. If a provisioning run fails completely or partly, messages also go to the administrators, and for runs resulting from an approved access request to the person who requested it as well.

Nova does not send a message when a target system goes offline; this shows up in the “Audit Log” and under “System Health”. Messages by email require email sending to be set up, see Requirements. The “Outbox” shows what was sent.

CHAPTER 12

Migration

12.1 Migrating from SAP IdM

SAP has announced the end of mainstream maintenance for SAP Identity Management at the end of 2027. For the move, Nova takes over the data from an export file using the “Migration Workbench” (Nova System Migration Workbench, NSMW). Nova needs no connection to SAP IdM for this, nor a fixed export format: an AI writes a suitable parser for each file. Administrators find the workbench under “Administration › Migration Workbench”.

What Nova takes over

In Nova	Matched by	Taken over
Identities	email address	name, department, phone, location, status, start date, salutation, title, manager; further fields as custom attributes
Entitlements	ID	name, description, colour, risk
Business roles	ID	name, description, colour, contained entitlements
Assignments	email and role	one entitlement or one business role per assignment, by ID or name

Nova displays custom attributes once they are defined under “Administration › Identity Management › Custom Fields”.

The workbench writes to Nova only. It creates no accounts in target systems and triggers no provisioning; newly created entitlements are not linked to any target system. It does not take over accounts, organisational units, approval workflows or passwords.

Procedure

1. **Upload.** Administrators upload the export file via “New migration”: CSV, Excel, JSON, XML or fixed-width text, at most 10 MB. Nova stores it with the migration.
2. **Generate the parser.** In the workbench chat, administrators describe the file or take over the suggestion “Read the file in and map all fields automatically.” The AI receives an excerpt and writes a parser for the whole file. Nova asks it for an explanation of the mapping, including a table of source column and Nova field, and shows it in the chat; “Parser script” shows the code. The instructions to the AI require multi-values that SAP IdM exports separated by “|” to be split per value.

3. **Dry run.** As soon as the AI delivers a parser, Nova applies it to the whole file and checks every record against the current data without changing identities, entitlements or assignments; “Dry-run” repeats this at any time. If the parser fails, Nova passes the error message back to the AI up to two times. The AI receives the result of the last dry run with the next message.
4. **Review the preview.** The “Dry-run preview” counts per object type what would be created, updated or skipped and what is faulty. Clicking a number filters the list. Skipped and faulty rows state their reason, such as a missing email address, an invalid date or an unknown role; for identities to be updated, the preview shows the changes field by field. Administrators request corrections in the chat or set them under “Options”.
5. **Run.** “Run” opens “Run migration” with the counts per object type; individual object types can be deselected. Nova then writes in the background in a single database transaction: if it fails, the data stays unchanged.
6. **Final report.** “Import finished” shows the counts per object type and, under “Not imported”, the faulty records with their reason.
7. **Roll back.** As long as the migration has the status “Activated”, its card offers “Roll back”. Nova first states how many records it will delete and restore. Using a journal, Nova deletes what the migration created, together with the assignments attached to these objects, and restores the stored previous values of changed records. Entitlements that the migration added to business roles that already existed stay there. This takes effect in Nova only. Afterwards the migration can be run again.

Options

- **“Mode”:** the default is “Create only” – Nova skips whatever already exists in Nova. “Create & update” updates existing records; for identities, empty fields in the file do not overwrite existing values.
- **“On error”:** the default is “Skip row” – Nova writes the records without errors. With “Abort migration”, Nova writes nothing as soon as one record is faulty.
- Further switches remove leading and trailing whitespace, treat empty cells as missing, convert role IDs to upper case and skip duplicates within the file; all are on by default. A date-format hint is passed to the AI.

AI and data

Only generating the parser requires AI (“Local” or “Cloud”); dry run, run and roll-back work without it. The AI receives an excerpt of the file, the chat messages and the result of the last dry run – in “Cloud” mode, this includes personal data from the first rows of the export. Nova runs the parser in a separate, restricted process; only Nova itself writes to the database, after its own checks. See Overview and modes.

Logging and retention

- **Audit log:** upload, changes to name and options, start and completion of the run with the counts per object type and the records written (up to 5,000 listed individually), roll-back with the records undone, deletion.

- **Provisioning log:** entries per affected identity, both for the run and for the roll-back.
- Every dry run and every run remains stored as a run of the migration, as does the chat with the AI.
- An activated migration can only be deleted once it has been rolled back. Until deletion, the uploaded file also remains stored in Nova.

The steps after the upload can also be started from the admin chat; running, rolling back and deleting require a confirmation there, see AI guardrails.

CHAPTER 13

News

13.1 Release notes

Nova is developed continuously. This page summarises the visible changes per month, newest first. Features from add-ons are marked as such; add-ons are switched off after installation and are enabled under “Administration” → “Plugins”. Limits that matter for a rollout are listed under Known limitations.

September 2026

- **Handbook:** The first edition of this handbook is published, in German and English.
- **Admin chat:** The Nova AI chat can carry out almost all functions of the user interface. Actions with far-reaching effects, such as deleting, are first shown as a preview and only carried out after confirmation in the next message. Answers can be downloaded as PDF, Excel, CSV, Markdown or text.
- **Outbox:** A new tab under “Monitoring” shows outgoing notifications with their status; sending can be set to “Active”, “Paused” or “Off”.
- **Mapping:** Attributes of source and target systems can also be mapped using expressions – concatenation, if-then rules, text functions – with an AI suggestion and a preview against a sample identity.
- **HR import:** Nova classifies every person the same way in all views: joiner, mover, update, leaver or unchanged.
- **SAP:** One SAP role can be granted to or revoked from many accounts at once. Passwords that Nova generates for new SAP accounts follow SAP's fixed rules on length and leading characters. Fixed: after a reconciliation, SAP composite roles appear as a direct assignment instead of a business role.
- **Generic LDAP:** Locking and unlocking use the directory's password policy; the password and group memberships are kept.
- **Change journal:** In addition to LDAP and Entra ID, the change journal now also records writes to SAP, SCIM and Keycloak.
- **Role mining and Migration Workbench:** All proposals of a layer can be accepted together, and applied mining runs can be deleted. On request, the Migration Workbench transfers only selected object types, and its final report lists the records that were not transferred.

August 2026

- **New user interface:** The new user interface is now the default; the previous one is still available under `/oldui` for the time being.

- **Dashboard:** Everyone can adjust the size, order and visibility of the tiles. A new tile shows the share of the entitlement catalogue that is contained in business roles.
- **Role mining:** Nova proposes business roles derived from direct assignments, arranged in layers from the base role to special functions. Proposals are reviewed, calculated in a trial run and only then applied. See Role mining.
- **Attribute history:** In an identity's "Changelog", Nova shows for each attribute which value applied from when to when, and where it came from.
- **User interface:** Connection errors of target systems are visible directly on the system; an identity's assignment overview is grouped by system. Several themes are available, and Nova can be installed as a web app.
- **AI settings:** The timeout and the maximum length of AI answers can be configured.
- **Add-on OIDC Identity Provider:** Other applications can use sign-in to Nova via OpenID Connect ("Login with Nova").

July 2026

- **Keycloak:** New connector; groups serve as entitlements. See Keycloak.
- **Recertification:** Campaigns with a worklist for reviewers and a completion step in which Nova removes denied entitlements on request; the result is available as CSV. See Recertification.
- **Segregation of duties:** SoD rules between roles and business roles. When an access request is made, Nova warns about conflicts; depending on the setting, Nova blocks the request instead or requires an additional review step. As an add-on, Nova also checks against the rule set of SAP GRC Access Control. See Segregation of duties (SoD).
- **Approvals:** Deputy rule for absent approvers; the job "Approval Escalation" sends reminders about overdue steps and escalates to the administrators. See Approval workflows.
- **Lifecycle:** A daily job runs the leaver routine once the leave date has passed. For rehires, administrators decide whether earlier entitlements are kept. See Joiners, movers, leavers.
- **AI agents:** AI agents are identities of their own, with a responsible person, a purpose and a review date. See AI agents as identities.
- **Operations:** Nova regularly checks that the target systems can be reached, retries failed provisioning runs automatically and notifies the administrators when a background job fails. See Monitoring and jobs.
- **Traceability:** Writes to LDAP and Entra ID appear in the "Provisioning Log"; administrators can undo them one by one. The filtered audit log can be exported as CSV or JSON.
- **Validity:** For target systems without validity dates of their own, Nova holds back future-dated assignments and pushes them on their start date.
- **Add-on Native MFA (TOTP):** a second factor with single-use backup codes for signing in with a password.

June 2026

- **Notifications:** Nova sends messages by email, for example about access requests and failed provisioning. With the Microsoft Teams add-on, access requests can also be approved in Teams.
- **Joiner, mover and leaver routines:** Custom steps are created in an editor – with conditions, calculated values and a preview of what each step does in Nova and in the target systems.
- **Signing in to Nova:** Password history, optional password expiry and a lockout after repeated failed attempts; Nova logs sign-in events such as failed attempts and lockouts.
- **Governance analysis:** Nova assesses a requirements questionnaire – partly by checking the current data, partly with AI – and presents the result as a scorecard.
- **Migration Workbench:** Exports from legacy systems can be read in any format. The AI writes a parser for them, a trial run shows in advance what would be transferred, and transfers can be rolled back.
- **Fixed:** The HR import no longer creates duplicate identities for one person.

13.2 Known limitations

This page lists limits that matter when planning a rollout. Status: September 2026.

Operation

- **A single instance.** Nova runs as exactly one instance, because the scheduler and running jobs live in its process. Several instances side by side – for load balancing or active-active high availability – are not supported; Nova scales through threads. See Requirements.
- **Time zone of schedules.** Jobs run according to the Europe/Berlin time zone; no other time zone can be set for them.
- **Transferring configuration.** The configuration export contains, among other things, target systems, roles, business roles and approval workflows; currently only the SoD rules can be imported. It therefore covers the move from a test to a production environment only in part. See Updates and backup.
- **Validity periods.** Only SAP stores the start and end of an assignment itself. In the other target systems, Nova applies the dates through background jobs; they take effect there only with the next run. The start is pushed by the daily job “Activate Pending Assignments”. The end is handled by “Cleanup Expired Assignments”: this job is not created during installation and only passes the removal on when “Auto-provision backend systems” is switched on. See Monitoring and jobs.

SAP

Nova connects over RFC and needs the SAP NetWeaver RFC SDK for this. The SDK is SAP software, is not part of the standard image and has to be provided by the customer. Without the SDK, neither SAP target systems nor SAP organisational data as a source system nor the SAP add-ons can be used. See SAP.

Active Directory and LDAP

- **Locking.** Except in Active Directory, Nova locks accounts through the password policy's lock attribute (`pwdAccountLockedTime`). This only takes effect if the directory enforces a password policy and the account Nova works with may write this attribute.
- **Nested groups.** Nova only writes groups within groups for DN-based groups; the OpenLDAP profile with `posixGroup` does not support nesting. Generic directories additionally need the `memberOf` attribute in the schema, with write permission.

See Active Directory and LDAP.

Microsoft Entra ID

In dynamic groups, membership results from rules; in groups synchronised from an on-premises Active Directory, it results from the synchronisation. Microsoft Graph does not allow changes to these memberships, so Nova cannot provision such groups. Security groups with assigned members that are managed in Entra ID itself are suitable as entitlements. See Microsoft Entra ID.

SCIM

SCIM applications implement the standard differently. Nova expects the endpoints under `‹base URL›/scim/v2`, looks up accounts and groups with filters on `userName` and `displayName`, changes accounts and memberships with `PATCH` and reads groups page by page. For authentication, Nova supports basic authentication and bearer tokens. Whether an application fits has to be checked per provider. See SCIM.

Last sign-in

Nova can only read when an account was last used if the target system provides it:

Target system	Source
SAP	user master record (`USR02`)
Active Directory	`lastLogonTimestamp` or `lastLogon`
OpenLDAP	`authTimestamp`, if the directory maintains it
generic LDAP	no default source
Entra ID	sign-in activity from Microsoft Graph; requires the `AuditLog.Read.All` permission
Keycloak	the realm's stored login events; if the realm stores none, only active sessions
SCIM	no standard attribute; the attribute entered under “Last-logon attribute (SCIM)”, or `lastLogin` or `lastLogon`

Status of acceptance testing

The connectors are tested against a fixed acceptance catalogue with real target systems. Status September 2026:

- **SAP** (ABAP) and **generic LDAP** (tested with ApacheDS) have passed the catalogue in two complete rounds; individual sub-requirements were explicitly excluded.
- **Microsoft Entra ID:** acceptance has not been completed yet.
- **Active Directory, Keycloak and SCIM** have not yet been run through the catalogue.