EGKITS SaaS V11

by IDSYSEG · Complete module guide · Version 1.1

Updated Nov 2026

Digital Virtual Office Complete Module Guide

An organization of AI employees that works under human authority

20 AI employees | Human approvals | Tenant isolation

EGKits SaaS V11

A complete reference for business owners, administrators, implementation teams, auditors and engineers. It explains how the Digital Virtual Office (DVO) lets a company run departments of AI employees who plan, delegate, use EGKits tools, search approved knowledge and report results, while every sensitive action waits for a person, every record stays inside its own tenant, and every cost and decision is measured and audited.

Field	Value
Document version	1.1
Publication date	21 November 2026
Applies to	EGKits SaaS V11, app code dvo
Audience	Owners, administrators, implementers, auditors and engineers
Distribution	Public. No tenant data, prompts, credentials or production logs.
English file	`EGKits_DVO_Digital_Virtual_Office_Guide_EN.pdf`
Arabic file	`EGKits_DVO_Digital_Virtual_Office_Guide_AR.pdf`

Enforced control

AI employees never approve their own requests, never change their own instructions, permissions, model or security settings, and never see another tenant's data. These rules are enforced by the server and covered by automated security tests, not by prompts.

Organization

Executive Office and seven departments: 20 ready-made AI roles with job descriptions, skills and reporting lines.

Human authority

Autonomy levels L1 to L3. High-risk work always waits in the Approval Center, bound to the exact operation.

Measured work

Budgets that block before a model call, a Cost Center, seven KPIs, a full audit trail and telemetry.

Contents

Select a heading to jump to that section.

1. How to read this guide 3

2. Executive summary 3

- What makes it safe for a business 3
- The 3D workplace 4
- What it is not 4

3. Core concepts 4

4. Architecture 5

- Design decisions 5
- What is reused and what is added 5
- Solution projects 6

5. Tenant isolation and security 6

- Eight isolation rules 6
- Secrets 7
- Platform administrators and support access 7
- Prompt injection and untrusted content 7
- Voice, meetings and privacy 7

6. Enabling and setting up 8

- Permissions 8
- Office settings 9

7. The organization 10

- Departments and their boundaries 10
- Reporting rules 10

8. AI employees 11

- Configuration versions 11
- The twenty standard roles 12

9. Autonomy and effective authority 12

10. Tools 13

- Policies and grants 13
- Built-in office tools 13
- EGKits module tools 14

11. Work requests 14

- Steps, checkpoints and recovery 15
- Limits per request 15
- Delegation and questions 15

12. The Approval Center 15

- Who may decide 16

13. Executive plans and the Command Center 16

- The Command Center 17

14. Tasks and projects 17

- Board, list, calendar and timeline 18
- Projects 18

15. Communications 18

16. The Knowledge Center 18

17. Memory 19

18. Avatars and voice 19

- Avatars 20
- Voice conversation 20

19. Meetings 21

- Scheduling 21
- The live room 21
- Meetings of AI employees only 22

20. Transcripts, minutes, decisions and action items 22

21. The Virtual Office and Meeting Analytics 23

- Meeting Analytics 24

22. Budgets, usage and the Cost Center 24

- The Cost Center 24

23. The KPI Center 24

24. Runtime operations 25

25. Audit, telemetry and retention 26

- Telemetry 26
- Retention 26

26. Notifications, reports and the home dashboard 27

- Reporting Center 27
- Home dashboard 27

27. The demo scenario 27

28. Configuration reference 28

29. Page reference 29

30. Troubleshooting and runbooks 30

31. Implementation checklist 31

32. Glossary 31

1. How to read this guide

The guide follows the order in which a company adopts the module: understand it, enable it, build the organization, equip the AI employees, give them work, decide on their requests, and then measure and operate the office. Each chapter names the page, the permission and the statuses exactly as the product shows them.

Reader	Start with	Then read
Business owner or CEO	2. Executive summary	7, 11 to 13, 19, 21, 22 and 23
Administrator	6. Enabling and setting up	7 to 10, 16, 18, 19, 24 and 28
Manager who assigns work	11. Work requests	12 to 21
Auditor or security officer	5. Tenant isolation and security	9, 12, 18, 20, 25 and 30
Engineer or operator	4. Architecture	24, 25, 28 and 30

Four kinds of note appear in the text. They are not interchangeable.

Note	Meaning
Enforced control	A rule the server enforces whatever the user, the prompt or the model does. Each one has automated tests.
Configuration	A value the tenant or the platform can change. The default is stated.
Known limit	A documented boundary of the current release.
Development only	Available only on a development host, never in production.

2. Executive summary

The Digital Virtual Office is an EGKits application that gives each company its own organization of AI employees. They sit in departments, report to managers, hold job descriptions and skills, and work through the same EGKits services a person uses: CRM, Sales, HR, General Ledger, CMS, Newsletter, the document archive and the office's own tasks, knowledge and channels. A person gives an instruction; the AI employee plans, may delegate parts to its team, asks when information is missing, and submits a result that the person can review, rate and act on.

What makes it safe for a business

- Human authority. An AI employee can do only what its grants, its autonomy level and the requesting person's live permissions all allow. High-risk and irreversible operations always wait for a person's approval, bound to the exact operation by a SHA-256 fingerprint.
- Tenant isolation. Every DVO record carries its tenant and is filtered by the data model. The tenant always comes from the signed-in workspace, never from the request body, and background jobs re-resolve it before touching data.
- Real work only. Results come from a real model call and real tool calls. In production a route to the stub provider fails the work with a clear error instead of reporting a simulated success.
- Measured cost. Budgets for the office, a department or an employee are checked before each model call. Usage is recorded per call, compared with invoices in the Cost Center and summarized as seven KPIs.
- Full traceability. Each operation writes an audit event with actor, employee, target, authorization result and correlation id, and one trace follows a person's instruction through every delegation to the final result.

The 3D workplace

Each AI employee has a 3D avatar and a voice in English and Arabic. People talk to an employee by voice, meet several of them in a 3D meeting room, and see the whole office on a live 3D floor. Every avatar is labelled as an AI employee, the microphone opens only when a person presses it, and every spoken answer is real work. After a meeting an AI employee drafts the minutes; people confirm decisions and approve action items, which become tracked tasks.

What it is not

- It is not a chat window. The AI Assistant answers one person; the DVO runs durable work with owners, statuses, approvals, costs and history.
- It does not replace managers. People set direction, approve sensitive actions, verify memory and rate results. AI memory is never treated as an authoritative business record.
- It does not work in the background by itself. Work starts only when a person asks, starts a task, starts plan steps or asks a question in a channel, and the page always shows the real status of that work.

A typical day in the office

1. The CEO asks the Chief of Staff for a plan: "Prepare the Release 1.0 launch."
2. The plan proposes four steps, each assigned to an AI employee. The CEO picks three and starts them; each becomes a task on the board.
3. The CTO reviews release readiness, delegates the regression check to QA and adds a rollback-runbook subtask for DevOps.
4. The Sales specialist prepares a CRM follow-up activity. Writing to CRM needs approval, so it waits in the Approval Center for a second person.
5. Results arrive with sources and costs. Managers rate them, and the KPI Center and Cost Center update.

3. Core concepts

These terms are used throughout the module and this guide.

Concept	What it is	Where you see it
AI employee	A named role with a department, a manager, a job description, skills and a versioned configuration.	AI employees, Organization
Configuration version	An immutable snapshot of instructions, model, autonomy, data scopes, memory scope and tool categories.	Employee profile
Autonomy level	L1 read only, L2 low-risk changes, L3 may request high-risk work through approval.	Employee profile
Work request	One execution of an instruction by one AI employee, with steps, usage, result and transcript.	Work requests
Delegation	A child work request given to a member of the employee's team, within depth, size and cost limits.	Work request detail
Approval	A person's decision on one exact operation, identified by its fingerprint. Used at most once.	Approval Center
Executive plan	Steps proposed by a planning AI employee. Only the requester can start the chosen steps.	Plans, Command Center

Concept	What it is	Where you see it
Task and project	Board work with an owner, status, assignee, dependencies and deliverables, grouped in projects.	Tasks, Projects
Channel	General, department, project, task or direct conversation with recorded events.	Communications
Knowledge document	A versioned file or text that AI employees may search only when an access rule allows it.	Knowledge Center
Memory entry	A note kept for an employee, department, project or the office. Context, never a business record.	Memory
Avatar and voice	The 3D character and the English and Arabic voices of an AI employee, chosen from an approved catalog.	Employee profile, voice conversation
Meeting	A scheduled or ad hoc session of people and AI employees with an agenda, limits and a transcript.	Meetings, meeting room
Meeting turn	One AI employee's answer in a meeting, run as a real work request.	Meeting room, Work requests
Minutes and action items	Drafted by an AI employee after a meeting; a person approves the minutes and each action item before it becomes a task.	Meeting record
Budget	A monthly token or cost limit for the office, a department or an employee.	Budgets and usage
Runtime	The tenant's logical AI runtime: a queue, limits, health and lifecycle.	Runtime
Work item	A durable queue row that runs AI work or indexing, with leases, retries and dead-lettering.	Runtime

4. Architecture

The DVO is a native EGKits module, not a separate product. It reuses the platform's tenancy, permissions, AI provider stack, background jobs, notifications, audit trail and reporting, and adds only what an organization of AI employees needs.

Request path



Design decisions

Topic	Decision
Runtime	Logical isolation inside the EGKits server behind ITenantAiRuntimeHost. A worker process or per-tenant container host can replace it without changing business code.
Orchestration	A native .NET orchestrator composed from the model router, chat providers, the [AiTool] tool registry, usage quotas and the provider circuit breaker. No Python or Node.js sidecar.
Persistence	A dedicated DvoDbContext in schema Dvo. Every entity is auditable and multi-tenant.
API	Endpoints under /api/v1/Dvo, grouped by area, each behind the DVO entitlement filter and a permission.
User interface	A lazily loaded Razor library with pages under /digital-virtual-office, in 12 cultures with real English and Arabic text and full right-to-left support.
External agent SDKs	Evaluated and not adopted: they need a second runtime and secret path, and their tools would bypass the mediator's permission, validation, audit and tenant checks. OpenAI models remain available through the provider layer.

What is reused and what is added

Platform capability	Existing EGKits component	DVO addition
AI providers and keys	Tenant provider configuration with encrypted keys	None. DVO never reads keys.
Models	Model configuration and router tiers	Model, tier, temperature and output limit per configuration version
Tools	[AiTool] commands with risk tiers	Tenant tool policies and per-employee grants
Audit	Automatic entity change trail	Operation-level DVO audit events
Usage	AI usage log	Usage by employee, department, task and billing month
Workflow	Core workflow engine	An approval workflow with its own store
Jobs	Hangfire with isolated queues	Queue dvo with dispatcher, reaper, health and reconcile jobs
Notifications	Notification center and live hub	Topics for approvals, budgets, failures, questions and meetings; live topics for work and meetings
Speech	AI provider secrets	Azure AI Speech (provider azure-speech) for recognition tokens and server-side speech
Calls	Platform call service	Voice calls among the people in a meeting room
Entitlement	App catalog and subscription	The dvo app plus a DVO entitlement filter on every route

Solution projects

Project	Purpose
`Dvo.Domain`	Entities and domain rules: hierarchy, lifecycles, the task state machine and approval hashing.
`Dvo.Application`	Abstractions, contracts, commands, validators and metadata.
`Dvo.Infrastructure`	Database context, repositories, handlers, runtime host, orchestrator, jobs and endpoints.
`Dvo.Tests`	Unit, integration and security tests with scripted and stub providers only.
EGKitsDVO	Pages, the org chart, the 3D avatar, meeting room and office floor components (Three.js), voice and dialogs.

5. Tenant isolation and security

Each company's AI employees, sessions, conversations, knowledge, tasks, memory, history, tools and configuration are private to that company. Isolation is enforced in data, scheduling and authorization at every layer.

Eight isolation rules

How a record stays inside its tenant

1. Tenant identity comes from the signed-in workspace. No DVO request accepts a tenant id from the client.
2. Every DVO table has a tenant column and a tenant query filter, including bulk updates and deletes.
3. Saving stamps the tenant on insert, and a row of another tenant is refused and logged as critical.
4. Background jobs never trust the tenant stored on a job: each job reloads the tenant, opens a fresh scope and reads its work through the tenant filter.
5. Each tenant has one runtime. Queue partitions, leases, concurrency and fair-share quanta are computed per tenant.

6 Provider credentials are resolved by the tenant-aware provider stack on the server. DVO never reads, stores or returns them.

7 Platform permissions cover only enable, suspend, provision and quota overrides. Platform administrators get no tenant content.

8 AI employees cannot change themselves: configuration, prompt, permission, model and grant changes refuse any caller running as an AI employee.

Enforced control
Tested by TenantIsolationTests, CredentialLeakTests, SecurityGateTests and isolation cases in every feature suite, including tampered job arguments and foreign rows.

Secrets

- Provider keys are stored encrypted per tenant by the AI Assistant module. DVO calls the model router, which resolves credentials on the server.
- Keys never appear in DVO configuration, DTOs, the generated client, Blazor, browser storage, prompts, work item payloads, audit details or DVO logs. A test checks every contract for credential-shaped properties.
- Error text that reaches tenant-visible records is redacted before it is stored: key assignments, bearer tokens, JWTs and common provider key formats are masked.
- To rotate a key, replace it in the tenant's AI provider settings. Running work picks up the new route on its next model call.

Platform administrators and support access

The platform workspace enables the app, provisions and suspends runtimes and overrides quotas on the AI Runtime Operations page. It sees only its own DVO data. A cross-tenant support session is refused on content endpoints unless the tenant turned on support access in its DVO settings, optionally with an expiry, and every such request writes an audit event.

Prompt injection and untrusted content

Knowledge passages, tool results, delegate results and dependency results are wrapped in delimited untrusted-data blocks and never override an employee's instructions. Even when planted text convinces the model, the server still refuses ungranted tools, refuses approvals from an AI context and keeps the employee's prompt, grants and configuration unchanged.

Enforced control
PromptInjectionTests run five golden scenarios through tool output and knowledge. No ungranted or unknown tool runs, nothing is approved, and planted text stays inside one untrusted block.

Voice, meetings and privacy

- Avatars are labelled as AI employees on every 3D view and are never presented as people.
- The microphone and camera start only after the person acts and the browser asks. Audio is never stored; only recognised text is written, and only when the transcript is kept.
- Joining a meeting needs consent, recorded with its time. Recording needs the workspace policy, a kept transcript and everyone's consent, and stores only metadata.
- Meeting speech is business context, not authorization. It reaches the model as untrusted data and cannot grant tools or approve anything.

- Live meeting updates use the tenant group of the platform notification hub; each connection joins only the group of the tenant in its own sign-in token.

Enforced control
Covered by MeetingTests, SpeechTests, AvatarTests and the office tests, including meetings that never cross tenants or reach people who were not invited, appearances that never cross tenants, consent before joining and transcript removal.

6. Enabling and setting up

The dvo app is off for every tenant until the platform enables it in the tenant's subscription. Until then, the backend entitlement filter refuses DVO requests from that tenant and background jobs skip it.

From subscription to first result

1. Platform: enable the dvo app in the tenant's subscription. The reconcile job provisions the tenant's runtime within its schedule (every 15 minutes by default).
2. Tenant administrator: assign the DVO permissions to the right roles. They are seeded into the tenant Admin role.
3. Configure AI providers and models in the AI Assistant settings. DVO uses the same model router and has no provider settings of its own.
4. Open Settings: choose the CEO user, the default autonomy, support access, retention and your notification choices.
5. Open Organization: preview and apply the EGKits standard template, or build departments yourself.
6. Review each AI employee, adjust the configuration and activate it. Template hires start as drafts.
7. Grant tools on each profile, set tool policies on Tools and create budgets on Budgets and usage.
8. For voice, save an Azure AI Speech key in the AI provider settings. For meetings, set the meeting rules in Settings and each employee's Meetings tab.
9. Start the first work request, or ask the Chief of Staff for a plan from the Command Center.

Permissions

Each page and action is protected by a permission resource. A person sees a menu item only when they hold its permission, and every server operation checks it again.

Permission resource	Protects
`Dvo`	Overview of the office.
`DvoOrganization`	Organization, departments, reporting lines and templates.
`DvoEmployees`	AI employees, profiles, configuration versions, avatars and voices, and the Virtual Office floor.
`DvoExecutions`	Work requests, the Command Center and executive plans.
`DvoApprovals`	The Approval Center. Approve and Reject are separate actions.

Permission resource	Protects
DvoTasks	Tasks, the board and projects.
DvoCommunications	Channels, messages and voice conversations.
DvoMeetings	Meetings, the meeting room and Meeting Analytics. View, Create, Update, Approve (needed to invite employees restricted to approvers) and Export are separate actions.
DvoKnowledge	The Knowledge Center and memory.
DvoTools	Tool catalog, policies and grants.
DvoBudgets	Budgets and usage, budget alerts and the Cost Center.
DvoReports	The KPI Center and DVO reports in the Reporting Center.
DvoRuntime	Runtime status, queue, suspend, resume and requeue.
DvoSettings	Office settings and, in development, the demo scenario.
DvoAudit	The audit log and full transcripts of other people's work.
DvoPlatform	Root workspace only: AI Runtime Operations for all tenants.

Office settings

Setting	Effect
CEO user	The person department heads report to when their manager is the CEO.
Default autonomy	The level new AI employees start with. Each configuration can change it.
Support access	Lets a platform support session read DVO content, with an optional expiry. Off by default; every use is audited.
Execution and audit retention	Days to keep finished execution history and audit events, from 7 to 3650.
Voice	Whether Azure AI Speech is set up, its region, the languages people can speak and the speech minutes used this month. Test speech checks the key and region.
Meetings	Keep transcripts, allow recording with everyone's consent and allow meetings with AI employees only (all off by default); create tasks from approved low-risk action items without a second step; AI answers per turn (3, at most 5); idle end (30 minutes, 5 to 240); and the turn and token limits of AI-only meetings (12 and 200,000).
My notifications	Each person turns each notification kind on or off for themselves.

Configuration

Subscription limits come from Dvo:PlanDefaults:{Plan} per subscription plan, then an optional per-tenant override set only by the platform, then Dvo:Fallback. No limit is hardcoded.

Limit	What it controls
Employees and departments	Seats counted on every hire, restore and template apply. A refused change writes nothing.
Concurrent executions	How many work items the tenant may run at once.
Monthly tokens and cost	Subscription quotas checked before model calls, on top of the tenant's own budgets.
Allowed providers and models	A configuration that names a provider or model outside the subscription is refused.
Knowledge storage	Megabytes of knowledge files the tenant may store.
Monthly speech minutes	Spoken answers per month. When used up, speech stops and conversations continue in text.
Concurrent meetings and participants	How many meetings may be live at once and how many may take part in one. New meetings are refused; live ones continue.
Retention days	The longest history the plan allows.

Enforced control

Reaching a quota never deletes data and never stops running work. New work is refused or blocked before its next model call, and existing work finishes or resumes when room is available.

7. The organization

The Organization page shows the org chart and the departments panel. Applying the EGKits standard organization template (egkits-standard-v1) creates the Executive Office and seven departments with twenty AI employees, in English and Arabic.



Departments and their boundaries

Department	Focus	Boundary written into every member's instructions
Executive Office	Turns the CEO's direction into delegated work.	Cannot override departmental authority or approve restricted actions.
Technology	Architecture, code review, testing, deployment planning.	Production deployment, infrastructure changes and secrets need explicit human authorization.
Finance	Analysis, billing review, reconciliation support, budgeting.	Never authorizes payments, transfers or journal postings.
Sales	CRM analysis, leads, opportunities, proposals.	External communication and commitments need approval.
Marketing	Plans, content, SEO, campaigns, newsletters.	Publishing and external campaigns need permission and approval.
Operations	Support, coordination, SLA monitoring.	Customer communication and operational changes use authorized workflows only.
Product	Strategy, requirements, roadmap, research.	No roadmap dates or promises to customers without approval.
HR	HR administration, internal documentation.	HR data is confidential and used only for the task that needs it.

Reporting rules

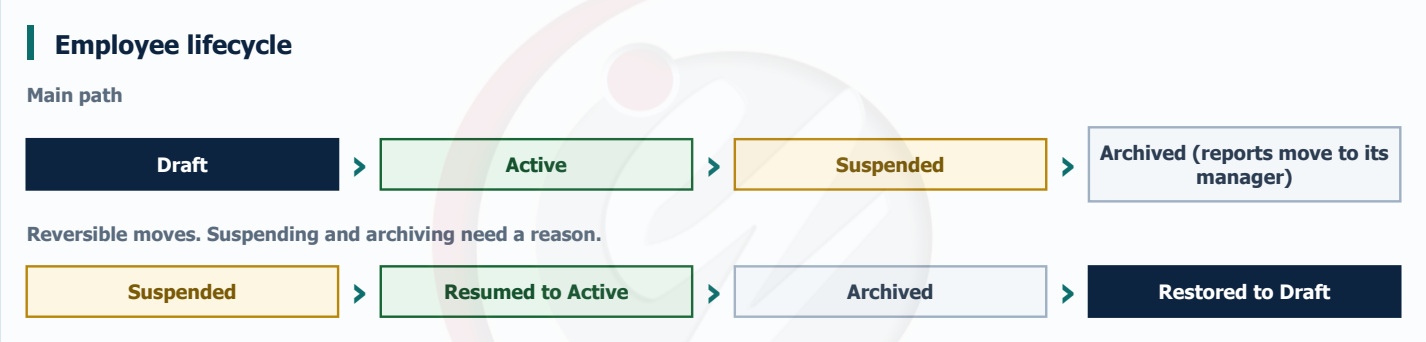
- Every AI employee has exactly one solid-line manager: another AI employee or a person. Without a named person, it is the CEO from Settings.

- Reporting cycles are refused. Up to five dotted lines add coordination without authority.
- A department head must be a member of that department.
- Archiving an employee moves its direct reports to its own manager and ends its dotted lines.
- Dragging a card onto a new manager in the chart is a real move: the server re-checks every rule. The chart supports arrow keys and Enter, and mirrors in right-to-left.
- Applying a template adds only what is missing, keeps renamed and archived employees, checks the limits for the whole set first and saves once.

Every organization change is written to the organization change log with version numbers and field names only, never prompt text. Configuration versions and change log rows are append-only.

8. AI employees

The AI employees directory lists every employee with code, department, manager, status and autonomy. The profile holds the job description, responsibilities, restrictions and skills, the configuration with its versions, the granted tools, performance, recent work and memory, the avatar and voice (chapter 18) and the meeting settings (chapter 19).



Configuration versions

Field	Meaning	Template default
Instructions	Role, responsibilities and restrictions, including the three common rules.	Generated per role and language
Model tier, provider, model	Routing through the model router. A provider or model outside the subscription is refused.	standard for heads, economy for others
Temperature and output limit	Sampling and the maximum output tokens per call.	0.3 and 4,000
Autonomy	L1, L2 or L3. See chapter 9.	L2 for Chief of Staff, QA, Support and Implementation; L1 for others
Data scopes	The EGKits apps the employee may read or act in. Must be known app codes.	Per role, for example gl, ar, ap for Finance
Memory scope and retention	None, Task, Employee or Department, and days to keep entries.	Employee, 90 days
Tool categories and names	What the employee may request. A person must still grant each tool.	Per role

- Saving the same content creates no version. A new version needs the expected current version number, so two people cannot overwrite each other.
- Versions are identified by a content hash and can be compared side by side on the profile.
- Only people edit configuration. Any request running as an AI employee, a system actor or a support session is refused.

The twenty standard roles

Code	Role	Department	Autonomy	Data scopes
DVO-01	Chief of Staff	Executive Office	L2	dvo
DVO-02	Chief Technology Officer	Technology	L1	dvo, documentarchive
DVO-03	Backend & Architecture Engineer	Technology	L1	dvo, documentarchive
DVO-04	Frontend & UX Engineer	Technology	L1	dvo
DVO-05	QA & Testing Engineer	Technology	L2	dvo
DVO-06	DevOps & Infrastructure Engineer	Technology	L1	dvo
DVO-07	Chief Financial Officer	Finance	L1	gl, ar, ap, fixedassets, einvoices, ereceipts
DVO-08	Accounting & Billing Specialist	Finance	L1	gl, ar, ap, einvoices, ereceipts
DVO-09	Chief Commercial Officer	Sales	L1	crm, sales, tradingpartners
DVO-10	Business Development Researcher	Sales	L1	crm, sales
DVO-11	Sales & Proposal Specialist	Sales	L1	crm, sales
DVO-12	Chief Marketing Officer	Marketing	L1	cms, newsletter, crm
DVO-13	SEO & Content Specialist	Marketing	L1	cms, newsletter
DVO-14	Chief Operating Officer	Operations	L1	helpdesk, services, crm
DVO-15	Customer Support Specialist	Operations	L2	helpdesk, crm
DVO-16	Implementation & Project Coordinator	Operations	L2	services, helpdesk
DVO-17	Chief Product Officer	Product	L1	dvo, crm, helpdesk
DVO-18	Business Analyst & Product Researcher	Product	L1	dvo, crm, helpdesk
DVO-19	Chief Human Resources Officer	HR	L1	hr
DVO-20	HR & Administrative Specialist	HR	L1	hr

Enforced control

Three rules are added to every template employee's instructions: never approve your own requests, never change your own instructions, permissions, model access or security settings, and never access data outside your scope or another company's data. The server enforces the same rules independently of the prompt.

9. Autonomy and effective authority

What an AI employee may do in a given moment is the intersection of four things: the tools granted to it, the requesting person's live permissions, the tenant's app entitlement and data scopes, and its autonomy level.

Level	Read tools	Low-risk changes	High-risk and irreversible
L1 Read only	Yes	No	No
L2 Low-risk changes	Yes	Yes, when explicitly granted	No
L3 Request high-risk work	Yes	Yes, when granted	Only through a person's approval

Writes in a restricted area, such as money, external communication and deployment, and any tool that needs a Delete permission, require L3 and approval whatever the tool's own risk tier.

The authority check before every tool call

1. The tenant's tool policy is enabled.
2. The employee's configuration requests the tool or its category.
3. A person granted the tool to this employee, and the grant is active.
4. The tool declares the permissions it needs.
5. The tool's app is enabled for the tenant and inside the employee's data scopes.
6. The requesting person holds those permissions now, checked live.
7. The autonomy level allows the risk tier; otherwise the call needs approval or is refused.

Decision shown on the step	Meaning
Allowed	Every check passed and the tool ran.
Needs approval	The work waits in the Approval Center.
Not granted, Not configured, Disabled by policy	A person must grant, configure or enable the tool.
Permission missing	The requesting person lacks a permission the tool needs.
App not enabled, Outside data scope	The tool's app is off for the tenant or outside the employee's scopes.
Autonomy too low	Raise the level in a new configuration version, or keep the tool out of reach.
Unknown tool, No declared permissions	The model asked for something that is not a usable tool.

10. Tools

AI employees act only through tools. A tool is an existing EGKits command or query marked as an AI tool, so the mediator runs the same permission checks, validation, audit and tenant scoping as for a person. Agents never read database tables directly.

Policies and grants

- Tool policy (Tools page): enable or disable a tool for the whole office, set its timeout and audit level (Minimal, Standard or Detailed) and raise its risk tier. A policy can raise a risk tier but never lower it.
- Tool grant (employee profile): gives one tool to one employee. A person can grant only tools they could use themselves.
- Collaboration tools are governed by the org chart and have no policy; the Tools page shows them with a dash instead of a grant count.

Built-in office tools

Tool	What it does
`dvo_delegate`	Gives part of the work to a member of the employee's team, found from the org graph.
`dvo_request_info`	Asks the requesting person a question; the work waits for the reply.

Tool	What it does
dvo_escalate	Raises a risk or a decision to a person with the facts and a recommendation.
dvo_submit_result	Submits the final result of the work.
dvo_create_task	Creates a task owned by the human requester: a draft at L1, waiting for that person's approval at L2. It never starts by itself.
dvo_search_knowledge	Searches the Knowledge Center within the employee's access rules.
dvo_remember	Writes an unverified memory entry within the configured memory scope.
dvo_recall	Recalls active memory entries; also needs the requester's knowledge permission.

EGKits module tools

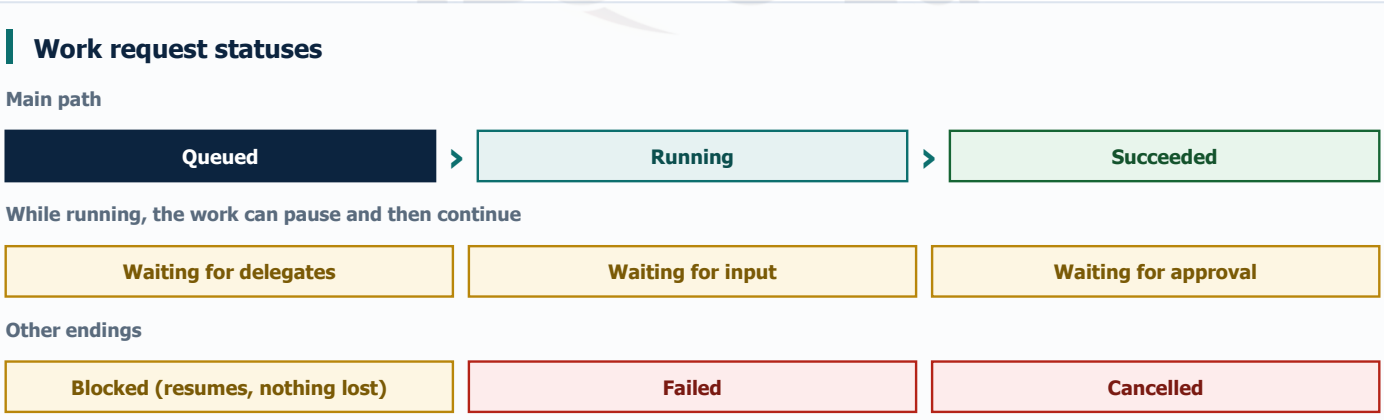
Module	Tools	Risk
CRM	Accounts, contacts and activities; create a follow-up activity.	Read; low-risk change
Sales	Leads and opportunities.	Read
General Ledger	Dashboard and journals.	Read
HR	Departments.	Read
Newsletter	Campaigns and segments.	Read
CMS	Pages.	Read

Every module tool requires the module's own view or create permission, and list results are capped at 50 rows per call.

11. Work requests

A work request is one AI employee working on one instruction. Start it from Work requests, from an employee's profile, by starting a task or plan step, or by asking an AI employee in a channel. The request runs with the authority of the person who started it.

Work request statuses



Status	What happened	What you can do
Queued	The work item waits for a free slot in the tenant's runtime.	Wait, or cancel.
Running	Model and tool calls are in progress, each recorded as a step.	Watch the steps; cancel if needed.
Waiting for delegates	Part of the work was delegated to team members.	Open the delegations.
Waiting for input	The employee asked a question or escalated.	Reply in the box on the request page.
Waiting for approval	A tool call needs a person's decision.	Decide in the Approval Center.

Status	What happened	What you can do
Blocked	Stopped before a model call by a budget, quota or configuration problem.	Fix the cause, then Resume.
Succeeded	The employee submitted a result.	Review, rate, act on it.
Failed	A coded failure, such as provider.failed. No result is invented.	Read the reason; start again.
Cancelled	A person cancelled; delegated children were cancelled too.	Nothing; history is kept.

Steps, checkpoints and recovery

- Every model call, tool call, delegation, question, escalation, submission, reply, delegate result and stop is a checkpointed step. After a restart the work resumes from the last completed step and reuses its history.
- A tool interrupted mid-call is re-run only if it is declared idempotent. Otherwise the step is marked Interrupted and the model is told the outcome is unknown.
- Step statuses: Running, Succeeded, Failed, Denied (refused before anything ran), Interrupted, Skipped and Awaiting approval.
- Provider errors are retried with backoff, then the work fails with provider.failed. Model replies are never made up.

Limits per request

Each request is bounded by model calls, steps, tokens, estimated cost and a 30-minute wall-clock deadline that spans retries and is extended while the work waits for a person or a delegate. Delegation is bounded by depth and by the size of the whole tree. The request page shows usage against each limit.

Configuration
Dvo:Orchestration sets the steps, model calls, minutes, tokens, cost, delegation depth, tree size, tool timeout, restricted tool categories and the approval tier that needs a second person.

Delegation and questions

- An employee delegates only to members of its own team: its reports, its department, or an explicit dotted line. Self-delegation, cycles, unavailable targets and trees that are too deep or too large are refused.
- Delegated results return to the manager as untrusted data, never as instructions. The request page lists each delegation and links to its own page.
- When information is missing, the employee asks the requester. The question appears on the request page, on the Command Center and as a notification; the reply continues the work.
- Cancelling flags the execution; the orchestrator finishes the run itself, cancels delegated children, cancels open approvals and tells the parent.

Enforced control
In production a route to the stub provider fails the work with a clear error. Simulated results are never reported as success outside development and tests.

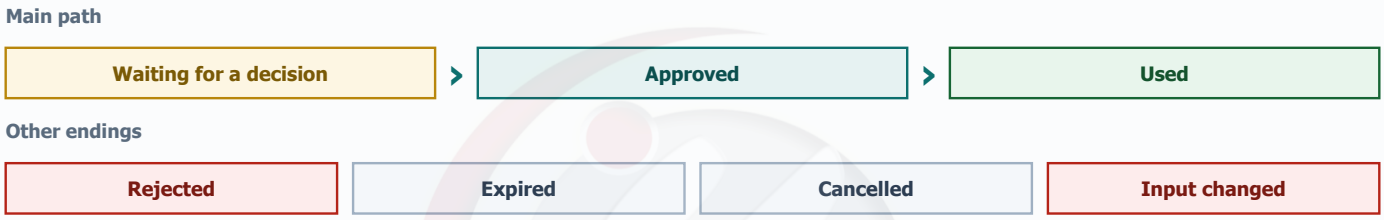
12. The Approval Center

When a tool needs approval, the work pauses and a request appears in the Approval Center. A person decides; the operation runs at most once, and only if it is still exactly what was approved.

From request to execution

1. The orchestrator stores a request with a SHA-256 fingerprint of tenant, execution, employee, tool and canonical arguments, and leaves the tool call unanswered.
2. The work shows Waiting for approval. Approvers with the right permissions are notified.
3. A person opens the approval, reads the operation, its exact input and the history, and approves or rejects with a comment.
4. After approval the call goes through the full authority check again and runs only if the fingerprint still matches and the approval is unused.
5. A rejection or expiry reaches the model as a refusal with the reason, and the work continues without that action.

Approval statuses



Who may decide

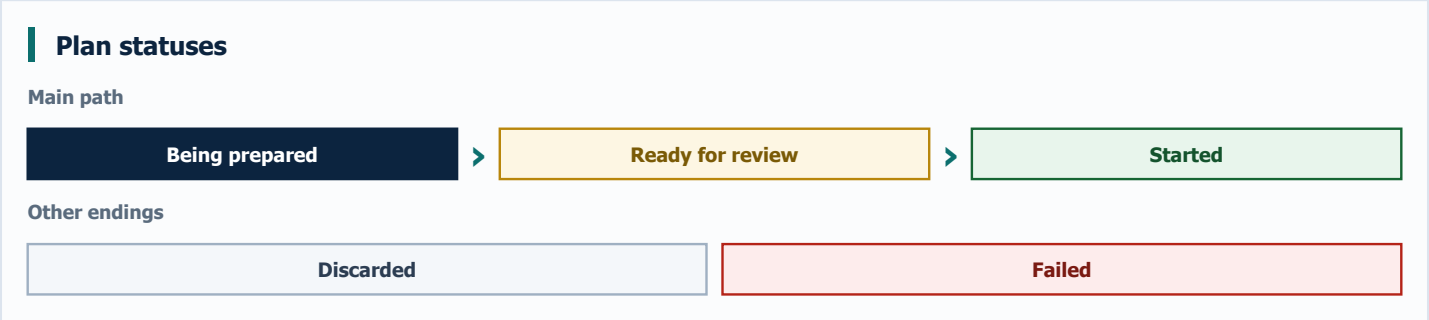
- A person with DvoApprovals Approve or Reject who also holds every permission the operation uses.
- An AI execution context can never decide, including on its own request.
- Irreversible operations need someone other than the person who requested the work.
- Each decision checks the approval's version, so a stale screen gets a conflict instead of a second decision.
- The approval page explains why the viewer cannot decide when that is the case.

Configuration
Pending approvals expire after Dvo:Orchestration:ApprovalExpiryHours, 72 hours by default, checked by a recurring job. Cancelling the work cancels its approvals.

The list filters by status, by what I can decide and by my requests, and lets people approve or reject in place. Each action is appended to the approval history and the audit log, and the approval also runs as an instance of the Core workflow engine stored in the approval's own transaction.

13. Executive plans and the Command Center

A person describes a goal in the New plan dialog. A planning execution of the Chief of Staff, or a planner the person chooses, may only propose a plan or ask a question. No other tool is offered, and any other call is refused with planning.only.



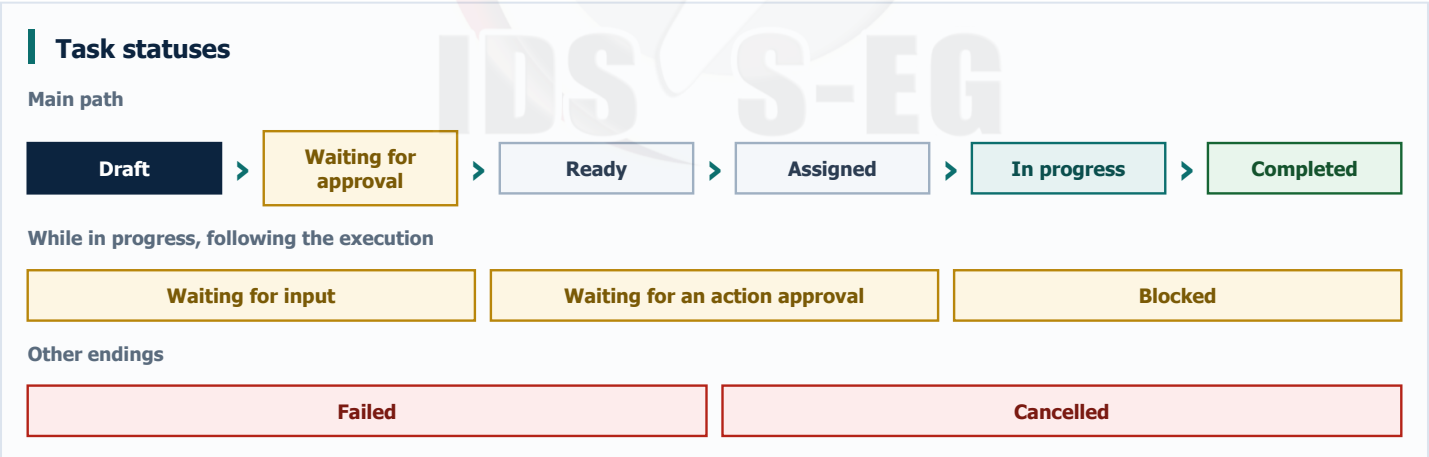
- The plan lists up to Dvo:Orchestration:MaxPlanSteps steps, each assigned to an available employee, with the tools that may need approval.
- Only the requester can pick and start steps or discard the plan. Starting the same steps again does nothing new.
- Each started step becomes a normal work request with the requester's authority and a board task linked to the plan, so dependencies can be added between those tasks.

The Command Center

The Command Center is the executive view of the office. It shows live counts of running, queued, waiting and blocked work, today's results, the month's tokens and cost against the office budget, approvals and questions waiting for the viewer, recent plans and an activity feed. It shows only what the viewer may see and refreshes live.

14. Tasks and projects

Tasks are the office's board work. A task runs only when a person starts it, either now or once every task it depends on is completed. From then on its status follows the real execution, which is queued in the same save that moves the task to In progress, so a task never shows as running without work behind it.



- A task that needs approval is approved by a person other than its author. Every change checks the signed-in person's live permissions.
- Tasks carry assignments with history, dependencies (cycles refused, results passed to the next employee as data), comments, deliverables, subtasks, priorities from Low to Urgent, dates and daily, weekly or monthly recurrence.
- Deliverables are the submitted result, text, a link or a Knowledge Center document.
- The next occurrence of a recurring task is created automatically and waits for a person to start it.
- Task numbers are per tenant, for example #16, with a retry when two people create a task at the same moment.

Board, list, calendar and timeline

Board moves go through the same operation as the matching buttons. Each task carries the moves the viewer may make, computed by the same rules, so the board offers only valid drop targets. The card menu offers the same moves for keyboard users. The calendar follows the culture's first day of the week and mirrors in right-to-left. Lists load up to 200 items and show how many match.

> **Enforced control**
> A task created by an AI employee belongs to its human requester, who must be allowed to create tasks. At L1 it stays a draft, at L2 it waits for that person's approval, and it never starts by itself.

Projects

Projects group tasks under one owner, with members (people and AI employees, as Lead or Member), a status of Active, On hold, Completed or Archived, and a shared project channel. The project page shows progress, its tasks and its members.

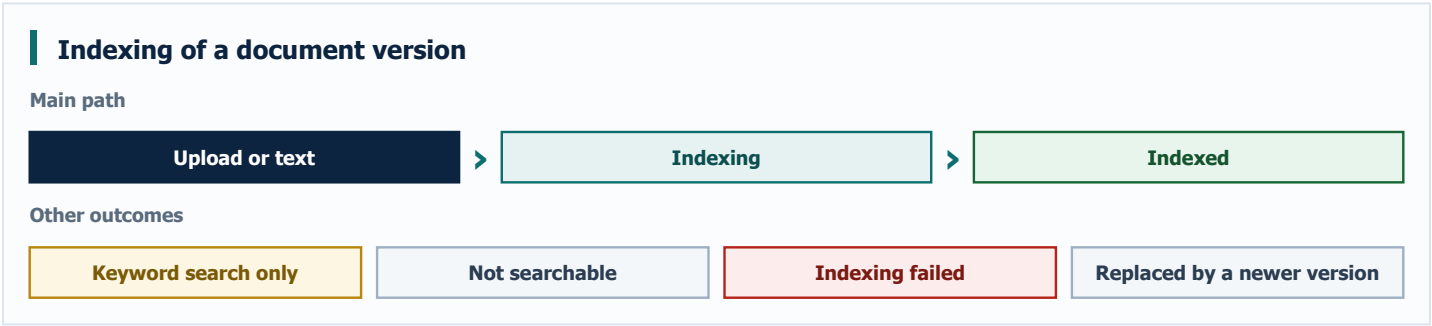
15. Communications

Channel	Who reads it
General	The whole office.
Department	People with access to the department's work.
Project	The project's members and managers.
Task	People who may see the task.
Direct	One person and one AI employee. Private to that person and DVO auditors.

- Recorded events such as task created, status changes, assignments, handoffs, approval requests, work started or failed are saved with the change they describe and link to the real record. They are stored as an event code with arguments and translated by the interface.
- Asking an AI employee in a channel starts real work with the asker's permissions, and it counts against budgets.
- The answer is posted only when that work finishes. Failed work is reported as failed, never as an answer.
- AI answers are shown only to people who may view AI employees' work.

16. The Knowledge Center

The Knowledge Center holds the company documents AI employees may consult: policies, product notes, playbooks, price lists, technical guides. Each document has a category (eleven standard ones), visibility, tags, an owner, an optional project and versions.



- Plain text, Markdown, CSV, JSON, XML, YAML, HTML and Word (.docx) are split into passages and indexed on the runtime. Other files are stored but not searchable.
- Files are kept in the platform's document archive storage under a per-tenant prefix, with paths built from server-side ids only. Storage counts against the plan's knowledge storage limit.
- Only the current version is searched. Archived documents are never retrieved.
- Ranking uses embeddings when an embedding model is configured, with keyword matching as fallback. Embedding calls are metered and skipped when a budget or quota has no room.

> **Enforced control**
> An AI employee retrieves a document only when an access rule names the document's audience as that employee, its department or all employees, and the person it works for may read the document. Candidates are narrowed to allowed sources before anything is compared.

The search tester on the Knowledge Center applies a chosen employee's access rules, so an administrator can see exactly which passages that employee would receive. The document page shows versions, upload and download, AI access rules, re-indexing and archive.

17. Memory

Memory keeps short notes that help AI employees work consistently: a customer's preferred format, a decision taken last week, a project convention. It is context, never a business record.

Area	Written by	Recalled by, when the memory scope allows it
One AI employee	That employee, or a person	That employee
Department	An employee of the department, or a person	Employees of that department
Project	An employee during work that belongs to the project, or a person	Employees during work on that project
Whole office	People only	Employees whose scope includes the office memory

- Entries written by an AI employee stay unverified until a person checks them. People can inspect, correct, verify and delete entries; deletes are recorded in the audit log.
- An AI employee writes and recalls only within the memory scope of its configuration, and recall also needs the requester's knowledge permission.
- Entries expire. Expired entries are kept for inspection and never recalled.
- Each AI employee has a cap on active entries (Dvo:Memory:MaxEntriesPerEmployee). At the cap it must ask a person to review its memory; nothing is deleted automatically.

> **Enforced control**
> Memory is never treated as authoritative business data. When a fact matters, it must come from the EGKits module that owns it, through a tool, under the requester's permissions.

18. Avatars and voice

Each AI employee has a 3D avatar and a voice in English and Arabic. People can talk to an employee by voice on its own page and hear its real answers with moving lips. The avatar is always labelled as an AI employee and is never presented as a person.

Avatars

- 22 Microsoft Rocketbox characters (MIT licence) with facial expressions and 18 shared animations such as idle, listen, talk and agree. The largest avatar is about 1 MB and all of them together about 12 MB.
- Each of the twenty standard roles (DVO-01 to DVO-20) has its own avatar and default voices. Other employees get a stable default from their id.
- The Avatar and voice tab of the profile shows a live preview. A person may choose another catalog avatar or voice; the change is versioned, written to the change history and audited as avatar.changed.
- The portrait is shown instead of 3D when the device has no WebGL, the file fails to load, the person asks for reduced motion or the page is hidden.

Enforced control
Only catalog avatars and voices are accepted, and a request running as an AI employee cannot change any employee's appearance or voice.

Voice conversation

Talk by voice on the profile or the Virtual Office opens /digital-virtual-office/voice/{id}. It is the same private conversation as in Communications: each question starts real work with the asker's permissions, and the answer is spoken only when that work finishes.

Avatar states



Part	How it works
Speech provider	Azure AI Speech. The key is saved with the other AI provider keys (provider azure-speech) and stays on the server.
Recognition	Runs in the browser only after the person presses Speak and allows the microphone, for one question. The browser gets a short-lived token, never the key. Audio is not recorded.
Speech	Text to speech runs on the server and returns audio with a viseme timeline for lip sync. Markdown, links and code are not read aloud, and long answers are split at sentence ends.
Languages	English (en-US) and Arabic (ar-EG), detected automatically or chosen by the person.
Quota	MonthlySpeechMinutes, metered once per spoken answer. A replayed clip is not billed again. When the quota runs out, speech stops and the conversation continues in text.
Settings	The Voice card in Settings shows the provider, region, languages, minutes used this month and a speech test.

Enforced control
Speech markup is escaped, so an answer cannot change the voice. Only answers the person may read are spoken, the speech token is audited without its value, and the key is only ever sent to an Azure host.

Development only
Without an Azure AI Speech key, a development host uses a silent demo voice: the lips move, nothing is spoken or recognised, every page says so and it is never billed.

19. Meetings

People meet AI employees by voice or text in a 3D room: one to one, with a few of them, or with all of them. AI employees take part as AI, never as people. Each of their answers is a real work request started with the permissions of the person who spoke, and each answers from its own role.

Meeting	What happens
One to one	A person and one AI employee. The employee answers every turn.
A few AI employees	A person names or picks who should answer; otherwise the most relevant ones answer, within the responder cap.
All AI employees	Invite all AI employees. Only the most relevant ones answer each question.
People together	People in the room can start a voice call among themselves through the platform call service.
AI employees only	Scheduled by a person, run by an AI moderator in rounds, with an agenda and hard limits. Needs the workspace policy.

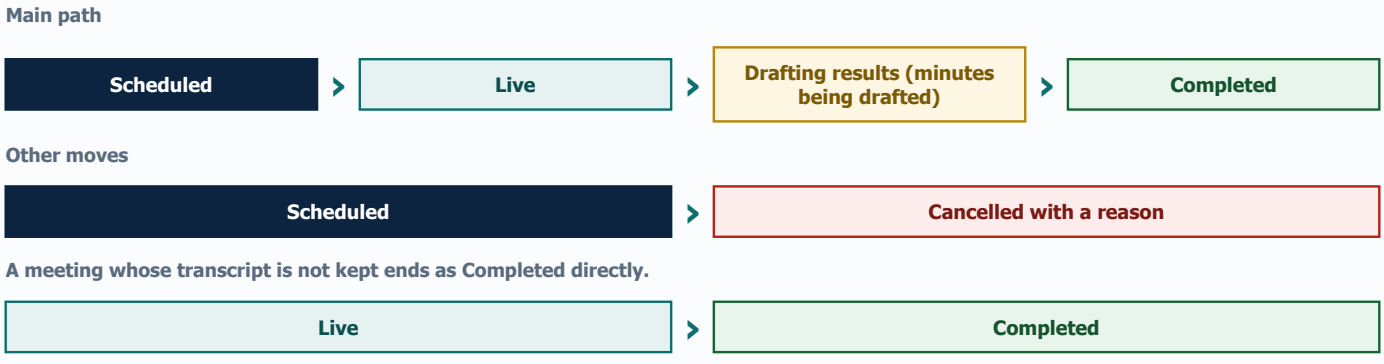
Scheduling

The Meeting Center (/digital-virtual-office/meetings) shows Live now, Upcoming, My meetings and Past, and lists the meetings the viewer may open. New meeting opens a form with the type, objective, start and duration, AI employees, people, moderator, agenda with AI turns per topic, privacy, language, transcript, recording, limits and repetition. The form reads the workspace policy and disables a choice the policy forbids, with the reason, instead of failing on save.

Field	Rule
Type	Ad hoc, Standup, Planning, Review, Brainstorm, Decision, Department, Executive, All hands or One-on-one.
AI employees	At least one. Only active employees whose meeting profile allows it can be invited.
Privacy	Participants only, or everyone in the workspace who may view meetings. Auditors may always open a meeting.
Agenda	Topics in order, each with a number of AI turns. The moderator moves on when the turns are used.
Limits	AI turns and tokens for the meeting (60 and 600,000 by default for meetings with people) and the AI answers per turn.
Repeats	None, daily, weekly or monthly.

Each AI employee has a Meetings tab on its profile: whether it may take part and moderate, the longest answer in words (30 to 400), whether anyone who may create meetings or only approvers may invite it, and preparation instructions.

Meeting lifecycle



The live room

- The room (/digital-virtual-office/meetings/{id}/room) is laid out face to face, at a small table, in a large room or as an auditorium, by the number of participants. Devices without 3D show portraits.
- A person speaks only after joining with consent. They speak through the microphone or type up to 4,000 characters, and may pick the AI employees who should answer.
- Explicit addressing comes before names and roles, and the responder cap (3 by default, at most 5) always applies. "… is preparing an answer" shows each turn being worked on.
- An answer appears only when its work has really finished. When it fails, the transcript says the AI employee did not answer.
- Read answers aloud plays each answer in the employee's voice. Interrupt stops answers still being prepared. The owner or moderator moves to the next topic, ends the meeting and turns recording on when it is allowed.
- State is kept in the database, so a restarted server or a reconnecting browser continues the meeting. The meeting job runs every minute, recovers stranded turns and ends a meeting after 30 idle minutes by default.

Enforced control
What is said in a meeting is a question, not an authorization: it cannot grant permissions, approve requests or change an AI employee's configuration. In meetings AI employees use only read tools and propose actions instead of doing them. Meeting speech reaches the model as untrusted data, and the moderator never assumes the AI employees agree.

Meetings of AI employees only

A meeting where only AI employees talk needs Allow meetings with AI employees only in Settings, an AI moderator and an agenda. It starts by itself at its time and runs for the person who scheduled it, with that person's permissions. Rounds stop at the workspace limits: 12 AI turns and 200,000 tokens by default.

Configuration
The ConcurrentMeetings (2) and MeetingParticipants (24) quotas come from the plan like the other DVO limits; the values shown are the fallback. Scheduling is refused, with the reason, when a limit is reached. Live meetings are never ended because of a quota.

20. Transcripts, minutes, decisions and action items

The meeting record (/digital-virtual-office/meetings/{id}) turns what was said into reviewed results. AI employees only propose; a person confirms decisions, approves minutes and approves action items.

From the meeting to tracked tasks

1. The transcript is written as people and AI employees speak, with system events such as joins and topic changes.
2. When the meeting ends, one AI participant drafts the minutes as real work. The output is checked against the minutes schema.
3. The owner and the human moderator are told that minutes wait for review. They edit and approve one version.
4. A person confirms the decisions that were really taken, in their own words.
5. A person approves action items. Each approved item for an AI employee becomes one task, linked to the meeting.
6. The Follow-up tab tracks the tasks and suggests a follow-up meeting while items are open.

Result	Rules	Statuses
Transcript	Append-only, searchable, exportable as plain text (audited as meeting.exported). When the workspace does not keep transcripts, the lines are removed when the meeting ends and no minutes are drafted.	—
Recording	Off unless the workspace allows it. Needs a kept transcript and everyone's consent on joining; the room shows the indicator. Only metadata is stored, never audio.	On, Off
Minutes	Summary, topics, each participant's own position, proposed decisions, open questions and risks. Output that fails the schema is not saved and the meeting says why. Draft again or write them by hand. A new draft replaces the old one; an approved version never changes.	Draft, Approved, Replaced
Decision	Exists only when a person confirms it, from a transcript line, a proposal or their own text (up to 1,000 characters). Withdrawn decisions stay in the history.	Confirmed, Withdrawn
Action item	Proposed by AI or added by a person, then reviewed. Approving needs an assignee: an AI employee or a person. An approved item for an AI employee becomes a task through the normal task service, at most once (key meeting:{meetingId}:action:{itemId}): right away when the workspace creates meeting tasks automatically, otherwise when the reviewer chooses Create task.	To review, Approved, Rejected, Task created

Enforced control

Decisions are never inferred from AI suggestions, and nothing in the minutes is a decision until a person confirms it. Items marked as restricted (financial, external, production or destructive) create tasks that wait in the Approval Center before any work starts.

Meeting notification	Sent to	Sent
Meeting invitation	The people invited, except the owner	Once per meeting
Minutes to review	The owner and the human moderator	Once per minutes version
Meeting action item done	The owner	When the task's real work completes

21. The Virtual Office and Meeting Analytics

The Virtual Office (/digital-virtual-office/office, second in the DVO menu) is a 3D floor with one zone per department and each AI employee at a desk, next to a list that can do everything the floor can. It needs the DvoEmployees view permission.

Activity shown under each avatar

Available | Waiting to start | Working | Waiting for a person | Blocked | In a meeting | Not active

- Selecting someone on the floor or in the list opens the same card, with Open profile, Talk by voice, Open the work and Enter room.
- Pointing at an avatar shows its name, code and activity. The list shows the same for keyboard and screen-reader users, and can be searched and filtered by activity.
- The floor mirrors for right-to-left languages, so the first department is on the right.
- The page follows work and meetings live and reloads once per burst of changes, with a two-minute check for changes nothing announces, such as an employee being activated.

Enforced control

Activity comes only from real executions that are not finished and from live meetings the viewer may see. It is never simulated. Work is shown only to people allowed to view work, meetings only to people allowed to view meetings, and the page says when something is hidden. The office is read-only.

Meeting Analytics

Meeting Analytics (/digital-virtual-office/meetings/analytics, opened from Meetings) shows, for a chosen period: meetings, completed, cancelled and AI-only, with the average length; AI answers given and failed and their tokens; minutes versions and approvals, confirmed decisions, and action items proposed, approved and rejected; tasks created and completed; and each AI employee's meetings and answers. Everything is counted only from the meetings the viewer may see: auditors see all of them, anyone else sees workspace-wide meetings, meetings they own and meetings they were invited to.

Enforced control
The seven meeting reports in the Reporting Center read only the meeting turns table, which holds no words. Meetings, agenda, transcripts, minutes, decisions, action items, participants, recordings and meeting profiles are never reportable, because reports are scoped to the tenant and would otherwise show private meetings to everyone who may run them.

22. Budgets, usage and the Cost Center

Every model call is recorded twice: in the platform's AI usage log and in DVO usage, with the employee, department, task, execution and billing month. Budgets and quotas are checked before each call, so spending stops before it happens, not after.

Budget element	Options
Applies to	Whole office, a department or an AI employee.
Unit	Tokens or USD per month.
Warning threshold	A percentage that raises a warning alert.
When used up	New model calls are blocked; the work shows Blocked with the code and resumes when the limit is raised.

- Budget alerts: every 15 minutes, month-to-date usage is compared with each active budget and with the subscription's monthly token and cost quotas. One alert is raised per budget, month and level (warning, then used up), and people with budget management permission are notified.
- Alerts inform only. Blocking stays with the budget guard before model calls, and nothing is stopped or deleted. Marking an alert as seen is audited.
- The Budgets and usage page shows monthly usage by employee, department and model next to each limit.

The Cost Center

For each month and provider, the Cost Center shows the estimated cost (recorded usage times the configured model rates) next to the amount a person entered from the provider's invoice, and the variance. Chat and embedding tokens are shown separately.

Known limit
Costs are estimates from configured rates until a person enters the invoiced amount for that month and provider.

23. The KPI Center

Seven measures are computed from recorded work only, never estimated by a model, for the whole office, each department and each AI employee.

Measure	Definition
Success rate	Share of finished work that succeeded.
On-time completion	Share of completed tasks with a due date that were completed by it.
Rework rate	Share of completed tasks that needed more than one attempt.
Human intervention rate	Share of finished work that had to wait for a person to answer or approve.
Approval rejection rate	Share of decided approval requests that people rejected.
Quality score	Average rating people gave finished work, from 1 to 5.
Cost per successful result	Estimated model cost divided by the work that succeeded.

How a result is judged

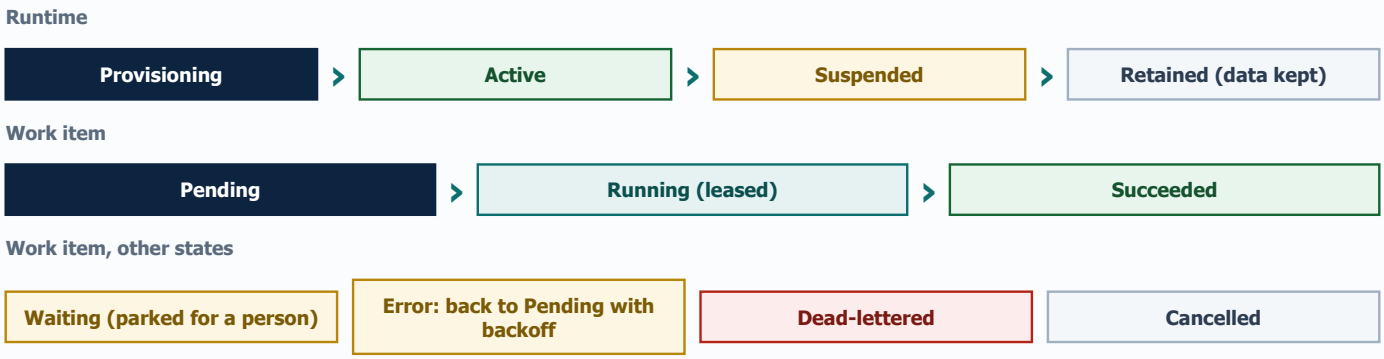
On target | At risk (within the tolerance) | Off target | No target | Not enough data

- Targets start from Dvo:Analytics:DefaultTargets. A tenant can change, disable or clear them; only people can, and every change is audited.
- A result shows no status when it has no target or fewer samples than the minimum sample size.
- Results are stored per month. A job recomputes the current month every hour, and the previous month during the first three days; a month nobody computed yet is computed when opened.
- Ratings: the requester or owner, and people who audit or supervise executions or manage tasks, rate finished work and tasks. One rating per person per piece of work. AI employees cannot rate anything, including their own work.

24. Runtime operations

Each tenant has one logical runtime: a database partition, a work queue and a configuration version. Work runs in process on Hangfire workers of the isolated dvo queue, so AI work cannot starve maintenance or integration jobs. No operating-system process is started per task.

Runtime and work item states



- Submission is idempotent per tenant and idempotency key, so a double click never starts work twice.
- The dispatcher visits tenants with pending work in round-robin order. A fair scheduler gives each tenant a quantum bounded by its free concurrency slots and the global slot budget, so a busy tenant cannot starve others.
- Claims are atomic conditional updates. The running worker extends its lease by heartbeat; a stale worker cannot overwrite the outcome.

- The lease reaper returns expired leases to the queue (counting the attempt) or dead-letters them after the last attempt. A person requeues dead-lettered work from the Runtime page, and the action is audited.
- Health snapshots are written from real queue metrics. The Runtime page shows queued, running and expired counts and can run an isolation verification for the tenant.
- Suspending a runtime keeps queued work queued, returns work being claimed without using an attempt and refuses new submissions. A runtime the platform suspended can only be resumed by the platform.
- The reconcile job provisions runtimes for tenants with the dvo app and suspends, never deletes, runtimes whose tenant lost it. Deprovisioning sets the runtime to Retained until the retention period ends.

The platform's AI Runtime Operations page lists every tenant's runtime with status, health and limits, and lets platform operators provision, suspend, resume and set quota overrides with a reason. It never shows tenant content.

> **Known limit**
> The runtime is logical and in process, so a handler shares its node with the rest of the server. Isolation is enforced in data, scheduling and authorization, not by operating-system boundaries. A worker-process or container host can replace it behind the same interface.

25. Audit, telemetry and retention

Every DVO operation writes an audit event with the actor (person, AI employee, system, support or platform), the employee, the runtime, the target, the authorization result, any approval reference, the outcome (Succeeded, Denied or Failed) and the correlation id. The Audit page searches these events, and each target links to its execution, approval, plan, employee, task, project, document or channel.

Telemetry

The meter and activity source EGKits.Dvo publish metrics and traces. Metric tags are codes, providers and outcomes only: there is no tenant id, user id or text, so cardinality stays bounded. Spans never carry instructions, prompts, results or tool arguments.

Instrument	Measures
`egkits.dvo.work_items`	Work item outcomes, with queue wait and duration histograms.
egkits.dvo.executions	Executions by status, failure code and whether delegated.
`egkits.dvo.model_calls`	Model calls and their duration per provider.
egkits.dvo.tokens	Input and output tokens per provider.
egkits.dvo.cost.estimated	Estimated cost in USD per provider.
`egkits.dvo.tool_calls`	Tool calls by kind and outcome.
`egkits.dvo.budget_blocks`	Budget and quota blocks by code.
egkits.dvo.approvals	Approval actions by risk tier.
egkits.dvo.leases.recovered	Recovered leases: requeued or dead-lettered.

One trace runs from the request that gave the instruction to the final result. Delegated and follow-up work inherits the trace, and each run continues it. Traces are exported when Telemetry:OtlpEndpoint is set.

Retention

- A daily job applies each tenant's own execution and audit retention (7 to 3650 days).
- A whole execution tree is removed only when every execution in it ended before the cutoff, no open task points at it, no

- plan is still being prepared or reviewed from it and no approval in it is pending.
- The cutoff never reaches into a month the KPI job can still recompute.
- Usage records, tasks, projects, plans, channel messages, KPI results, ratings, budget alerts, knowledge and memory are kept regardless of age.
- Each run that removes anything writes one retention.purged audit event with the cutoffs and counts.

26. Notifications, reports and the home dashboard

Notification	Sent to	Not repeated within
Approval requested	People who may decide	60 minutes
Approval expired	Approvers	60 minutes
Budget alert	Budget managers	12 hours
Work failed	The requester	30 minutes
Question from an AI employee	The requester	60 minutes
Meeting invitation	The people invited	Once per meeting
Minutes to review	The owner and the human moderator	Once per minutes version
Meeting action item done	The meeting owner	Once per task

Notifications go through the platform notification center, and each person can turn each kind off in Settings. Each notice is written in the language of the recipient's profile. Delivery history and choices belong to the tenant.

Reporting Center

Forty DVO reports and two dashboards (Cost and Usage, and Workforce Performance) are seeded at four levels. The workspace has the monthly cost trend with a three-month forecast, tokens, model calls, cost by model, provider and usage kind, provider invoices, budget policies and alerts, and the KPI scorecard. Each department has its cost, usage summary, KPI results and status, and staffing. Each AI employee has its usage and cost, the top spenders, KPI results, quality ratings and meeting answers. Meetings have outcomes, the answer rate, a weekly trend and the reasons AI employees were given the turn. Employee and department names replace identifiers, KPI tiles compare a date range with the period before, and each report carries a summary, highlights and suggested actions. They require the DvoReports view permission. Reports are limited to usage, budgets, alerts, KPI results and targets, ratings, provider costs, meeting turns and the employee and department names that label them.

Enforced control
Conversations, prompts, configurations, tool calls, knowledge, memory, audit records, meetings, transcripts, minutes, decisions and action items are never reportable. The report schema policy fails closed.

Home dashboard

The EGKits home dashboard shows DVO attention signals: pending approvals (a warning when one expires within four hours), the viewer's work waiting for an answer, the viewer's work that failed in the last 24 hours and open budget alerts (critical when a budget is used up).

27. The demo scenario

> **Development only**
> The demo scenario exists only on a host running in the Development environment with Dvo:Demo:Enabled. Elsewhere its endpoints return 404, nothing routes to the scripted provider, and the Settings card and the page banner stay hidden.

Load demo scenario on the Settings page (needs DvoSettings Update) works only in an empty office with an active runtime whose limits can hold it. It takes about a minute and builds a complete working company, Retail Cloud:

- The standard template, with all 20 roles in 8 departments hired, activated and given their tools.
- 9 knowledge documents indexed for real, 5 projects, 15 tasks, channel messages and budgets, one of them deliberately tiny.
- A kickoff that starts real work on the real orchestrator for the person who loaded it.

What runs	What it shows
CEO weekly briefing	Delegation to three managers and a merged result.
Release review in Release 1.0	Task-based work that creates a rollback-runbook subtask in the same project.
CRM follow-up	A write that waits for a second person's approval in the Approval Center.
Data migration	A question to the requester and completion after the reply.
Onboarding	Work blocked by its budget, as designed.
Sales channel question	An answer posted by the researcher when its work finishes.
Launch plan	A proposed four-step plan ready for review.
Meetings	Each demo employee answers in meetings from the knowledge its role may read. The demo minutes contain only what was said: action items come only from lines starting with "Action:", and decisions are never inferred.
Voice	Without a speech key, the silent demo voice moves the lips and says on every page that nothing is spoken.

Only the model replies are scripted. The demo employees use the dvo-demo provider, which every page labels Demo (scripted), and every scripted result starts with [Demo]. Remove demo scenario cancels the demo's running work, deletes only the rows the scenario recorded and writes an audit event; embedding usage and notification delivery records are kept because they record real consumption.

28. Configuration reference

Server settings live in the Dvo configuration section. Tenant-level choices are on the Settings page; limits per plan are in Dvo:PlanDefaults.

Setting	Default	Guidance
`Dvo:AllowStubProviders`	false	Must stay false in production. An execution routed to the stub provider then fails.
`Dvo:Demo:Enabled`	false	Ignored outside Development.
`Dvo:JobsEnabled`	true	false stops all DVO background work on a node without losing queued work.
`Dvo:GlobalMaxConcurrency`	8	Most work items one process runs at once across tenants. Size to workers and provider rate limits.
`Dvo:DispatchBatchSize`	16	Work items claimed per dispatcher pass.
`Dvo:DefaultLeaseSeconds`	300	How long a worker holds an item without a heartbeat.
`Dvo:DefaultMaxAttempts`	5	Attempts before an item is dead-lettered.

Setting	Default	Guidance
Dispatch, reaper, health, reconcile, approval expiry, knowledge	1, 2, 5, 15, 10, 15 min	Recurring job schedules, in UTC.
`Dvo:Analytics:KpiCron`	hourly at :20	KPI recomputation.
`Dvo:Analytics:BudgetAlertCron`	every 15 min	Budget and quota alert scan.
`Dvo:RetentionCron`	daily 02:30	Retention purge.
`Dvo:Orchestration:ApprovalExpiryHours`	72	Pending approvals expire after this.
`Dvo:MeetingCron`	every minute	Starts AI-only meetings at their time, keeps rounds going, recovers stranded turns and ends idle meetings.
Dvo:Meetings:DefaultMaxAiTurns / DefaultMaxTokens	60 / 600,000	Default limits of meetings with people. AI-only meetings take theirs from Settings.
`Dvo:Meetings:DurationGraceMinutes`	15	How long a meeting may run past its expected length before AI turns stop.
`Dvo:Speech:RecognitionLanguages`	en-US, ar-EG	Languages the browser listens for, with automatic detection.
Dvo:Speech:MaxSpokenCharacters / SegmentCharacters	2,400 / 320	Characters of one answer that are spoken, and the size of each spoken part. The full text stays in the conversation.
Dvo:Speech:SynthesisCostPerMillionCharacters / RecognitionCostPerHour	16 / 1 USD	Rates used to estimate speech cost in usage records.
`Dvo:PlanDefaults:{Plan}`	none	Employees, departments, concurrency, monthly tokens and cost, storage, retention, providers and models per plan.
`Dvo:Fallback`	small limits	Used when neither the plan nor a tenant override defines a limit.

Deployment: apply the ten DvoDbContext migrations (schema, organization, orchestration, approvals, work, analytics, demo, avatars, speech and meetings). Every table is in the Dvo schema, or named Dvo_<Table> on databases without schemas. List the dvo queue in HangfireOptions:Queues on every node that should run AI work. Serve .glb files as model/gltf-binary so 3D avatars load.

29. Page reference

Every page of the module, its address and the permission resource that shows it in the menu. Detail pages open from their list.

Page	Address	Permission
Overview	`/digital-virtual-office`	`Dvo`
Virtual Office	`/digital-virtual-office/office`	`DvoEmployees`
Command Center	`/digital-virtual-office/command-center`	`DvoExecutions`
Approval Center and approval	`/digital-virtual-office/approvals`	`DvoApprovals`
Plans and plan	`/digital-virtual-office/plans`	`DvoExecutions`
Tasks and task	`/digital-virtual-office/tasks`	`DvoTasks`
Projects and project	`/digital-virtual-office/projects`	`DvoTasks`
Communications	`/digital-virtual-office/communications`	`DvoCommunications`
Voice conversation	`/digital-virtual-office/voice/{id}`	`DvoCommunications`
Meetings	`/digital-virtual-office/meetings`	`DvoMeetings`
Meeting record	`/digital-virtual-office/meetings/{id}`	`DvoMeetings`
Meeting room	`/digital-virtual-office/meetings/{id}/room`	`DvoMeetings`

Page	Address	Permission
Meeting Analytics	`/digital-virtual-office/meetings/analytics`	`DvoMeetings`
Knowledge Center and document	`/digital-virtual-office/knowledge`	`DvoKnowledge`
Memory	`/digital-virtual-office/memory`	`DvoKnowledge`
Organization	`/digital-virtual-office/organization`	`DvoOrganization`
Department workspace	`/digital-virtual-office/departments/{id}`	`DvoOrganization`
AI employees and profile	`/digital-virtual-office/employees`	`DvoEmployees`
Work requests and request	`/digital-virtual-office/executions`	`DvoExecutions`
Tools	`/digital-virtual-office/tools`	`DvoTools`
Budgets and usage	`/digital-virtual-office/budgets`	`DvoBudgets`
Cost Center	`/digital-virtual-office/costs`	`DvoBudgets`
KPI Center	`/digital-virtual-office/kpis`	`DvoReports`
Runtime	`/digital-virtual-office/runtime`	`DvoRuntime`
Settings	`/digital-virtual-office/settings`	`DvoSettings`
Audit	`/digital-virtual-office/audit`	`DvoAudit`
AI Runtime Operations (platform)	`/admin/dvo-platform`	`DvoPlatform`

30. Troubleshooting and runbooks

Symptom	Likely cause	What to do
"The Digital Virtual Office is not available."	The dvo app is not in the tenant's subscription.	Ask the platform to enable it; the runtime is provisioned on the next reconcile.
Work stays Queued	No node runs the dvo queue, the runtime is suspended, or the tenant's concurrency is full.	Check the Runtime page counts and the node's Hangfire queues.
Work shows Blocked	A budget, a quota or a configuration problem, named by the code.	Raise the budget or ask for a quota override, then Resume.
Failed with provider.failed	The AI provider was unavailable after all retries, or the key is wrong.	Check the provider status and the AI provider settings, then start again.
A step shows Not granted	The employee was not given the tool.	Grant it on the profile if you hold its permissions.
A step shows Permission missing	The requesting person lacks a permission the tool uses.	Ask a person who holds it to start the work, or adjust roles.
I cannot approve	You lack DvoApprovals Approve or a permission of the operation, or you requested irreversible work yourself.	The approval page states the reason. Ask another approver.
Approval shows Input changed	The model changed the arguments after approval.	Decide on the new request; the old approval cannot be reused.
A document is never found	No access rule names the employee, its department or all employees; the file is not searchable; or it is archived.	Use the search tester with that employee and fix the rule or the format.
Dead-lettered work	All attempts failed; the error is redacted and audited.	Fix the cause, then requeue from the Runtime page.
A portrait shows instead of the 3D avatar	No WebGL, reduced motion is on, or the server does not serve .glb files as model/gltf-binary.	Check the device settings and the static file content types.
Voice shows Not set up	No Azure AI Speech key is saved for the workspace.	In AI settings choose Azure AI Speech, set https://<region>.api.cognitive.microsoft.com/, save the key, then Test speech.

Symptom	Likely cause	What to do
Nobody answers in a meeting	The speaker may not give work to AI employees, a turn or topic limit was reached, or no employee found the topic relevant.	Pick an employee to ask directly, or raise the limits for the next meeting.
No minutes after a meeting	The transcript was not kept, or the draft failed the schema check.	Press Draft again, or write the minutes with Edit.
Suspected cross-tenant access	Investigation needed.	Run the isolation verification, search the audit by correlation id; isolation violations dead-letter the work and are logged as critical.
A key may have leaked	A provider key was pasted into text or echoed by an error.	Rotate the key in AI provider settings, then search audit details and errors for its prefix.

> **Tip**
> Work waiting for a person, an approval or a delegate is parked by design and resumes when that happens. It is not stuck.

31. Implementation checklist

- The ten DVO migrations are applied and the dvo queue is listed on the nodes that run AI work.
- Dvo:AllowStubProviders is false and Dvo:Demo:Enabled is false in production.
- Plan limits are defined in Dvo:PlanDefaults for every subscription plan that includes the app.
- The tenant's AI providers and models are configured and tested in the AI Assistant settings.
- DVO permissions are assigned to roles; approvers hold both DvoApprovals and the module permissions they will approve.
- The CEO user, default autonomy, retention and support access are set in Settings.
- The organization is applied and reviewed; each employee's configuration, data scopes and autonomy are checked before activation.
- Tools are granted on purpose; high-risk tools have policies with the right audit level.
- Office, department and employee budgets exist, with warning thresholds.
- Knowledge documents have access rules, and the search tester returns what each role should see.
- For voice, an Azure AI Speech key is saved and Test speech passes; the monthly speech minutes fit the plan.
- Meeting rules are set on purpose: transcripts, recording, AI-only meetings, answers per turn and idle end. Each employee's Meetings tab says who may invite it.
- A pilot group runs real work for two weeks; results are rated and KPI targets are adjusted.
- Alerts on dead letters, queue wait, provider failures, budget blocks and approval backlog are in place.

32. Glossary

Term	Meaning
Action item	A proposed follow-up from a meeting that becomes one task only after a person approves it.
AI moderator	An AI employee that opens, steers and closes a meeting of AI employees only. It never assumes agreement.
Autonomy level	L1, L2 or L3: how far an AI employee may act on its own.
Canonical arguments	Tool input normalized so the same operation always gives the same fingerprint.

Term	Meaning
Correlation id	The id that links a request, its work items, audit events and trace.
Data scope	An EGKits app an employee may read or act in.
Dead letter	A work item that failed all its attempts and waits for a person.
Dotted line	A coordination relationship without authority.
Effective authority	Grants, live permissions, entitlement and autonomy combined.
Fingerprint	SHA-256 of tenant, execution, employee, tool and canonical arguments.
Idempotent	Safe to run twice with the same result.
Lease	A worker's time-limited claim on a work item, extended by heartbeat.
Minutes version	One draft of a meeting's minutes. A newer draft replaces it; an approved version never changes.
Restricted area	A tool category such as money, external communication or deployment that always needs L3 and approval.
Runtime	A tenant's logical AI runtime with its queue, limits and health.
Solid line	The single manager who directs and reviews an employee's work.
Stub provider	A test-only provider refused in production.
Untrusted data block	Delimited content the model must treat as data, never as instructions.
Viseme	A mouth shape timed to speech so the avatar's lips follow the spoken answer.

EGKits publishes this guide to describe the module as shipped in EGKits SaaS V11. It contains no tenant data, prompts, credentials or production logs. Screens and limits can change in later releases; the release notes list each change.