



Myra AI Workspace

Administration Manual

Version History

Version	Date	Reason for Change
V2.5	08/2026	Updated sections • Home • Getting access
V2.4	08/2026	New sections • Updated sections New sections: • Frequently asked questions • API Endpoints Updated sections: • Views • Configurations • FAQ & Error messages
V2.3	08/2026	Updated sections • New sections Updated sections: • Product description New sections: • Technical description
V2.2	08/2026	New sections • Playground Updated sections • Home page and GUI layout • Agents • Members tab • Guardrails section • Managing gateways • Configuring guardrails • Request Logs • Gateways
		Updated sections

Version	Date	Reason for Change
V2.1	08/2026	• Views
V2.0	08/2026	New sections • Role concept • Notational conventions Updated sections • Getting Started • Views • Configurations • FAQ & Error messages
V1.3	08/2026	Updated sections • FAQ & Error messages
V1.2	08/2026	Updated sections • Initial setup • Configuring guardrails • Managing routing rules • Creating a gateway • Enabling a model on a gateway • Access • Chat New sections • Creating an auth token • Editing an auth token • Revoking an auth token • Deleting a provider key
V1.1	08/2026	Updated sections • Views • Configurations
V1.0	08/2026	Initial version

Contents

1. Product description	20
1.1 Functions	22
1.2 Limitations	25
1.3 Target groups	27
1.4 Functions of your role	28
1.5 Permissions	29
1.6 Technical description	33
1.6.1 Functional overview	34
1.6.2 System architecture	35
1.6.3 Technical details	37
1.6.4 Error handling	40
1.6.4.1 Limits and billing	40
1.6.4.2 Access and security	40
1.6.4.3 Request and model	41
1.6.4.4 Provider and operations	42
1.6.5 Use cases	44
2. Getting started	45
2.1 Getting access	45
2.1.1 Acknowledging the note about MYRA AI	49
2.1.2 Walkthrough of the interface	51
2.1.3 Signing out	54
2.2 Role concept	55

2.2.1 Assigning roles	56
2.2.2 Custom roles	57
2.3 Initial setup	58
2.3.1 Step 1: Create a tenant	60
2.3.2 Step 2: Create a gateway	62
2.3.3 Step 3: Enable a model	64
2.3.4 Step 4: Set a budget and a rate limit	66
2.3.5 Step 5: Set up the guardrails	68
2.3.6 Step 6: Examine the roles	70
2.3.7 Step 7: Create the user accounts	71
2.3.8 Step 8: Create the groups	73
2.3.9 Step 9: Issue a token	75
2.3.10 Step 10: Send a test request and examine it	77
3. Views	79
3.1 Home page and GUI layout	79
3.2 Workspace	85
3.2.1 Chat	85
3.2.1.1 Add files, photos, or options menu	88
3.2.1.2 Choose a model dialog	90
3.2.1.3 Actions on a response	92
3.2.1.4 Report this response dialog	94
3.2.2 Projects	97
3.2.2.1 Detail view of a project	99
3.2.2.2 Overview tab	101
3.2.2.3 Files tab	102
3.2.2.4 Members tab	104
3.2.2.5 Conversations tab	105

3.2.2.6 Feed tab	106
3.2.3 Workflows	108
3.2.3.1 Builder of a workflow	110
3.2.3.2 Workflow copilot dialog	112
3.2.3.3 Versions dialog	113
3.2.3.4 Runs of a workflow	115
3.2.4 Agents	117
3.2.4.1 List of the agents	118
3.2.4.2 Detail view of an agent	121
3.2.4.3 Agent catalog dialog	123
3.2.5 Playground	125
3.2.6 Library	127
3.2.6.1 Versions dialog	129
3.2.7 Scheduled	131
3.3 Overview	133
3.3.1 Dashboard	133
3.3.2 Recent chats	136
3.3.3 Analytics	139
3.3.3.1 Cost Analytics tab	140
3.3.3.2 Cost Analytics breakdowns	142
3.3.3.3 Detail panel of a tenant	144
3.3.3.4 Live Monitor tab	146
3.3.3.5 Health tab	148
3.3.4 Feedback	150
3.3.4.1 Inbox tab	151
3.3.4.2 Edit feedback dialog	153
3.3.4.3 Sessions tab	154
3.3.4.4 Reports tab	156
3.4 Settings	158
3.4.1 Approvals	158

3.4.1.1	Personal tab	159
3.4.1.2	Approve this workflow run? page	161
3.4.1.3	Agents tab	163
3.4.1.4	Tool approvals tab	165
3.4.1.5	Approve review dialog and Deny review dialog	167
3.4.1.6	Configurations tab	168
3.4.2	Preferences	171
3.4.2.1	My preferences tab	172
3.4.2.2	My Commands tab	174
3.4.2.3	Command – / dialog	177
3.4.2.4	Version history – / dialog	178
3.4.2.5	Share / dialog	179
3.4.2.6	My PII configuration tab	180
3.4.3	Connections	182
3.4.3.1	MCP Connectors tab	183
3.4.3.2	Connector directory dialog	184
3.4.3.3	Chat Integrations tab	186
3.4.4	Request Logs	189
3.4.4.1	Request details area	191
3.4.4.2	Request Trace area	192
3.4.5	Gateways	195
3.4.5.1	Detail view of a gateway	197
3.4.5.2	Gateway section	199
3.4.5.3	Provider Keys section	201
3.4.5.4	Auth Tokens section	202
3.4.5.5	Guardrails section	203
3.4.5.6	Routing Rules section	205
3.4.5.7	Circuit Breaker section	207
3.4.6	Providers	208
3.4.7	Security	210
3.4.8	Costs	212
3.4.8.1	User Costs tab	213
3.4.8.2	Provider Costs tab	215

3.4.8.3 Entries in the New Model Price and Edit dialogs	217
3.5 User Management	219
3.5.1 Users	219
3.5.1.1 Detail view of a user account	221
3.5.1.2 Tokens section	222
3.5.2 Groups	224
3.5.2.1 Group dialog	225
3.5.3 Roles & permissions	228
3.5.4 Tenants	230
3.5.4.1 Detail view of a tenant	232
3.5.4.2 Overview tab	233
3.5.4.3 Gateways tab	234
3.5.4.4 API Keys tab	235
3.5.4.5 Commands tab	236
3.5.4.6 Prompt Examples tab	238
3.6 Account	240
3.6.1 My Account	240
3.6.1.1 Entries in the New Token dialog	243
3.6.1.2 Entries in the Token Created dialog	244
3.6.1.3 Effects in the Delete Account dialog	245
3.6.1.4 Effects in the Erase all my data dialog	246
3.6.2 Plans	248
3.6.2.1 Entries in the Edit plan dialog	250
3.6.3 Vouchers	252
3.6.3.1 Entries in the New voucher dialog	253
4. Configurations	255
4.1 Account	255
4.1.1 Personal preferences	256
4.1.1.1 Entries in the General area	257

4.1.1.2	Entries in the Personal AI preferences area	257
4.1.1.3	Effects of the changes	258
4.1.2	Managing personal commands	260
4.1.2.1	Creating a personal command	260
4.1.2.1.1	Entries in the New Command dialog	261
4.1.2.1.2	Entries in the Edit Command dialog	262
4.1.2.2	Deleting a personal command	264
4.1.2.2.1	Effects of the deletion of a command	265
4.1.3	Managing PII entries	266
4.1.3.1	Entries for the PII entry	267
4.1.3.2	Effects of the deletion of a keyword	268
4.1.4	Deleting the account	269
4.1.4.1	Deleting the account	269
4.1.4.1.1	Effects of the deletion of the user account	271
4.1.4.2	Erasing all your data	271
4.1.5	Managing tokens	275
4.1.5.1	Creating a token	275
4.1.5.1.1	Data in the New Token dialog	277
4.1.5.2	Revoking a token	278
4.1.5.2.1	Effects of the revocation	280
4.2	Workspace	281
4.2.1	Managing projects	282
4.2.1.1	Creating a project	282
4.2.1.1.1	Data in the New Project dialog	284
4.2.1.1.2	Access tiers	285
4.2.1.2	Editing a project	286
4.2.1.2.1	Data in the Edit project form	287
4.2.1.3	Deleting a project	288
4.2.1.3.1	Effects of the deletion	290
4.2.2	Managing agents	291
4.2.2.1	Creating an agent	291
4.2.2.1.1	Entries in the New Agent dialog	293
4.2.2.2	Editing an agent	295

4.2.2.2.1 Entries in the Edit Agent dialog	296
4.2.2.3 Deleting an agent	298
4.2.2.3.1 Effects of the deletion	299
4.2.3 Managing workflows	300
4.2.3.1 Creating a workflow	300
4.2.3.1.1 Entries in the New workflow dialog	301
4.2.3.2 Editing a workflow	302
4.2.3.2.1 Entries for the workflow	303
4.2.3.2.2 Entries per step type	304
4.2.3.2.3 Error handling of a step	304
4.2.3.2.4 Entries in the Schedule dialog	305
4.2.3.2.5 Entries in the Webhook dialog	306
4.2.3.2.6 Entries in the Form dialog	306
4.2.3.3 Deleting a workflow	307
4.2.3.3.1 Effects of the deletion	308
4.2.4 Managing scheduled tasks	310
4.2.4.1 Creating a scheduled task	310
4.2.4.1.1 Entries in the New scheduled task dialog	311
4.2.4.2 Editing a scheduled task	313
4.2.4.2.1 Entries in the Edit scheduled task dialog	314
4.2.4.3 Deleting a scheduled task	315
4.2.4.3.1 Effects of the deletion	316
4.2.5 Managing the prompts of the library	317
4.2.5.1 Creating a prompt of the library	317
4.2.5.1.1 Entries in the New prompt dialog	318
4.2.5.2 Editing a prompt of the library	319
4.2.5.2.1 Entries in the Edit prompt dialog	320
4.2.5.3 Deleting a prompt of the library	322
4.2.5.3.1 Effects of the deletion	323
4.3 Connections	324
4.3.1 Managing MCP connectors	325
4.3.1.1 Creating an MCP connector	325
4.3.1.1.1 Entries in the New MCP Connector dialog	326
4.3.1.2 Editing an MCP connector	329

4.3.1.2.1 Entries in the Edit MCP Connector dialog	330
4.3.1.3 Deleting an MCP connector	332
4.3.1.3.1 Effects of the deletion	333
4.3.2 Managing chat integrations	335
4.3.2.1 Creating a chat integration	335
4.3.2.1.1 Entries in the Add chat connection dialog	337
4.3.2.2 Editing a chat integration	339
4.3.2.2.1 Effects of the changes	340
4.3.2.3 Rotating the credentials of a chat integration	341
4.3.2.3.1 Entries in the Rotate bot token dialog	342
4.3.2.4 Deleting a chat integration	344
4.3.2.4.1 Effects of the deletion	345
4.4 Gateway	347
4.4.1 Managing gateways	348
4.4.1.1 Creating a gateway	348
4.4.1.1.1 Entries in the New Gateway dialog	349
4.4.1.2 Enabling a model on a gateway	351
4.4.1.2.1 Entries in the Add Model dialog	354
4.4.1.3 Deleting a provider key	356
4.4.1.3.1 Effects of the deletion of a key	359
4.4.1.4 Editing a gateway	360
4.4.1.4.1 Entries in the Edit Gateway dialog	363
4.4.1.5 Creating an auth token	366
4.4.1.5.1 Entries in the Create Auth Token dialog	370
4.4.1.6 Editing an auth token	371
4.4.1.6.1 Entries in the Edit Token dialog	375
4.4.1.6.2 Effects of the changes	376
4.4.1.7 Revoking an auth token	377
4.4.1.7.1 Effects of the revocation	380
4.4.1.8 Deleting a gateway	381
4.4.1.8.1 Effects of the deletion	382
4.4.2 Configuring guardrails	383
4.4.2.1 Adding a guardrail detector	383
4.4.2.1.1 Detector types	386

4.4.2.1.2 Data common to each detector	388
4.4.2.1.3 Data for the Regex / Pattern detector	388
4.4.2.1.4 Data for the Keyword detector	389
4.4.2.1.5 Data for the Jailbreak detector	390
4.4.2.1.6 Data for the Presidio (NLP) detector	390
4.4.2.1.7 Data for the Prompt Guard detector	391
4.4.2.1.8 Data for the Prompt Injection detector	392
4.4.2.1.9 Data for the PII Protector detector	392
4.4.2.1.10 Data for the Custom PII Blacklist detector	393
4.4.2.1.11 Execution plan	393
4.4.2.2 Changing the sequence of the guardrail detectors	394
4.4.2.3 Removing a guardrail detector	396
4.4.2.3.1 Effects of the removal	397
4.4.3 Managing routing rules	399
4.4.3.1 Creating a routing rule	399
4.4.3.1.1 Data in the New Routing Rule dialog	402
4.4.3.1.2 Conditions	403
4.4.3.1.3 Action: Direct route	403
4.4.3.1.4 Action: Load balance	404
4.4.3.2 Editing a routing rule	404
4.4.3.3 Changing the sequence of the routing rules	406
4.4.3.4 Deleting a routing rule	408
4.4.3.4.1 Effects of the deletion	409
4.4.4 Managing governance templates	410
4.4.4.1 Creating a governance template	410
4.4.4.1.1 Entries in the New governance template dialog	411
4.4.4.1.2 Entries in the Edit governance template dialog	413
4.4.4.2 Applying a governance template	413
4.4.4.2.1 Entries in the Preview & apply dialog	414
4.4.4.3 Deleting a governance template	415
4.4.4.3.1 Effects of the deletion	416
4.5 User Management	417
4.5.1 Managing groups	418
4.5.1.1 Creating a group	418

4.5.1.1.1 Data in the New group dialog	419
4.5.1.2 Editing a group	420
4.5.1.2.1 Data in the dialog of the group	422
4.5.1.3 Deleting a group	423
4.5.1.3.1 Effects of the deletion	424
4.5.2 Managing user accounts	425
4.5.2.1 Creating a user account	425
4.5.2.1.1 Data in the New User dialog	426
4.5.2.2 Editing a user account	428
4.5.2.2.1 Data in the Edit dialog	431
4.5.2.3 Deleting a user account	433
4.5.2.3.1 Effects of the deletion	435
4.5.3 Managing roles	436
4.5.3.1 Creating a role	436
4.5.3.1.1 Data in the New role dialog	437
4.5.3.2 Editing a role	439
4.5.3.2.1 Data in the Edit NAME dialog	440
4.5.3.3 Deleting a role	441
4.5.3.3.1 Effects of the deletion	442
4.5.4 Managing tenants	444
4.5.4.1 Creating a tenant	444
4.5.4.1.1 Entries in the New Tenant dialog	446
4.5.4.2 Editing a tenant	448
4.5.4.2.1 Entries in the Branding tab	451
4.5.4.2.2 Entries in the Data & Retention tab	452
4.5.4.2.3 Entries in the Routing tab	453
4.5.4.2.4 Entries in the Identity & SSO tab	453
4.5.4.3 Deleting a tenant	455
4.5.4.3.1 Effects of the deletion	457
4.6 Billing	459
4.6.1 Managing the billing	460
4.6.1.1 Entries in the billing settings	461
4.6.1.2 Effects of the changes	462

4.6.2 Managing model prices	463
4.6.2.1 Creating a model price	463
4.6.2.1.1 Effects of the new model price	464
4.6.2.2 Editing a model price	465
4.6.2.2.1 Effects of the changed prices	466
4.6.2.3 Deleting a model price	467
4.6.2.3.1 Effects of deleting a model price	468
4.6.3 Managing plans	470
4.6.3.1 Effects of the plan change	471
4.6.4 Managing vouchers	473
4.6.4.1 Creating a voucher	473
4.6.4.1.1 Effects of the new voucher	474
4.6.4.2 Deactivating a voucher	475
4.6.4.2.1 Effects of the deactivation	476
5. FAQ & Error messages	477
5.1 Sign-in	477
5.1.1 Messages	478
5.1.2 That code is incorrect or has expired	479
5.1.2.1 Cause	479
5.1.2.2 Solution	479
5.1.3 Too many attempts	480
5.1.3.1 Cause	480
5.1.3.2 Solution	480
5.1.4 Could not reach the server	481
5.1.4.1 Cause	481
5.1.4.2 Solution	481
5.1.5 Frequently asked questions	482
5.1.5.1 How long is a one-time code valid?	482
5.1.5.2 How often can I enter a code?	482
5.1.5.3 Why do I get no email with the code?	482
5.1.5.4 How long do I stay signed in?	482

5.1.5.5 Can I sign in with a password?	482
5.2 Gateway	483
5.2.1 Messages	484
5.2.2 Rate limit exceeded	485
5.2.2.1 Cause	485
5.2.2.2 Solution	485
5.2.3 The usage budget of this workspace is used up	486
5.2.3.1 Cause	486
5.2.3.2 Solution	486
5.2.4 Request blocked by content policy	487
5.2.4.1 Cause	487
5.2.4.2 Solution	487
5.2.5 The prompt is too long for this model	488
5.2.5.1 Cause	488
5.2.5.2 Solution	488
5.2.6 The request took too long to complete	489
5.2.6.1 Cause	489
5.2.6.2 Solution	489
5.2.7 The requested model is not available on this gateway	490
5.2.7.1 Cause	490
5.2.7.2 Solution	490
5.2.8 This gateway requires EU model hosting	491
5.2.8.1 Cause	491
5.2.8.2 Solution	491
5.2.9 The model provider returned an error	492
5.2.9.1 Cause	492
5.2.9.2 Solution	492
5.2.10 Frequently asked questions	493
5.2.10.1 Which budget is used up?	493
5.2.10.2 How high is my rate limit?	493
5.2.10.3 Why are other models available to me than to other users?	493

5.2.10.4 What does EU model hosting mean?	493
5.3 Gateway token	494
5.3.1 Messages	495
5.3.2 Access denied for this gateway	496
5.3.2.1 Cause	496
5.3.2.2 Solution	496
5.3.3 Missing or invalid gateway token	497
5.3.3.1 Cause	497
5.3.3.2 Solution	497
5.3.4 Frequently asked questions	498
5.3.4.1 Where do I find my tokens?	498
5.3.4.2 Does a token apply to more than one gateway?	498
5.3.4.3 Does a token expire?	498
5.3.4.4 What do I do if a token gets into the wrong hands?	498
5.4 Your own objects	499
5.4.1 Messages	500
5.4.2 The agent could not be saved	501
5.4.2.1 Cause	501
5.4.2.2 Solution	501
5.4.3 The workflow could not be created	502
5.4.3.1 Cause	502
5.4.3.2 Solution	502
5.4.4 Changed in another tab	503
5.4.4.1 Cause	503
5.4.4.2 Solution	503
5.4.5 Duplicate step	504
5.4.5.1 Cause	504
5.4.5.2 Solution	504
5.4.6 Save the connector first, then test	505
5.4.6.1 Cause	505

5.4.6.2 Solution	505
5.4.7 No gateways available	506
5.4.7.1 Cause	506
5.4.7.2 Solution	506
5.4.8 Frequently asked questions	507
5.4.8.1 Why does the editor show Changed in another tab?	507
5.4.8.2 Why can an MCP connector only be tested after it is saved?	507
5.4.8.3 Who can edit my agents, workflows, and connectors?	507
5.5 Administration	508
5.5.1 Messages	509
5.5.2 Internal gateway error	510
5.5.2.1 Cause	510
5.5.2.2 Solution	510
5.5.3 Gateway configuration error	511
5.5.3.1 Cause	511
5.5.3.2 Solution	511
5.5.4 No API key is configured for the selected provider	512
5.5.4.1 Cause	512
5.5.4.2 Solution	512
5.5.5 The gateway couldn't authenticate with the model provider	514
5.5.5.1 Cause	514
5.5.5.2 Solution	514
5.5.6 You do not have the permission	516
5.5.6.1 Cause	516
5.5.6.2 Solution	516
5.5.7 Another plan edit is in progress	517
5.5.7.1 Cause	517
5.5.7.2 Solution	517
5.5.8 Frequently asked questions	518
5.5.8.1 What is a gateway configuration error?	518
5.5.8.2 Where do I store the API key of a provider?	518

5.5.8.3 Why does Myra AI Workspace show that I do not have a permission?	518
5.5.8.4 Why can a plan not be edited at the moment?	518
5.6 Interface	519
5.6.1 Something went wrong	520
5.6.1.1 Cause	520
5.6.1.2 Solution	520
5.6.2 Frequently asked questions	521
5.6.2.1 The message stays after a reload – what now?	521
5.6.2.2 How do I report an unsuitable response of the AI?	521
5.6.2.3 Why do I not see a view that this manual describes?	521
6. Reference	522
6.1 Notational conventions	522

1. Product description

Myra AI Workspace is a secured AI working environment. It combines the security and the control of an AI gateway with the tools of an AI platform for daily work, so that every person in a company can work with AI without confidential information leaving the company.

Below the interface, the product is a multi-tenant reverse proxy on the global Myra Security CDN. It sits between the customer's applications and the programming interfaces of the AI providers, enforces the policies there, manages the credentials, and records every request. Policy enforcement – authentication, rate limiting, cache, and content inspection – runs in the same process as the request. No additional network hop occurs, and the latency stays consistently low even under high load.

The product is used through two access paths that pass through the same processing and are subject to the same policies:

- web interface for people who work with AI every day
- OpenAI-compatible programming interface for software that calls AI models

21 AI providers are connected behind one common interface, so that changing a provider or adding a second one requires no change to the customer's code.

A gateway is the named installation of a tenant. It sets which providers and models are available, which rules apply to the content, which budget is available, and what Myra AI Workspace logs. This keeps the use of the AI traceable, limitable, and bound to the requirements of the organization.

As a security-as-a-service solution, Myra AI Workspace requires no installation on the customer side when it is operated as a managed service. Myra Security sets up the instance and maintains it. The service runs on the certified EU infrastructure of Myra Security and addresses the following risks:

- **Uncontrolled outflow of personal information to providers outside the EU:** Personal, financial, or otherwise regulated information is detected and masked inside the EU, or converted into reversible placeholders, before a request leaves the certified infrastructure.
- **Prompt injection and jailbreak attempts:** A two-tier guardrail chain detects instruction overrides, role takeovers, and injected system messages, and evaluates requests and responses semantically across 14 categories.

- **Uncontrolled spending on AI:** Hard cost limits apply per access token, per tenant, and per gateway, and are enforced before a call goes out.
- **Disclosure of the providers' credentials:** The provider keys are stored in a vault encrypted with AES-256 and are never disclosed to the calling application.
- **Missing evidence:** Every request produces a structured log entry with identity, routing, token count, cost, latency per phase, and the verdict of every guardrail, sealed in the audit log by a chain of checksums.
- **Provider lock-in and provider outages:** Fallback chains, weighted load distribution, and a circuit breaker keep the traffic running when a provider fails or slows down.
- **Data leaving the EU through an unnoticed detour:** On a gateway with EU routing enforced, the endpoint actually observed is evaluated after the response, and an endpoint outside the EU is rejected before a single byte reaches the caller.

1.1 Functions

The following functions are available:

- **Unified provider interface:** Connects 21 AI providers – among them OpenAI, Anthropic, Google Gemini, Vertex AI, AWS Bedrock, Azure OpenAI, Mistral, Cohere, Cloudflare Workers AI, HuggingFace, and Myra – behind one OpenAI-compatible endpoint. Every model name goes to `/compat/chat/completions`, and the gateway derives the provider from the model name itself.
- **Multi-tenancy:** Divides the platform into tenants and gateways. Every gateway is a self-contained policy scope with its own provider keys, its own authentication tokens, its own rate limits, and its own routing rules, separated at the storage layer and addressed through the URL path `/v1/{tenant}/{gateway}/...`.
- **Two-tier guardrails:** Enforces eleven types of detectors in two tiers. Tier 1 runs in-process below one millisecond and comprises jailbreak detection with 18 prepared attack phrases, keyword matching, and named pattern sets of regular expressions for card numbers, credentials, and information under HIPAA and GDPR. Tier 2 runs as a separate service in the certified Myra infrastructure and comprises NLP PII detection across more than 50 types of personal information, prompt guard across 14 semantic categories, and reversible PII protection.
- **PII protection:** Converts personal information into reversible placeholders before the request leaves the EU, and restores the original values in the response. The model therefore processes no real information and still answers coherently. A tenant blocklist and a personal list per account extend the detection.
- **EU data residency:** Rejects every path that would leave the EU before the request reaches a provider, evaluates the endpoint actually observed after the response, and records a residency zone (`eu`, `non_eu`, `unknown`) per request segment as evidence for data localization reporting.
- **Budgets:** Enforces hard cost limits per access token, per tenant, and per gateway, with automatic reset daily at midnight UTC, monthly, or never. An exhausted budget is answered with status `429` and a message that names the exhausted level.
- **Rate limiting:** Enforces a sliding-window limit per gateway and per token before a call goes out, and returns the headers `X-RateLimit-Limit`, `X-RateLimit-Remaining`, and `Retry-After`.

- **Routing and failover:** Evaluates an ordered rule chain that rewrites provider and model, splits traffic by weights, and appends fallback chains. A failed request is retried on a further provider, and the circuit breaker automatically removes a persistently disturbed provider from the routing.
- **Key vault (BYOK):** Stores the API keys of the providers encrypted with AES-256 and decrypts them only for the associated gateway at request time. The calling application authenticates with its own token and never sees a provider key.
- **Response cache:** Answers a verbatim repetition from the cache and marks it with the header `X-AIG-Cache: HIT`. This lowers cost and latency.
- **Chat and daily work:** Provides a persistent multi-turn conversation with model selection, attachments, read-aloud and dictation through a voice operated in the EU, reusable commands, a prompt library for the whole installation, web search with sources, and a ghost mode that stores nothing and logs nothing.
- **Projects and knowledge:** Groups conversations under shared instructions, a knowledge base, and a fixed model. Knowledge files, scanned PDF files, and images are evaluated by text recognition, and connected sources such as SharePoint and Confluence are synchronized including their permissions.
- **Agents and automation:** Provides reusable assistants made of system prompt, model, tools, and bound knowledge, multi-step workflows in the graphical editor, scheduled tasks with delivery into the chat, and access to an agent from Mattermost, Slack, or Microsoft Teams.
- **Code interpreter and generation:** Runs Python on an attached file in an isolated environment and returns charts and files, generates and edits images on a model operated in the EU, and creates artifacts that can be exported as an Office or OpenDocument file.
- **Governance and audit:** Maps regulatory obligations onto the settings through governance templates, lets a second administration approve a change to the prompt library or to an agent, seals every entry of the audit log with a chain of checksums, enforces retention and deletion periods per tenant, project, and account, and marks every response as AI-generated, as the EU Artificial Intelligence Act requires.
- **Observability:** Records a structured request log with identity, routing, status, cache state, token count, cost, latency per phase, and the verdict of every guardrail, and shows usage and spending by tenant, gateway, provider, model, agent, and account, together with a live monitor and a view of the providers' health. Counters and latency distributions are available at `/metrics`.
- **SIEM integration:** Sends one structured event per request to Splunk, Elasticsearch, Vector, or syslog.

- **Access management:** Signs in without a password through a six-digit one-time code by e-mail or through single sign-on with OIDC and SAML 2.0, takes over accounts and groups through SCIM, restricts access to defined address ranges through an IP allowlist, and assigns rights through the roles `admin`, `tenant_admin`, `ki_manager`, `member`, `finance`, `viewer`, and `demouser`.
- **Self-registration and billing:** Provides the **Starter** and **Pro** plans with a monthly or yearly billing period, a monthly message quota, invoices in the application, and cancellation through a public page without signing in.

1.2 Limitations

The following limitations apply:

- **Tier 2 guardrails default to fail open:** If the guardrail service of the certified Myra infrastructure is unreachable, the request passes as if the guardrail had not triggered. The setting must be switched off for every tier 2 check as soon as the processing is subject to regulatory requirements.
- **The IP allowlist accepts IPv4 ranges only:** IPv6 addresses cannot be entered, and the `ip_allowlist` field is maintained exclusively through the administration interface – it does not appear in the web interface.
- **A provider's own web search never counts as EU-clean:** A model that runs the search on its own server is blocked before the call only on a gateway with EU routing enforced. On a gateway without an enforced floor, the search axis of such a request is always `non_eu`.
- **The residency zone stays unknown for opaque paths:** OpenRouter, a cross-region Bedrock profile without evidence for the EU-27, a per-gateway base URL to an endpoint that cannot be identified, and an Azure path without a region are reported as not EU-clean.
- **The response cache covers non-streamed responses only:** A response is stored only with status 200, only if it is not transferred as a stream, and only if `cache_ttl` is greater than 0.
- **The viewer role cannot make inference requests:** The **Chat** and **Playground** views stay reachable but refuse the creation of tokens and the sending of messages.
- **The context-window estimate is approximate:** The gateway calculates with about 3.5 characters per token and a fixed amount per attachment, so that a request with many attachments can exceed the context window even when its visible text is short.
- **The value of a token is shown once:** The plain text appears only in the response to the creation and cannot be restored. A lost token must be deleted and created again.
- **Self-registration is a property of the respective installation:** If it is switched off, the registration pages are unavailable, and no reference to the registration appears on the sign-in page.



For internal information only. The following limitations are not intended to be passed on to customers.

- **The platform runs on a single server per environment:** Rate limiting, cache, and budget counters are held in the memory of the process. A distributed state backend is not connected, so operation across several servers would let the counters drift apart.
- **Encryption at rest is not attested:** For the database and the stored files, no evidence of encryption at rest exists. The provider keys are not affected by this, they are stored encrypted with AES-256 in every case.
- **The file store is switched off by default in the production environment:** Attachments and knowledge files become available only after the store has been explicitly switched on for the respective installation.

1.3 Target groups

Myra AI Workspace addresses organizations that use language models under regulated conditions. On the customer side, it is used by departments that use AI in daily work, development teams that connect applications to the programming interface, administrations that set up tenants, gateways, budgets, and policies, and the security, data protection, and compliance areas that keep the evidence of the processing. On the Myra Security side, sales and onboarding set up the instance, and support maintains it.

1.4 Functions of your role

The following options are available:

Function	Limitation
Hold conversations with the AI and create projects, agents, workflows, scheduled tasks, and MCP connectors	
Manage the prompt library, the governance rules, and the approvals of the tenant	
Manage the gateways, their models, routing rules, budgets, and API tokens	
Manage user accounts, groups, and roles	
Create tenants, manage their settings, and delete them permanently	This function stays reserved for the Administrator role. Your own role does not take up the Create tenants and Purge tenants permissions.
Analyse the cost by tenant, gateway, provider, model, and user, and manage the billing	
Manage the model prices, the plans, the vouchers, and the pool of system keys	These functions stay reserved for the Administrator role. Your own role does not take up the permissions of the Platform module.
Process the operational state, the audit log, the feedback, and the maintenance runs of the platform	These functions stay reserved for the Administrator role.

1.5 Permissions

The **Administrator** role has every permission of the catalogue. The permissions of the modules are available for selection in the editor of the **Roles & permissions** view when you put your own roles together.

The following options are available:

Module	Permission	Description
Roles	Manage roles within the tenant	Creates your own roles of the tenant and changes their permissions.
Roles	Manage system (Myra-curated) roles	Manages the system roles maintained by Myra.
Agents	Create and manage own agents	Creates and manages your own agents.
Agents	Review, approve, and moderate agents	Reviews, approves, and moderates agents.
Agents	Curate the tenant agent catalog	Curates the agent catalogue of the tenant.
Workflows	Create and run own workflows	Creates and starts your own workflows.
Workflows	View all workflow runs in the tenant	Shows all workflow runs of the tenant.
Workflows	Enable or disable workflows for the workspace	Switches the workflows for the workspace on or off.
Scheduled tasks	Create and manage own scheduled tasks	Creates and manages your own scheduled tasks.
Projects	Create projects	Creates projects.
	View model prices	Shows the model prices.

Module	Permission	Description
Model catalogue		
Connectors	Create and manage own MCP connectors	Creates and manages your own MCP connectors.
Connectors	Manage tenant connector-sync and chat-bridge	Manages the connector synchronization and the chat integrations of the tenant.
Groups	Manage groups and memberships	Manages groups and their memberships.
Governance	Manage governance rules, prompts, and approvals	Manages the governance rules, the prompts, and the approvals.
Analytics	View analytics dashboards	Shows the analyses.
Analytics	View cost breakdowns	Shows the cost breakdown.
Analytics	View per-user cost	Shows the cost per user.
Analytics	View cross-tenant cost	Shows the cost across all tenants.
Analytics	View request logs and traces	Shows the request logs and the related traces.
Users	Manage tenant users	Manages the users of the tenant.
Users	Manage users across all tenants	Manages the users across all tenants.
SSO	Manage SSO, SAML, and SCIM	Manages SSO, SAML, and SCIM.
Workspace	Manage tenant settings	Manages the settings of the tenant.
	Create tenants	Creates tenants.

Module	Permission	Description
Workspace		
Workspace	Purge tenants	Deletes tenants permanently.
Gateways	Manage gateways and routing	Manages the gateways and their routing.
Gateways	Manage gateway API tokens	Manages the API tokens of the gateways.
Approvals	Approve agent egress requests	Approves the outbound requests of agents.
Billing	Manage billing settings	Manages the billing settings.
Platform	Manage model prices and deprecations	Manages the model prices and the deprecations.
Platform	Manage plan and entitlement config	Manages the plans and their entitlements.
Platform	Manage vouchers	Manages the vouchers.
Platform	Manage the system key pool	Manages the pool of system keys.
Platform	Triage feedback, reports, and client errors	Triages the feedback, the reports, and the client errors.
Platform	View the platform health dashboard	Shows the operational state of the platform.
Platform	View the platform audit log	Shows the audit log of the platform.
Platform	Run maintenance and retention sweeps	Runs the maintenance and retention sweeps.



The permissions of the **Platform** module stay reserved for the **Administrator** role. The **Manage system (Myra-curated) roles**, **Manage users across all tenants**, **View cross-tenant cost**, **Create tenants**, and **Purge tenants** permissions stay reserved as well. Your own role does not take them up.

1.6 Technical description

This section describes the technical structure of Myra AI Workspace, the processing of a request, the interfaces, and the error handling.

1.6.1 Functional overview

Myra AI Workspace is a multi-tenant reverse proxy between the customer's applications and the programming interfaces of the AI providers. The structure knows two levels. A tenant is the top-level account – usually a company, an application, or a team – and a gateway is a named installation within a tenant, for example `production` or `staging`. Every gateway carries its own provider keys, its own authentication tokens, its own rate limits, and its own routing rules, fully separated from every other gateway.

Every inference request carries the tenant and the gateway in the URL path, either as `/v1/{tenant}/{gateway}/compat/chat/completions` for the unified endpoint or as `/v1/{tenant}/{gateway}/{provider}/chat/completions` for the native path of a provider. The gateway resolves both slugs per request and loads the matching gateway configuration. A change to the configuration takes effect within seconds.

Each request passes three phases. Each phase either stops and answers the request immediately, or it hands the request on to the next step:

- **Access phase:** Runs before the body of the request is read, which makes a rejection inexpensive. It covers the authentication against the stored token hash, the sliding window rate limit of the gateway and of the token, the quota check against the budgets of the token, the tenant, and the gateway, and the IP allowlist.
- **Content phase:** Runs with the complete body available. It covers the cache lookup, the two-tier guardrail chain on the outgoing request, the evaluation of the routing rules, the call to the provider, the guardrail chain on the incoming response, the cost calculation, and the cache store.
- **Logging phase:** Runs after the response has been sent, so an error here no longer affects the caller. It writes the structured log record and the metrics.

The routing rules are evaluated in their configured order, and the first matching rule wins. A rule overrides the provider, the model, and the fallback chain. If no rule matches, the provider and the model of the request itself apply:

- On a native endpoint both come from the URL path.
- On the unified endpoint the gateway derives the provider from the field `model` over the prefix of the model name, falling back to `OpenRouter`.

If the provider answers with an error of class `5xx`, the gateway retries the call up to `retry_count` times – the default value `2` yields three attempts in total – and then works through the fallback chain. Responses of class `4xx` are returned to the caller immediately and without a retry.

1.6.2 System architecture

The gateway runs as a single OpenResty process – nginx with LuaJIT – that carries the complete policy chain inside the worker itself. There is no sidecar and no additional network hop for authentication, rate limiting, caching, or the tier 1 guardrails. The product consists of the following components:

- **Gateway process (nginx/OpenResty, Lua):** The multi-tenant reverse proxy. It hooks into the three nginx phases access, content, and log, and runs an ordered middleware chain in which every step either answers the request immediately or enriches a per-request context object for the next step. One process serves many tenants and gateways with hard separation between them.
- **Web interface (React, TypeScript, Vite):** The surface for daily work and for administration. The navigation bar carries the groups **WORKSPACE** (Chat, Projects, Workflows, Agents, Library, Scheduled) and **ÜBERSICHT** (Dashboard, Recent chats, Analytics, Feedback), plus the three settings areas **EINSTELLUNGEN**, **BENUTZERVERWALTUNG**, and **KONTO**. Which entries appear depends on the permissions of the account. The navigation bar only hides entries. Every view and every interface checks the permission again on the server.
- **Configuration and log database (MySQL/MariaDB):** Holds the durable configuration – tenants, gateways, provider keys, authentication tokens, routing rules, users, roles, and model prices – and the structured request log. The configuration of a gateway is read from the database and held in shared memory for the time of `config_cache_ttl`, so a request needs no database read of its own.
- **Shared memory (nginx shared dictionaries):** The hot state of the process – response cache, sliding window counters of the rate limit, gateway configuration, decrypted provider keys with a lifetime of 60 seconds, and the counters for the metrics.
- **myrapod (LiteLLM proxy in front of the Myra GPU cluster):** Serves every request to the provider `myra` and hosts the tier 2 services – the Presidio analyzer and anonymizer for the detection of personal data, the Llama Guard 3 classification for Prompt Guard, and the document conversion. Every model operated by Myra is reached exclusively over this path.
- **Blob store:** A file store per owner on the host, into which attachments, knowledge files, uploads, and generated images are moved out of the database.
- **Supporting services:** An SMTP relay for the one-time codes of the login, the push services for the mobile clients, a search interface for the web search, an OpenTelemetry collector for the tracing, and the SIEM target of the customer.

- **Upstream providers:** The connected AI providers, called over HTTPS during the content phase.

The production environment is reachable at the following addresses:

Address	Meaning
ai.myra.eu	Web interface
ai-api.myra.eu	Inference interface under /v1/
ai-api-admin.myra.eu	Administration interface under /admin/
ai-docs.myra.eu	Documentation



For internal information only. The deployment is one server per environment. The complete hot state lives in the shared memory of that one process. The distributed state backend that a setup across several nodes would need is not implemented, and the gateway refuses to start with any other state backend. The integration environment runs on the same machine as the production environment – production over IPv6, integration over IPv4.

1.6.3 Technical details

Interfaces. The product exposes two separate interfaces. The inference interface under `/v1/` carries the requests to the AI providers and authenticates with an opaque token per gateway. The administration interface under `/admin/v1/` manages tenants, gateways, users, tokens, provider keys, routing rules, model prices, statistics, logs, and guardrail events, and authenticates with a browser session held in the signed JWT cookie `aig_admin`.

Authentication. The gateway accepts a token in three headers and takes the first one present:

- `x-aig-token`
- `Authorization: Bearer <token>`
- `x-api-key` for compatibility with the SDK of Anthropic

Through this order the gateway takes the place of the interface of OpenAI without any change to an existing client. A token is hashed with SHA-256 before it is stored, is valid for exactly one gateway, and carries an optional expiry, an own rate limit, an own budget, and a binding to an account. The role of that account decides what the token may do.

Roles. The platform knows the roles `admin`, `tenant_admin`, `ki_manager`, `member`, `finance`, `viewer`, and `demouser`. The role is enforced during authentication, before the body of the request is read. Deleting a user invalidates every authentication token of that user immediately.

Login. The web interface authenticates without a password over a six-digit one-time code sent by e-mail. After five failed attempts per address the check answers with `429`. A tenant additionally sets up OpenID Connect or SAML 2.0. Both methods sign in existing accounts only and never create one. A temporary fault never invalidates a session and never leads to a login – the check always fails to the safe side with `503`.

Request headers. Payload logging is switched off per gateway with `log_payloads: false` and per request with `x-aig-collect-log-payload: false`. The header `x-aig-collect-log: false` drops the log record entirely. The header `x-aig-byok-alias` selects a specific provider key. On an unknown alias the gateway never falls back to the default key.

Response headers. A cached response carries `X-AIG-Cache: HIT`. An error response carries the persistent code in `X-AIG-Error`. Every response carries `X-RateLimit-Limit` and `X-RateLimit-Remaining`. A rejection additionally carries `Retry-After`.

Data separation. The URL prefix `{tenant_slug}/{gateway_slug}` is resolved per request, and an unknown slug is answered with `404 tenant_not_found`. The whole internal state is named separately per tenant and gateway, all data is strictly assigned to the owning tenant on the storage layer, and rate limits and budgets are kept per gateway and never shared.

Streaming. With `"stream": true` the gateway forwards the chunks of the response as they arrive. If an error occurs after the stream has been opened, the HTTP status can no longer be changed. The response carries the status `200` on the wire, and the error arrives as an event in the stream. A streaming client therefore evaluates both paths – the HTTP status before the first event and the field `error` in the events after it. The log records the actual status, not the `200` of the wire.

Cryptography and key handling. A provider key is encrypted with AES-256-CBC and stored as a pair of initialization vector and ciphertext. The key is decrypted on first use and held in shared memory for 60 seconds, then injected into the call to the provider. The master key comes from the environment of the process. The session of the web interface is a stateless JSON web token with HMAC-SHA256, held in the cookie `aig_admin` with the attributes `HttpOnly` and `SameSite=Strict` and a lifetime of eight hours by default.



For internal information only. The derivation of the master key is global today. A derivation per tenant is planned but not implemented.

Rate limit and budget. The rate limit is a sliding window over two counters. The counter of the previous window is weighted by the share of the current window that has not yet elapsed and added to the counter of the current window. Costs are held as micro-dollars, checked before the request against the limit of the token, of the tenant, and of the gateway, and increased after the response of the provider.

Roles and permissions. A role is a set of permissions, and every gate resolves against a permission instead of a role name. The seven roles named above are held as system roles maintained by Myra. A tenant additionally defines its own roles from the same set of permissions. A role of a tenant never carries a permission of the platform class, which the data model itself enforces. Single sign-on assigns the default role `member` and never sets the platform role. Groups of the identity provider are mapped onto internal groups over an allowlist per tenant.

Response cache. Beside the exact match over a checksum of provider, model, and the canonical body of the request, a gateway optionally switches on a semantic cache. On a miss the prompt is embedded and compared with the stored embeddings of the same gateway and the same model. Above the configured similarity – `0.95` by default – the

stored response is returned with the headers `X-AIG-Cache: SEMANTIC_HIT` and `X-AIG-Similarity` . Responses transferred as a stream are not cached semantically.

1.6.4 Error handling

Every error response of the gateway carries an HTTP status and a persistent error code. The code names the cause precisely and does not change when the wording of the message is revised. A calling application therefore evaluates `error.code` and never the text `error.message`. The code arrives both in the JSON body and in the header `X-AIG-Error`. The web interface shows a translated message for every code. The errors fall into four groups.

1.6.4.1 Limits and billing

Error	Cause and remedy
429 rate_limited	The sliding window limit is exhausted. Wait for the period stated in <code>Retry-After</code> and raise the limit of the gateway or of the token.
429 quota_exceeded	A budget of the token, the tenant, or the gateway is reached. Raise the budget or reset the counter.
402 subscription_inactive	The subscription is not active. Reactivate the subscription or store a valid payment method.
402 trial_expired	The trial period has expired. Reactivate the subscription or store a valid payment method.

1.6.4.2 Access and security

Error	Cause and remedy
401 unauthorized	The token is missing, unknown, expired, revoked, or scoped to a different gateway. The wording of the message distinguishes the three cases.
403 forbidden	The source address is not in the IP allowlist, or the role does not permit the action.
403 data_residency_blocked	The selected model would run over a provider or a region outside the EU. Select an EU model or a gateway without the armed floor.

Error	Cause and remedy
400 guardrail_blocked	A guardrail rejected the content. Rephrase the content, or adjust the pattern or the action of the guardrail in case of a false positive.
403 processing_restricted	Processing is restricted for the account, for example after a request under Article 18 GDPR.

1.6.4.3 Request and model

Error	Cause and remedy
400 invalid_request	The body is not valid JSON, or a mandatory field is missing.
400 context_length_exceeded	The prompt exceeds the context window. Shorten the input, switch on context compaction, or move to a model with a larger context window.
413 context_overflow	The prompt exceeds the context window. The response carries <code>estimated</code> , <code>limit</code> , <code>suggested_model</code> , and <code>can_compact_now</code> .
413 request_too_large	No provider accepted the request, usually because of oversized attachments.
400 model_capability_mismatch	The selected model does not support the requested tools.
400 web_search_not_supported	The selected model does not support the web search.
404 tenant_not_found	The tenant or the gateway in the path does not exist.

1.6.4.4 Provider and operations

Error	Cause and remedy
502 provider_error	The provider answered with a server error after every retry.
502 all_providers_failed	Every provider of the routing chain failed. Check the request log, the status pages of the providers, and the keys of every provider in the chain.
424 provider_key_missing	No key is stored for the resolved provider on this gateway.
504 request_timeout	The request did not finish within the time budget. Simplify the request or switch off the web search.
503 service_starting	The gateway is starting up. Repeat the request after the time stated in <code>Retry-After</code> .
500 configuration_error	To be resolved by Myra Security. Support needs the request ID and the gateway slug from the request log.
500 internal_error	To be resolved by Myra Security. Support needs the request ID and the gateway slug from the request log.

Circuit breaker. If the circuit breaker is switched on and a provider reaches the configured failure threshold within the configured window, the gateway skips that provider and moves to the next entry of the chain. After the cool-down the circuit breaker allows one probe request. Only an attempt that actually returns a response closes it. The state per provider is read with `GET /admin/v1/gateways/{id}/circuit-breaker`.

Guardrail service unreachable. If a tier 2 service cannot be reached, the setting **Fail-open** of the affected guardrail decides the outcome:

- Switched on – the default – the request passes unchecked.

- Switched off, the request is blocked.

The request log shows whether a tier 2 verdict is missing on a request that passed.

1.6.5 Use cases

A public sector body and a regulated enterprise face the same conflict. The most capable models – OpenAI, Anthropic, Google Gemini, AWS Bedrock – are operated by companies from the United States, are subject to the law there, and stand in data centers there. Sending confidential information to those providers directly raises questions under the GDPR and under the requirements of the respective industry:

- Who accesses that information?
- Under which law?
- Under whose supervision?

At the same time, teams that are denied AI start using consumer services on their own, which is exactly the outflow the organization wanted to prevent.

Myra AI Workspace places the certified EU infrastructure of Myra Security between the organization and every one of those providers. All requests and responses run through the network of Myra, and a provider from the United States only ever receives what the policies of the customer explicitly permit. The services that inspect the content and detect personal data run inside that certified infrastructure, so confidential content is checked, masked, or replaced by placeholders inside the EU.

Three typical scenarios illustrate the use of the product:

- **The legal department of an insurer** works with contracts that contain personal data. The department uses a project of the access level **PII protection required**. Personal data is replaced by reversible placeholders before the request leaves the EU and restored in the answer. The department budget is capped at 2,000 USD per month.
- **A development team** integrates a customer-facing chatbot. The application authenticates with its own token against the gateway `production` and never sees a provider key. The token carries a budget of 50 USD per month and its own rate limit. A fallback chain and the circuit breaker keep the chatbot answering when a provider fails.
- **The compliance officer** answers the question of an auditor after a security incident: what did the model receive, and what did it answer? The request log holds the identity, the routing, the token counts, the cost, the latency per phase, the verdict of every guardrail, and the residency zone of every segment, sealed against change by a chain of checksums.

2. Getting started

2.1 Getting access

Myra AI Workspace does not use a password. You request a six-digit one-time code and receive it by email.

The following requirements are met:

- ☒ Your user account exists in Myra AI Workspace.
- ☒ You have access to the mailbox of the registered email address.

Proceed as follows to sign in to Myra AI Workspace:





Figure 1: Sign-in

- ▶ Click on the **Continue with Email code** button.
 - ↳ The prompt for the email address appears.



The image shows a sign-in form for Myra AI Workspace. At the top is the MYRA logo. Below it is the title "Sign in to Myra AI Workspace". Under the title is a label "Email address" followed by a text input field containing "admin@example.com". Below the input field is a checkbox labeled "Stay logged in for 30 days on this device". Below the checkbox is a large blue button labeled "Send code". At the bottom of the form is a link labeled "← Back".

Figure 2: Sign-in with the email address

- ▶ Enter the email address of your user account.
- ▶ If necessary, select **Stay logged in for 30 days on this device**.
- ▶ Click on the **Send code** button.
 - ↳ The prompt for the one-time code appears.




Figure 3: Sign-in with the one-time code

- ▶ Enter the code from the email.
- ▶ Click on the **Sign in** button.
- Myra AI Workspace starts with the navigation bar of your role.

The sign-in asks for two entries:

- The **Email address** text field takes the address of your user account.
- The **6-digit code** text field takes the one-time code that you received.

The **Stay logged in for 30 days on this device** checkbox keeps the session open on that device for 30 days.



If your tenant has an identity provider, the **Continue with SSO** button also appears. The **or sign in with a one-time code** hint stands below it. Sign-in through OIDC or SAML 2.0 only signs in existing user accounts. It does not create a new user account.



If a one-time code has expired, the sign-in shows the **That code is incorrect or has expired. Please request a new one.** message. In that case, click on the **Didn't get the code? Resend** link. The link becomes available again 30 seconds after the send. Until then it shows the remaining time.

After the first sign-in, Myra AI Workspace shows two dialogs that each appear only once:

- note about MYRA AI
- walkthrough of the interface

Which of the two appears first depends on the view that your role opens after the sign-in. The **Member** and **Viewer** roles land in the chat and see the walkthrough there first. The other roles land in the dashboard and acknowledge the note there first.

2.1.1 Acknowledging the note about MYRA AI

The **A note about MYRA AI** dialog explains that a language model generates the responses, how you report a response, and where your entries go. It appears when you first open a view outside the chat and cannot be bypassed. The **Esc** key and a click next to the dialog do not close it, and the dialog has no close button.

Proceed as follows to acknowledge the note:




Figure 4: A note about MYRA AI dialog

- Read the note.
- Click on the **Got it** button.
- The dialog closes and does not appear again on this device.

The dialog states three points:

Paragraph	Statement
Generated	MYRA AI uses large language models. Output can be inaccurate, incomplete, biased, or otherwise inappropriate. Treat it as a starting

Paragraph	Statement
responses	point, not a final answer – especially for medical, legal, financial, or safety-critical decisions.
Report of a response	You report an offensive, unsafe, or wrong response directly in the conversation with the Report action on the assistant message.
Processing of your entries	Your entries and all uploaded files go to the model provider that is configured for your gateway. The privacy policy link opens the policy in a new tab.



Myra AI Workspace stores the acknowledgement in the browser of the device you use, not in the user account. After the deletion of the browser data, and in a different browser or on a different device, the note appears again. It does not appear in the chat.

2.1.2 Walkthrough of the interface

A walkthrough leads through the most important views. It appears once per user account, for example in the chat or in the dashboard. While the note about MYRA AI is open, the walkthrough waits.




Figure 5: Walkthrough, first step

- ▶ Read the text of the step.
- ▶ Click on the **Next** button until the last step appears.
- ▶ Click on the **Get started** button.
- The walkthrough closes and does not appear again.

The walkthrough contains the following controls:

Element	Description
Try it →	Opens the view of the step and ends the walkthrough. The link is missing when you are already in that view.
Dots	

Element	Description
	Show the number of steps. The filled dot is the current step. A click on a dot goes directly to that step.
Step X of Y	States the current step and the number of steps.
Skip	Ends the walkthrough immediately.
Back	Goes to the previous step. The button is missing in the first step.
Next	Goes to the next step.
Get started	Ends the walkthrough. In the last step the button replaces the Next button.

Which steps the walkthrough shows depends on your role:

Step	Title	View	Roles
1	Welcome – a quick look around	Chat	all roles with a walkthrough
2	Organize with Projects	Projects	all roles with a walkthrough
3	Prompt Library	Prompts	Administrator and Tenant Administrator
4	You're in control of your data	My PII configuration	all roles with a walkthrough
5	Bring your team	Users	Administrator and Tenant Administrator

The **Administrator** and **Tenant Administrator** roles therefore go through five steps, the other roles through three.



Myra AI Workspace stores the completed walkthrough in the user account. The **Skip** and **Get started** buttons and the **Try it** → link complete it in the same way. There is no way to start it again. The **Demo User** role only reaches the chat and receives neither the note nor the walkthrough.

2.1.3 Signing out

The menu under your name at the bottom of the navigation bar contains the **Log out** entry.

Proceed as follows to log out of Myra AI Workspace:

- ▶ On the start page, click on your name at the bottom of the navigation bar.
 - ↳ The user menu opens.




Figure 6: User menu

- ▶ Click on the **Log out** entry.
- Myra AI Workspace ends the session and shows the sign-in.

2.2 Role concept

Myra AI Workspace has seven system roles. Each user account has exactly one of them. The role decides which entries the navigation bar shows and which requests the server accepts.

The permissions of a role come from a catalog of 38 permissions. Each permission belongs to a module and to a class. The class decides which roles get the permission. The navigation bar hides entries. Each view and each interface examines the permission again on the server.

The following information and options are available in a table:

Role	Permissi ons	Description
Administrator	38	Manages the platform and all tenants.
Tenant Administrator	25	Manages one tenant completely.
AI Manager	10	Controls the agents, the governance, and the groups of a tenant.
Member	5	Works with the AI every day and creates own objects.
Finance	3	Examines the costs and manages the billing settings.
Viewer	0	Examines the workspace without working in it.
Demo User	0	Runs a limited demonstration without a user account.

2.2.1 Assigning roles

The user who manages the users assigns a role. Which roles are available depends on the role of that user.

The following information is available:

Own role	Assigns the roles
Administrator	All seven roles.
Tenant Administrator	AI Manager, Member, Finance, and Viewer.
AI Manager	None.
Member	None.
Finance	None.
Viewer	None.
Demo User	None.

2.2.2 Custom roles

The **Administrator** and **Tenant Administrator** roles also compose custom roles. To do this, they select individual permissions from the catalog in the **Roles & permissions** view. Then they assign the completed role to users in the workspace.

The permissions of the platform are not available for selection. They stay reserved for the **Administrator** role.

2.3 Initial setup

After the installation, Myra AI Workspace is empty. The following objects do not exist yet:

- tenant
- gateway
- provider key

Only when a gateway carries at least one enabled model does Myra AI Workspace accept a request. The following sections take you through the initial setup in the sequence in which the steps build on each other.

The following requirements are met:

- ☒ Your user account has the **Manage gateways and routing** permission.
- ☒ Your user account has the **Manage tenants** permission if you create a tenant.
- ☒ An API key of the applicable provider is available, or you use a platform key from Myra.

The following steps are necessary:

Step	Result
1. <u>Create a tenant</u>	The tenant exists and carries a plan.
2. <u>Create a gateway</u>	The gateway has an endpoint URL.
3. <u>Enable a model</u>	The gateway forwards requests to the provider.
4. <u>Set a budget and a rate limit</u>	The spend and the load are capped.
5. <u>Set up the guardrails</u>	The gateway examines the content of the requests and the responses.
6. <u>Examine the roles</u>	The permissions agree with the structure of your organization.
7. <u>Create the user accounts</u>	The users get their invitation.
8. <u>Create the groups</u>	

Step	Result
	The access to agents, projects, and connectors is grouped.
9. <u>Issue a token</u>	Applications outside the interface reach the gateway.
10. <u>Send a test request and examine it</u>	The setup is demonstrably complete.

2.3.1 Step 1: Create a tenant

A tenant keeps the data, the users, and the billing of an organization separate from all others. Each gateway belongs to exactly one tenant.

Proceed as follows to create a tenant:





Figure 7: Tenants view

- ▶ In the **User Management** area, open the **Tenants** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Click on the **New Tenant** button.
 - ↳ The **New Tenant** dialog opens.



The image shows a 'New Tenant' dialog box with the following fields and options:

- Name ***: Text input with 'my-company'. Below it, a note says 'Lowercase letters, numbers, and hyphens only.'
- Plan**: Dropdown menu with 'Enterprise' selected.
- Tenant Budget (USD)**: Text input field. Below it, a note says 'Caps total spend across all of this tenant's gateways.'
- Budget Period**: Dropdown menu with 'Monthly' selected.
- Trial ends (demo accounts)**: Text input with '9/30/2026'. Below it, a note says 'Last day of the trial. After this the tenant is blocked from inference until upgraded. The trial budget is the Tenant Budget (USD) above – set the period to Total for a one-off trial cap. Leave empty for no trial.'
- Default user type**: Dropdown menu with 'Consumer' selected. Below it, a note says 'Chat persona for /easy – picks chip labels, example prompts, and system prompts.'
- Assistant name**: Text input with 'Myra AI'. Below it, a note says 'How the assistant identifies itself in chat (e.g. "Myra AI", "Bundestag Assistant"). Leave blank to use the platform default.'

At the bottom right are two buttons: 'Cancel' and 'Create Tenant'.

Figure 8: New Tenant dialog

- ▶ Enter the name, select the plan, and set the budget of the tenant.
- ▶ Click on the **Create Tenant** button.
- The tenant appears in the list.

The **Creating a tenant** section describes which data the dialog expects and how the data takes effect.


2.3.2 Step 2: Create a gateway

The following requirements are met:

- ☑ Step 1 is complete, see [Create a tenant](#).

The gateway is the endpoint to which the applications send their requests. It sets which models are available, which budget applies, and which examinations the content goes through.

Proceed as follows to create a gateway:



Gateways

Tenant: `muster-gmbh`

Search gateways by name...

New Gateway

Select Tenant


Filter tenants by name...

`muster-gmbh` `beispiel-ag` `probe-kg`

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	–	0s	–	–	7/19/2026

Figure 9: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- In the **Select Tenant** selection list, select the tenant from step 1.
- Click on the **New Gateway** button.



The 'New Gateway' dialog box contains the following fields and options:

- Name ***: A text input field containing 'prod'. Below it, a note states: 'Lowercase letters, numbers, and hyphens only.'
- Gateway Budget (USD)**: A dropdown menu set to 'unlimited'. Below it, a note states: 'Applies to this gateway only. The tenant's overall spend is capped separately by its own budget.'
- Cache TTL (s)**: A text input field containing '300'. Below it, a note states: '0 = disabled'.
- Retry Count**: A text input field containing '2'. Below it, a note states: 'Default 2 – a failed upstream call is retried up to twice. Set 0 to disable retries.'
- Timeout (ms)**: A text input field containing '120000'. Below it, a note states: 'How long to wait for the upstream before failing. The default suits most models; raise it only for very long generations.'
- Require auth token**: A checked checkbox.
- Guardrails (0)**: A section with several buttons: 'Regex / Pattern', 'Keyword', 'Jailbreak', 'Presidio (NLP)', 'Prompt Guard', 'Prompt Injection', 'PII Protector', and 'Custom PII Blacklist'.
- Below the guardrails, a note states: 'No guardrails configured. Add one above to start scanning requests or responses.'
- At the bottom, there are two buttons: 'Cancel' and 'Create Gateway'.

Figure 10: New Gateway dialog

- In the **Name** field, enter an identifier made of lowercase letters, numbers, and hyphens, for example **prod**.
- Set the **Require auth token** check box so that only applications with a token reach the gateway.
- Click on the **Create Gateway** button.
- The **Gateways** view shows the new gateway.

The **Creating a gateway** section describes which data the dialog expects and how the data takes effect.

2.3.3 Step 3: Enable a model

The following requirements are met:

- ☑ Step 2 is complete, see [Create a gateway](#).

Without an enabled model, the gateway refuses every request. Together with the model, you store the key with which Myra AI Workspace calls the provider.

Proceed as follows to enable a model:





Figure 11: Detail view of a gateway

- ▶ Open the gateway from step 2 with the **Open** → row action.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
- ▶ In the **Provider Keys** section, click on the **Add Model** button.



The 'Add Model' dialog box is centered on the screen. It has a title bar 'Add Model'. Below the title bar are two tabs: 'Use My Own Key' (selected) and 'Use Platform Key'. Below the tabs is the text 'Enable a provider's models on this gateway.' Below this text are two input fields: 'Provider' with a dropdown menu showing 'anthropic' and a downward arrow, and 'Alias' with a text box containing 'default'. Below these fields is the 'API Key *' section, which includes a text box with the placeholder 'Paste API key...' and an eye icon to its right. Below the text box is the text 'Stored encrypted. Replaces existing key for this provider/alias.' At the bottom right of the dialog are two buttons: 'Cancel' and 'Store Key'.

Figure 12: Add Model dialog

- ▶ If you have your own contract with the provider, select the method **Use My Own Key**, otherwise the method **Use Platform Key**.
- ▶ Select the **Provider**, enter the **API Key** for your own key, and click on the **Enable** button at the model you want.
- Myra AI Workspace shows **Key stored and encrypted successfully**.

Do the step again for each further model. The **Enabling a model on a gateway** section describes which data the dialog expects.



After this step, Myra AI Workspace accepts requests. Steps 4 to 9 are not mandatory for the operation, but a production system needs them.


2.3.4 Step 4: Set a budget and a rate limit

The following requirements are met:

- ☑ Step 3 is complete, see [Enable a model](#).

The budget caps the spend, the rate limit caps the load. Both prevent a single defect in an application from using up the cost or the quotas of the provider.

Proceed as follows to set a budget and a rate limit:



Edit Gateway: produktion

BUDGET & LIMITS

Gateway Budget (USD)	Budget Period
250	Month
Cache TTL (s)	Retry Count
300	2
Timeout (ms)	Rate Limit (req)
120000	600
Rate Window (s)	
60	

☒ Require auth token
☒ Log request/response payloads

PROVIDER BASE URLS


CIRCUIT BREAKER

Figure 13: Edit Gateway dialog

- In the detail view of the gateway, click on the **Edit** button.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- In the **Gateway Budget (USD)** field, enter the spend cap.

2.3.4 Step 4: Set a budget and a rate limit

- In the **Rate Limit (req)** and **Rate Window (s)** fields, enter the permitted number of the requests per time window.
- Click on the **Save changes** button.
- The **Budget** and **Rate Limit** statistics of the detail view show the new values.



Gateways

Gateway: production Edit Delete

BUDGET

€342.00 / €460.00 74% used

This month
= converted at today's rate; the cap is stored and enforced in USD.

CACHE TTL
5 m

RETRIES
2

TIMEOUT
60 s

AUTH
required

RATE LIMIT
60/60 s

LOG PAYLOADS
yes

WEBHOOK
all events

TRACING
on

WEB SEARCH
on

IMAGE GEN
off

PROMPT CACHING
on

COMPACTION
on

COMPACT ERR
32k tok

Endpoint URL
anthropic
https://ai-gateway.beispiel.de/v1/register-gmh/production/anthropic/chat/completions Copy

Provider Keys Add Model

PROVIDER	ALIAS	ADDED
anthropic	Editorial Anthropic	5/14/2026
openai	Analysis OpenAI	6/2/2026

Auth Tokens Generate

An authentication token is a secret key an application sends with each request to prove it may use this gateway – like a password for software. Create one per app or integration as you can rotate it individually.

LABEL	HASH (FIRST 16)	RATE LIMIT	SPEND CAP	EXPIRES	CREATED
Editorial system	a7b2c3d4e5f6d7f8...	60/60 s	€50.00	12/31/2026, 11:59:00 PM	4/2/2026
Analysis	9f8e7d6c5b4a3d21...	10/60 s	–	never	6/11/2026

Figure 14: Detail view of a gateway

The **Editing a gateway** section describes which data the dialog contains and how the data takes effect.

2.3.5 Step 5: Set up the guardrails

The following requirements are met:

- ☑ Step 4 is complete, see [Set a budget and a rate limit](#).

Guardrails examine the requests and the responses before they reach the provider or go back to the users. Without a guardrail, all content leaves the gateway unchanged.

Proceed as follows to set up the guardrails:




Figure 15: Guardrails section

- In the detail view of the gateway, open the **Guardrails** section.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- With the **Pii Protector** button, add a detector that removes personal data from the requests.

- ▶ With the **Jailbreak** button, add a detector that blocks attempts to get around the instructions.
- ▶ Click on the **Save Guardrails** button.
- The **Saved** entry appears, and the **Execution plan** section shows the sequence of the examinations.

The **Adding a guardrail detector** section describes which detector types are available and which data each type expects.



If more than one provider is available and you want to distribute the requests by conditions, create the applicable rules in the **Routing Rules** section. See the **Creating a routing rule** section.


2.3.6 Step 6: Examine the roles

The following requirements are met:

- ☑ Step 5 is complete, see [Set up the guardrails](#).

A role groups the permissions that a user account gets. Myra AI Workspace supplies prepared roles. Create your own role only if none of them agrees with your structure.

Proceed as follows to examine the roles:



ROLE	TYPE	PERMISSIONS	USERS	ACTIONS
Tenant Administrator	System	38	2	Edit
Member	System	12	18	Edit
Editorial	Custom	7	6	Edit

Figure 16: Roles & permissions view

- ▶ In the **User Management** area, open the **Roles & permissions** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Examine which permissions the available roles carry.
- ▶ If you need your own role, click on the **New role** button.
- The role is available for selection when you create a user account.

The **Creating a role** section describes which permissions are available.


2.3.7 Step 7: Create the user accounts

The following requirements are met:

- ☑ Step 6 is complete, see [Examine the roles](#).

Each person who works with Myra AI Workspace needs a user account of their own. The role assigned decides which views the person sees.


Proceed as follows to create a user account:



EMAIL	NAME	ROLE	TENANT	LAST LOGIN	CREATED
alex.muster@beispiel.de	Alex Muster	Tenant Administrator	muster-gmbh	8/16/2026	1/12/2026
kim.probe@beispiel.de	Kim Probe	Member	muster-gmbh	8/15/2026	3/4/2026
sam.beispiel@beispiel.de	Sam Beispiel	Viewer	beispiel-ag	never	6/21/2026

Figure 17: Users view

- ▶ In the **User Management** area, open the **Users** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Click on the **New User** button.



The image shows a 'New User' dialog box with the following fields and controls:

- Tenant ***: A dropdown menu with 'muster-gmbh' selected.
- Email ***: A text input field containing 'alice@example.com'. Below this field is the text: 'We'll email them an invitation to set up their account.'
- Name**: A text input field containing 'Alice'.
- Role**: A dropdown menu with 'Member' selected.
- At the bottom right, there are two buttons: 'Cancel' and 'Create User'.

Figure 18: New User dialog

- ▶ Enter the name and the email address.
- ▶ Select the tenant from step 1 and the role from step 6.
- ▶ Click on the **Create User** button.
- Myra AI Workspace sends the invitation to the email address entered.

The **Creating a user account** section describes which data the dialog expects. You change the role of an available account with the **Edit** row action. See the **Editing a user account** section.

2.3.8 Step 8: Create the groups

The following requirements are met:

- ☑ Step 7 is complete, see [Create the user accounts](#).

A group collects user accounts. With a group, you give the access to agents, projects, and connectors together instead of giving it to each person.

Proceed as follows to create a group:





Figure 19: Groups view

- ▶ In the **User Management** area, open the **Groups** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Click on the **New group** button.
 - ↳ The **New group** dialog opens.



The image shows a 'New group' dialog box. It has a title 'New group' at the top. Below the title, there are two input fields. The first is labeled 'Name' and contains the text 'e.g. Engineering'. The second is labeled 'Description' and contains the text 'Optional'. At the bottom right of the dialog box, there are two buttons: 'Cancel' and 'Create'.

Figure 20: New group dialog

- ▶ Enter the name and click on the **Create** button.
- ▶ Add the user accounts from step 7 to the group.
- The group is available for selection when you share agents, projects, and connectors.

The **Creating a group** section describes which data the dialog expects.


2.3.9 Step 9: Issue a token

The following requirements are met:

- ☑ Step 8 is complete, see [Create the groups](#).

An authentication token identifies an application to the gateway. Without a token, a gateway whose **Require auth token** check box is set refuses every request.

Proceed as follows to issue a token:



The image shows a 'New Token' dialog box with the following fields and options:

- Gateway ***: A dropdown menu with 'muster-gmbh/produktion' selected.
- Label**: A text input field containing 'dev laptop'.
- Expires At**: A date/time input field with the placeholder 'mm/dd/yyyy --:--'. Below it, a note says 'Leave blank = never expires'.
- Budget (USD)**: A text input field containing 'unlimited'.
- Rate limit (req/window)**: Two adjacent input fields. The first contains 'none' and the second contains '60s'. Below them, a note says 'requests / window_sec (blank = gateway default)'.
- At the bottom right, there are two buttons: 'Cancel' and 'Create Token'.

Figure 21: New Token dialog

- ▶ In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- ▶ In the **Tokens** section, click on the **New Token** button.
- ▶ Select the gateway from step 2 and enter a label that names the purpose, for example **ci-pipeline**.
- ▶ Click on the **Create Token** button.

→ Myra AI Workspace creates the token and shows it one time.




Figure 22: Token Created dialog

Copy the token immediately. It cannot be shown again later. See the **Creating a token** section.


2.3.10 Step 10: Send a test request and examine it

The following requirements are met:

- ☑ Step 9 is complete, see [Issue a token](#).

The initial setup is complete only when a request has demonstrably gone through the gateway.

Proceed as follows to send a test request and examine it:



Request Logs
3 entries

Export CSV Export XLSX Refresh

Tenant: All tenants ▾ Provider: All providers ▾ Model: e.g. gpt-4o Status: All ▾ Blocked: All ▾ Guardrail outcome: All ▾ Limit: 50 ▾

Since: mm/dd/yyyy Until: mm/dd/yyyy

TIME	PROVIDER	MODEL	STATUS	IN	OUT	CACHE	COST	LATENCY	UPSTREAM	GUARDRAIL	DETECTORS	FLAGS
4/2/2026, 9:41:02 AM	anthropic	claude-sonnet-4-5	200	1,240	310	–	\$0.03	381 ms	352 ms	29 ms	pii-protect	scrubbed >
4/2/2026, 9:40:55 AM	openai	gpt-4o-mini	200	620	148	80% hit • 12k	\$0.01	298 ms	271 ms	–	–	cached >
4/2/2026, 9:38:11 AM	anthropic	claude-sonnet-4-5	400	0	0	–	–	12 ms	–	8 ms	prompt_guard	blocked >

Figure 23: Request Logs view

- ▶ In the detail view of the gateway, copy the endpoint URL with the **Copy** button.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
 - ↳ The address has the form `https://<host>/v1/<tenant>/<gateway>/compat/chat/completions`. The selection list is set to **compat**. This endpoint finds the provider from the name of the model.
- ▶ Send a request to this address and give the token from step 9 in the Authorization: Bearer header.
- ▶ In the **Settings** area, open the **Request Logs** group.
- The request appears in the list with the provider, the model, the number of the tokens, the cost, and the latency.

The request has the following form:

```
curl -s -X POST \
  "https://<Host>/v1/<Mandant>/<Gateway>/compat/chat/completions" \
  -H 'Content-Type: application/json' \
  -H 'Authorization: Bearer <Token>' \
  -d '{
    "model": "gpt-4o",
    "messages": [{"role": "user", "content": "Hello!"}]
  }'
```

If the request does not appear, examine the message that the call returned. The sections of the **FAQ & Error messages** chapter name the causes.

3. Views

3.1 Home page and GUI layout




Figure 24: Layout of the interface

Myra AI Workspace shows the navigation bar at the left edge. It contains the header with the logo at the top (1), the entries of the current area below it (2), and your name with your role at the bottom (3). The content area with the selected view is next to it, on the right (4).

The following areas are available:

N o.	Area	Description
1	Header	Shows the logo and holds the Dark mode , Documentation , Notifications , Send feedback , and Collapse buttons.
2	Navigation area	Lists the entries of the current area, grouped by section.
3		Shows your name with your role and opens the user menu.

N o.	Area	Description
	User block	
4	Content area	Shows the selected view.

The application separates the **Workspace** area from the areas of the administration. The workspace contains the daily work with the AI. The areas of the administration contain the settings. You move between them through your name at the bottom of the navigation bar. The **Back to workspace** link at the top of the navigation bar goes from every area of the administration back to the workspace.

The navigation bar shows only the entries for which your role has the permission. An entry that this guide does not name does not appear for your role.




Figure 25: Navigation bar

The following options are available:

Entry	Description
Chat	Opens a conversation with the AI.
Projects	Holds related chats, files, and its own instructions together, so that the AI uses the same context for one field of work.
Workflows	Creates multi-stage workflows from your agents – with branches, approvals, and delivery.
Agents	Lists the agents saved with their instructions, model, tools, schedules, and run history.
Playground	Runs a prompt against one or more models and presents the answers side by side.
Library	Shows in the Prompts tab the central prompt templates of the organization with categories, search, and versioning.
Scheduled	Runs a prompt on a schedule. The results appear in a conversation and, if you want, by email as well.
Dashboard	Shows the real-time metrics of Myra AI Workspace.
Recent chats	Shows all conversations so far, the most recent first.
Analytics	Shows in the Cost Analytics , Live Monitor , and Health tabs the analyses of the cost and the operation.
Feedback	Shows in the Inbox , Sessions , and Reports tabs the feedback of the users.

Proceed as follows to open an administration area:

- On the start page, click on your name at the bottom of the navigation bar.
 - ↳ The user menu opens.




Figure 26: User menu

- Click on the entry of the area you want.
- The navigation bar shows the selected area.




Figure 27: Navigation bar in the Settings area

The **Settings** area is divided into groups.

The following options are available:

Group	Description
Approvals	Processes the open approvals for workflows, agents, tools, and configurations.
Preferences	Sets the language, the currency, the chat behaviour, the personal commands, and the masking of personal data.
Connections	Manages the MCP connectors and the chat integrations.
Request Logs	Shows the request logs and the related traces.
Gateways	Manages the gateways, their routing, and their API tokens.
Providers	Shows the AI providers supported with their configuration and their availability status in real time.
Security	Manages the governance rules of the workspace.
Costs	Shows the cost per user and the model prices.

The **User Management** area is divided into groups.

The following options are available:

Group	Description
Users	Manages the users of the tenant and their roles.
Groups	Manages groups and their memberships, to give access by group.
Roles & permissions	Puts your own roles together from a curated set of permissions and assigns them to the users in the workspace.
Tenants	Creates tenants and manages their settings.

You assign each of the seven roles:

- **Administrator**
- **Tenant Administrator**

- **AI Manager**
- **Member**
- **Finance**
- **Viewer**
- **Demo User**

The **Account** area is divided into groups.

The following options are available:

Group	Description
Profile	Shows your profile, your tokens, and the data export.
Plans	Manages the self-serve plans with their models, budget, degrade configuration, and Stripe price IDs.
Vouchers	Manages the codes that can be redeemed to give new registrations a one-time trial credit.

3.2 Workspace

3.2.1 Chat




Figure 28: Chat view

The **Chat** view opens a conversation with the AI. It greets you with your first name and the time of day, for example **Good morning,** , and carries the **Your privacy-friendly AI assistant – just ask in your own words.** subtitle below. As long as no name is available, the greeting is **How can I help today?**

The following options are available in the header:

Option	Description
New chat button	Starts a new conversation. The previous conversation remains available under Recent chats .
Ghost mode button	Keeps the next conversation private. Myra AI Workspace does not log it. The button appears before the first message only, because ghost mode belongs to the conversation and cannot be switched later. In the Demo user role it does not appear.
Memories button	Sets whether the conversation uses your saved memories. The button carries the state as on or off and appears before the first message only.

The following options are available in the input area:

Option	Description
Ask a question or describe a task... input field	Takes your input. In a running conversation the input field carries the text Write a message....
Add files, photos, or options button	Opens the menu with the attachments, the web search, the project assignment, and the answer styles.
Dictate button	Records your voice and writes what you say into the input field. While recording, the button carries the note Recording – tap to stop , afterwards Transcribing....
Privacy label	Shows how the next message is protected: Local model , Local + masking , Privacy on , or Privacy off . A click switches the masking.
Preview button	Shows what goes to the model after the masking. The button appears on a gateway with an active privacy filter only.
Choose a model button	Opens the Choose a model dialog. The button carries the name of the selected model, Auto by default. A dot in front of it shows the protection level of the route.
Send message button	Sends the request to the selected model. While a dictation is running, the button carries the label Stop & Send .
Stop generation button	Cancels the running answer. The button takes the place of Send message while the answer arrives.

Before the first message, the following tiles stand below the input field:

Tile	Description
Prompt Examples	Opens the prompt examples of the tenant and puts the selected example into the input. If none is stored, the tile shows No prompt examples yet.
Personal Commands	Opens your own slash commands and puts the selected command into the input. If none exists, the tile shows No personal commands yet.

Tile	Description
------	-------------

Shared Command	Opens the commands others have shared with you. If none exists, the tile shows No shared commands yet.
-----------------------	---

s



The first prompt examples are also available as buttons directly above the tiles. If the suggestions cannot be loaded, the view shows **Couldn't load suggestions. Try again in a moment.**



Whether the Enter key sends or starts a new line is set in the preferences. If the Enter key starts a new line, the input area carries the hint **Ctrl/Cmd+Enter to send • Enter for newline.**



You attach 10 files per message at most, each file 104 MB at most. Myra AI Workspace reports "**<name> is too large to attach (limit: MB).**, You can **attach up to files per message – remove one first.**, or **Couldn't read <name> – try a different file.** You also drag a file onto the view with the mouse.



If no gateway is set up for your tenant yet, the view shows **Chat isn't available yet – your admin needs to set up a gateway for your tenant.** With the **Manage gateways** permission it shows **No gateway is configured yet. Create one to enable chat for everyone on this tenant.** instead and offers the **Create a gateway** link.

The following sections describe the view in detail:

- **Add files, photos, or options menu**
- **Choose a model dialog**
- **Actions on a response**
- **Report this response dialog**

3.2.1.1 Add files, photos, or options menu




Figure 29: Add files, photos, or options menu

The menu collects everything you give a message in addition.

The following options are available:

Option	Description
Add files or photos entry	Opens the file selection and attaches the selected files to the message.
Web search entry	Lets the AI search the web in addition.
Recency drop-down list	Limits the web search in time: Any time , Past 24 hours , Past week , Past month , or Past year .
Add to project entry	Opens the Select project list and assigns the conversation to a project.
New project name input field	Creates a new project with the Create button and assigns the conversation to that project.
Use style entry	Opens the Select style list with the answer style of the AI.

Option	Description
Back entry	Leads from a sub-list back into the main menu.

The following answer styles are available:

Style	Description
Personal style	Applies your profile writing style.
Concise	Short, direct answers.
Explanatory	Step-by-step walkthroughs.
Formal	Professional register, no contractions.
Casual	Friendly, conversational tone.



If you assign the conversation to a project, Myra AI Workspace asks **Move this conversation to ?** When you take it out, Myra AI Workspace asks **Remove this conversation from ?**. If the project prescribes its own model, the entry carries the note **uses** and the confirmation adds **Switching will change the model from to .**



If no project exists yet, the list shows **No projects yet**. While loading it shows **Loading projects...**, on an error **Couldn't load projects**.

3.2.1.2 Choose a model dialog




Figure 30: Choose a model dialog

The dialog lists the models available on your gateway, grouped by provider.

The following information is available per model:

Information	Description
Provider	Heading of the group, for example Myra , Anthropic , OpenAI , or Other .
Name	Name of the model and, below it, a short description of its strengths.
Processing location	Icon in front of the name with the note Hosted by Myra (Germany) , Data processed in the EU , Data processed in the US , Data processed in China , or Processing region varies .
Included	Label for a model provided by Myra, with the note Provided by Myra – usage is billed to your Myra budget (no own API key needed) .

The following options are available:

Option	Description
Search models... input field	Filters the list by the name of the model. If no model matches, the dialog shows No matches for "" .
Auto entry	Leaves the choice to Myra AI Workspace and carries the description Pick the right model for the task . This is the default value.
entry	Binds the conversation to this model. The selected model carries a check mark.
Adaptive thinking toggle	Lets the model think longer about complex tasks and carries the description Think harder for complex tasks .



While loading, the dialog shows **Loading models...**, on an error **Couldn't load models**, and with no model available **No models available**.



If a model does not support adaptive thinking, the toggle carries the note **This model does not support adaptive thinking**. If a model always reasons, it carries the note **This model always reasons and can't turn it off**.



As soon as the first message is sent, the model is locked for that conversation. The button then carries the note **Model is locked for this chat (data security)**. With the **New chat from here** row action on a reply you switch the model. If a project prescribes a model, it carries the note **uses**. A slash command with a fixed model locks the **Auto** entry with the note **Not available for commands**.



On a gateway with the **EU-only** setting, the dialog shows the **EU-only gateway – models are limited to EU-hosted, approved providers**. hint and lists only models of EU-hosted providers.

3.2.1.3 Actions on a response




Figure 31: Actions on a response

Every response of the AI carries the **AI-generated** marker above the text and the **Message actions** toolbar below it. The marker stands on every response and cannot be switched off.

The following actions are available:

Action	Description
Copy	Copies the text of the response to the clipboard. The button briefly changes to Copied! . It is available during the response as soon as the first text exists.
Read aloud	Reads the response aloud. During the playback the button carries the Stop label.
Edit	Corrects the text of the response in the display. The changed response carries the edited marker and the Revert to original action. The correction does not go to the model. It only applies to the display.
Regenerate	Generates the response again with the same request. The action only appears on the last response of the conversation.
▼ next to Regenerate	Opens the Regenerate with reason menu.
Report	Opens the Report this response dialog.
Comment	Opens the Comments area for this message. You write a comment, mention other members, and set a comment to Resolve .

Action	Description
New chat from here	Branches a new conversation. Myra AI Workspace copies the conversation and the files up to this response. The model stays switchable in the new conversation.
Download ▾	Opens the menu with the file formats.

The **Regenerate with reason** menu contains the following entries:

Entry	Effect
Try again	Generates the response again without a further requirement.
Too long	Requests a shorter response. The new response carries the Adjusted • shorter marker.
Too short	Requests more detail. Marker Adjusted • more detailed .
Too formal	Requests a more casual, conversational tone. Marker Adjusted • more casual .
Too casual	Requests a more formal, professional tone. Marker Adjusted • more formal .
Off topic	Requests a response to what was actually asked. Marker Adjusted • on-topic .
Inaccurate	Requests a more careful and more precise response. Marker Adjusted • more accurate .

The **Download** menu offers the following formats in three groups:

Group	Formats
Finished file	Download as PDF
Microsoft Office	Download as Word , PowerPoint, Download as Excel
OpenDocument and text	OpenDocument Text, OpenDocument Presentation, OpenDocument Spreadsheet, TXT, Download as CSV



If your tenant has a letterhead template, the entry for Word additionally carries the **Word (using letterhead)** version. During the creation the entry states **Generating PDF....**



Which actions appear depends on the state of the response. During the response only **Copy** is available. **Regenerate** only appears on the last response, **Edit** only on a settled response of plain text without an image and without a generated file, and **Download** only if the response contains text or an image. On a response whose conversation exceeds the context window of the model, **Regenerate** does not appear.

3.2.1.4 Report this response dialog

The dialog reports a response to Myra Security GmbH for review. It carries the **Help us improve our models and content filters. Your report goes to Myra Security GmbH for review.** hint.

Proceed as follows to report a response:





Figure 32: Actions on a response

- ▶ On the response, click on the **Report** action.
 - ▷ On the start page, open the **Chat** view in the **Workspace** work area.
 - ↳ The **Report this response** dialog opens.



The dialog box is titled "Report this response" and contains the following elements:

- A subtitle: "Help us improve our models and content filters. Your report goes to Myra Security GmbH for review."
- A "Reason" section with a dropdown menu currently showing "Offensive or hateful content".
- A "Notes (optional)" section with a text area containing the placeholder text "Anything else we should know?".
- At the bottom, two buttons: "Cancel" and "Submit report".

Figure 33: Report this response dialog

- ▶ In the **Reason** select field, select the matching reason.
- ▶ If necessary, describe the case in the **Notes (optional)** text field.
- ▶ Click on the **Submit report** button.
- Myra AI Workspace states **Report submitted**.




Figure 34: Report submitted dialog

The following entries are available in the dialog:

Entry	Description
Reason	Type of the problem. Offensive or hateful content , Unsafe or dangerous advice , Inaccurate or misleading , and Other are available.
Notes (optional)	Free text with the Anything else we should know? placeholder. Describe here what is wrong with the response.
Cancel button	Closes the dialog without a report.
Submit report button	Sends the report. During the send the button carries the Sending... label.



If the report cannot be submitted, the dialog states **Failed to submit report** and stays open with your entries. The reported response appears for the administration in the **Reports** view.

3.2.2 Projects




Figure 35: Projects view

The **Projects** view bundles related chats, knowledge files, and custom instructions so the AI shares the same context across a body of work. The view lists the projects that you created and the projects that other users added you to.

The following information is available:

Column	Description
Project	Name of the project with its icon and its color. If a project has no name, the Untitled entry appears.
Description	Short description of the project. The column stays empty if no description was entered at creation.
Members	Number of the users who have access to the project.
Files	Number of the knowledge files of the project. The column counts only the files you can see. Some may be private to their creator until a chat is shared.
Role	Your role in this project: Owner , Editor , or Viewer .
Updated	Time of the last change to the project.

The following options are available:

Option	Description
New Project button	Opens the New Project dialog. See the Creating a project section.
Search projects... field	Limits the list to the projects whose name or description contains the text entered.
All filter	Shows all projects that you have access to.
Owner filter	Shows only the projects in which you have the Owner role.
Editor filter	Shows only the projects in which you have the Editor role.
Viewer filter	Shows only the projects in which you have the Viewer role.
Sort by list	Sorts the list by Recent Activity , Last edited , or Date created .
Open button	Opens the detail view of the project in the row.
Create your first project link	Appears only for an empty list and opens the New Project dialog.



If no project exists yet, the view shows **No projects yet**. If the search finds no project, the view shows **No projects match your search**.

The following sections describe the view in detail:

- [Detail view of a project](#)
- [Overview tab](#)
- [Files tab](#)
- [Members tab](#)
- [Conversations tab](#)
- [Feed tab](#)

3.2.2.1 Detail view of a project




Figure 36: Detail view of a project

The detail view lists the content of a project in tabs. If the project no longer exists, the view shows **Project not found**.

The following information is available:

Tab	Description
Overview	Shows the description, the key figures, the access tier, and the project instructions of the project.
Files	Lists the knowledge files that the model reads on demand. The counter of the tab counts only the files you can see. Some may be private to their creator until a chat is shared.
Members	Lists the users of the project with their project role.
Conversations	Lists the conversations that are assigned to the project.
Feed	Lists, under the Project Feed title, the conversations that project members shared.

The following options are available:

Option	Description
Back button	Returns to the list of the projects.
Start Chat button	Opens a new conversation in the context of the project.
Continue Last Chat button	Opens the conversation of the project used last.
Edit project button	Opens the form with the data of the project. See the Editing a project section.
Delete project button	Opens the Delete Project? dialog. See the Deleting a project section.
Share with organization button	Releases the project in the Members tab to all users of the organization. After the release, the button shows Shared with organization – unshare .

Next to the name of the project, the view shows your project role as a badge.



The **Edit project** button appears for the **Owner** and **Editor** project roles and for the **Administrator** role. The **Delete project** button, the administration of the members, and the **Conversation retention (days)** field appear only for the **Owner** project role and for the **Administrator** role.



If your organization switched off sharing with the whole organization, the view shows **Sharing with the whole organization is switched off for your organization. Ask an administrator to enable it**.

3.2.2.2 Overview tab




Figure 37: Overview tab

The **Overview** tab summarizes the data of the project.

The following information is available:

Item	Description
Description	Short description of the project. The item appears only if a description is entered.
Access tier	Access tier of the project: User configurable , PII protection required , or Local only .
Members	Number of the members. A click on the tile opens the Members tab.
Files	Number of the knowledge files you can see. A click on the tile opens the Files tab.
Conversations	Number of the conversations of the project. A click on the tile opens the Conversations tab.
Project Instructions	Instructions that apply to every conversation of the project. The item appears only if instructions are entered.

3.2.2.3 Files tab




Figure 38: Files tab

The **Files** tab contains the files available for the model to read on demand:

- plain text up to 5 MB
- PDF, DOCX, XLSX, and PPTX up to 20 MB
- images (PNG, JPEG, WEBP – text extracted via OCR) up to 20 MB

The following information is available:

Column	Description
File	Name of the file. The Only you addition marks a file that only its creator can see until that person shares the related chat into the project.
Size	Size of the file.
Tokens (est.)	Estimated number of the tokens that the content of the file occupies in the context of the model.
Uploaded	Time of the upload. During the processing, the Processing... status appears. After an error, the Extraction failed status appears.

The following options are available:

Option	Description
Upload File button	Opens the file selection. As an alternative, drag the files onto the Drop files here, or click to browse area.
Download file icon	Downloads the file in its original form.
Copy reference URL icon	Copies the address that references the file in a conversation.
Move to folder icon	Opens the Move "<name>" dialog with the Folder field. An empty field moves the file to the top level. Use "/" to nest folders.
Remove file icon	Removes the file from the project after a confirmation.
Share this chat button	Shares the chat that generated the file into the project. Its files and messages become visible to project members.



If no file was uploaded yet, the tab shows **No knowledge files yet**. If a file exceeds the limit, the tab shows **NAME exceeds LIMIT limit**. If a file is already in the project, the tab shows **NAME is already in this project**.

3.2.2.4 Members tab




Figure 39: Members tab

The **Members** tab lists the users who have access to the project.

The following information is available:

Column	Description
Name	Name of the member.
Email	Email address of the member.
Role	Project role of the member: Owner , Editor , or Viewer .
Joined	Time at which the member was added to the project.

In addition, the **Groups with access** area lists the groups with the **Group**, **Members**, and **Role** columns.

The following options are available:

Option	Description
Invite Member button	Opens the Invite Member dialog with the Email address field, the Role list, and the Share my existing conversations with the project check box.

Option	Description
Edit role icon	Opens the Change role – NAME dialog and changes the project role of the member.
Remove button	Removes the member from the project after the Remove this member? confirmation.



A project must have at least one owner. If you change the role of the last owner, the dialog shows **A project must have at least one owner.** If the project lists no further users yet, the tab shows **No members yet.**



The **Share my existing conversations with the project** check box is cleared by default. If you select it, Myra AI Workspace adds all your current private conversations of this project to the **Feed** tab. The conversations are then visible to each member of the project, not only to the person that you invite. You can unshare each conversation later in the **Feed** tab. If the application adds the member but the sharing fails, it shows **Member added, but sharing your existing conversations failed. You can share them individually from the Feed.**

3.2.2.5 Conversations tab



Overview	Files	Members	Conversations	Feed
----------	-------	---------	----------------------	------

Conversations			New Conversation 
TITLE	MODEL	UPDATED	
Offer Muster GmbH	claude-opus-4-6	8/15/2026	
Check price tiers	claude-sonnet-4-6	8/14/2026	
Shorten scope of work	gpt-4.1	6/10/2026	

Figure 40: Conversations tab

The **Conversations** tab lists the conversations of the project.

The following information is available:

Column	Description
Title	Title of the conversation.
Model	Model that the conversation uses.
Updated	Time of the last message.

The following options are available:

Option	Description
New Conversation button	Starts a conversation in the context of the project.

 If the project lists no conversation yet, the tab shows **No conversations in this project yet.**

3.2.2.6 Feed tab




Figure 41: Feed tab

The **Feed** tab lists, under the **Project Feed** title, the conversations shared by project members.

The following information is available:

Column	Description
Title	Title of the conversation shared.

Column	Description
Model	Model that the conversation uses.
Shared by	Users who shared the conversation.
Shared at	Time of the release.



If no conversation was shared yet, the tab shows **No conversations shared yet. Open a conversation in this project and click the team icon to share it here.**

3.2.3 Workflows




Figure 42: Workflows view

The **Workflows** view builds multi-step workflows from your agents – with branches, approvals and delivery. The list shows the workflows of the selected access tier.

The following information is available:

Column	Description
Name	Name of the workflow.
Status	State of the workflow: Draft , Published or Disabled . Only a published workflow runs.
Updated	Time of the last change, as a distance from the present.

The following options are available:

Option	Description
New workflow button	Opens the New workflow dialog. The button is locked while no access tier is selected.
Access tier drop-down list	Selects the access tier whose workflows the list shows. The drop-down list appears only if more than one access tier exists. A hint about the selected tier is shown below it.
Open row action	Opens the builder of the workflow.
	Creates a copy of the workflow and opens its builder.

Option	Description
Duplicate row action	
Enable row action	Sets a disabled workflow to Published .
Disable row action	Sets the workflow to Disabled . Triggers no longer start a run.
Delete row action	Opens the Delete this workflow and all its runs? prompt.
Create example workflow button	Creates the example workflows and opens the first of them in the builder. The button appears only in the empty list.



If no workflow exists yet, the view shows **No workflows yet** with the text **Create your first workflow to chain agents, add conditions and deliver results**. If no access tier exists, the view shows this instead of the list. If the list cannot be loaded, the view shows **Could not load workflows**.



The **Workflows** entry appears only if your administration has enabled the workflows for the workspace. The permission for this is called **Enable or disable workflows for the workspace**.

The following sections describe the view in detail:

- [Builder of a workflow](#)
- [Workflow copilot dialog](#)
- [Versions dialog](#)
- [Runs of a workflow](#)

3.2.3.1 Builder of a workflow




Figure 43: Builder of a workflow

The builder assembles the sequence from steps. The chain of the steps is on the left, the settings of the selected step are on the right.

The following options are available in the toolbar:

Option	Description
All workflows link	Leads back to the Workflows view.
Name input field	Renames the workflow.
Run budget (EUR) input field	Limits the cost of a single run. If the field stays empty, no limit applies. The field then shows no cap .
Save-state marker	Shows Saved , Saving... , Unsaved changes , Changed in another tab with the Reload button, or Save failed with the Retry button.
Runs button	Opens the Runs view of the workflow.
AI copilot button	Opens the Workflow copilot dialog.
Versions button	Opens the Versions dialog.
Publish button	Publishes the draft as a new version. While the draft is published, the button shows the label Publishing...
Run now button	Starts a run. The button is locked while the workflow is not published.
Unpublished changes marker	Appears if the draft differs from the published version. The run still uses the last published version.

The **Add step** button inserts the following steps:

Step	Description
Trigger	Starts the workflow: manually, on a Schedule , through a Webhook or through a Form . The trigger is the first step of the chain.
Agent step	Runs one of your agents against an input.
Condition	Compares a value and runs the If true branch or the Otherwise branch.
Loop	Runs the For each item branch once per item of the selected list, up to ten items per run.
Approval	Holds the run until the named approver decides in the My Approvals view.
Deliver	Delivers the result through Mattermost , through a Webhook or by Email .

The following options are available on a step:

Option	Description
Add step button	Opens the Choose a step type selection and inserts the selected step at this position.
Delete step button	Removes the step. The Step deleted message offers the Undo button.
Insert a value button	Opens the Insert a value from an earlier step selection. If no earlier step produces a value yet, the selection shows No earlier step produces a value yet .
Test this step button	Runs an agent step on its own against sample input. The test run is not saved and does not appear in the run history.



If a step is incomplete, the builder shows the message **Fix steps before publishing** above the chain. Each entry jumps to the affected step. If the workflow was created with a newer builder, the view shows **Unsupported layout** and locks the editing.

3.2.3.2 Workflow copilot dialog




Figure 44: Workflow copilot dialog

The **Workflow copilot** dialog builds and checks the sequence with a model.

The following tabs are available:

Tab	Description
Generate	Describe a process in plain language and get a proposed draft workflow.
Fix	Repairs the current draft using its validation errors.
Explain	Summarizes the current workflow in plain language.

The following options are available in the dialog:

Option	Description
Describe the process input field	Takes the description of the process. On the Fix tab, the field carries the label Additional guidance (optional) .
Model drop-down list	Selects the model the copilot works with. If no model is available, the dialog shows No model is available to run the copilot .
Ask the copilot button	Starts the copilot. While the copilot works, the button shows the label Working....

3.2.3.3 Versions dialog




Figure 45: Versions dialog

The **Versions** dialog lists the published versions of the workflow.

The following information is available:

Column	Description
Version	Number of the published version. The current version carries the Current marker.

Column	Description
Published by	User account that published the version.
When	Time of the publication.
Runs	Number of the runs of this version.
Compare	Checkbox that selects the version for the comparison. You can select two versions. The selection of a third version replaces the older of the two selected ones.

The following options are available in the dialog:

Option	Description
Roll back row action	The row action is not available on the current version. It opens the Publish version as the new current version? Your current draft is not changed; future runs will use this version. prompt.
Compare button	Shows the differences between the two selected versions. If both versions agree, the dialog shows No configuration differences between these versions.



If no version is published yet, the dialog shows **No published versions yet.**
If a version uses a deleted agent or approver, the roll back shows **This version references an agent or approver that no longer exists.**

3.2.3.4 Runs of a workflow




Figure 46: Runs view of a workflow

The **Runs** view lists the runs of the workflow. A click on a row opens the detail view of the run with the history per step.

The following information is available:

Column	Description
Status	State of the run.
Reason	Reason of the failure. For a successful run, the column shows —.
Trigger	Trigger that started the run: Manual , Schedule , Webhook or Form .
Cost	Cost of the run in euro.
Started	Time of the start.
Duration	Run time of the run.

The following options are available:


Option	Description
Back to builder link	Leads back to the builder of the workflow.
Run now button	Starts a run. If the workflow is not published, the view shows Publish the workflow before running it.

Option	Description
Cancel row action	Ends a run that is still going. The row action appears only on a run that has not finished.



If no run exists yet, the view shows **No runs yet** with the text **Run this workflow to see its history, per-step cost and delivery here.**

3.2.4 Agents



Agents

Saved agents: instructions, model, tools, schedules and run history.

Agent catalog **New Agent +**

All Active Scheduled Draft

Search agents by name... ☐ Shared with me Any tool ▼

NAME	MODEL	TOOLS	USED	SCHEDULE	VERSION	UPDATED
Daily briefing tagesbriefing	claude-sonnet-4-5	web_search	128	in 3 hours	v3	8/15/2026
Minutes assistant protokoll-assistent	gpt-4o-mini	files	46	—	v1	8/11/2026
Repo research repo-recherche	mistral-large	fetch, mcp×2	12	in 4 days	v2	7/28/2026

Figure 47: Agents view

The **Agents** view lists the saved agents. The subtitle reads **Saved agents: instructions, model, tools, schedules and run history.**

The following sections describe the view in detail:

- [List of the agents](#)
- [Detail view of an agent](#)
- [Agent catalog dialog](#)

3.2.4.1 List of the agents




Figure 48: List of the agents

The following information is available:

Column	Description
Name	Name of the agent. The shared label stands for an agent that another person shared with you.
Model	Model the agent runs on. If no model is set, then no model appears there.
Tools	Tools the agent may use: Web search , Knowledge files , Fetch , and MCP connectors .
Used	Number of the previous invocations of the agent.
Schedule	Next run as well as the number of the active schedules in the active schedules form. If a run is due, then Due now appears there.
Version	Current version of the agent.
Updated	Time of the last change.

The following options are available:

Option	Description
New Agent button	Opens the New Agent dialog.
Buttons below Start from a template:	Open the New Agent dialog with the instructions of the selected template: Daily Brief , Meeting Notes , Web Researcher , PII-Safe Document Q&A , Repo Researcher , Docs Assistant , or GitHub Triage .
Access level selection list	Gateway whose agents the list holds. The hint below the selection list describes the selected level.
Agent catalog button	Opens the Agent catalog dialog.
All , Active , Scheduled , and Draft tabs	Restrict the list to the respective status.
Search agents by name... input field	Restricts the list to a search term in the name.
Shared with me check box	Restricts the list to the agents that other people shared with you.
Filter by tool selection list	Restricts the list to one tool: Web search , Knowledge files , Fetch , or MCP connectors . The Any tool entry removes the restriction.
Filter by tag selection list	Restricts the list to one tag. The Any tag entry removes the restriction.
Filter by usage selection list	Restricts the list to Used or Never used . The Any usage entry removes the restriction.
Sort by selection list	Order of the list: Newest , Recently updated , Name (A-Z) , or Most used .
Open row action	Opens the detail view of the agent.



If no agent exists yet, then the view reports **No agents here yet.** with the **Create your first agent** button and the **Name & instructions, Pick model & tools**, and **Run or schedule** steps. If no agent matches the filters, then the list reports **No agents match the current filters.** If the loading fails, then the view reports **Could not load agents.**

The following options are available in the header of the detail view:

Option	Description
← Back button	Returns to the list of the agents.
Edit button	Opens the Edit Agent dialog with the entries of the agent.
Share button	Opens the Share "" dialog. The Who can see this agent selection offers Private (only me) , Specific people , A group , and Everyone in the organization . The hint reads Your agent's instructions become visible to anyone you share with. The Save sharing button applies the selection.
Duplicate button	Opens the Duplicate Agent dialog and creates your own copy.
Delete button	Opens the Delete agent ""? prompt.

3.2.4.2 Detail view of an agent



Figure 49: Detail view of an agent

The detail view opens when you select the **Open** row action in the list. The header names the name of the agent and the access level in the **Access tier: ()** form.

The following tabs are available:

Tab	Description
Configuration	Shows Slug , Provider , Max tokens , Tools , Version , Instructions (system prompt) and, for an imported agent, Imported from . If a provider is pinned, then the entry reads Pinned: runs on , otherwise Chosen at run time . The Call this agent from your code area holds a code example with the Copy button and the Manage API tokens link.
Try it	Runs the agent with an input. The input field carries the Input label, and the Run button starts the run. Below the result, the view names the trace in the Trace: form.
Schedules	Lists the schedules with the Name , Cadence , Enabled , Delivery , and Next run columns.
Webhooks	Lists the webhook triggers with the Name , URL , Enabled , and Last fired columns. The introduction reads Let an external system trigger this agent by POSTing to a unique URL. Each trigger has its own secret; the run executes as the agent's owner.

Tab	Description
Runs	Lists the previous runs with the Status , When , HTTP , Version , and Delivery columns. The status is Running , Succeeded , or Failed .
Versions	Lists the versions with the Version , Name , Model , and Created columns. The running version carries the Current label.

The following options are available in the **Runs** tab:

Option	Description
Entry of a run	Opens the Run detail panel with the Result , Steps , Cost & tokens areas – with the Cost , Input tokens , Output tokens , and Cache tokens values – as well as Agent config at run time and the Delivery entry.



The delivery of a run is **Pending**, **Delivered**, **Delivery failed**, **Skipped (empty)**, **Blocked (PII)**, **Blocked (unsafe URL)**, **Blocked (not allowlisted)**, **Blocked (destination changed)**, **Awaiting approval**, **Delivering**, **Denied**, or **Approval expired**. If the PII policy blocks the delivery, then the view names the possible remedies under **What you can do:**. If the agent holds no model, then the view reports **This agent has no model configured and cannot run. Edit it and choose a model.**

3.2.4.3 Agent catalog dialog




Figure 50: Agent catalog dialog

The **Agent catalog** dialog lists the agents that were shared across your organization. The subtitle reads **Agents shared across your organization. Clone one to make your own editable copy.**

The following information is available:

Column	Description
Name	Name of the agent. The Featured label stands for a highlighted agent.
Model	Model the agent runs on.
Tools	Tools the agent may use.
Rating	Average rating in the () form. If no rating is available yet, then Not rated appears there.
Actions	Holds the actions of the row.

The following options are available in the dialog:

Option	Description
Search the catalog by name... input field	Restricts the list to a search term in the name.
Filter by category selection list	Restricts the list to one category. The All categories entry removes the restriction.
Icons of the stars	Rate the agent.



Cloning copies an agent's instructions, model and tools. Connectors and knowledge sources are not copied – add your own after cloning. If no agent has been shared organization-wide yet, then the dialog reports **No agents have been shared organization-wide yet.**

3.2.5 Playground




Figure 51: Playground view

The **Playground** view runs a prompt against one or more models and presents the answers side by side. The header names the **Tenant**, the **Gateway**, and the **token active** state.

The following information is available per model panel:

Entry	Description
Response	Answer of the model. While the run lasts, Running... appears there. Without a result, (empty response) appears there.
Latency, Input tokens, Output tokens, Cache write tokens, Cache read tokens, and Estimated cost	Metrics of the run.

The following options are available:

Option	Description
Add Model button	Adds another model panel. The selection list names the number of the models as in catalog and marks the runnable models with runnable .
System prompt button	Shows the input field of the system prompt. After that, the button carries the Hide system label. The restore default link resets the text.
Parameters button	Shows the Temperature slider. After that, the button carries the Hide params label.
Web Search button	Enables the web search for the run. After that, the button carries the Web Search ON label.
Message input field	Message that is sent to the models. The Tip: ⌘↵ / Ctrl↵ to run hint names the key combination.
► Run button	Starts the run. During the run, the button carries the Running... label.
Show raw button	Shows the answer unformatted. After that, the button carries the Show rendered label.
Remove panel icon	Removes the model panel.



If you open a link that the model produced, then the playground reports **You are leaving this tool** with the **This link was generated by an AI model. You are solely responsible for any site you choose to visit.** hint and the **Open link** → button. If no gateway is selected, then the view reports **No active gateway token. Select a gateway first.**

3.2.6 Library



Figure 52: Prompt Library (Organization) view

The **Prompt Library (Organization)** view contains the central prompt templates that the administration maintains and that the organization reuses in the chat. The templates carry categories, are searchable, and are versioned.

The following information is available:

Column	Description
Name	Name of the template.
Category	Category of the template. If no category is set, the column shows —.
Description	Short description of the template. If no description is set, the column shows —.
Template	Start of the prompt text. The column shortens long texts.
Version	Sequential number of the current version, for example v3 .

The following options are available:

Option	Description
New prompt button	Opens the New prompt dialog. The button is disabled while no organization is selected.
Organization drop-down list	Selects the organization whose library the list shows. The drop-down list appears for the Administrator role only. All other roles see the library of their own organization.

Option	Description
Search prompts... input field	Searches the name, the description, and the prompt text of the templates.
Filter by category drop-down list	Limits the list to one category. The drop-down list shows the categories of the listed templates and the All categories entry.
Use in chat row action	Opens the chat with the template in the input field. The template stays editable there and is not sent immediately.
Edit row action	Opens the Edit prompt dialog.
Versions row action	Opens the Versions – dialog.
Delete row action	Opens the Delete the prompt ""? confirmation.



If the library is empty, the view shows **No prompts in the library yet**. If a search finds nothing, the view shows **No matches for ""**. If no organization is selected yet, the view shows **Choose an organization first**.

3.2.6.1 Versions dialog




Figure 53: Versions dialog

The **Versions** – dialog shows the earlier versions of a template.

The following information is available:

Column	Description
Version	Sequential number of the version, for example v2 .
Name	Name of the template at the time of this version.
Template	Prompt text of this version.

The following options are available in the dialog:

Option	Description
Restore row action	Opens the Restore version v? This creates a new version with the restored content. confirmation. The earlier version stays available.

 If no earlier version exists, the dialog shows **No versions yet.**



Changes to the prompt library need the approval of a second person with the **Manage governance rules, prompts, and approvals** permission. The requesting person never approves their own change. If a change is submitted for approval, the view shows **Change submitted for approval.** and the list stays unchanged until the approval.

3.2.7 Scheduled




Figure 54: Scheduled tasks view

The **Scheduled tasks** view runs a prompt on a schedule. The results arrive in a conversation, and by email if you like. Each task is shown as its own area in the list.

The following information is available:

Entry	Description
Active scheduled es	Number of the tasks that currently run.
Next run	Time of the next run, as a distance from the present. If no run is due, the tile shows —.
Name of the task	Name you assigned when you created the task.
Schedule and model	Schedule of the task, for example Daily at 06:00, Monday, 06:00 or Every 2 h , followed by the model. If the task also delivers by email, the line carries the addition and email .
State marker	Scheduled before the first run, Last run OK after a successful run, Last run failed or Retrying after an error, and Paused for a task that is held.
Error message	For a failed run, the reason is shown below the schedule.

The following options are available:

Option	Description
New scheduled task button	Opens the New scheduled task dialog. The button is locked if no gateway is available or if the limit of your plan is reached.
Edit row action	Opens the Edit scheduled task dialog with the entries of the task.
Pause row action	Holds the task. The task then carries the Paused marker and no longer runs.
Resume row action	Resumes a paused task.
Delete row action	Opens the Delete the scheduled task ""? It will stop running and this cannot be undone. prompt.
× button	Closes the view and leads back to the conversation.



If no task exists yet, the view shows **No scheduled tasks yet. Create one to get a prompt answered on a schedule.** If the list cannot be loaded, the view shows **Scheduled tasks could not be loaded.**



Your plan limits the number of the active tasks. If the limit is reached, the view shows **Limit reached: at most active scheduled tasks on your plan. Pause or delete one to add another.**

3.3 Overview

3.3.1 Dashboard




Figure 55: Dashboard view

The **Dashboard** view shows the real-time metrics of Myra AI Workspace. The header carries the **Real-time Myra AI Workspace metrics** subtitle.

The following metric areas are available:

Metri c	Description
Requ ests	Number of the requests of the selected period. The subline names the cache hit rate as % cache hit rate and, if requests were limited, the rate limited addition. If no request is available, then the subline reads No requests yet .
Cost	

Metri c	Description
	Cost of the selected period. The subline names the saving as saved via cache . If no saving is available, then the subline reads No cache savings yet .
Guar drail Hits	Number of the hits of the content rules. The subline breaks the hits down as blocked • scrubbed • flagged . If no hit is available, then the subline reads No guardrail hits .

The following information is available:

Area	Description
Input Token Cachi ng Efficie ncy	Share of the input tokens that came from the cache, with the Served from cache (reads) , Processed fresh (cache writes) , and New tokens (not cached) bars. Below that, the area names the Cached input cost line with the (reads at 0.1x rate) addition, the Uncached input cost line with the (writes + new tokens) addition, and the Compaction saved line. The area appears only if tokens are in the cache.
Usage by Tenan t –	Usage per tenant with the Tenant , Requests , Input Tokens , Output Tokens , and Cost columns.
Top Model s –	Most used models with the Provider , Model , Requests , Cost , and Avg Latency columns.
Recen t Reque sts	Last requests with the Time , Tenant , Gateway , Provider , Model , Status , Tokens , Cost , Latency , and Flags columns. A flag is cached , blocked , or rate limited .
Recen t Guard rail Event s	Last hits of the content rules with the Time , Tenant , Gateway , Outcome , Detector , Reason , and Latency columns. You can expand a row, which then shows the Prompt sent to LLM (after PII substitution) and the LLM response (before PII token restoration) .

The following options are available:

Option	Description
Today, Yesterday, Last 7 days, Last hour, and Last minute tabs	Period of the evaluation. The period applies to the metric areas and to the Usage by Tenant and Top Models areas.
Export CSV button	Downloads the metrics of all periods and the model breakdown of the selected period as a CSV file.
Export XLSX button	Downloads the same data as an Excel file.

Above the view, a notice banner appears when a budget reaches its limit. The banner names the scope as **Tenant, Gateway, Token, Plan allowance**, or **Rate limit**, carries the **Manage token** link for a token, and the **Manage in gateway settings** link for a service token. The **Dismiss** button hides the banner.



In a new workspace, the view reports **No data yet** with the **This is a new workspace – your dashboard will fill in with real metrics as you start using the applications.** hint and the **Use Applications** button. If the loading fails, then the view reports **Couldn't load your usage stats.** and offers the **Retry** button. For the roles without an analytics permission, the view carries the **My Usage** title with the **Your personal API usage** subtitle and shows your personal usage only. The export buttons as well as the **Usage by Tenant, Top Models**, and **Recent Guardrail Events** areas do not appear then.

3.3.2 Recent chats




Figure 56: Recent chats view

The **Recent chats** view shows all your previous conversations, newest first. You continue a conversation or open its project.

The following information is available:

Column	Description
Conversation	Title of the conversation as a link. A click on the link opens the conversation. If the conversation does not carry a title yet, the column shows Untitled chat .
When	Time of the last change. On the same day the column shows Today , on the previous day Yesterday , within the last seven days d ago , and before that the date.
Model	Model the conversation last used. If no model is stored, the column shows — .
Project	Project the conversation belongs to. If the conversation belongs to no project, or the name of the project cannot be resolved, the column shows — .
Turns	Number of exchanges in the conversation. As long as the list does not supply the number, the column shows — .
Actions	Row actions of the conversation. The following table lists them.

The following options are available:

Option	Description
New chat button	Opens a new conversation. If you opened the view from a project, the button carries the label New project chat and the new conversation belongs to that project.
Search conversations... input field	Searches the conversations. In the Active view, Myra AI Workspace searches on the server from two characters on and therefore also in the content of the conversations. In the Archived view, the search filters the rows already loaded by title, model, and project.
All projects drop-down list	Shows the conversations of the selected project only. The drop-down list carries the projects of the loaded conversations.
Any model drop-down list	Shows the conversations of the selected model only. The drop-down list carries the models of the loaded conversations.
Status drop-down list	Switches between Active and Archived . The Archived view carries its own list with its own load-more. The Archive row action is replaced by Restore there.
Continue row action	Opens the conversation and continues it.
Share to project feed row action	Makes the conversation visible in the feed of its project. Myra AI Workspace reports Shared to project feed . The row action appears for a conversation that belongs to a project only, and not in the Archived view.
Remove from project feed row action	Takes the conversation back out of the feed of its project. Myra AI Workspace reports Removed from project feed . The row action takes the place of Share to project feed as soon as the conversation is shared.
Rename row action	Replaces the title in the row with an input field. The title holds 255 characters at most. The Enter key applies the title, the Esc key discards the change.
Archive row action	Takes the conversation out of the Active view. It is available in the Archived view afterwards.

Option	Description
Restore row action	Brings the archived conversation back into the Active view. The row action appears in the Archived view only.
Delete row action	Opens the Delete this conversation? This can't be undone. confirmation. After the confirmation, Myra AI Workspace removes the conversation from every list.
Load more button	Loads the next page of the list. The button appears only while further conversations are available. While loading, it carries the label Loading...



A dot behind the title marks a new scheduled result that you have not opened yet. The dot carries the note **New scheduled result** and disappears as soon as you open the conversation.



If no conversation exists yet, the view shows **No conversations yet.** In the **Archived** view it shows **No archived conversations.** If a search or a filter returns no result, the view shows **No matches.**



If the list cannot be loaded, the view shows **Couldn't load your chats.** and offers the **Retry** button. If the search fails, the view shows **Search failed. Please try again.** with the **Retry** button as well. If the load-more fails, the button carries the label **Couldn't load more – retry.**

3.3.3 Analytics




Figure 57: Cost Analytics view

The **Analytics** view shows the evaluations of the cost and the operation.

The following tabs are available:

Tab	Description
Cost Analytics	Breaks the spend down by tenant, gateway, provider, model, and user.
Live Monitor	Shows the running requests in real time.
Health	Shows the operational state across the failure surfaces of the last 24 hours. An empty cell means that this surface is healthy.

The following sections describe the view in detail:

- Cost Analytics tab
- Cost Analytics breakdowns
- Detail panel of a tenant
- Live Monitor tab
- Health tab

3.3.3.1 Cost Analytics tab




Figure 58: Cost Analytics tab

The **Cost Analytics** tab breaks the spend down by tenant, gateway, provider, model, channel, agent, and user. The header holds the selection list of the currency and the tabs of the period.

The following metric areas are available:

Metric	Description
--------	-------------

Cost of all requests of the selected period.

Metric	Description
Total Spend	
Cache Savings	Amount that the hits in the prompt cache saved.
Total Requests	Number of the requests of the selected period.
Error Rate	Share of the failed requests. The subline names the number as errors .
Top Spender	Tenant with the highest cost. The subline names its cost with the this period addition. If no data is available, then the area shows No data .
Budget Warnings	Number of the tenants that reached 80% of their budget. The subline reads tenants ≥ 80% of budget .

Below the metric areas, the tab shows the **30-Day Overview** chart with the cost and the requests per day, and the latency percentiles **p50**, **p95**, and **p99** with the **Latency percentiles (non-blocked)** label.

The following activity counters are available:

Counter	Description
Conversations	Number of the chat sessions of the period.
Messages	Number of the messages sent.
Active users	Number of the people with at least one request.
New agents	Number of the agents newly created.
New projects	Number of the projects newly created.
New documents	Number of the documents newly created.

The following options are available:

Option	Description
Currency selection list	Currency in which the view states the amounts.
Today, Last 7 days, and Last 30 days tabs	Period of the evaluation. The period applies to the metric areas, the activity counters, and all breakdowns.
Group selection list	Restricts the evaluation to one group. The All (workspace) entry removes the restriction.
Anonymized (aggregate only) check box	If the check box is selected, then the breakdown by user shows no email addresses, and the export contains no personal rows. The check box is selected by default.
Filter rows input field	Restricts the rows of the selected breakdown to a search term. The of display names the number of the hits, and the Clear button removes the filter.
Export CSV button	Downloads the rows of the selected breakdown as a CSV file.
Export XLSX button	Downloads the rows of the selected breakdown as an Excel file.

3.3.3.2 Cost Analytics breakdowns




Figure 59: Tenant breakdown

The following breakdowns are available:

Breakdown	Columns
By Tenant	Tenant, Requests, Cost, Share, Cache, Errors, Blocked, Avg Latency, and Budget. If the tenant holds no budget, then the Budget column shows the value unlimited .
By Gateway	Gateway, Tenant, Requests, Cost, Share, Cache, Errors, Blocked, and Avg Latency.
By Provider	Provider, Models, Requests, Cost, Share, and Avg Latency.
By Model	Model, Provider, Requests, Input tokens, Output tokens, Cost, Share, and Avg Latency.
By Channel	Channel, Requests, Input tokens, Output tokens, Cost, Errors, and Avg Latency. The channel is UI (chat), API key, or Agent.
By Agent	Agent, Invocations, Conversations, Distinct users, and Cost. Above the table, the breakdown shows the Calls over time chart.
By User	User, Requests, Cost, Share, Cache, Errors, Blocked, and Avg Latency.

The **Share** column presents the share of the row in the total cost of the breakdown as a bar.

3.3.3.3 Detail panel of a tenant




Figure 60: Detail panel of a tenant

The detail panel opens when you click a row in the **By Tenant** breakdown. The header names the tenant as well as its plan and its budget with the budget period.

The following information is available:

Entry	Description
Request s, Cost, and Avg Latency	Metrics of the tenant in the selected period.
Budget Utilization	Share of the budget consumed. The area appears only for a tenant with a budget.
Cost – Last 30 Days	Course of the cost of the last 30 days.

Entry	Description
Top Models	Most used models with the Model , Requests , Cost , and Avg Latency columns.
Monthly Spend	Cost per month with the Month and Spend columns.
Anthropic Usage (API)	Usage of the Anthropic API with the Input Tokens , Output Tokens , Cache Read , Web Searches , and Total Cost metrics as well as the Date , Model , Tier , Input , Output , Cache Read , and Cost table columns. The Authoritative label stands for numbers from the provider account, the Estimated label for calculated numbers.
Entry Last Synced:	Time of the last synchronization with the Anthropic API. Without a synchronization it states never .

The following options are available in the detail panel:

Option	Description
Close icon	Closes the detail panel.



If no usage data of the Anthropic API is available yet, then the area reports **No usage data yet. Syncs hourly once an admin key is configured.**

3.3.3.4 Live Monitor tab



Figure 61: Live Monitor tab

The **Live Monitor** tab shows the running requests in real time and refreshes itself. The header names the time of the last refresh as well as the state of the auto-refresh as **on** or **paused**.

The following information is available:

Area	Description
Requests/ min	Course of the requests per minute over the last samples. The now addition names the current value.
Last minute, Last hour, and Today	Metrics of the respective period: Requests, Cached, Blocked, Scrubbed, Flagged, Cost, Saved, Avg latency, Provider ms, Input tokens, and Output tokens.
Tenants – Today	Usage per tenant with the Tenant, Requests, In Tokens, Out Tokens, Cost, and Quota Left columns. If the tenant holds no budget, then the Quota Left column shows the value no limit .

Area	Description
Last 10 Requests	Last requests with the Time, Tenant, Provider, Model, Status, In, Out, Cost, Saved, Latency, and Prov ms columns. The blocked label stands for a blocked request, the cached label for a response from the cache.
Recent Guardrail Events	Last hits of the content rules.

The following options are available:

Option	Description
Tenant selection list	Restricts the display to one tenant. The All tenants entry removes the restriction.
Interval selection list	Interval of the auto-refresh: 1 s, 2 s, 3 s, 5 s, 10 s, or 30 s.
Pause button	Stops the auto-refresh. After that, the button carries the ► Resume label.
↺ Refresh button	Loads the values again immediately.



If no request is available yet, then the tab reports **No requests yet.**

3.3.3.5 Health tab




Figure 62: Health tab

The **Health** tab shows the operational state across the failure surfaces of the last 24 hours. An empty cell means that this surface is healthy. If no surface is affected, then the tab reports **All clear – every surface is green.**

The following information is available:

Area	Description
Synthetic probes	Results of the probe runs with the Probe, Pass, Fail, Consecutive fails, Last run, and Status columns. The status is ok, flaky, or down.
Request log non-2xx (24h, excluding 499)	Failed responses with the Status, Count, and Distinct tenants columns.
Client-aborted requests (24h, status 499)	Requests aborted by the client with the Count, Distinct tenants, and Distinct users columns.

Area	Description
Client errors phoned home (24h)	Errors reported by the application with the Message prefix and Count columns.
Provider health	State of the providers with the Provider , Status , Latency , Checked , and Message columns.
Open queues	The Open content reports and Scrub-preview failures (24h) counters.
Recent probe runs (60 most recent, 24h)	Last probe runs with the Time , Probe , Status , and Failure columns.



The tab is available to platform administrators only. For other roles the view reports **Platform admins only**. If no values are available, then the tab reports **No data**.



If no data is available for the selected period, then the view reports **No data for this period**. If the loading fails, then the view reports **Failed to load analytics**.

3.3.4 Feedback



Figure 63: Feedback Inbox view

The **Feedback** view shows the feedback of the users.

The following options are available:

Tab	Description
Inbox	Lists the feedback received.
Sessions	Lists the feedback on individual chat sessions.
Reports	Lists the reports that users submitted about inappropriate or inaccurate model output.



The tabs of this view need the **Triage feedback, reports, and client errors** permission of the **Platform** module. This permission stays reserved for the **Administrator** role. A role of your own does not take it.

The following sections describe the view in detail:

- [Inbox tab](#)
- [Edit feedback dialog](#)
- [Sessions tab](#)
- [Reports tab](#)

3.3.4.1 Inbox tab



Figure 64: Inbox tab

The tab carries this note: **Bug reports and feature suggestions submitted by users.** The heading carries the **new** label while unprocessed feedback exists.

The following information is available:

Column	Description
Date	Time at which the feedback came in.
Type	Bug, Feature, or Other.
Summary	One-line description that the person who reported entered.
User	Email address of the person who reported. Without an address, the column shows —.
Processed	Row action that sets the processing state.

The following options are available:

Option	Description
All, Bug, Feature, Other buttons	Limit the list to one type.
Row of the table	Expands the details of the feedback and collapses them again.

Option	Description
Mark as processed row action	Marks the feedback as processed. For feedback processed already, the row action shows Mark as unprocessed .
Edit row action	Opens the Edit feedback dialog. The row action is in the row expanded.
Delete row action	Opens the Delete the feedback entry ' This cannot be undone. > confirmation. The row action is in the row expanded.

The row expanded contains the following information:

Entry	Description
Descr iption	Detailed text of the feedback. Without a text, the entry shows No additional details .
Page:	Address of the page on which the feedback was written. The entry appears only if an address is available.
Screen shot:	Images that the person who reported attached. A click on an image enlarges it and another click reduces it again. If an image cannot be loaded, the row shows Could not load the image .



If no feedback exists, the tab shows **No feedback submissions yet**. If the load, the save, the deletion, or the marking fails, the tab shows **Failed to load feedback.**, **Failed to save the feedback entry.**, **Failed to delete the feedback entry.**, or **Failed to update the feedback entry**.

3.3.4.2 Edit feedback dialog




Figure 65: Edit feedback dialog

You change the following entries in the dialog:

Entry	Effect of the change
Type drop-down list	Changes the classification of the feedback: Bug , Feature , or Other . The Type column and the filter follow the change.
Summary input field	Changes the text in the Summary column, at most 255 characters. The entry is required.
Description input field	Changes the detailed text in the row expanded, at most 4,000 characters. An empty field removes the text.
Page URL input field	Changes the address behind the Page: entry, at most 1,024 characters. An empty field removes the address.

The following options are available in the dialog:

Option	Description
Save button	Applies the change. While the change is saved, the button has the Loading... label. The button is disabled while the Summary is empty.
Cancel button	Closes the dialog without applying the change.

 You do not change the images attached in the dialog. They come from the **Send Feedback** form with which the person who reported submitted the feedback.

3.3.4.3 Sessions tab




Figure 66: Sessions tab

The tab carries this note: **Star ratings submitted by users at the end of a chat session.** The heading carries the **new** label while unprocessed ratings exist.

The following information is available:

Colum n	Description
Date	Time at which the rating was submitted.

Column	Description
Project	Project of the conversation. Without a project, the column shows —.
User	Email address of the person who rated. Without an address, the column shows —.
Rating	Grade from 1 to 5, where 5 is the best grade. Myra AI Workspace highlights a high grade in green and a low grade in red.
Comment	Text of the rating, shown truncated. Without a text, the column shows —.
Actions	Row action that sets the processing state.

The following options are available:

Option	Description
All, Unprocessed, Processed buttons	Limit the list to one processing state.
Row of the table	Expands the complete comment and collapses it again. A row without a comment does not expand.
Mark as processed row action	Marks the rating as processed. For a rating processed already, the row action shows Mark as unprocessed .



If no rating exists, the tab shows **No session feedback yet**. If the load or the marking fails, the tab shows **Failed to load feedback**. or **Failed to update the feedback entry**.

3.3.4.4 Reports tab




Figure 67: Reports tab

The tab carries this note: **User-submitted reports of inappropriate or inaccurate model output.** The heading carries the **open** label while open reports exist.

The following information is available:

Column	Description
When	Time at which the report came in.
Report er	Email address of the person who reported. Without an address, the column shows —.
Reason	Offensive, Unsafe advice, Inaccurate, or Other.
Status	Open, Triaged, or Dismissed.
Preview	First 120 characters of the message reported. A longer text ends with three dots.

The following options are available:

Option	Description
All, Open, Triaged, Dismissed buttons	Limit the list to one status.
Row of the table	Expands the details of the report and collapses them again.
Mark triaged row action	

Option	Description
	Sets the status to Triaged . The row action does not appear for a report triaged already.
Dismiss row action	Sets the status to Dismissed . The row action does not appear for a report dismissed already.
Reopen row action	Sets the status to Open . The row action does not appear for a report that is open already.

The row expanded contains the following information:

Entry	Description
Reporter notes	Text that the person who reported entered. The entry appears only if a note is available.
Reported message	Complete text of the model output reported. Without a text, the entry shows (empty) .
ID:	ID of the report. After it, Message: and Conversation: show the IDs of the message and the conversation, if they are available.



If no report exists, the tab shows **No reports yet**. With a filter set, it shows **No reports with status ""**. If the load or the change fails, the tab shows **Failed to load reports**. or **Failed to update report**.

3.4 Settings

3.4.1 Approvals




Figure 68: My Approvals view

In the **Approvals** group, you decide on the runs, deliveries, and changes that wait for an approval. You approve, deny, or annotate each review individually.

The following options are available:

Tab	Description
Personal	Shows the workflow runs that wait for your decision.
Agents	Shows the scheduled agent runs whose delivery requires human approval before it is sent. Myra AI Workspace never delivers denied or expired holds.
Tool approvals	Shows the pending governance and moderation reviews. A tool review holds an unattended agent run at a gated tool. A flagged prompt or output is an interaction that a guardrail marked for review.
Configurations	Shows the changes to the prompt library and to the agents of the organisation that require an approval before they take effect.

The following sections describe the view in detail:

- [Personal tab](#)
- [Approve this workflow run? page](#)
- [Agents tab](#)
- [Tool approvals tab](#)
- [Approve review dialog and Deny review dialog](#)
- [Configurations tab](#)

3.4.1.1 Personal tab




Figure 69: Personal tab

The tab carries this note: **Workflow runs waiting for your decision.** It lists each waiting run as an entry of its own.

The following information is available for each entry:

Entry	Description
Name of the workflow	Workflow whose run waits for your decision. The name is in the first line of the entry.
Request ed by	Person who started the run. Without a name, the email address appears, otherwise —.

Entry	Description
Decide by	Time left until the approval expires, for example in 3 hours . The point in time appears as a hint when you point at the entry. Without a deadline, the entry shows — .
Cost so far	Costs that the run caused up to the approval step.
Step	Name of the approval step in the workflow. The technical key appears as a hint when you point at the entry.
Run	Short form of the run ID with a leading number sign, for example #a1b2c3d4 . The complete ID appears as a hint when you point at the entry.

The following option is available:

Option	Description
Entry of the list	Opens the Approve this workflow run? page.



If no approval waits for you, the view shows **You have no approvals waiting**. If the list cannot be loaded, the view shows **Could not load your approvals**.



The list contains only the runs in which you are the named approver. A run for which another person decides does not appear here.

3.4.1.2 Approve this workflow run? page

Back to approvals

Approve this workflow run?

Workflow

Angebot prüfen

Requested by

Kim Probe

Run reference

#a1b2c3d4

Approval step

Freigabe Preis

Cost so far

€0.42

Decide by

in 2 days

This run masked personal data; delivery stays blocked until it is cleared.

What was produced so far

Angebot entwerfen

Dear Ms Muster, we are pleased to offer you the following terms for the maintenance of your plant in 2026.

Preis prüfen

List price €4,800.00 • Discount 8% • Final price €4,416.00

► Technical details

Approve

Reject

Figure 70: Approve this workflow run? page

The page collects the information that you need for the decision.

The following information is available:

Entry	Description
Workflow	Name of the workflow from which the run comes.
Requested by	Person who started the run.
Run reference	Short form of the run ID, for example #a1b2c3d4. The complete ID appears as a hint when you point at the entry.
Approval step	Name of the step at which the run waits.
Cost so far	Costs that the run caused up to this step.
Decide by	Time left until the approval expires.

Entry	Description
What was produced so far	Output of the steps that ran already, each under the name of its step.
Technical details	Collapsible area with the complete state of the run in JSON format.

The following options are available:

Option	Description
Back to approvals link	Returns to the list without a decision.
Approve button	Continues the run and shows Approved – the run will continue . While the decision is sent, the button has the Submitting... label.
Reject button	Shows the Reason for rejection input field.
Reason for rejection input field	Takes the reason for the record, at most 1,000 characters. The entry is required.
Confirm rejection button	Ends the run and shows Rejected – the run was ended . The button is disabled while no reason is entered.
Cancel button	Discards the reason and hides the input field again.



If the run masked personal data, the page shows **This run masked personal data; delivery stays blocked until it is cleared**.




If no output is available yet, the page shows **No step output to show yet**. If the output cannot be prepared for reading, the page shows **No readable preview – see technical details below**.



If another approver decided already, Myra AI Workspace shows **This run was already decided.** and returns to the list. Further messages are **You are not the approver for this run., This approval is no longer available to you.,** and **Please enter a rejection reason (up to 1000 characters).**

3.4.1.3 Agents tab



The screenshot shows the 'Agents' tab with a sub-tab 'Agent Approvals'. It displays a table of scheduled agent runs requiring human approval. The table has columns: REQUESTED, AGENT, ACTION, STATUS, DELIVERY, and EXPIRES. There are three rows: one pending, one approved, and one denied. A filter bar at the top shows '1 pending'.

REQUESTED	AGENT	ACTION	STATUS	DELIVERY	EXPIRES
8:12 AM	Wochenbericht	Email to team@beispiel.de	Pending	–	in 21 h
yesterday	Marktlage	Webhook to ops.beispiel.de	Approved	Delivered	–
6/12/2026	Angebot	Mattermost to #vertrieb	Denied	–	–

Figure 71: Agents tab

The tab carries this note: **Scheduled agent runs whose delivery requires human approval before it is sent. Denied or expired holds are never delivered.**

The following information is available:

Column	Description
Request ed	Time at which the run requested the approval.
Agent	Name of the agent. Without a name, the technical key of the agent appears.
Action	Short description of the delivery, for example the dispatch by email, webhook, or Mattermost.
Status	Pending, Approved, Denied, Expired, or Completed.
Delivery	

Column	Description
	State of the delivery after the decision. While no state is available, the column shows —.
Expires	Time at which the approval expires. For an approval decided already, the column shows —.

The following options are available:

Option	Description
All, Pending, Approved, Denied, Expired, Completed buttons	Limit the list to one status.
Approve row action	Approves the delivery. The scheduler delivers it on its next run and re-checks each egress gate.
Deny row action	Denies the delivery. Myra AI Workspace never delivers it.



The row actions are available for a pending approval only. A pending approval that nobody decides within 24 hours expires and counts as denied.



The heading carries the **pending** label while pending approvals exist. If no approval exists, the tab shows **No delivery approvals yet**. With a filter set, it shows **No approvals with status ""**. If the load or the decision fails, the tab shows **Could not load approvals**. or **Could not record the decision**.

3.4.1.4 Tool approvals tab




Figure 72: Tool approvals tab

The tab carries this note: **Pending governance and moderation reviews. A tool review holds an unattended agent run at a gated tool; a flagged prompt or output is an interaction a guardrail marked for review. Approve, deny, or annotate each.**

The following information is available:

Column	Description
Requested	Time at which the review was requested.
Kind	Tool gate, Flagged prompt, or Flagged output.
Agent	Name of the agent. Without an agent, the column shows —.
Subject	For a tool gate, this column gives the names of the gated tools. For a flagged prompt or output, it gives the trace reference of the interaction.
Status	Pending, Approved, Denied, Expired, or Completed.
Note	Note that the deciding person left. Without a note, the column shows —.
Expires	Time at which the review expires. For a review decided already, the column shows —.

The following options are available:

Option	Description
All, Pending, Approved, Denied, Expired, Completed buttons	Limit the list to one status.
All kinds, Tool gate, Flagged prompt, Flagged output buttons	Limit the list to one kind.
Approve row action	Opens the Approve review dialog.
Deny row action	Opens the Deny review dialog.



For a flagged prompt or output, the list never shows the text of the interaction, only its trace reference. The reference carries this note: **Interaction trace reference – correlate it in Request Logs**. You see the content in the **Request Logs** view, which has its own access controls.



If no review exists, the tab shows **No reviews yet**. With a filter set, it shows **No reviews match the current filters**. If the load or the decision fails, the tab shows **Could not load reviews**. or **Could not record the decision**.

3.4.1.5 Approve review dialog and Deny review dialog




Figure 73: Approve review dialog

Both dialogs take the decision together with an annotation for the record.

The following information is available in the dialog:

Entry	Description
Subject of the review	For a tool gate Gated tool(s) , for a flagged interaction Flagged interaction: .

The following options are available in the dialog:


Option	Description
Annotation (optional) input field	Takes a note for the record, at most 4,096 characters. The note then appears in the Note column.
Approve button	Approves the review. For a tool gate, the agent run runs the tool on the next sweep.

Option	Description
Deny button	Denies the review. For a tool gate, the tool gets a denial result and the run continues.
Cancel button	Closes the dialog without a decision.



A pending tool gate that nobody decides within 24 hours expires. The agent run then ends failed.

3.4.1.6 Configurations tab



Config Approvals 2 pending

Changes to the organisation's prompt library and agents that require a second administrator's approval before they take effect. The requester can never approve their own change; agent changes are re-checked against current residency rules at approval time.

REQUESTED	CHANGE	ITEM	STATUS	EXPIRES
10:02 AM	Update - Prompt library	Angebotsbaustein	Pending	in 20 h
8:47 AM	Create - Agent	Wochenbericht	Pending	in 19 h
yesterday	Delete - Agent schedule	—	Applied	—

Figure 74: Configurations tab

The tab carries this note: **Changes to the organisation's prompt library and agents that require a second administrator's approval before they take effect. The requester can never approve their own change; agent changes are re-checked against current residency rules at approval time.**

The following information is available:

Column	Description
--------	-------------

Time at which the change was submitted.

Column	Description
Requested	
Change	Kind of the change and object affected, for example Update • Prompt library . As the change, Create , Update , Delete , and Restore appear. As the object, Prompt library , Agent , Agent schedule , or Project .
Item	Name of the entry affected. Without a name, the column shows —.
Status	Pending , Applying , Applied , Denied , Apply failed , or Expired .
Expires	Time at which the approval expires. For an approval decided already, the column shows —.

The following options are available:

Option	Description
All , Pending , Applying , Applied , Denied , Apply failed , Expired buttons	Limit the list to one status.
Approve row action	Applies the change. Myra AI Workspace applies it in one step and sets the status to Applied .
Deny row action	Discards the change. The entry stays unchanged.
Your request label	Marks a row that you submitted yourself. The Approve row action does not appear there.



Myra AI Workspace enforces the four-eyes principle server-side. The requester never approves their own change. The **Deny** row action stays available to them so that they can withdraw the request.



Myra AI Workspace re-checks an agent change against the current residency rules at approval time. If the change cannot be applied, the row carries the **Apply failed** status.



If no approval exists, the tab shows **No config approvals yet**. With a filter set, it shows **No approvals with status ""**. If the load or the decision fails, the tab shows **Could not load approvals.** or **Could not record the decision.**

3.4.2 Preferences




Figure 75: Preferences view

In the **Preferences** group you set how Myra AI Workspace behaves for you. The settings apply to your user account only. Below the heading, the **How the workspace behaves for you – language, currency, chat behaviour, and your default model.** note appears.

The following options are available:

Tab	Description
My preferences	Sets the language, the currency, the chat behaviour, and your default model.
My Commands	Shows the slash-commands available in chat – your own plus any shared with you.

Tab	Description
My PII configuration	Masks the words and phrases you list here in your messages before they leave Myra, in addition to the detectors for personal data configured at the gateway.

The following sections describe the view in detail:

- [My preferences tab](#)
- [My Commands tab](#)
- [Command – / dialog](#)
- [Version history – / dialog](#)
- [Share / dialog](#)
- [My PII configuration tab](#)

3.4.2.1 My preferences tab




Figure 76: My preferences tab

The tab carries the **General** and **Personal AI preferences** areas.

The following options are available in the **General** area:

Option	Description
Language drop-down list	Language of the interface: English or Deutsch . Myra AI Workspace changes the interface immediately and saves the choice in the user account, so that emails and notifications also follow this language.
Display currency drop-down list	Currency in which the app shows budgets and costs: USD \$ or EUR € . The drop-down list carries this note: Budgets and costs are shown in this currency across the app. Independent of your language, and a display setting only – amounts are stored and enforced in USD.
Press Enter to send messages checkbox	Behavior of the Enter key in chat. If the checkbox is selected, the Enter key sends the message. If it is not selected, the Enter key adds a line break and the Ctrl+Enter key combination sends.
Default model button	Opens the model picker. Without a choice, the button shows Auto (cheapest available) . The button carries this note: New chats start with this model unless a project or the message model picker overrides it. Falls back to the cheapest available model when unset or unavailable.

The following options are available in the **Personal AI preferences** area:

Option	Description
What should I call you? input field	Name that the AI uses for you, at most 64 characters.
What best describes your work? drop-down list	Field of work, for example Product Management , Engineering , or Legal . Without a choice, the drop-down list shows – Not set –.
How should the AI respond to you? input field	Permanent instructions for the style of the replies, at most 5,000 characters. Below the field, the / 5000 counter appears.

Option	Description
Save preferences button	Saves the three entries and shows Saved. New chats use these preferences. The button is disabled while no entry is changed or the instructions are too long.




The **General** area saves each entry immediately. The **Personal AI preferences** area saves only with the **Save preferences** button. The note of the area is: **These notes are added to every chat – they're not a question. Tell us about you or how you'd like replies written.**



From 4,000 characters, the counter adds the **long instructions may crowd out conversation context** note. From 5,000 characters, the area shows **Instructions exceed 5000 characters.** and does not save.

3.4.2.2 My Commands tab



The screenshot shows the 'My Commands' tab in the MYRA interface. At the top, there are three tabs: 'My preferences', 'My Commands' (selected), and 'My PII configuration'. Below the tabs, the title 'My Commands' is displayed, followed by the subtitle 'Slash-commands available in chat – your own plus any shared with you.'.

A paragraph explains: 'Commands are reusable slash-command templates (like /translate) that you type into the chat to insert a prepared prompt. They're different from the prompt examples on the chat start screen, which are one-tap suggestions for starting a conversation.'

Under the heading 'Personal commands', there is a 'New command' button. Below it, a note states: 'Your own slash-command shortcuts. Visible only to you; you can edit and delete them.' There is a search bar labeled 'Search commands...' and a dropdown menu labeled 'All categories'.

A table lists personal commands:

COMMAND	CATEGORY	DESCRIPTION	TEMPLATE	
/uebersetzen	Editorial	Translates a paragraph	Translate to {{language}}: {{text}}	
/zusammenfassen	Sales	Summarises a note	Summarise in five points: {{text}}	

Below the table, under the heading 'Shared commands', a note states: 'Slash commands defined for your tenant. Read-only here – ask a tenant admin to change them.'

A table lists shared commands:

COMMAND	DESCRIPTION	TEMPLATE	
/angebot	Writes a quote paragraph	Write a quote for {{customer}}	

Figure 77: My Commands tab

The tab carries this note: **Commands are reusable slash-command templates (like /translate) that you type into the chat to insert a prepared prompt.** It contains the **Personal commands** and **Shared commands** sections.

The following information is available in both sections:

Column	Description
Command	Name of the slash-command with a leading slash, for example /translate .
Category	Category of the command. Without a category, the column shows — . The column appears in the Personal commands section only.
Description	Short description that appears in the picker of the chat. Without a description, the column shows — .
Template	Prompt of the command. The column truncates long templates. The View row action shows them in full.

The following options are available:

Option	Description
New command button	Opens the New Command dialog.
Search commands... input field	Searches in the personal commands.
Filter by category drop-down list	Limits the list to one category. The drop-down list contains the categories of the shown commands and the All categories entry.
View row action	Opens the Command — / dialog. The row action is in each row of both sections.
Edit row action	Opens the Edit Command dialog. The row action is available for your own command only.
Versions row action	Opens the Version history — / dialog. The row action is available for your own command only.

Option	Description
Share row action	Opens the Share / dialog. The row action is available for your own command only.
Delete row action	Opens the Delete this command? confirmation. The row action is available for your own command only.
Manage tenant commands button	Opens the tenant view in which the shared commands are maintained. The button appears for the Tenant administrator role only.
Shared with you label	Marks a row that another person shared with you. Such a row carries the View row action only.



The **Personal commands** section carries this note: **Your own slash-command shortcuts. Visible only to you; you can edit and delete them.** The **Shared commands** section carries this note: **Slash commands defined for your tenant. Read-only here – ask a tenant admin to change them.**



If no own command exists yet, the section shows **No commands yet. Create one to use /commandname shortcuts in chat.** If the search finds nothing, the section shows **No matches for ""**. If no shared command exists, the section shows **No shared commands defined for this tenant yet.**

3.4.2.3 Command – / dialog




Figure 78: Command dialog

The dialog shows a command in full, including a shared one.

The following information is available in the dialog:

Entry	Description
Description	Short description of the command. Without a description, the entry shows —.
Template	Complete prompt of the command without truncation.
Variables detected:	Placeholders of the form {{variable}} that the chat asks for on each call. The entry appears only if the template contains a placeholder.

The following option is available in the dialog:

Option	Description
Close button	Closes the dialog.

3.4.2.4 Version history – / dialog



Figure 79: Version history dialog

The dialog lists the earlier versions of your own command.

The following information is available in the dialog:

Column	Description
Version	Number of the version in the form v .
Command	Name of the command at the time of this version.
Template	Prompt of this version, shown truncated.

The following option is available in the dialog:

Option	Description
Restore row action	Opens the Restore version v? This creates a new version with the restored content. confirmation. After the confirmation, Myra AI Workspace shows Version restored.



If no earlier version exists yet, the dialog shows **No version history yet.** The restore deletes no version. It creates a new version with the old content.

3.4.2.5 Share / dialog




Figure 80: Dialog for sharing a command

The dialog gives selected users access to your own command. It carries this note:
Share this prompt with selected users in your organization. They can use it read-only; only you can edit, restore, or delete it.


The following options are available in the dialog:

Option	Description
Add a user by email input field	Takes the email address of the person. The Add member button puts the person on the list.
Remove button	Takes the person off the list.
Save sharing button	Applies the list. While the list is saved, the button has the Saving... label.
Cancel button	Closes the dialog without changing the sharing.



If the command is not shared yet, the dialog shows **Not shared with anyone yet**. If the email address belongs to no user account, the dialog shows **User not found. They must have an account first**.

3.4.2.6 My PII configuration tab



My preferences My Commands **My PII configuration**

My PII configuration

Words and phrases you list here will be masked in your messages before they leave Myra, in addition to the gateway's pre-configured PII detectors.

Add a keyword

e.g. ProjectAtlas, ClientXYZ, ...

Add

☒ Case-sensitive
☒ Whole word only

ADD A KEYWORD	CASE-SENSITIVE	WHOLE WORD ONLY	
ProjectAtlas	✓	✓	✕
ClientXYZ	—	✓	✕

Figure 81: My PII configuration tab

The tab carries this note: **Words and phrases you list here will be masked in your messages before they leave Myra, in addition to the gateway's pre-configured PII detectors.**

The following options are available:

Option	Description
Add a keyword input field	Takes the word or the phrase, at most 200 characters.
Case-sensitive checkbox	Masks the keyword only in the spelling that you entered.

Option	Description
Whole word only checkbox	Masks the keyword only as an independent word, not as part of a longer word.
Add button	Puts the keyword on the list.
Add a list input field	Takes up to 20 words, one word per line. The field appears for the Tenant administrator role only.
Add list button	Applies the entered words. While the words are applied, the button has the Adding... label.
Upload .txt file button	Reads a text file with one word per line. Myra AI Workspace scans the file for malware before it reads the file.
Delete row action	Removes the keyword from the list.

The following information is available:

Column	Description
Add a keyword	Word or phrase that Myra AI Workspace masks.
Case-sensitive	Check mark if the spelling must agree, otherwise —.
Whole word only	Check mark if only the whole word is masked, otherwise —.



If no keyword exists yet, the tab shows **No personal keywords yet. Anything you add will be replaced with [MYRA-CUSTOM:...] tokens before the model sees your message.**



The list takes bare words only – no spaces, commands, or regular expressions. The punctuation @ . - _ is permitted. Myra AI Workspace reports each rejected entry with the reason, for example **too long** or **contains a disallowed character**.

3.4.3 Connections




Figure 82: MCP Connectors view

The following options are available:

Tab	Description
MCP Connectors	Connects servers for the Model Context Protocol (MCP) so the AI calls external tools in the chat.
Chat Integrations	Connects Mattermost, Slack, or Microsoft Teams to an agent so your team speaks to the agent from these platforms.

The following sections describe the view in detail:

- [MCP Connectors tab](#)
- [Connector directory dialog](#)
- [Chat Integrations tab](#)

3.4.3.1 MCP Connectors tab




Figure 83: MCP Connectors tab

The tab lists the connectors of the tenant. It carries the note **Connect to Model Context Protocol servers so the AI can call external tools during chat.**

The following information is available:

Column	Description
Name	Name of the connector. For a connector from the directory, the icon of the service comes first.
Server URL	Address of the MCP server that the connector calls.
Auth	Method with which the connector signs in to the server: Bearer , Header , or None .
Created	Date on which the connector was created.

The following options are available:

Option	Description
Add from directory button	Opens the Connector directory dialog.
New connector button	Opens the New MCP Connector dialog.
Edit row action	Opens the Edit MCP Connector dialog.
Delete row action	

Option	Description
--------	-------------

Opens the **Delete this MCP connector? The model will no longer have access to its tools.** confirmation.



If no connector exists yet, the tab shows **No MCP connectors yet. Add one to give the AI access to external tools.**

3.4.3.2 Connector directory dialog




Figure 84: Connector directory dialog

The **Connector directory** dialog pre-fills a connector from a catalogue of known services. The dialog carries the note **Pick a service to pre-fill a connector. You'll register an OAuth app for it and your users authorize it once.**

The following information is available for each service:

Entry	Description
Name and icon	Service that the connector addresses, for example GitHub , Confluence , or Google Drive .
Description	Short description of what the AI does with this service.
Examples	Two example tasks that work in the chat with this connector.
Label	No sign-in needed for a service without a sign-in, Connects automatically for a service that registers itself on the first sign-in, and One-time admin setup for a service for which an OAuth app is to be registered.

The following options are available in the dialog:

Option	Description
Configure button	Opens the New MCP Connector dialog with the entries of the service filled in.
Configure existing button	Opens the Edit MCP Connector dialog for the connector of this service that already exists.
Cancel button	Closes the dialog without creating a connector.



The dialog gives the callback URL for the OAuth app: **Use this callback URL in your OAuth app:** Enter this URL at the provider before the first person signs in.

3.4.3.3 Chat Integrations tab




Figure 85: Chat Integrations tab

The tab lists the chat connections of the tenant. It carries the note **Let your team talk to an agent from Mattermost, Slack or Microsoft Teams. Paste the platform's credentials, pick the agent that answers, and enable it.**

The following information is available:

Colum n	Description
Platfor m	Chat platform of the connection: Mattermost , Slack , or Microsoft Teams .
Worksp ace	Identifier of the workspace on the platform, for example the Mattermost team ID, the Slack team ID, or the Entra tenant ID.
Agent	Agent that answers the messages from the chat.
Gatew ay	Gateway through which the connection bills and enforces its requests.
Status	Enabled if the bot answers, otherwise Disabled .
Web search	On if the bot may search the web for its answers, otherwise Off .
All tools	On if the bot may use the full toolset of its agent, otherwise Off .
	Date on which the connection was created.

Colum n	Description
Create d	

The following options are available:

Option	Description
Add connection button	Opens the Add chat connection dialog.
Enable row action	Switches the bot on. The row action appears while the status is Disabled .
Disable row action	Switches the bot off. The bot then stops answering, but the connection stays.
Enable web search row action	Lets the bot search the web for its answers.
Disable web search row action	Takes the web search away from the bot.
Enable all tools row action	Gives the bot the full toolset of its agent, that is running code, fetching URLs, MCP connectors, and image generation.
Disable all tools row action	Takes the tools away from the bot, except the web search.
Rotate token row action	Opens the Rotate bot token dialog.
Delete row action	Opens the Delete this chat connection? The bot will stop answering and its access token is revoked. confirmation.



If no chat connection exists yet, the tab shows **No chat connections yet.**



The tools run under the credentials of the person who owns the bot, and anyone in the chat workspace triggers them. Enable **All tools** only if you trust every member of this channel.

3.4.4 Request Logs




Figure 86: Request Logs view

The **Request Logs** view records the requests to the gateways and shows the matching trace for every request. The trace names the providers involved, the model used, the tokens, the cost, and the latency. The subtitle names the number of the entries.

The following information is available:

Column	Description
Time	Time at which the gateway received the request.
Provider	Provider that answered the request, for example anthropic .
Model	Model that answered the request, for example gpt-4o .
Status	HTTP status code of the response, for example 200 , 401 , 429 , or 500 .
In	Number of the input tokens of the request.
Out	Number of the output tokens of the response.
Cache	Share of the tokens read from the prompt cache, for example 80% hit . The note of the column names the cache write and cache read tokens.
Cost	Cost of the request in US dollars.
Latency	Total duration of the request in the gateway in milliseconds.

Column	Description
Upstream	Duration of the request at the provider in milliseconds. If no duration is available, then the column shows — .
Guardrail	Duration of the guardrail check in milliseconds and its verdict.
Detectors	Detectors that fired. From the third detector on, the column summarizes the remaining ones as + .
Flags	Labels of the request: cached , blocked , scrubbed , fallback , and attempts .

The following filters are available:

Filter	Description
Tenant selection list	Restricts the list to one tenant. The All tenants entry removes the filter.
Provider selection list	Restricts the list to one provider. The All providers entry removes the filter.
Model input field	Restricts the list to one model, for example gpt-4o .
Status selection list	Restricts the list to one status code: 200 OK , 400 , 401 , 429 , or 500 .
Blocked selection list	The Blocked only entry restricts the list to the blocked requests.
Guardrail outcome selection list	Restricts the list to one outcome of the guardrails: Any hit , Blocked , Scrubbed , or Flagged .
Limit selection list	Maximum number of the loaded entries: 20 , 50 , 100 , or 200 .
Since and Until input fields	Restrict the list to a period.

The following options are available:

Option	Description
Refresh button	Loads the list again with the filters set.
Export CSV button	Downloads the filtered entries as a CSV file. During the export, the button carries the Exporting... label.
Export XLSX button	Downloads the filtered entries as an Excel file.
Row of the list	Shows or hides the prompt and the response of the request.
View trace icon	Opens the Request Trace area below the list. The icon appears only for a request with a trace.

The following sections describe the view in detail:

- [Request details area](#)
- [Request Trace area](#)

3.4.4.1 Request details area



TIME	PROVIDER	MODEL	STATUS	IN	OUT	CACHE	COST	LATENCY	UPSTREAM	GUARDRAIL	DETECTORS	FLAGS
4/2/2026, 9:41:02 AM	anthropic	claude-sonnet-4-5	200	1,240	310	–	\$0.03	381 ms	352 ms	29 ms	pii-protect	scrubbed
Prompt (PII scrubbed) Summarize the contract of [MYRA-PII:PERSON] and name the notice period. Contact: [MYRA-PII:EMAIL]												
LLM response (PII tokens visible if echoed) The contract runs for twelve months and renews for twelve months at a time. The notice period is three months before the end of the term.												
4/2/2026, 9:40:55 AM	openai	gpt-4o-mini	200	620	148	80% hit · 12k	\$0.01	298 ms	271 ms	–	–	cached

Figure 87: Request details area

The area opens when you click a row of the list, and it shows the stored wording of the request.

The following information is available:

Entry	Description
Prompt	Wording of the request. If Myra AI Workspace replaced personal data, then the entry carries the Prompt (PII scrubbed) label.

Entry	Description
LLM response	Wording of the response. If Myra AI Workspace replaced personal data, then the entry carries the LLM response (PII tokens visible if echoed) label.

3.4.4.2 Request Trace area




Figure 88: Request trace area

The area opens via the **View trace** icon and shows the path of the request through the gateway as a sequence of steps.

The following information is available:

Item	Description
Item below the heading	First eight characters of the trace ID, followed by ...
Badge next to the heading	Status of the trace: done , blocked , or error .

The following steps are available:

Step	Description
Receive d	The gateway accepted the request. The step names the model, the number of the messages, and the stream addition for a streaming request.
Transfo rmed	The gateway translated the request into the format of the provider.
Routin g	The gateway selected the target. The step names the number of the fallbacks.
Guardr ail	The guardrails checked the request. The step names the verdict, for example safe .
→ Provide r	The gateway sent the request to the provider. The step names the attempt, for example attempt 2 .
← Provide r	The provider responded.
Provide r Error	The provider responded with an error.
Deliver ed	The gateway delivered the response. The step names the size of the response or the streaming addition.

The following options are available in the area:

Option	Description
Row of a step	Shows and hides the recorded data of the step as JSON.
Close icon	Closes the area.



If no request matches the filters, then the view reports **No logs match your filters**. If the loading fails, then the view reports **Failed to load logs**. If the request logging is disabled for the selected tenant, then the view reports **Request logging is disabled for tenant ""**. **No new request log entries are being stored for its traffic**. If the export exceeds the maximum number of the rows, then the view reports **Export truncated at rows**. **Narrow the date range to capture the rest**. If no trace is available, then the area reports **Trace not found**, and without recorded steps **No steps recorded**.

3.4.5 Gateways




Figure 89: Gateways view

The **Gateways** view manages the named installations of a tenant. A gateway sets which providers and models are available, which rules apply to the content, which budget applies, and which authentication tokens get access. Above the list, the **Tenant:** entry names the tenant whose gateways the view lists.

The following information is available:

Colum n	Description
Name	Identifer of the gateway. It contains lowercase letters, numbers, and hyphens, and is part of the endpoint URL.
Auth	Shows whether a request to the gateway must carry an authentication token: required or open .
Budget	Spend cap of the gateway. If no budget is set, the — entry appears.
Cache	Lifetime of the cache in seconds. The value 0s means that the cache is disabled.
Rate Limit	Permitted requests per time window, in the form /s . If no rate limit is set, the — entry appears.

Column	Description
Detect ors	Number of guardrail detectors that are active on the gateway. If no detector is active, the — entry appears.
Create d	Time at which the gateway was created.

The following options are available:

Option	Description
Select Tenant selection list	Selects the tenant whose gateways the view lists. The Filter tenants by name... input field narrows the selection.
Search gateways by name... input field	Narrows the list to the gateways whose name contains the text you enter.
New Gateway button	Opens the New Gateway dialog. See the Creating a gateway section.
Open → row action	Opens the detail view of the gateway in the row.
Delete row action	Opens the Delete gateway ""? This cannot be undone. prompt. See the Deleting a gateway section.
Quick start guide link	Appears only when the list is empty and opens the quick start documentation.



If no gateway exists for the tenant yet, the view shows **No gateways for yet**. If the search finds no gateway, the view shows **No gateways match "<search text>"**. If the filter finds no tenant, the selection shows **No tenants match "<search text>"**.

The following sections describe the view in detail:

- [Detail view of a gateway](#)
- [Gateway section](#)
- [Provider Keys section](#)

- [Auth Tokens section](#)
- [Guardrails section](#)
- [Routing Rules section](#)
- [Circuit Breaker section](#)

3.4.5.1 Detail view of a gateway




Figure 90: Detail view of a gateway

The detail view lists the settings and the access of a gateway in sections. You can collapse and expand each section with the icon to the right of its heading.

The following information is available:

Section	Description
Gateway:	Lists the statistics of the configuration and the endpoint URL.
Provider Keys	Lists the stored API keys of the providers.
Auth Tokens	Lists the tokens that get access to this gateway.

Section	Description
Guardrails	Lists the detectors for content and their recent events.
Routing Rules	Lists the rules that override the default routing.
Circuit Breaker	Lists the state of the providers whose calls have failed repeatedly.

The following options are available:

Option	Description
← Gateways link	Returns to the list of gateways.
Edit button	Opens the Edit Gateway: dialog. See the Editing a gateway section.
Delete button	Opens the Delete gateway ""? This cannot be undone. prompt. While the gateway is deleted, the button shows the Deleting... label.
Reset Spend button	Opens the Reset the budget spend counter to \$0? prompt and resets the spend counter after your confirmation. The button appears only if a budget is set.
Endpoint URL selection list	Selects the provider for which the view builds the endpoint URL.
Copy button	Copies the endpoint URL to the clipboard. After the copy, the button shows the Copied! label.
Collapse section icon	Collapses the section. When the section is collapsed, the icon shows the Expand section label.

3.4.5.2 Gateway section




Figure 91: Gateway section

The following statistics are available:

Statistic	Description
Budget	Spend and cap of the gateway in the selected budget period (This month , Today , or All time), with the % used entry and a breakdown per model. If no budget is set, the No spend cap entry appears. If the spend cannot be determined, the Spend unavailable entry appears.
Cache TTL	Lifetime of the cache in seconds.
Retries	Number of retries of a failed call to the provider. The default is 2 .
Timeout	Time to wait for the provider before the request fails. The default is 120000 ms .

Statistic	Description
Auth	required or open .
Rate Limit	Permitted requests per time window. If no rate limit is set, the – entry appears.
Log Payloads	yes or no .
Web hook	The forwarded events or the all events entry. If no webhook URL is set, the – entry appears.
Tracing	on or off . If the trace contains the message bodies, the + bodies entry appears as well.
Web Search	on or off .
Image Gen	on or off .
Prompt Caching	The configured lifetime of the prompt cache, or the off entry.
Compaction	on or off .
Compact err	Threshold in the form k tok , or the off entry.
Endpoint URL	Address to which an application sends its requests. The selection list sets the provider for which the view builds the address.



If the currency of the workspace is euro, the **≈ converted at today's rate; the cap is stored and enforced in USD.** hint appears below the budget.

3.4.5.3 Provider Keys section



Figure 92: Provider Keys section

The following information is available:

Column	Description
Provider	Provider to which the key belongs, for example anthropic or openai .
Alias	Label of the key. The alias distinguishes several keys of the same provider.
Added	Time at which the key was stored.

The following options are available:

Option	Description
Add Model button	Opens the Add Model dialog. See the Enabling a model on a gateway section.
Delete row action	Opens the Delete key for /? Routing to this provider will fail. prompt.



If no key is stored, the section shows **No keys stored. Add one to enable provider routing.**

3.4.5.4 Auth Tokens section




Figure 93: Auth Tokens section

The following information is available:

Column	Description
Label	Label of the token, for example ci-pipeline .
Hash (first 16)	The first 16 characters of the hash value. They identify a token without disclosing it.
Rate Limit	Permitted requests per time window for this token.
Spend Cap	Amount above which Myra AI Workspace blocks this token.
Expires	Expiry time of the token. If no expiry is set, the never entry appears. A token close to its expiry carries the Expires soon marker, an expired one the Expired marker.
Created	Time at which the token was generated.

The following options are available:

Option	Description
Generate button	Opens the Create Auth Token dialog.
Edit row action	Opens the Edit Token – dialog.
Revoke row action	Opens the Delete this token? Requests using it will fail. prompt.



If no token exists yet, the section shows **No tokens yet. Generate one to allow access to this gateway.**

3.4.5.5 Guardrails section




Figure 94: Guardrails section

The following information is available:

Entry	Description
blocked marker	Number of requests that a detector blocked in the last 24 hours.
scrubbed marker	Number of requests from which a detector removed content.
flagged marker	Number of requests that a detector flagged but did not change.
avg ms guardrail latency • last 24h entry	Average time that the check of a request needs in addition.
Time column	Time of the event.
Outcome column	Outcome of the check: blocked , scrubbed , or flagged .
Detectors column	Detectors that triggered.
Reason column	Reason that the detector reported.
Guardrail ms column	Duration of the check in milliseconds.
Total ms column	Total duration of the request in milliseconds.
Guardrails () entry	Quantity of the detectors that this gateway has.
Area of a detector	Type, name, and the • summary of a detector. When it is expanded, the area shows all data of the detector.
Execution plan section	Shows the detectors in the sequence in which the gateway examines them. See the Execution plan section.

The following options are available:

Option	Description
Regex / Pattern, Keyword, Jailbreak, Presidio (NLP), Prompt Guard, Prompt	Adds a detector of the applicable type. See the Adding a guardrail detector section.

Option	Description
Injection, PII Protector, and Custom PII Blacklist buttons	
Switch in the area of a detector	Turns the detector on and off. A detector that is off stays in the list but is not evaluated. Its area is dimmed and carries the DISABLED badge.
Area of a detector	Expands and collapses the data of the detector. The Move up and Move down icons move the detector, the icon to remove it deletes it. See the Changing the sequence of the guardrail detectors and Removing a guardrail detector sections.
Save Guardrails button	Applies the detectors. After the save, the Saved entry appears.
Show recent activity button	Shows the table of the recent guardrail events. When the table is shown, the button carries the Hide recent activity label.



If no event has been recorded yet, the section shows **No guardrail events recorded for this gateway yet.**

3.4.5.6 Routing Rules section



ORDER	CONDITIONS	ACTION	STATUS
1	header:x-customer-tier eq premium	anthropic / claude-opus-4-6 openai/gpt-4.1	on
2	provider eq anthropic	load balance - weighted_random anthropic/claude-sonnet-4-6 70% openai/gpt-4.1 30%	on
3	any	openai / gpt-4.1-mini	off

Figure 95: Routing Rules section

The following information is available:

Column	Description
Order	Position of the rule in the evaluation. Myra AI Workspace evaluates the rules from top to bottom.
Conditions	Conditions that must all match. If the rule carries no condition, the any entry appears.
Action	Target of the rule: a direct route to / , or the load balance • entry with the targets and their shares.
Status	on or off .

The following options are available:

Option	Description
New Rule button	Opens the New Routing Rule dialog.
Move up (evaluate earlier) row action	Moves the rule up by one position.
Move down (evaluate later) row action	Moves the rule down by one position.
Edit row action	Opens the Edit Rule dialog.
Del row action	Opens the Delete this routing rule? prompt.



If no rule exists, the section shows **No routing rules. Add rules to override default provider routing.**

3.4.5.7 Circuit Breaker section




Figure 96: Circuit Breaker section

The following information is available:

Colum n	Description
Provid er	Provider whose calls the circuit breaker monitors.
State	closed (requests run), open (requests are not forwarded), or half-open (a probe request tests the provider).
Failure s	Number of failures in the sliding window.
Opene d at	Time at which the circuit breaker opened.

The following options are available:

Option	Description
Refresh button	Loads the state of the providers again.



Below the table, the **Threshold: failures in s • Cooldown:** entry names the configured values. If no circuit breaker is open, the section shows **All providers are healthy (no open breakers)**.

3.4.6 Providers




Figure 97: Providers view

The **Providers** view shows the supported AI providers with their configuration and their availability status in real time. The list shows the configured providers first and the remaining providers alphabetically after them.

The following information is available:

Column	Description
Provider	Name of the AI provider, for example Anthropic or OpenAI .
Supported	The ✓ Supported label. Myra AI Workspace connects the provider.
Configured	State of the provider key: The Configured label stands for at least one stored key, the Not configured label for no stored key, and the No key needed label for a provider without a key.
Live Status	Availability according to the status page of the provider: Operational , Degraded , or Outage . If the provider has no public status page, then the column shows – with the No public status page note.
Last Checked	Time since the last query of the status page, for example 4m ago .

The following metric areas are available:

Metric	Description
Supported providers	Number of the providers that Myra AI Workspace connects.
Configured (any gateway)	Number of the providers for which at least one gateway holds a key.
Operational (status page)	Number of the operational providers in relation to the providers with a status page, for example 6 / 8 .
Status last refreshed	Time of the last successful query. The area appears only after the first successful query.

The following options are available:

Option	Description
Refresh button	Queries the availability status again. During the query, the button carries the Refreshing... label.
Add key → button	Opens the Gateways view to store an API key for this provider. The button appears only with the Not configured label.



Myra AI Workspace polls the status from the status pages of the providers every 5 minutes and only for providers with a known status page endpoint. If no provider data is available, then the view reports **No provider data available**.

3.4.7 Security




Figure 98: KI-Governance view

The **KI-Governance** view contains rule templates that map compliance obligations to gateway settings. You show a preview of a template for a gateway and apply it.

The following information is available:

Column	Description
Name	Name of the template. A template you created carries the Non-authoritative label in addition.
Key	Template key followed by the version, for example v1 .
Controls	Target controls of the template in the form =on or =off . If the template carries no control, the column shows No controls .

The following options are available:

Option	Description
New template button	Opens the New governance template dialog. The button is disabled while no tenant is selected.
Tenant drop-down list	Selects the tenant whose templates the list shows. The drop-down list appears for the Administrator role only. All other roles see the templates of their own tenant.
	Opens the Preview & apply – dialog.

Option	Description
Preview & apply row action	
Edit row action	Opens the Edit governance template dialog. For an authoritative template the row action is disabled.
Delete row action	Opens the Delete the template "?" confirmation. For an authoritative template the row action is disabled.



If no template exists yet, the view shows **No governance templates yet**. If no tenant is selected yet, the view shows **Select a tenant to manage its governance templates**.



An authoritative template is read-only. The disabled row actions carry the note **Authoritative templates are read-only**. The Data Protection Officer signs off authoritative content.



Myra AI Workspace always applies the residency minimum. A template does not lower it.

3.4.8 Costs




Figure 99: Finance view

The following tabs are available:

Tab	Description
User Costs	Shows the cost attribution and the billing settings of the tenant.
Provider Costs	Shows the model prices that the cost attribution uses.

The following sections describe the view in detail:

- [User Costs tab](#)
- [Provider Costs tab](#)
- [Entries in the New Model Price and Edit dialogs](#)

3.4.8.1 User Costs tab




Figure 100: User Costs tab

The **Costs** view shows the cost attribution and the billing settings of the tenant in the **User Costs** tab. The header carries the **Cost attribution and billing settings for your tenant.** subtitle.

The following information is available:

Area	Description
Total cost	Cost of the tenant in the selected period. Below that, the area shows the daily cost trend as a curve.
Users with spend	Number of the users to whom cost is attributed in the selected period. The subline reads in the selected range.
Billing settings	Sets the address to which Myra AI Workspace sends the invoices, and names the plan of the tenant.
Cost per user	Cost per user with the User , Requests , and Cost columns. The table is sorted by the cost in descending order.

The following options are available:

Option	Description
Time range selection list	Period of the evaluation: Last 7 days or Last 30 days .
Currency selection list	Currency in which the view states the amounts.
Invoice recipient input field	Email address invoices are addressed to, for example billing@example.com . If you clear the field, then Myra AI Workspace removes the recipient.
Plan entry	Plan of the tenant. You cannot change the entry in this view.
Save button	Applies the invoice recipient. While it saves, the button carries the Saving... label, and after that the view reports Saved .



The **Billing settings** area appears only for a tenant for which the operator enabled the billing. If the address entered is not a valid email address, then the view reports **Enter a valid email address**. If the saving fails, then the view reports **Could not save. Please try again**. If no cost occurred in the period, then the table reports **No cost recorded in this range**.

3.4.8.2 Provider Costs tab




Figure 101: Provider Costs tab



The **Provider Costs** tab lists the model prices that the cost attribution uses. The subtitle names the number of the entries in the **entries – used for cost attribution** form.

The following information is available:

Column	Description
Provider	Provider of the model.
Model	Identifier of the model. The per query label stands for a row that prices the web search per query. The Deprecated label stands for a model that the provider decommissioned.
Display Name	Name under which the model appears in the interface.
Hosted in	Region in which the model runs.
Input \$/1K tokens	Price per 1,000 input tokens.
	Price per 1,000 output tokens.

Column	Description
Output \$/1K tokens	
Cache Read \$/ 1K	Price per 1,000 tokens read from the cache.
Updated	Time of the last change of the price.

The following options are available:

Option	Description
Sync from Providers button	Applies the published prices of the providers. While it syncs, the button carries the Syncing... label. After that, the view reports Synced:  , All models up to date , or Sync completed with errors –  .
New Price button	Opens the New Model Price dialog.
Buttons of the filter	Restrict the table to one provider. The All button removes the restriction.
Show deprecated () button	Shows the decommissioned models. After that, the button carries the Hide deprecated () label.
Edit row action	Opens the Edit: / dialog.
Delete row action	Opens the Delete pricing for /? prompt. After the confirmation, Myra AI Workspace no longer bills new requests to this model.

3.4.8.3 Entries in the New Model Price and Edit dialogs

New Model Price

Provider

anthropic

Model *

claude-sonnet-4-6

Input \$/1K tokens *

0.003

Output \$/1K tokens *

0.015

Prices are per 1,000 tokens. Providers usually publish per 1,000,000 – divide their figure by 1,000.

Cache Write 5m \$/1K (optional)

—

Cache Write 1h \$/1K (optional)

—

Cache Read \$/1K (optional)

—

Cancel

Add Price

Figure 102: New Model Price dialog

You set the following entries:

Entry	Description
Provider	Provider of the model.
Model *	Identifier of the model, for example claude-sonnet-4-6 . The entry is required. In the Edit dialog, you cannot change the identifier.
Input \$/1K tokens *	Price per 1,000 input tokens, for example 0.003 . The entry is required.
Output \$/1K tokens *	Price per 1,000 output tokens, for example 0.015 . The entry is required.
Cache Write 5m \$/1K (optional)	Price per 1,000 tokens that are written to the cache for five minutes.

Entry	Description
Cache Write 1h \$/1K (optional)	Price per 1,000 tokens that are written to the cache for one hour.
Cache Read \$/1K (optional)	Price per 1,000 tokens read from the cache.

The following options are available in the dialogs:

Option	Description
Add Price button	Creates the model price. In the Edit dialog, the button carries the Save Changes label.
Cancel button	Closes the dialog without applying the entries.



The prices are per 1,000 tokens. Providers usually publish per 1,000,000 – divide their figure by 1,000. Web search is priced per query, not per token. The **Input \$/1K tokens** input field is USD per 1,000 queries (10 = \$0.01 per query). The value of the **Output \$/1K tokens** field stays 0.

3.5 User Management

3.5.1 Users




Figure 103: Users view

The **Users** view administers the user accounts of the tenant and their roles. The subtitle of the view gives the number of the user accounts.

The following information is available:

Colum n	Description
Email	Email address of the user account. A deleted account carries the Deleted badge, a disabled account the Disabled badge.
Name	Name of the user.
Role	Role of the user account, for example Member or Tenant Administrator .
Tenan t	Tenant that the user account belongs to.
Last Login	Time of the last login. If the account never logged in, the never entry appears.
Create d	Time at which the user account was created.

The following options are available:

Option	Description
Search users by name or email... field	Limits the list to the user accounts whose name or email address contains the text entered. The field appears as soon as the list holds at least one user account.
Show deleted check box	Shows the deleted user accounts with the Deleted badge.
All tenants list	Limits the list to one tenant.
Export CSV button	Downloads the list as a CSV file with the Email, Name, Role, Tenant, Last Login, Created, and Status columns.
Export XLSX button	Downloads the same data as an XLSX file.
New User button	Opens the New User dialog. See the Creating a user account section.
Column heading	Sorts the list by the column. A further click reverses the sort direction.
Open → row action	Opens the detail view of the user account.
View as user row action	Shows the work area in the view of this user account.
Disable row action	Blocks the user account. The account then carries the Disabled badge.
Enable row action	Releases the block of a disabled user account.
Restore row action	Restores a deleted user account.



The **All tenants** list appears only for the **Administrator** role. If no user account exists yet, the view shows **No users yet. Create the first user to get started.** If the search finds no user account, the view shows **No users match "<search text>".**

The following sections describe the view in detail:

- Detail view of a user account
- Tokens section

3.5.1.1 Detail view of a user account




Figure 104: Detail view of a user account

The detail view lists the data of a user account and its tokens.

The following information is available:


Item	Description
Email	Email address of the user account.
Name	Name of the user.
Role	Role of the user account.
Tenant	Tenant that the user account belongs to.

Item	Description
User ID	Identifier of the user account.
Created	Time at which the user account was created.
Last Login	Time of the last login or the never entry.

The following options are available:

Option	Description
← Users button	Returns to the list of the user accounts.
Resend Invite button	Sends the invitation to set up the account by email again. After the sending, the button shows Invite Sent ✓.
Edit button	Opens the Edit: EMAIL dialog. See the Editing a user account section.
Delete User button	Deletes the user account after a confirmation. See the Deleting a user account section.
New Token button	Opens the New Token for EMAIL dialog. See the Creating a token section.

3.5.1.2 Tokens section



Tokens							New Token
LABEL	GATEWAY	SCOPES	EXPIRES	BUDGET	RATE LIMIT	CREATED	
Editorial script	muster-gmbh/redaktion	all	12/31/2026	\$50	60/60s	4/2/2026	⛔
Analysis	muster-gmbh/produktion	chat	never	\$10	10/60s	6/11/2026	⛔

Figure 105: Tokens section

The **Tokens** area lists the API tokens of the user account.

The following information is available:

Column	Description
Label	Label of the token.
Gateway	Gateway that the token applies to, in the tenant/gateway form.
Scopes	Scopes of the token. A token without a restriction carries the all badge.
Expires	Expiry date of the token or the never entry.
Budget	Budget of the token in USD. Without a budget, the – entry appears.
Rate Limit	Requests per time window. Without a rate limit, the – entry appears.
Created	Time at which the token was created.

The following options are available:

Option	Description
Revoke row action	Revokes the token after the Revoke this token? confirmation. Requests with this token then fail.



If the user account holds no token yet, the area shows **No tokens yet**. For the **Viewer** role, the area shows **Viewer users cannot make inference requests**.

3.5.2 Groups




Figure 106: Groups view

The **Groups** view administers groups and their memberships. Through a group, you grant access to a knowledge space to several users at once. The subtitle of the view gives the number of the groups or the tenant whose groups the list holds.

The following information is available:

Column	Description
Name	Name of the group.
Description	Short description of the group. The column stays empty if no description is entered.
Members	Number of the users in the group.
Created	Time at which the group was created.

The following options are available:

Option	Description
Tenant list	Selects the tenant whose groups the list holds. Without a selection, the view shows Select a tenant to manage its groups...
New group button	Opens the New group dialog. See the Creating a group section.
Row of the group	Opens the dialog with the data, the members, and the grants of the group.

Option	Description
--------	-------------

Delete row action	Deletes the group after the Delete the group "<name>"? Its memberships and grants are removed. confirmation.
--------------------------	---

 If no group exists yet, the view shows **No groups yet. Create one to grant access by group.**

3.5.2.1 Group dialog




Figure 107: Dialog of a group

The dialog carries the name of the group and administers its data, its members, and its grants.

The following options are available in the dialog:

Option	Description
--------	-------------

Name field	Changes the name of the group.
-------------------	--------------------------------

Option	Description
Description field	Changes the short description of the group.
Save button	Applies the changed data.
Cancel button	Closes the dialog without applying the changes.

The **Members** area lists the users of the group with their email address and gives their number.

Option	Description
Add member (pick from the list or type an email) list	Selects the users that the group takes in.
Add member button	Takes the users in. After the addition, the dialog shows EMAIL added to the group.
Remove row action	Removes the member after the Remove NAME from this group? They lose the access this group grants. confirmation.

The **Access grants** area lists the knowledge spaces that the group has access to and gives the role granted for each.

Option	Description
Grant access to a knowledge space list	Selects the project that the group receives access to.
Role list	Selects the role of the grant: Viewer – can read and chat, Editor – can add knowledge files, or Owner – full control.
Add grant button	Grants the access to all members of the group.
Remove row action	Revokes the access after the Revoke this group's access to NAME? Every member of the group loses it. confirmation.



If the group holds no member yet, the dialog shows **No members yet.** If the group holds no grant yet, the dialog shows **No grants yet.**

3.5.3 Roles & permissions

Roles & permissions

New role

ROLE	TYPE	PERMISSIONS	USERS	ACTIONS
Tenant Administrator	System	38	2	Edit
Member	System	12	18	Edit
Editorial	Custom	7	6	Edit

Figure 108: Roles & permissions view

The **Roles & permissions** view composes custom roles from a curated set of permissions and assigns them to users in the work area.

The following information is available:

Column	Description
Role	Name of the role with its description.
Type	Origin of the role: System for the roles delivered by Myra, Custom for the roles composed by yourself.
Permissions	Number of the permissions that the role grants.
Users	Number of the user accounts that the role is assigned to.
Actions	Holds the actions of the row. A role of the System type carries the Managed by Myra note here.

The following options are available:

Option	Description
New role button	Opens the New role dialog. See the Creating a role section.
Edit row action	Opens the Edit NAME dialog. See the Editing a role section.

Option	Description
Delete row action	Deletes the role after the Delete the role NAME? Its permissions and every user assignment are removed with it. confirmation.

The editor groups the permissions into modules:

- **Agents**
- **Analytics**
- **Approvals**
- **Billing**
- **Model catalog**
- **Gateways**
- **Governance**
- **Groups**
- **Connectors**
- **Projects**
- **Roles**
- **Scheduled tasks**
- **SSO**
- **Workspace**
- **Users**
- **Workflows**



A permission that your own user account does not hold cannot be granted. The editor shows it as not grantable and gives **You do not hold KEY, so you cannot grant it.** Roles of the **System** type cannot be edited or deleted. If no custom role exists yet, the view shows **No custom roles yet. Create one to grant a tailored set of permissions.**

3.5.4 Tenants




Figure 109: Tenants view

The **Tenants** view creates tenants and manages their settings. A tenant groups the gateways, the user accounts, and the budget of an organization. Below the title, the view names the number of tenants in the form **tenants**.

The following information is available:

Column	Description
Name	Identifier of the tenant. It contains lowercase letters, numbers, and hyphens.
Plan	Plan of the tenant: Free , Standard , or Enterprise .
Budget	Spend cap of the tenant with its period. If no budget is set, the unlimited entry appears.
Created	Time at which the tenant was created.

The following options are available:

Option	Description
New Tenant button	

Option	Description
	Opens the New Tenant dialog. See the Creating a tenant section. The button appears only if your user account may create tenants.
Filter tenants by name... input field	Narrows the list to the tenants whose name contains the text you enter.
Open → row action	Opens the detail view of the tenant in the row.



If no tenant exists yet, the view shows **No tenants yet. Create your first tenant to get started.** If the filter finds no tenant, the view shows **No tenants match "<search text>".**



The view needs the **Create tenants** and **Purge tenants** permissions. These permissions stay reserved for the **Administrator** role. A role of your own does not take them.

The following sections describe the view in detail:

- [Detail view of a tenant](#)
- [Overview tab](#)
- [Gateways tab](#)
- [API Keys tab](#)
- [Commands tab](#)
- [Prompt Examples tab](#)

3.5.4.1 Detail view of a tenant




Figure 110: Detail view of a tenant

The detail view lists the settings of a tenant in tabs.

The following information is available:

Tab	Description
Overview	Lists the statistics of the tenant.
Gateways	Lists the gateways of the tenant.
API Keys	Lists the admin API keys of the tenant.
Commands	Lists the shared commands that are available to all users of the tenant in Chat.
Prompt Examples	Lists the prompt examples that appear on the home screen.

The following options are available:

Option	Description
--------	-------------

Returns to the list of tenants.

Option	Description
← Tenants link	
Edit button	Opens the Edit: dialog. See the Editing a tenant section.
Add user button	Opens the dialog that creates a user account in this tenant. After the creation, the view shows User added to this tenant.
Export data button	Exports the data of the tenant. While the export runs, the button shows the Exporting... label.
Delete Tenant button	Opens the Delete tenant ""? This will also delete all their gateways, keys, and tokens. prompt. See the Deleting a tenant section.

3.5.4.2 Overview tab




Figure 111: Overview tab

The following statistics are available:

Statistic	Description
Name	Identifier of the tenant.

Statistic	Description
Plan	Plan of the tenant: Free , Standard , or Enterprise .
Tenant Budget Limit	Spend cap with its period (Monthly , Daily , or Lifetime). If no budget is set, the unlimited entry appears.
Trial ends	Last day of the trial with the active or expired marker. The statistic appears only for a tenant with a trial.
Gateways	Number of the gateways of the tenant.
Tenant ID	Internal identifier of the tenant.
Created	Time at which the tenant was created.
Conversation retention	Period in days after which Myra AI Workspace deletes inactive conversations. If no period is set, the Disabled entry appears.
Request-log retention	Period in days after which Myra AI Workspace deletes request logs. If no period is set, the Disabled entry appears.

3.5.4.3 Gateways tab

Overview

Gateways

API Keys

Commands

Prompt Examples

Gateways

All Gateways +

NAME	AUTH	BUDGET	CACHE TTL	CREATED	
produktion	required	\$460.00	300s	2/12/2026	
redaktion	required	\$120.00	60s	3/3/2026	
auswertung	open	—	0s	5/27/2026	

Figure 112: Gateways tab

The following information is available:

Column	Description
Name	Identifier of the gateway.
Auth	Shows whether a request must carry an authentication token: required or open .
Budget	Spend cap of the gateway.
Cache TTL	Lifetime of the cache in seconds.
Created	Time at which the gateway was created.

The following options are available:

Option	Description
All Gateways → button	Opens the Gateways view of this tenant.
Open → row action	Opens the detail view of the gateway in the row.



If no gateway exists yet, the tab shows **No gateways yet**. With the **All Gateways** button you create one.

3.5.4.4 API Keys tab




Figure 113: API Keys tab

The following information is available:

Column	Description
Provider	Provider to which the key belongs.
Alias	Label of the key.
Added	Time at which the key was stored.

The following options are available:

Option	Description
Add Key button	Opens the form for an Anthropic Admin Key . Myra AI Workspace stores the key encrypted and uses it to sync the usage data of the Anthropic Admin API.
Save Key button	Stores the key you entered. While the key is stored, the button shows the Saving... label.
Remove row action	Opens the Remove key ()? prompt.



If no key exists, the tab shows **No admin keys configured. Add an Anthropic Admin key (sk-ant-admin01-...) to enable authoritative usage reporting.**

3.5.4.5 Commands tab




Figure 114: Commands tab

The following information is available:

Column	Description
Command	Shortcut of the command in the form <code>/</code> .
Description	Short description that appears in the picker in Chat.
Template	Text that the command puts into the input field. Placeholders are written in double braces.

The following options are available:

Option	Description
Add Command button	Opens the Add Shared Command dialog. The dialog expects a Name , a Description , and a Template .
Edit row action	Opens the Edit / dialog.
Delete row action	Opens the Remove this command? prompt.



If the name is missing, the dialog shows **Name is required**. If the template is missing, the dialog shows **Template is required**. If no command exists yet, the tab shows **No shared commands yet. Commands defined here are available to all users in this tenant via the `/commandname` shortcut in Chat.**

3.5.4.6 Prompt Examples tab




Figure 115: Prompt Examples tab

The following information is available:

Column	Description
Label	Label of the example that appears on the home screen.
Model matchers	Ordered list of the preferred models. The first model available to the tenant wins.
Effort	Reasoning effort for reasoning models: Default , Low , Medium , High , Extra , or Max .
Enabled	Shows whether the example appears on the home screen: Yes or No .

The following options are available:

Option	Description
Add Example button	Opens the Add Prompt Example dialog with the Content , Steering , and Routing & behaviour sections.
Edit row action	Opens the Edit dialog.
Delete row action	Opens the Remove this prompt example? prompt.



If the label is missing, the dialog shows **Label is required**. If the prompt is missing, the dialog shows **Prompt is required**. If no example exists yet, the tab shows **No prompt examples yet. Examples appear under the Prompt Examples tile on the /easy home screen.**

3.6 Account

3.6.1 My Account




Figure 116: My Account view

The **My Account** view shows your profile, your sign-in, your tokens, and the data export. Below the heading, the **Your profile, sign-in, tokens, and data export.** note appears.

The **Profile** section shows the identity data of your user account.

The following information is available:

Entry	Description
Email	Email address of your user account. You sign in with this address. It is also the confirmation for the permanent erasure of your data. The address is not editable in this view.

Entry	Description
Role	Role of your user account as a colored label, for example admin , tenant_admin , member , viewer , or finance . The role controls which views and options are available to you. Only users with the Administrator role can change the role.



You do not set the interface language, the behavior of the Enter key, and the personal AI preferences here, but in the **Preferences** view in the **Settings** area.

The **Tokens** section lists your personal access tokens for the inference API. A token applies to one gateway and always carries the **inference** scope.

The following information is available:

Column	Description
Label	Name that you entered when you created the token. If you entered no name, the column shows — .
Gateway	Tenant and gateway of the token in the form / . If the gateway is no longer available, the column shows the first eight characters of the gateway ID.
Expires	Date on which the token becomes invalid. If the token does not expire, the column shows never .
Budget	Spend budget and consumption of the token. If the consumption is known, the column shows / • left . If the consumption is unknown, it shows • usage unknown . Without a budget, the column shows used . Without a budget and without a known consumption, it shows — .
Rate Limit	Number of requests and length of the time window in the form /s . Without a rate limit, the column shows — .

Column	Description
Created	Date on which the token was created.

The following options are available:

Option	Description
New Token button	Opens the New Token dialog. The button appears only for a role that can create tokens.
Revoke row action	Opens the Revoke this token? confirmation. After the confirmation, Myra AI Workspace makes the token invalid and removes it from the list.



If no token exists yet, the section shows **No tokens yet. Create one to use the inference API.** The **Learn more** link opens the documentation. While the list loads, the section shows **Loading...** If the list cannot be loaded, the section shows **Could not load. Please try again.** and offers the **Retry** button.



If no gateway is available, Myra AI Workspace shows the **No gateways available** warning when you click **New Token** and does not open the dialog.

The following sections describe the view in detail:

- **Entries in the New Token dialog**
- **Entries in the Token Created dialog**
- **Effects in the Delete Account dialog**
- **Effects in the Erase all my data dialog**

3.6.1.1 Entries in the New Token dialog

You set the following entries when you create the token:

Entry	Description
**Gateway **	Gateway that the token applies to. The drop-down list contains the gateways of all tenants that you are permitted to use, in the form /. The entry is mandatory. The first gateway of the list is preselected.
Label	Name of the token in the list, for example dev laptop . The entry is optional. Without a name, the Label column shows — later.
Expires At	Date and time on which the token becomes invalid. The field carries the Leave blank = never expires note.
Budget (USD)	Spend budget of the token in US dollars. The field carries the unlimited placeholder. Without an entry, the token does not limit the cost.
Rate limit (req/window)	Number of requests in the left field and length of the time window in seconds in the right field, at least 1 each. The right field is preset to 60 and stays disabled while the left field is empty. The field carries the requests / window_sec (blank = gateway default) note.

The following options are available in the dialog:

Option	Description
Create Token button	Creates the token and opens the Token Created dialog. While the token is created, the button has the Creating... label. The button is disabled while no gateway is selected.
Cancel button	Closes the dialog without creating the token.



If no gateway is selected, the dialog shows **Select a gateway**. The token always carries the **inference** scope. You cannot give it a different scope.

3.6.1.2 Entries in the Token Created dialog

The **Token Created** dialog shows the created token one single time.

The following options are available in the dialog:

Option	Description
Display of the token	Copies the token to the clipboard when you click it. The display carries the Click to copy note.
Copy button	Copies the token to the clipboard. After the copy, the button briefly has the Copied! label.
Done button	Closes the dialog.



The dialog gives this warning: **Copy this token now – it will not be shown again.** After you close the dialog, Myra AI Workspace does not show the token again. If you lost it, revoke it and create a new one.

The **Your data** section gives you a machine-readable copy of your account data under GDPR Article 15. The section carries this note: **Download a machine-readable copy of your account data (conversations, memories, agents, projects, usage statistics and more) under GDPR Article 15.**

The following option is available:

Option	Description
Export my data button	Creates the export and saves it as the myra-ai-datenexport-.json file. While the export is prepared, the button has the Preparing export... label and is disabled.



If the export fails, the section shows the error above the button. The export does not change your user account. You can repeat it as often as necessary.

The **Danger Zone** section contains the two ways to end your user account. The ways differ in whether your data stays available.

The following options are available:

Option	Description
Delete Account button	Opens the Delete Account dialog. The section carries this note: Deleting your account signs you out and locks you out of MYRA AI. Your data is retained per our privacy policy; only an administrator can restore the account.
Erase all my data button	Opens the Erase all my data dialog. The section carries this note: Permanently erase your account and all associated data (GDPR Article 17, right to erasure). This is irreversible and cannot be undone by you or an administrator.

3.6.1.3 Effects in the Delete Account dialog

The **Delete Account** dialog deactivates your user account without erasing your data.

The confirmation has the following effects:

Object	Effect
Sign-in	Myra AI Workspace signs you out immediately and does not accept a new sign-in.
Data	Your conversations and your other data stay available as specified in the privacy policy.
Restoration	Only users with the Administrator role can restore the user account. You cannot reactivate it yourself.

The following options are available in the dialog:

Option	Description
Delete Account button	Deactivates the user account and signs you out. While the account is deleted, the button has the Deleting... label.
Cancel button	Closes the dialog without deactivating the user account.
privacy policy link	Opens the privacy policy in a new tab.



The dialog gives the note that you erase your data permanently with the **Erase all my data** button on your profile instead.

3.6.1.4 Effects in the Erase all my data dialog

The **Erase all my data** dialog erases your user account and your data permanently under GDPR Article 17.

The confirmation has the following effects:

Object	Effect
User account	Myra AI Workspace erases the user account permanently. Neither you nor users with the Administrator role can restore it.
Content	Myra AI Workspace deletes your conversations, memories, agents, projects, uploaded files, and personal settings.
Record	Myra AI Workspace retains a tamper-evident deletion record without any of your content for compliance.
Tokens	Your access tokens become invalid. Calls to the inference API with these tokens fail afterwards.

You set the following entry in the dialog:

Entry	Description
Type your email () to confirm	Confirmation of the erasure. The Erase everything button stays disabled while the entry is not exactly the email address of your user account.

The following options are available in the dialog:

Option	Description
Erase everything button	Erases the user account and all associated data. While the data is erased, the button has the Erasing... label.
Cancel button	

Option	Description
	Closes the dialog without erasing anything. The confirmation entry is discarded.
	If the erasure fails, the dialog shows the error and you stay signed in. In this case, the user account stays unchanged.
	Myra AI Workspace cannot undo the erasure under GDPR Article 17. Export your data with the Export my data button before you confirm the erasure.

3.6.2 Plans



Figure 117: Self-serve Plans view

The **Self-serve Plans** view manages the plans of the self-registration. A plan defines the available models, the budget, the degrade configuration, and the Stripe price id. The view shows every plan as an area of its own. The subtitle names the number of the plans.

The following information is available per plan area:

Entry	Description
Area title	Name of the plan, for example free , standard , or enterprise .
Budget (USD / month)	Spending limit of a tenant of this plan per month.
Scheduled tasks cap	Maximum number of the scheduled tasks that a tenant of this plan holds.
Seat cap (max users)	

Entry	Description
	Maximum number of the user accounts of a tenant of this plan.
Models	Model ids that the plan releases. If no model id is stored, then the area shows —.
Efficient models (over-cap subset)	Model ids that remain available after the budget is reached.
Fallback (provider:model)	Model that Myra AI Workspace uses if the selected model does not respond.
Degrade provider and Degrade model	Provider and model that Myra AI Workspace degrades to above the budget.
Degrade multiplier	Multiple of the budget from which the degrade applies, for example 1.5x .
Stripe price id (monthly) and Stripe price id (annual)	Price ids with which Stripe bills the plan.
Last updated	Time of the last change of the plan.

The following options are available:

Option	Description
Edit row action	Opens the Edit plan: dialog.

3.6.2.1 Entries in the Edit plan dialog




Figure 118: Edit plan dialog

You change the following entries in the dialog:

Entry	Effect of the change
Budget (USD / month)	Changes the monthly spending limit. The value must be a number greater than or equal to 0.
Scheduled tasks cap	Changes the maximum number of the scheduled tasks. The value must be a whole number of at least 1.
Seat cap (max users)	Changes the maximum number of the user accounts. The value must be a whole number between 1 and 10000.
Models	Defines the released models, one model id per line. Removed models are no longer available to the tenants of this plan.
Efficient models (over-cap subset)	Defines the models that still respond above the budget, one model id per line.

Entry	Effect of the change
Fallback (provider:model)	Defines the replacement model. The value must be in the form provider:model , for example openai:gpt-4o-mini . If you leave the field empty, then the fallback is disabled.
Degrade multiplier	Defines the multiple of the budget from which Myra AI Workspace degrades. The value must be greater than 1.0 and at most 10.0. If you leave the field empty, then the degrade is disabled.
Degrade provider and Degrade model	Define the provider and the model of the degrade. Enter both entries or leave both empty.
Stripe price id (monthly) and Stripe price id (annual)	Change the price ids of the billing.

The following options are available in the dialog:

Option	Description
Save Changes button	Applies the changes. During the saving, the button carries the Saving... label.
Cancel button	Closes the dialog without applying the changes.



If another person is editing the same plan at that moment, then the dialog reports **Another plan edit is in progress. Please retry in a moment**. The view is available only to user accounts with the **Manage plan and entitlement config** permission.



The view needs the **Manage plan and entitlement config** permission of the **Platform** module. This permission stays reserved for the **Administrator** role. A role of your own does not take it.

3.6.3 Vouchers

Vouchers							New voucher
CODE	TYPE	CREDIT	REDEEMED	EXPIRES	ACTIVE	ACTIONS	
TEST-2026-.....	Unique (single use)	\$25	1 / 1	12/31/2026	Yes	Deactivate	
MESSE-.....	Generic (campaign)	\$10	42 / 100	10/15/2026	Yes	Deactivate	
ALT-.....	Generic (campaign)	\$10	8 / 50	Never	No	Reactivate	

Figure 119: Vouchers view

The **Vouchers** view manages redeemable codes. A code grants a new signup a one-time trial credit in US dollars.

The following information is available:

Column	Description
Code	Character string that the signup redeems, for example TRIAL-2026 .
Type	Kind of the code: Generic (campaign) for a code with several redemptions or Unique (single use) for a code with exactly one redemption.
Credit	Trial credit in US dollars that the redemption grants, for example \$50 .
Redeemed	Number of the redemptions in relation to the maximum number, for example 12 / 1000 .
Expires	Date up to which the code can be redeemed. If no date is stored, then the column shows Never .
Active	Redeemability of the code: Yes or No .
Actions	Row actions of the code.

The following options are available:

Option	Description
New voucher button	Opens the New voucher dialog.
Deactivate row action	Opens the Deactivate voucher ? It can no longer be redeemed. prompt. After the confirmation, the code shows No in the Active column.
Reactivate row action	Opens the Reactivate voucher ? prompt. After the confirmation, the code can be redeemed again. The row action appears only for a deactivated code.

3.6.3.1 Entries in the New voucher dialog




Figure 120: New voucher dialog

You define the following entries during the creation:

Entry	Description
Code *	Character string that the signup enters, for example TRIAL-2026 . The entry is required.
Type	Kind of the code: Generic (campaign) or Unique (single use) . Unique codes are single-use. The maximum number of the redemptions limits generic codes.
Credit (USD) *	Trial credit in US dollars, for example 50 . Myra AI Workspace applies the credit once in the first billing period of the new tenant. The entry is required.
Max redemptions	Maximum number of the redemptions of a generic code, for example 1000 . The value must be a whole number of at least 1.
Expires (optional)	Last day on which the code can be redeemed. If you leave the field empty, then the code does not expire.

The following options are available in the dialog:

Option	Description
Create voucher button	Creates the voucher. During the creation, the button carries the Creating... label.
Cancel button	Closes the dialog without creating the voucher.



If no voucher exists, then the view reports **No vouchers yet**. If the loading fails, then the view reports **Failed to load vouchers**.



The view needs the **Manage vouchers** permission of the **Platform** module. This permission stays reserved for the **Administrator** role. A role of your own does not take it.

4. Configurations

4.1 Account

The following configurations concern your own user account. They determine how Myra AI Workspace behaves for this account and apply only to the person who makes them.

The following sections describe these configurations:

- [Personal preferences](#)
- [Managing personal commands](#)
- [Managing PII entries](#)
- [Deleting the account](#)
- [Managing tokens](#)

4.1.1 Personal preferences

The personal preferences set how Myra AI Workspace behaves for your user account.

Proceed as follows to edit the personal preferences:




Figure 121: My preferences tab

- ▶ In the **Settings** area, open the **Preferences** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **My preferences** tab.
- ▶ In the **General** area, set the language, the **Display currency**, the behaviour of the **Enter key**, and your **Default model**.
 - ↳ Myra AI Workspace saves each of these immediately.
- ▶ In the **Personal AI preferences** area, enter what the AI is to call you and what describes your work.
- ▶ Enter how the AI is to respond to you.
- ▶ Click on the **Save preferences** button.

→ Myra AI Workspace shows **Saved. New chats use these preferences.**

4.1.1.1 Entries in the General area

You set the following entries in the area:

Entry	Effect of the change
Language drop-down list	Changes the interface to English or Deutsch . Myra AI Workspace applies the language immediately and saves it in the user account. Emails and notifications also follow it.
Display currency drop-down list	Shows budgets and costs in USD \$ or EUR € . The entry applies to the display only. Myra AI Workspace stores each amount in USD and enforces each budget in USD.
Press Enter to send messages checkbox	Sets whether the Enter key in chat sends the message or adds a line break. If the checkbox is not selected, the Ctrl+Enter key combination sends.
Default model button	Sets the model with which a new chat starts. A project or the message model picker overrides it. Without a choice and with an unavailable model, Myra AI Workspace uses the cheapest available model.



The area saves each entry immediately. There is no button to save. If the save fails, Myra AI Workspace restores the previous value and shows **Could not save your Enter-key preference. Please try again.** or **Your default model could not be saved. Please try again.**

4.1.1.2 Entries in the Personal AI preferences area

You set the following entries in the area:

Entry	Effect of the change
What should I call you? input field	Sets the name that the AI uses for you. The field takes at most 64 characters.
What best describes your work? drop-down list	Sets your field of work, for example Product Management , Engineering , or Legal . The – Not set – entry sends no field of work.

Entry	Effect of the change
How should the AI respond to you? input field	Sets permanent instructions for the style of the replies. The field takes at most 5,000 characters.

The following options are available in the area:

Option	Description
Save preferences button	Saves the three entries. While the entries are saved, the button has the Saving... label. The button is disabled while no entry is changed or the instructions exceed 5,000 characters.



If the instructions read like a question to the AI, the area asks back. The **Cancel** button returns to the input field, and the **Save anyway** button applies the text unchanged. The instructions are no chat – Myra AI Workspace adds them to each new conversation and does not answer them.



If the instructions contain personal data, the area names the types detected and asks back. The **Edit** button returns to the input field, and the **Save anyway** button applies the text unchanged. The instructions go to the model with each new conversation.

4.1.1.3 Effects of the changes

The changes have the following effects:

Object	Effect
New chats	A new chat starts with the default model and takes the name, the field of work, and the instructions into its context.
Ongoing conversations	A conversation already started keeps its model and its context. The changes apply from the next chat only.
Other users	The preferences apply to your user account only. No other user account and no shared project takes them over.

Object	Effect
Budgets and costs	The Display currency changes the display only. A budget remains stored in USD and is enforced in USD.

4.1.2 Managing personal commands

A personal command is available as a slash command in the chat.

4.1.2.1 Creating a personal command

Proceed as follows to create a personal command:





Figure 122: My Commands tab

- ▶ In the **Settings** area, open the **Preferences** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **My Commands** tab.
- ▶ In the **Personal commands** section, click on the **New command** button.
 - ↳ The **New Command** dialog opens.



New Command

Name *

Description

Category

Template *

Use {{variable}} for placeholders. Example: Translate to {{language}}: {{text}}

Figure 123: New Command dialog

- ▶ In **Name**, enter the name of the command.
- ▶ In **Template**, enter the prompt.
- A placeholder of the form `{{variable}}` is asked for when the command is used in the chat.
- ▶ Click on the **Create** button.
- Myra AI Workspace creates the command in the **Personal commands** section. It is available in the chat as a slash command.

4.1.2.1.1 Entries in the New Command dialog

You set the following entries at creation:

Entry	Description
Name	Name of the command without the leading slash, for example translate . Myra AI Workspace replaces spaces with hyphens. The entry is required.
Description	Short description that appears after the command in the picker of the chat. The entry is optional.

Entry	Description
Category	Category by which the list in the My Commands tab filters. The entry takes at most 64 characters and is optional.
Template	Prompt that the command inserts into the chat. A placeholder of the form <code>{{variable}}</code> is asked for on each call. The entry is required.

The following options are available in the dialog:

Option	Description
Variables detected: display	Lists the placeholders that Myra AI Workspace finds in the template. The display appears as soon as the template contains a placeholder.
Create button	Creates the command. While the command is saved, the button has the Saving... label. The button is disabled while Name or Template is empty.
Cancel button	Closes the dialog without creating the command.



The name is unique within your personal commands. If the name is taken already, Myra AI Workspace does not create the command and reports the conflict.

4.1.2.1.2 Entries in the Edit Command dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name	Changes the shortcut with which you call the command in chat. The previous shortcut is no longer available.
Description	Changes the text in the picker of the chat.
Category	Changes the assignment in the Filter by category drop-down list.

Entry	Effect of the change
Template	Changes the prompt that the command inserts. The change applies from the next call. Messages sent already remain unchanged. If you remove a placeholder, the chat no longer asks for it.
	Myra AI Workspace creates a new version with each save. The Versions row action opens the Version history – / dialog and restores an earlier version with the Restore row action. The restore creates a new version in turn.
	A command that another person shared with you carries the Shared with you label and cannot be edited. Contact the person who created it, or your tenant administration for a command of the tenant.

4.1.2.2 Deleting a personal command

Proceed as follows to delete a personal command:




Figure 124: My Commands tab

- ▶ In the **Settings** area, open the **Preferences** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **My Commands** tab.
- ▶ In the **Personal commands** section, delete the command.
 - ↳ The **Delete this command?** query opens.




Figure 125: Delete command query

- ▶ Confirm the query.
- Myra AI Workspace deletes the command. It is no longer available in the chat.

4.1.2.2.1 Effects of the deletion of a command


The deletion has the following effects:

Object	Effect
Command	Myra AI Workspace removes the command from the Personal commands section. A restore is not possible.
Version history	The earlier versions of the command are removed with it.
Sharing	The sharing ends. The command is no longer available to the persons with whom you shared it.
Chat	The <code>/<name></code> shortcut is no longer available in the picker. The name is free again for a new command.
Conversations	The messages that the command inserted remain in the conversations.

4.1.3 Managing PII entries

A PII entry masks a word or a phrase in your messages before Myra AI Workspace sends them to the provider. The entry applies in addition to the detectors configured on the gateway.

Proceed as follows to create a PII entry:



The screenshot shows the 'My PII configuration' tab with the following elements:

- Navigation tabs: My preferences, My Commands, **My PII configuration**
- Section title: **My PII configuration**
- Description: Words and phrases you list here will be masked in your messages before they leave Myra, in addition to the gateway's pre-configured PII detectors.
- Form area:
 - Label: Add a keyword
 - Input field: e.g. ProjectAtlas, ClientXYZ, ...
 - Buttons: Add (blue), Case-sensitive (checked), Whole word only (checked)
- Table of configured keywords:

ADD A KEYWORD	CASE-SENSITIVE	WHOLE WORD ONLY	
ProjectAtlas	✓	✓	🗑️
ClientXYZ	–	✓	🗑️

Figure 126: My PII configuration tab

- ▶ In the **Settings** area, open the **Preferences** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **My PII configuration** tab.
- ▶ In **Add a keyword**, enter the word or the phrase.
- ▶ If necessary, set the **Case-sensitive** and **Whole word only** switches.
- ▶ Click on the **Add** button.
- Myra AI Workspace replaces the keyword with a token of the form [MYRA-CUSTOM:...]
before it is sent.

4.1.3.1 Entries for the PII entry

You set the following entries at creation:

Entry	Description
Add a keyword	Word or phrase that Myra AI Workspace masks in your messages. The entry takes at most 200 characters and bare words only. The punctuation @ . - _ is permitted.
Case-sensitive	Sets that Myra AI Workspace masks the spelling entered only. Without the entry, Myra AI Workspace masks each spelling.
Whole word only	Sets that Myra AI Workspace masks the keyword as an independent word only. Without the entry, Myra AI Workspace also masks it as part of a longer word.

The entry has the following effects:

Object	Effect
New messages	Myra AI Workspace replaces the keyword with a token of the form [MYRA-CUSTOM:...] before it sends the message. The model does not see the keyword.
Gateway detectors	The entry adds to the detectors for personal data configured at the gateway. It does not replace them.
Messages sent already	The messages remain unchanged. The masking applies from the next message only.
Other users	The list applies to your user account only.



For the **Tenant administrator** role, the **Add a list** area takes up to 20 words at once – one word per line – or reads a .txt file. Myra AI Workspace scans the file for malware before it reads the file and reports each rejected entry with the reason.

4.1.3.2 Effects of the deletion of a keyword

The deletion has the following effects:

Object	Effect
Keyword	Myra AI Workspace removes the keyword from the list.
New messages	Myra AI Workspace no longer masks the word. The model sees it in clear text from the next message, unless a detector of the gateway applies.
Messages sent already	The masking remains in the messages sent already.

4.1.4 Deleting the account

The deletion of the user account ends your access to Myra AI Workspace. The conversations and the other data are retained per the privacy policy. Only users with the **Administrator** role restore the account.

4.1.4.1 Deleting the account

Proceed as follows to delete your user account:




Figure 127: Danger zone section

- ▶ In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- ▶ In the **Danger Zone** section, click on the **Delete Account** button.
 - ↳ The **Delete Account** dialog opens.




Figure 128: Delete Account dialog

- ▶ Click on the **Delete Account** button.
- Myra AI Workspace signs you out and shows the sign-in.




Figure 129: Sign-in

4.1.4.1.1 Effects of the deletion of the user account

The **Delete Account** dialog in the **View: My Account** section lists the effects of the deletion in detail – on the sign-in, on your data, and on the restore.

4.1.4.2 Erasing all your data

The erasure of all your data removes the user account and its content permanently.

The following requirements are met:

- ☒ You know the email address of your user account.

Proceed as follows to erase all your data:




Figure 130: Danger zone section

- ▶ In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- ▶ In the **Danger Zone** section, click on the **Erase all my data** button.
 - ↳ The **Erase all my data** dialog opens.




Figure 131: Erase all my data dialog

- ▶ Enter your email address to confirm.
- ▶ Click on the **Erase everything** button.
- Myra AI Workspace signs you out and shows the sign-in.




Figure 132: Sign-in



Myra AI Workspace does not undo the deletion under Article 17 GDPR. Users with the **Administrator** role also do not restore the data deleted. A tamper-evident deletion record without any of your content is retained for compliance.



The **Erase all my data** dialog in the **View: My Account** section lists the effects of the erasure in detail – on the user account, on your content, on the deletion record, and on your tokens. Export your data with the **Export my data** button before you confirm the erasure.

4.1.5 Managing tokens


A personal access token uses the inference API outside the interface.

4.1.5.1 Creating a token

The following requirements are met:

- ☒ At least one gateway exists for your tenant.

Proceed as follows to create a token:



My Account
Your profile, sign-in, tokens, and data export.

Profile

Email: alex.muster@beispiel.de Role: member

Tokens [New Token](#)

LABEL	GATEWAY	EXPIRES	BUDGET	RATE LIMIT	CREATED
Editorial script	redaktion	12/31/2026	\$12.40 / \$50.00 - \$37.60 left	60/60 s	4/2/2026
Analysis	produktion	never	\$10.00 - usage unknown	10/60 s	6/11/2026

Your data

Download a machine-readable copy of your account data (conversations, memories, agents, projects, usage statistics and more) under GDPR Article 15.

[Export my data](#)

Danger Zone

Deleting your account signs you out and locks you out of MYRA AI. Your data is retained per our privacy policy; only an administrator can restore the account.


[Delete Account](#)

Permanently erase your account and all associated data (GDPR Article 17, right to erasure). This is irreversible and cannot be undone by you or an administrator.

[Erase all my data](#)

Figure 133: My Account view

- In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- In the **Tokens** section, click on the **New Token** button.
 - ↳ The **New Token** dialog opens.



The 'New Token' dialog box is a light gray rectangular window with a white background. It contains several input fields and buttons. At the top, the title 'New Token' is in bold. Below it, the 'Gateway *' field is a dropdown menu showing 'muster-gmbh/produktion'. The 'Label' field is a text input containing 'dev laptop'. The 'Expires At' field is a date/time picker showing 'mm/dd/yyyy --:--', with a note below it: 'Leave blank = never expires'. The 'Budget (USD)' field is a text input containing 'unlimited'. The 'Rate limit (req/window)' field consists of two sub-fields: 'none' and '60s', with a note below them: 'requests / window_sec (blank = gateway default)'. At the bottom right, there are two buttons: 'Cancel' and 'Create Token'.

Figure 134: New Token dialog

- ▶ Select the gateway and enter the label.
- ▶ If necessary, set the expiry date, the budget, and the rate limit.
- ▶ Click on the **Create Token** button.
- Myra AI Workspace creates the token and shows it one time.




Figure 135: Token Created dialog



A token always applies to exactly one gateway. If no gateway exists, Myra AI Workspace shows **No gateways available** and creates no token.

4.1.5.1.1 Data in the New Token dialog

You set the following data at creation:

Item	Description
Gateway *	Gateway that the token applies to. A token always applies to exactly one gateway. The item is mandatory.
Label	Label of the token, for example dev laptop . The label appears in the Label column of the list. The item is optional.
Expires At	Date on which the token loses its validity. An empty field means: never expires.
Budget (USD)	Amount in USD that requests with this token may consume. An empty field means: unlimited.

Item	Description
Rate limit (req/window)	Number of the requests per window seconds. An empty field applies the default of the gateway.

The following options are available in the dialog:

Option	Description
Create Token button	Creates the token and shows it once. During the creation, the button shows Creating...
Cancel button	Closes the dialog without creating the token.



Myra AI Workspace shows the token only immediately after the creation. Copy it before you close the dialog. For the **Viewer** role, the area shows **Viewer users cannot make inference requests.**

4.1.5.2 Revoking a token

In the **My Account** view, a token cannot be edited. You revoke it and create a new one. The administration changes the spend cap, the rate limit, and the expiry of a token in the detail view of the gateway.

Proceed as follows to revoke a token:




Figure 136: My Account view

- In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- In the **Tokens** section, click on the **Revoke** button at the token.
 - ↳ The **Revoke this token?** query opens.




Figure 137: Revoke token query

- ▶ Confirm the query.
- The token loses its validity. The gateway answers requests with this token with **Missing or invalid gateway token.**

4.1.5.2.1 Effects of the revocation

The revocation has the following effects:

Object	Effect
Token	Myra AI Workspace removes the token from the list. The revocation cannot be undone.
Requests	The gateway answers requests with this token with Missing or invalid gateway token.
Budget and rate limit	The data of the token ends with the token. A new token receives its own data.

The following options are available in the confirmation:

Option	Description
Revoke button	Revokes the token.
Cancel button	Closes the confirmation without revoking the token.

4.2 Workspace

The workspace holds the content that is used in the chat:

- Projects bundle related chats, knowledge files, and instructions.
- Agents and workflows run recurring tasks.
- The library keeps the prompts available.

The following configurations create these objects, change them, and remove them again.

The following sections describe these configurations:

- **Managing projects**
- **Managing agents**
- **Managing workflows**
- **Managing scheduled tasks**
- **Managing the prompts of the library**

4.2.1 Managing projects

A project holds related chats, files, and its own instructions together.

4.2.1.1 Creating a project

The following requirements are met:

- ☒ Your user account has the **Create projects** permission.

Proceed as follows to create a project:





Figure 138: Projects view

- ▶ On the start page, open the **Projects** view in the **Workspace** work area.
- ▶ Click on the **New Project** button.
 - ↳ The **New Project** dialog opens.



The 'New Project' dialog box is a white rectangular window with a thin blue border. It contains several sections for configuring a new project. At the top, there's a title 'New Project'. Below it, the 'Icon' section shows a 3x3 grid of icons, with the first icon (a folder) selected. The 'Color' section shows a row of 10 colored circles, with the first one (blue) selected. The 'Name' section has a text input field with the placeholder 'e.g. Research Project'. The 'Description' section has a text input field with the placeholder 'Short description (optional)'. The 'Project Instructions' section has a larger text input field with the placeholder 'Custom system prompt for all conversations in this project (optional)'. The 'Access tier' section has a dropdown menu with 'User configurable' selected. Below this, there's a small note: 'All models available: people choose PII protection per chat (today's behaviour)'. The 'Preferred model' section has a dropdown menu with 'Automatic (no preference)' selected. Below this, there's another small note: 'Preselects this model for new chats in the project. Members can still switch model per chat.' At the bottom right, there are two buttons: 'Cancel' and 'Create Project'.

New Project

Icon

Color

Name *

e.g. Research Project

Description

Short description (optional)

Project Instructions

Custom system prompt for all conversations in this project (optional)

Access tier

User configurable

All models available: people choose PII protection per chat (today's behaviour).

Preferred model

Automatic (no preference)

Preselects this model for new chats in the project. Members can still switch model per chat.

Cancel Create Project

Figure 139: New Project dialog

- ▶ Enter the name of the project.
- ▶ If necessary, set the icon, the color, the description, the project instructions, and the access tier.
- ▶ Click on the **Create Project** button.
- Myra AI Workspace creates the project and opens it.




Figure 140: View of a project

4.2.1.1.1 Data in the New Project dialog

You set the following data at creation:

Item	Description
Icon	Icon of the project. The icon marks the project in the list and in the detail view.
Color	Color of the icon. The color applies to the display only.
Name *	Name of the project, for example Research Project . The item is mandatory.
Description	Short description of the project. The description appears in the Description column of the Projects view and in the Overview tab. The item is optional.
Project Instructions	Custom system prompt for all conversations in this project. The AI applies the prompt in every conversation of the project. The item is optional.
Access tier	Sets which models the project uses and whether PII protection can be turned off. The default is User configurable .

The following options are available in the dialog:

Option	Description
Create Project button	Creates the project and opens it. During the creation, the button shows Creating....
Cancel button	Closes the dialog without creating the project.



If the name field stays empty, the dialog shows **Name is required.**

4.2.1.1.2 Access tiers

The access tier applies to every conversation of the project.

Access tier	Effect
User configurable	All models available. People choose PII protection per chat.
PII protection required	All models available, but PII protection is always on and cannot be turned off in this project.
Local only	Only local Myra models. Data never leaves Myra infrastructure. Cloud models, web search, and external tools are disabled.

4.2.1.2 Editing a project

Proceed as follows to edit a project:




Figure 141: Projects view

- ▶ On the start page, open the **Projects** view in the **Workspace** work area.
- ▶ Open the project.
 - ↳ The view of the project opens.




Figure 142: View of a project

- ▶ Click on the **Edit project** button.
 - ↳ The **Overview** tab shows the data of the project as a form.
- ▶ Change the data.
- ▶ Click on the **Save changes** button.
- Myra AI Workspace applies the changes for every conversation of the project.




Figure 143: View of a project

4.2.1.2.1 Data in the Edit project form

You change the following data in the form:

Item	Effect of the change
Name	Changes the name in the list of the projects, in the detail view, and in the selection of the project in the chat.
Icon	Changes the icon of the project in the list and in the detail view.
Color	Changes the color of the icon.
Description	Changes the short description in the Description column and in the Overview tab.
Project Instructions	Changes the system prompt for all conversations of the project. The change applies from the next message, in existing conversations as well.
Access tier	Changes the models available and the PII protection of the project. If you select the Local only access tier, only local Myra models are available for selection.
Conversation retention (days)	Permanently deletes this workspace's conversations with no activity for this many days (30–3650). To disable the deletion, leave the field empty. A workspace period may only shorten retention versus the organization policy.

Item	Effect of the change
Preferred model	Preselects this model for new chats in the project. Members can still switch model per chat. With the Automatic (no preference) selection, the project sets no model.

The following options are available in the form:


Option	Description
Save changes button	Applies the changes. During the saving, the button shows Saving....
Cancel button	Closes the form without applying the changes.



The **Owner** and **Editor** project roles and the **Administrator** role open the form. The **Conversation retention (days)** field appears only for the **Owner** project role and for the **Administrator** role. The deletion is irreversible.

4.2.1.3 Deleting a project

Proceed as follows to delete a project:



Projects					
New Project + Last edited Date created					
PROJECT	DESCRIPTION	MEMBERS	FILES	ROLE	UPDATED
Quotes 2026	Templates and price blocks	4	12	Owner	8/12/2026
Handbook new edition	Chapters and terminology	3	27	Editor	8/9/2026
Market watch	Sources and summaries	2	8	Viewer	7/30/2026

Figure 144: Projects view

- On the start page, open the **Projects** view in the **Workspace** work area.
- Open the project.
 - ↳ The view of the project opens.




Figure 145: View of a project

- Click on the **Delete project** button.
 - ↳ The **Delete Project?** dialog opens with the **This cannot be undone.** hint.




Figure 146: Delete Project dialog

- Click on the **Delete project** button.
- Myra AI Workspace deletes the project.

4.2.1.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Project	Myra AI Workspace deletes the project permanently.
Members	Myra AI Workspace removes all members of the project.
Knowledge files	Myra AI Workspace removes all files of the project.
Conversations	Myra AI Workspace detaches the conversations from the project, but does not delete them.

The following options are available in the dialog:

Option	Description
Delete Project button	Deletes the project. During the deletion, the button shows Deleting....
Cancel button	Closes the dialog without deleting the project.



The **Owner** project role and the **Administrator** role delete the project. If the deletion fails, the dialog shows **Delete failed** with the cause.

4.2.2 Managing agents


An agent holds instructions, a model, tools, schedules, and the run history together.

4.2.2.1 Creating an agent

The following requirements are met:

- ☑ Your user account has the **Create and manage own agents** permission.
- ☑ At least one model is enabled on the gateway.

Proceed as follows to create an agent:



Agents

Saved agents: instructions, model, tools, schedules and run history.

Agent catalog **New Agent +**


All Active Scheduled Draft

Search agents by name... ☐ Shared with me Any tool ▾

NAME	MODEL	TOOLS	USED	SCHEDULE	VERSION	UPDATED
Daily briefing tagesbriefing	claude-sonnet-4-5	web_search	128	in 3 hours	v3	8/15/2026
Minutes assistant protokoll-assistent	gpt-4o-mini	files	46	–	v1	8/11/2026
Repo research repo-recherche	mistral-large	fetch, mcpv2	12	in 4 days	v2	7/28/2026

Figure 147: Agents view

- On the start page, open the **Agents** view in the **Workspace** work area.
- Click on the **New Agent** button.
 - ↳ The **New Agent** dialog opens.




The 'New Agent' dialog box contains the following fields and sections:

- Name:** A text input field.
- Instructions (system prompt):** A large text area for entering instructions.
- Model:** A dropdown menu currently showing 'claude-sonnet-4-5'. Below it, a note states: 'The model this agent runs on. Models your tenant can run are shown; toggle off "only runnable" in the picker to browse the full catalog.'
- Tools:** Two checkboxes: 'Web search' and 'URL fetch (agentic)'. Both are currently unchecked.
- Connectors this agent can use:** A section with a note: 'The agent runs these as its owner; each connector's tool permissions still apply.'
- Knowledge areas:** A section with a note: 'Select one or more knowledge areas. The agent searches across all of them and cites the source file, page and section.'
- Advanced:** A section header with a right-pointing arrow.
- Buttons:** 'Cancel' and 'Create' buttons at the bottom right.

Figure 148: New Agent dialog

- ▶ Enter the name and the slug and select the model.
- ▶ Set the instructions and the maximum number of the tokens.
- ▶ Click on the **Create** button.
- Myra AI Workspace creates the agent and opens its detail view.



The 'Daily briefing' agent detail view includes the following elements:

- Header:** A 'Back' button and the title 'Daily briefing'. On the right are 'Edit', 'Share', 'Duplicate', and 'Delete' buttons.
- Tabs:** 'Configuration' (active), 'Try it', 'Schedules', 'Webhooks', 'Runs', and 'Versions'.
- Configuration Panel:**
 - Slug:** tagesbriefing
 - Model:** claude-sonnet-4-5
 - Provider:** anthropic
 - Max tokens:** 4096
 - Tools:** web_search
 - Version:** v3
 - Instructions (system prompt):** Summarize the news of the day in five points. Give the source for each point.

Figure 149: Detail view of an agent

4.2.2.1.1 Entries in the New Agent dialog

You set the following entries when you create the agent:

Entry	Description
Name *	Name of the agent in the list of agents and in the detail view. The entry is mandatory.
Instructions (system prompt)	System prompt that controls each run of the agent. The text sets the task, the tone and the limits of the agent.
Model *	Model the agent runs on. The drop-down list shows the models your tenant can run. If you switch off the only runnable setting in the picker, you browse the full model catalog. The entry is mandatory.
Tools: Web search	The agent searches the web to support its answer.
Tools: URL fetch (agentic)	The agent fetches the content of the addresses it selects itself.
Tools: Include sources list in the answer	The agent appends the sources it used to its answer.
Connectors this agent can use	MCP connectors the agent calls. The agent runs these as its owner. The tool permissions of each connector still apply. If no connector exists yet, the dialog shows No connectors yet. and points to the MCP Connectors page.
Knowledge areas	Knowledge areas the agent searches. The agent searches across all selected areas and cites the source file, the page and the section. If no knowledge area exists yet, the dialog shows No knowledge areas yet. and points to the Projects page.
Advanced: Slug *	Stable handle in the invoke URL of the agent. The entry is mandatory.
Advanced: Max tokens	Upper limit of the tokens per answer of the agent. The value is a positive integer.

The following options are available in the dialog:

Option	Description
Advanced group	Expands and collapses the Slug and Max tokens entries.
Test run (preview) area	Runs the draft against a sample prompt before you save it. Model, instructions, web search, agentic fetch, connectors and knowledge areas take effect and run under your user account. The Run preview button starts the test run.
Create button	Creates the agent. While the agent is created, the button shows the label Saving....
Cancel button	Closes the dialog without creating the agent.



The dialog requires the name, the slug and the model. If an entry is missing, it shows **Name is required**, **Slug is required** or **Model is required**. If the **Max tokens** entry is not a positive integer, it shows **Max tokens must be a positive integer**.



If the selected model does not support tools, the dialog shows **The selected model does not support tools. Pick a tool-capable model to enable web search, URL fetch, connectors and knowledge areas**. If the Playground is switched off for the tenant, the **Test run (preview)** area shows the message **The preview requires the Playground, which is disabled for this tenant**.

4.2.2.2 Editing an agent

Proceed as follows to edit an agent:




Figure 150: Agents view

- On the start page, open the **Agents** view in the **Workspace** work area.
- Open the agent.
 - ↳ The detail view of the agent opens.



Figure 151: Detail view of an agent

- Click on the **Edit Agent** button.

- Change the instructions, the model, the tools, the schedules, or the webhook triggers.
- Click on the **Save** button.
- Myra AI Workspace creates a new version of the agent.



Figure 152: Detail view of an agent

4.2.2.2.1 Entries in the Edit Agent dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name	Changes the name in the list of agents and in the detail view.
Instructions (system prompt)	Changes the behaviour of the agent from the next run on. Completed runs stay unchanged.
Model	Changes the model the agent runs on. If the new model does not support tools, web search, URL fetch, connectors and knowledge areas are no longer available.
Tools	Extends or limits the tools of the agent from the next run on.
Connectors this agent can use	Changes the MCP connectors of the agent. Connectors you do not manage are kept when you save. The dialog shows a note about them.

Entry	Effect of the change
Knowledge areas	Changes the knowledge areas the agent searches. Bindings you do not manage are kept when you save. The dialog shows a note about them.
Advanced: Slug	The entry is locked. The slug cannot be changed after creation, so that the invoke URL of the agent stays valid.
Advanced: Max tokens	Changes the upper limit of the tokens per answer of the agent.

The following options are available in the dialog:

Option	Description
Test run (preview) area	Runs the changed draft against a sample prompt before you save it.
Save button	Applies the changes and creates a new version of the agent. While the agent is saved, the button shows the label Saving...
Cancel button	Closes the dialog without applying the changes.



Each save creates a new version. The **Versions** tab of the detail view lists the versions. The **Restore this version** row action creates a new version with the selected configuration and overwrites nothing.

4.2.2.3 Deleting an agent

Proceed as follows to delete an agent:




Figure 153: Agents view

- On the start page, open the **Agents** view in the **Workspace** work area.
- Open the agent.
 - ↳ The detail view of the agent opens.



Figure 154: Detail view of an agent

- Click on the **Delete** button.
 - ↳ The **Delete agent** ""? query opens.




Figure 155: Delete agent query

- ▶ Confirm the query.
- Myra AI Workspace deletes the agent with its schedules and webhook triggers.

4.2.2.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Agent	Myra AI Workspace removes the agent from the list of agents. The invoke URL of the slug no longer responds.
Schedules	Myra AI Workspace removes the schedules of the agent. Scheduled runs no longer take place.
Webhook triggers	Myra AI Workspace removes the webhook triggers of the agent. Calls to the trigger addresses fail from then on.
Runs	The Runs tab with the run history of the agent is no longer available.
Shares	The shares to people, groups or the organization end. The agent no longer appears in their list of agents.



If the deletion fails, the detail view shows **The agent could not be deleted.**

4.2.3 Managing workflows

A workflow chains your agents into a multi-stage sequence with branches, approvals, and delivery.

4.2.3.1 Creating a workflow

The following requirements are met:

- ☑ Your user account has the **Create and run own workflows** permission.
- ☑ Your administration has switched on the workflows for the workspace.

Proceed as follows to create a workflow:




Figure 156: Workflows view

- On the start page, open the **Workflows** view in the **Workspace** work area.
- Click on the **New workflow** button.
 - ↳ The **New workflow** dialog opens.


A modal dialog titled "New workflow" with a "Name" label and a text input field containing "e.g. Press review with approval". At the bottom right are "Cancel" and "Create" buttons.

Figure 157: New workflow dialog

- ▶ In the **Name** text field, enter the name.
- ▶ Click on the **Create** button.
- Myra AI Workspace creates the workflow and opens the editor.


A workflow editor interface. At the top, it shows "All workflows" with a dropdown menu set to "Review quote", followed by "Run budget (EUR)" with a dropdown set to "no cap", and a "Saved" status. On the right are tabs for "Runs", "AI copilot", "Versions", "Publish", and "Run now". The main area shows a single step named "Trigger" with a plus icon below it. A sidebar on the right titled "Trigger" shows "Type: Manual" and a description: "Manual only – no automatic trigger configured. Add a Schedule, Webhook, or Form below." Below this are buttons for "Schedule", "Webhook", and "Form".

Figure 158: Editor of a workflow

4.2.3.1.1 Entries in the New workflow dialog

You set the following entries when you create the workflow:

Entr y	Description
Name *	Name of the workflow in the list of the workflows, for example Press review with approval . The entry is mandatory.

The following options are available in the dialog:

Option	Description
Create button	Creates the workflow and opens its builder. The button is locked while the Name input field is empty.
Cancel button	Closes the dialog without creating the workflow.



Myra AI Workspace creates the workflow on the access tier that is selected in the **Workflows** view. An agent step uses only agents of the same access tier. If the creation fails, the dialog shows **The workflow could not be created.**



A new workflow carries the **Draft** status and does not run yet. Only the **Publish** button in the builder makes it runnable.

4.2.3.2 Editing a workflow

Proceed as follows to edit a workflow:




Figure 159: Workflows view

- ▶ On the start page, open the **Workflows** view in the **Workspace** work area.
- ▶ Open the workflow.
 - ↳ The editor of the workflow opens.

4.2.3 Managing workflows




Figure 160: Editor of a workflow

- ▶ Click on the **Add step** button and select the step type.
- ▶ Connect the steps and set their settings.
- The editor shows the save state as **Saved**.



Figure 161: Editor of a workflow



If the editor shows **Changed in another tab**, click on the **Reload** button. If you do not, the save operation writes over the change of the other tab.

4.2.3.2.1 Entries for the workflow

You change the following entries in the toolbar of the builder:

Entry	Effect of the change
Name	Changes the name in the list of the workflows. The builder saves the change automatically.
Run budget (EUR)	Limits the cost of a single run. If the run exceeds the budget, Myra AI Workspace stops it. If the field stays empty, no limit applies.

4.2.3.2.2 Entries per step type

You set the following entries on the selected step:

Step	Entries
Trigger	Type: Manual , Schedule , Webhook or Form . Form fields: label, key, type (Text , Number , Select , File upload) and Required for each field. For a file field also Allowed file types (PDF, PNG, JPEG, WebP, Text, Word (DOCX), Excel (XLSX)) and Max size per file (MB, 0 = default 10 MB) . At most five file fields are allowed per form.
Agent step	Agent: the agent that runs the step. Input for this step: the task the agent receives. It may reference earlier steps. The system prompt of the agent is set on the agent, not here.
Condition	Value , Comparison (equals , not equal , contains , greater than , less than , is empty) and Compare with . The result leads into the If true branch or into the Otherwise branch.
Loop	Iterate over: the list whose items run through the For each item branch once each, up to ten items per run.
Approval	Approver: the user account that decides. It finds the request in the My Approvals view. If not decided in time: End the run (nothing goes out) or Continue automatically (auto-approve) . Timeout (hours): period until the decision.
Deliverer	Channel: Mattermost with the Mattermost channel ID (26 characters), Webhook with the Webhook URL (https) or Email with a Subject . The email goes to the recipient configured for the workspace. Message: the text to deliver. An empty message delivers nothing.

4.2.3.2.3 Error handling of a step

You set the following entries in the **Error handling** area:

Entry	Effect of the change
If this step fails	

Entry	Effect of the change
	Stop the run (default) ends the run. Continue without this step's result continues the run without the result. Run the error branch runs the steps in the On error branch.
Attempts (1–5)	Number of the attempts before the step counts as failed.
Seconds between attempts (0–30)	Waiting time between two attempts.

4.2.3.2.4 Entries in the Schedule dialog

You set the following entries in the dialog:

Entry	Effect of the change
Trigger	Every fixed interval runs the workflow at the selected interval. Once a day (UTC) runs it at a set time.
Interval and Interval unit	Distance between two runs. The value lies between one minute and 366 days.
Time of day (UTC)	Time of the daily run in the format HH:MM in UTC. The dialog shows the matching local time below it.
Schedule enabled	Switches the automatic run on and off. A switched-off schedule starts no run. The manual run stays available.

The following options are available in the dialog:

Option	Description
Save schedule button	Applies the schedule. While the schedule is saved, the button shows the label Saving....
Remove schedule button	Deletes the schedule. The workflow then runs only manually or through another trigger.



If several scheduled runs fail in a row, Myra AI Workspace switches the schedule off and shows **This schedule was turned off automatically**. Fix the cause, switch the schedule back on and save – only the saving resumes the automatic runs.

4.2.3.2.5 Entries in the Webhook dialog

The following options are available in the dialog:

Option	Description
Create webhook button	Creates the webhook. The dialog then shows the Webhook URL and the Secret .
Secret entry	Is shown only once and cannot be retrieved later. Copy the secret immediately.
Delete webhook button	Deletes the webhook. Calls to the address no longer start a run.



An external system starts a run by POSTing to the webhook URL with the secret header. The JSON body becomes the output of the trigger step and is available to the following steps.

4.2.3.2.6 Entries in the Form dialog

You set the following entries in the dialog:

Entry	Effect of the change
Form enabled	Switches the public form page on and off. The form works only while the workflow is published.
Require anti-spam check (CAPTCHA)	Requires a check against automated input before the form is submitted.

The following options are available in the dialog:

Option	Description
Create form button	Creates the public form page and shows the Public form URL .
Regenerate link button	Assigns a new address. The previous address then leads nowhere.
Delete form button	Deletes the public form page.



Anyone with the link submits the form without signing in and thereby starts a run that drives the agents of this workflow with their full tool access. Only enable the form on a workflow you trust with anonymous input.

4.2.3.3 Deleting a workflow

Proceed as follows to delete a workflow:




Figure 162: Workflows view

- ▶ On the start page, open the **Workflows** view in the **Workspace** work area.
- ▶ Click on the **Delete** button at the workflow.
 - ↳ The **Delete this workflow and all its runs?** query opens.




Figure 163: Delete workflow query

- ▶ Confirm the query.
- Myra AI Workspace deletes the workflow with all its runs.

4.2.3.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Workflow	Myra AI Workspace removes the workflow from the list of the workflows. The builder is no longer available.
Runs	Myra AI Workspace deletes all runs of the workflow with their history and their cost.
Versions	The published versions are removed. A roll back to an earlier version is no longer possible.
Triggers	Schedule, webhook and public form are removed. Calls to the webhook address and to the form address no longer start a run.
Open approvals	Approvals that wait for a decision are removed. They no longer appear in the My Approvals view.
Agents	The agents of the steps stay. The deletion of the workflow deletes no agent.



If the deletion fails, the view shows **The workflow could not be deleted.**

4.2.4 Managing scheduled tasks


A scheduled task runs a prompt on a schedule.

4.2.4.1 Creating a scheduled task

The following requirements are met:

- ☒ Your user account has the **Create and manage own scheduled tasks** permission.


Proceed as follows to create a scheduled task:



Scheduled tasks					
NAME	PROMPT	SCHEDULE	MODEL	GATEWAY	STATUS
Daily briefing	Summarise today's news	Daily	claude-sonnet-4-5	produktion	Last run OK
Weekly review	Summarise the week in five points	Weekly	gpt-4o-mini	redaktion	Paused

Figure 164: Scheduled view

- ▶ On the start page, open the **Scheduled** view in the **Workspace** work area.
- ▶ Click on the **New scheduled task** button.
 - ↳ The **New scheduled task** dialog opens.



The image shows a 'New scheduled task' dialog box. It contains the following fields and options: 'Name' (text input), 'Prompt' (text area with example text 'e.g. Summarize today's AI news in five bullet points'), 'Model' (dropdown menu showing 'claude-sonnet-4-5'), 'Gateway' (dropdown menu showing 'produktion'), 'Schedule' (dropdown menu showing 'Daily'), 'Time' (text input showing '06:00'), and 'Also send by email (optional)' (text input showing 'you@example.com'). There are 'Cancel' and 'Save' buttons at the bottom right. Small text below the 'Time' field explains that the time is local and stored in UTC, and that the first run starts shortly after creation.

Figure 165: New scheduled task dialog

- ▶ Enter the prompt and set the schedule.
 - ▶ If necessary, switch on the delivery of the result by email.
 - ▶ Save the task.
- Myra AI Workspace runs the prompt on the schedule. The result appears in a conversation.

4.2.4.1.1 Entries in the New scheduled task dialog

You set the following entries when you create the task:

Entry	Description
Name *	Name of the task in the list of the scheduled tasks. The entry is mandatory.
Prompt *	Prompt Myra AI Workspace sends to the model on each run, for example Summarize today's AI news in five bullet points . The entry is mandatory.
Model *	Model that answers the prompt. The entry is mandatory.

Entry	Description
Gateway	Gateway the task runs on. The drop-down list appears only when you create the task and only if more than one gateway is available.
Schedule	Daily, Weekly or Interval. The selection decides which further entries the dialog shows.
Run every and Unit	Distance between two runs for the Interval selection, in minutes, hours or days . The value lies between one minute and 366 days.
Weekday	Weekday of the run for the Weekly selection.
Time	Time of the run in the format HH:MM for the Daily and Weekly selections. You enter your local time. Myra AI Workspace stores it in UTC. Across a daylight-saving change the run time shifts by one hour.
Also send by email (optional)	Email address Myra AI Workspace sends the result to in addition. If the field stays empty, the result arrives in the conversation only.

The following options are available in the dialog:

Option	Description
Save button	Creates the task. While the task is saved, the button shows the label Saving....
Cancel button	Closes the dialog without creating the task.



The first run starts shortly after creation. After that the task follows its schedule.



If an entry is missing or invalid, the dialog shows **Please enter a name., Please enter a prompt., Please choose a model., The interval must be at least 1 minute and at most 366 days.** or **Please enter a time as HH:MM.** If the model does not belong to your plan, the dialog shows **This model is not available on your plan.**

4.2.4.2 Editing a scheduled task

Proceed as follows to edit a scheduled task:

Scheduled tasks

New scheduled task

NAME	PROMPT	SCHEDULE	MODEL	GATEWAY	STATUS
Daily briefing	Summarise today's news	Daily	claude-sonnet-4-5	produktion	Last run OK
Weekly review	Summarise the week in five points	Weekly	gpt-4o-mini	redaktion	Paused

Figure 166: Scheduled view

- ▶ On the start page, open the **Scheduled** view in the **Workspace** work area.
- ▶ Open the task.
 - ↳ The **Edit scheduled task** dialog opens.

Edit scheduled task

Name

Press review

Prompt

e.g. Summarize today's AI news in five bullet points

Model

claude-sonnet-4-5

Schedule

Daily

Time

06:00

Your local time. Stored in UTC – across a daylight-saving change the run time shifts by one hour.

Also send by email (optional)

you@example.com

Leave empty to deliver into the conversation only.

Cancel

Save

Figure 167: Edit scheduled task dialog

- ▶ Change the prompt, the schedule, or the delivery.

- Save the task.
- Myra AI Workspace uses the change from the next run.

4.2.4.2.1 Entries in the Edit scheduled task dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name	Changes the name in the list of the scheduled tasks.
Prompt	Changes the prompt from the next run on. Results that were already delivered stay unchanged.
Model	Changes the model that answers the prompt.
Schedule	Changes the kind of the schedule: Daily , Weekly or Interval . Myra AI Workspace calculates the next run again.
Run every, Unit, Weekday and Time	Change the time of the next run.
Also send by email (optional)	Changes or removes the additional delivery by email. If you clear the field, the result arrives in the conversation only.

The following options are available in the dialog:

Option	Description
Save button	Applies the changes. Myra AI Workspace uses them from the next run on.
Cancel button	Closes the dialog without applying the changes.




The gateway of an existing task cannot be changed. The **Gateway** drop-down list appears only when you create the task. If the saving fails, the dialog shows **The task could not be saved – please check the fields.** or **Saving failed. Please try again.**



The **Pause** row action holds the task without changing it. The **Resume** row action takes it up again. A paused task does not count against the limit of your plan. If the switch fails, the view shows **Could not update the task. Please try again.**

4.2.4.3 Deleting a scheduled task

Proceed as follows to delete a scheduled task:



Scheduled tasks						New scheduled task
NAME	PROMPT	SCHEDULE	MODEL	GATEWAY	STATUS	
Daily briefing	Summarise today's news	Daily	claude-sonnet-4-5	produktion	Last run OK	
Weekly review	Summarise the week in five points	Weekly	gpt-4o-mini	redaktion	Paused	

Figure 168: Scheduled view

- ▶ On the start page, open the **Scheduled** view in the **Workspace** work area.
- ▶ Delete the task.
 - ↳ The **Delete the scheduled task ""? It will stop running and this cannot be undone.** query opens.




Figure 169: Delete task query

- ▶ Confirm the query.
- Myra AI Workspace does not run the task any more. A run that is still due does not occur.

4.2.4.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Scheduled task	Myra AI Workspace removes the task from the list. The deletion cannot be undone.
Schedule	A run that is still due is removed. Myra AI Workspace no longer runs the prompt.
Results	The results that were already delivered stay in their conversations.
Limit of the plan	The deletion frees a slot. You can then create another task.



If the deletion fails, the view shows **Could not delete the task. Please try again.**

4.2.5 Managing the prompts of the library

A prompt of the library is available to the whole organization in the chat.

4.2.5.1 Creating a prompt of the library

The following requirements are met:


- ☑ Your user account has the **Manage governance rules, prompts, and approvals** permission.

Proceed as follows to create a prompt:



Figure 170: Library view

- ▶ On the start page, open the **Library** view in the **Workspace** work area.
- ▶ Click on the **New prompt** button.
 - ↳ The **New prompt** dialog opens.



The image shows a 'New prompt' dialog box with the following fields and buttons:

- Name (required)**: Text input field with placeholder 'e.g. Complaint reply'.
- Category**: Text input field with placeholder 'e.g. Support'.
- Description**: Text input field with placeholder 'Short description (optional)'.
- Template (required)**: Large text area with placeholder 'The prompt text...'.
- Buttons**: 'Cancel' and 'Create' buttons at the bottom right.

Figure 171: New prompt dialog

- ▶ Enter the name, the category, and the prompt text.
- ▶ Save the prompt.
- Myra AI Workspace creates the prompt and submits it for approval.

4.2.5.1.1 Entries in the New prompt dialog

You set the following entries when you create the template:

Entry	Description
Name (required)	Name of the template in the list of the library, for example Complaint reply . The name holds a maximum of 255 characters.
Category	Category of the template, for example Support . The category holds a maximum of 64 characters and supplies the Filter by category drop-down list. The entry is optional.
Description	Short description of the template. The description appears in the list of the library. The entry is optional.

Entry	Description
Template (required)	Prompt text the users move into the chat with the Use in chat row action.

The following options are available in the dialog:

Option	Description
Create button	Creates the template. While the template is saved, the button shows the label Saving....
Cancel button	Closes the dialog without creating the template.



If a mandatory entry is missing, the dialog shows **Name is required.** or **Template is required.**

4.2.5.2 Editing a prompt of the library

Proceed as follows to edit a prompt:



Figure 172: Library view

- ▶ On the start page, open the **Library** view in the **Workspace** work area.
- ▶ Open the prompt.
 - ↳ The **Edit prompt** dialog opens.




Figure 173: Edit prompt dialog

- Change the entries and save the prompt.
- Myra AI Workspace creates a new version and submits it for approval.



A change to the prompt library becomes effective only after the approval of a second person. The person who makes the request never approves their own change. The changes that are open are in the **Settings** area, in the **Approvals** group, in the **Configurations** tab.

4.2.5.2.1 Entries in the Edit prompt dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name (required)	Changes the name in the list of the library and in the Versions dialog.
Category	

Entry	Effect of the change
	Assigns the template to a different category. If you clear the field, the Category column shows — and the template appears under All categories only.
Description	Changes the description in the list. If you clear the field, the Description column shows —.
Template (required)	Changes the prompt text. Each save creates a new version. The Version column counts up. Texts that were already used in the chat stay unchanged.

The following options are available in the dialog:

Option	Description
Save button	Applies the changes as a new version. While the template is saved, the button shows the label Saving...
Cancel button	Closes the dialog without applying the changes.



The **Restore** row action in the **Versions** dialog puts an earlier version back in place. Myra AI Workspace creates a new version with the restored content. The versions in between stay available. A restore also needs the approval of a second person.

4.2.5.3 Deleting a prompt of the library

Proceed as follows to delete a prompt:



Figure 174: Library view

- ▶ On the start page, open the **Library** view in the **Workspace** work area.
- ▶ Delete the prompt.
 - ↳ The **Delete the prompt ""?** query opens.




Figure 175: Delete prompt query

- ▶ Confirm the query.
- Myra AI Workspace shows **Prompt deleted**.

4.2.5.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Template	Myra AI Workspace removes the template from the library. The users no longer find it in the chat.
Versions	Myra AI Workspace removes all versions of the template. A restore through the Versions dialog is no longer possible.
Conversations	The prompts that were already sent stay in the conversations.



A deletion also needs the approval of a second person. Until the approval, the view shows **Change submitted for approval.** and the template stays in the list.

4.3 Connections

Connections integrate external systems into Myra AI Workspace. A connector for the Model Context Protocol (MCP) provides the AI in the chat with the tools of an external server, and a chat integration opens the chat in a messenger.

The following sections describe these configurations:

- [Managing MCP connectors](#)
- [Managing chat integrations](#)

4.3.1 Managing MCP connectors

A connector for the Model Context Protocol (MCP) makes the tools of an external server available to the AI in the chat.

4.3.1.1 Creating an MCP connector

The following requirements are met:

- ☑ Your user account has the **Manage tenant settings** permission.
- ☑ The address of the MCP server and the related credentials are available.

Proceed as follows to create an MCP connector:




Figure 176: MCP Connectors view

- ▶ In the **Settings** area, open the **Connections** group and its **MCP Connectors** tab.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Click on the **New connector** button.
 - ↳ The **New MCP Connector** dialog opens.

New MCP Connector

Name *

Credential scope

Tenant-shared Per-user

A single auth value is stored on this connector and used by every user in the tenant. Right for internal MCPs that don't have per-user auth.

Server URL *

JSON-RPC 2.0 endpoint of the MCP server.

Authentication

None Bearer token Custom header

Tool permissions

Choose whether tools from this connector are allowed by default, then block or allow individual tools as exceptions. Blocking is enforced on the server — a blocked tool can never run.

Allow by default Block by default

CancelCreate

Figure 177: New MCP Connector dialog

- ▶ Enter the name and select under **Credential scope** whether the connector is **Tenant-shared** or **Per-user**.
- ▶ Enter the address of the server and give the credentials.
- ▶ Save the connector.
- Myra AI Workspace creates the connector. The AI calls its tools in the chat.



The test of a connector needs a saved connector. Before that, the dialog shows **Save the connector first, then test.**

4.3.1.1.1 Entries in the New MCP Connector dialog

You set the following entries when you create the connector:

Entry	Description
Name	Name of the connector in the list, for example My Tools Server . The entry is mandatory.

Entry	Description
Credential scope	Tenant-shared stores a single auth value on the connector that every user of the tenant uses. That is the right scope for an internal MCP server without a sign-in per person. Per-user asks every person for their own credential. That is the right scope for a service such as Gmail or GitHub whose token represents a specific person. The users then connect themselves in the My MCP Connectors view.
Microsoft Entra tenant ID	Directory identifier of your organization as a GUID from the Entra admin center. The entry appears for a tenant-scoped Microsoft 365 connector only, and when you create it only. Myra AI Workspace derives the server URL and the sign-in addresses from it. The Server URL is then read-only. The entry is mandatory.
Server URL	JSON-RPC 2.0 endpoint of the MCP server, for example https://my-mcp-server.example.com/mcp . The entry is mandatory.
Authentication	Method with which the connector signs in to the server. For Tenant-shared , None , Bearer token , and Custom header are available. For Per-user , None and Bearer token are available. For Bearer token , you enter the token, for Custom header , the header in the Header-Name: header-value format, for example X-API-Key: your-secret-value .
Tool permissions	Basic rule for the tools of this connector: Allow by default or Block by default . Myra AI Workspace enforces the rule on the server. A blocked tool never runs.

If you select **Per-user** and **Bearer token**, the **OAuth (advanced)** area gives you these entries in addition:

Entry	Description
Authorize URL	Address at which the users sign in to the provider.
Token URL	Address at which the provider issues the tokens.
	Identifier of the OAuth app at the provider. For a Microsoft 365 connector, the entry is mandatory.

Entry	Description
OAuth Client ID	
OAuth Client Secret	Secret of the OAuth app. Myra AI Workspace stores it encrypted and never shows it again. PKCE is always used. For a public client, the field stays blank.
Scopes	Space-separated permissions that the sign-in requests. For Microsoft Entra, they include offline_access and the resource scope so that access is refreshed without a new sign-in.

The following options are available in the dialog:

Option	Description
Test connection button	Calls the server and shows Connected ✓ or Connected ✓ – tool(s) found . The button appears for Tenant-shared only, because only there the credential is in the dialog. It is disabled while the Server URL is empty.
Create button	Creates the connector. While it stores, the button carries the Saving... label.
Cancel button	Closes the dialog without creating the connector.



If a mandatory entry is missing, the dialog shows **Name is required, Server URL is required, Microsoft Entra tenant ID is required.** or **OAuth client ID is required for a Microsoft 365 connector.**



On a test, the dialog shows **Save the connector first, then test.** while the connector is not stored yet, **Connect required – add your credential first.** while no credential is available, and **Error:** on an error.

4.3.1.2 Editing an MCP connector

Proceed as follows to edit an MCP connector:




Figure 178: MCP Connectors view

- ▶ In the **Settings** area, open the **Connections** group and its **MCP Connectors** tab.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the connector.
 - ↳ The **Edit MCP Connector** dialog opens.




Figure 179: Edit MCP Connector dialog

- ▶ Change the data.
- ▶ Save the connector.
- Myra AI Workspace applies the change.



The field for the stored secret stays empty. Leave it empty to keep the stored secret.

4.3.1.2.1 Entries in the Edit MCP Connector dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name	Changes the name in the list and in the selection of the tools in the chat.
Cred ential	Switches between a shared and a personal access. After a change to Per-user , every person connects themselves in the My MCP Connectors view. Until then, the AI does not call the tools for that person. After a change to

Entry	Effect of the change
scope	Tenant-shared , the value stored on the connector applies to every user of the tenant.
Server URL	Points the connector at a different server. The AI then calls the tools of this server. If the credential is not valid there, the calls fail. For a Microsoft 365 connector, the entry is read-only.
Authentication	Changes the sign-in method. On a change, Myra AI Workspace discards the value entered so far. Store the credential again afterwards.
OAuth Client Secret	Replaces the stored secret. If you leave the field blank, Myra AI Workspace keeps the stored secret.
Tool permissions	Changes which tools the AI calls. The change applies immediately to each new call and to every user of the tenant. A blocked tool never runs.

The following entry is no longer changeable after the connector is created:

Entry	Description
Connector ID	Identifier of the connector. You address the connector by this identifier from the API. The Copy button puts the identifier on the clipboard.



The per-tool exceptions are available in the **Edit MCP Connector** dialog only, because they need a stored connector. While the tools are not loaded, the **Tool permissions** area carries the note that you load this connector's tools with the **Test connection** button and then set per-tool exceptions. For each tool, you then set **Allow** or **Block**.

4.3.1.3 Deleting an MCP connector

Proceed as follows to delete an MCP connector:




Figure 180: MCP Connectors view

- ▶ In the **Settings** area, open the **Connections** group and its **MCP Connectors** tab.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Delete the connector.
 - ↳ The **Delete this MCP connector? The model will no longer have access to its tools.** query opens.




Figure 181: Delete MCP connector query

- Confirm the query.
- Myra AI Workspace deletes the connector.

4.3.1.3.1 Effects of the deletion

The deletion has the following effects:

Objec t	Effect
Conn ector	Myra AI Workspace removes the connector from the list. It cannot be restored.
Crede ntials	Myra AI Workspace deletes the credential stored on the connector and, for a connector with the Per-user scope, the tokens of every user as well. The tokens at the provider stay. Revoke them there if they are no longer to be valid.
Tool permi ssions	The basic rule and the per-tool exceptions go with the connector.
Chat	The AI no longer calls the tools of this server for any person of the tenant. A task that needs such a tool then fails.

Object	Effect
Agents	An agent that uses this tool continues to work without the tool. Check the agents concerned before the deletion.
API	Calls that address the connector by its Connector ID no longer find it.

4.3.2 Managing chat integrations

A chat integration connects Mattermost, Slack, or Microsoft Teams to an agent. Your team then talks to the agent from the chat platform.

4.3.2.1 Creating a chat integration

The following requirements are met:

- ☑ Your user account has the **Manage tenant connector-sync and chat-bridge** permission.
- ☑ You have the credentials of the chat platform.
- ☑ The agent that answers exists.

Proceed as follows to create a chat integration:





Figure 182: Chat Integrations view

- ▶ In the **Settings** area, open the **Connections** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Chat Integrations** tab.
- ▶ Click on the **Add connection** button.
 - ↳ The **Add chat connection** dialog opens.



Add chat connection

Platform
 Mattermost

Gateway
 Select a gateway...

Agent
 Select an agent...

Workspace (team id)
 the Mattermost team id

Mattermost URL
 https://mattermost.example.com

Bot access token

Relay secret

A high-entropy shared secret the relay presents on every event. Store it only in the relay config.

☐ Allow web search
 Let this bot search the web to answer. Other external tools (fetching URLs, running code) stay disabled on the chat channel.

☐ Allow all tools
 Give this bot its agent's full toolset – running code, fetching URLs, MCP connectors, image generation. These run under the bot owner's credentials, and anyone in the chat workspace can trigger them. Leave off unless you trust every member of this channel.

Cancel Save

Figure 183: Add chat connection dialog

- ▶ Under **Platform**, select the chat platform.
- The credential fields depend on this selection.
 - ▶ Select the **Gateway** and the **Agent** that answers.
 - ▶ Enter the identifier of the workspace and the credentials of the platform.
 - ▶ If necessary, set the **Allow web search** and **Allow all tools** switches.
 - ▶ Click on the **Save** button.
- Myra AI Workspace creates the connection with the **Enabled** status. Your team talks to the agent from the platform selected.



Allow all tools gives the bot its agent's full toolset – running code, fetching URLs, MCP connectors, image generation. These run under the bot owner's credentials, and anyone in the chat workspace triggers them. Set the switch only if you trust every member of the channel.

4.3.2.1.1 Entries in the Add chat connection dialog

You set the following entries for every platform:

Entry	Description
Platform	Chat platform of the connection: Mattermost , Slack , or Microsoft Teams . The selection decides which credential fields the dialog gives you next.
Gateway	Gateway through which the connection bills and enforces its requests. The drop-down list gives the gateways of the tenant. The entry is mandatory.
Agent	Agent that answers the messages from the chat. The drop-down list gives the agents of the selected gateway and is disabled while no gateway is selected. The entry is mandatory.
Workspace	Identifier of the workspace on the platform. For Mattermost, that is the team ID, for Slack, the team ID in the T0123ABCD form, for Microsoft Teams, the GUID of the Entra tenant. The entry is mandatory.
Allow web search	Lets the bot search the web for its answers. Other external tools – fetching URLs, running code – stay switched off on the chat channel. The check box is available when you create the connection only.
Allow all tools	Gives the bot the full toolset of its agent: running code, fetching URLs, MCP connectors, and image generation. The check box is available when you create the connection only.

For the **Mattermost** platform, these entries come in addition:

Entry	Description
Mattermost URL	Address of your Mattermost server, for example https://mattermost.example.com . The entry is mandatory.
Bot access token	Access token of the bot account from Mattermost. The entry is mandatory.
Relay secret	High-entropy shared secret that the relay presents on every event. Store it in the configuration of the relay only. The entry is mandatory.

For the **Slack** platform, these entries come in addition:

Entry	Description
Bot token (xoxb-...)	Bot token of the Slack app. The entry is mandatory.
Signing secret	Request-signing secret from Basic Information → App Credentials in Slack. Store it here only. The entry is mandatory.

For the **Microsoft Teams** platform, these entries come in addition:

Entry	Description
Bot app ID	Application identifier from your Azure Bot registration. The entry is mandatory.
Bot app type	Multi-tenant for an app that serves several Entra tenants, or Single-tenant for an app in your own tenant.
Home tenant ID (Entra GUID)	GUID of the Entra tenant that holds the bot registration. The entry is mandatory.
Service URL	Messaging endpoint from the bot registration, for example https://smba.trafficmanager.net/emea/ . The entry is mandatory.
Client secret	Secret of the bot registration. The entry is mandatory.

The following options are available in the dialog:

Option	Description
Save button	Creates the connection. While it stores, the button carries the Loading... label.
Cancel button	Closes the dialog without creating the connection.



Myra AI Workspace never gives the credentials out again. They appear neither in the list nor in the dialog. To change them, use the **Rotate token** row action.

4.3.2.2 Editing a chat integration

The platform, the workspace, and the agent that answers stay unchanged on an existing chat integration. You can change the status and the two tool switches in the row of the connection.

Proceed as follows to edit a chat integration:




Figure 184: Chat Integrations view

- ▶ In the **Settings** area, open the **Connections** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Chat Integrations** tab.
- ▶ In the row of the connection, select **Disable**, **Enable web search**, or **Enable all tools**.
- Myra AI Workspace applies the change immediately and shows it in the **Status**, **Web search**, and **All tools** columns.

4.3.2.2.1 Effects of the changes

The changes in the row have the following effects:

Change	Effect
Disable	The bot stops answering. The connection and its credentials stay. The Status column then shows Disabled .
Enable	The bot answers again. The Status column then shows Enabled .
Enable web search	The bot searches the web for its answers. The Web search column then shows On .
Disable web search	The bot answers from the knowledge of the model and the sources of its agent only. The Web search column then shows Off .
Enable all tools	The bot uses the full toolset of its agent. The All tools column then shows On .
Disable all tools	The bot uses no external tool other than the web search. The All tools column then shows Off .

The following entries are no longer changeable after the connection is created:

Entry	Description
Platform	The chat platform of the connection. For a different platform, create another connection.
Workspace	The identifier of the workspace on the platform.
Agent	The agent that answers.
Gateway	The gateway that bills.



The changes apply immediately to the next message. Myra AI Workspace does not stop an answer that is already running.

4.3.2.3 Rotating the credentials of a chat integration

You can change the credentials of the chat platform without creating the connection again.

Proceed as follows to rotate the credentials:





Figure 185: Chat Integrations view

- ▶ In the **Settings** area, open the **Connections** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Chat Integrations** tab.
- ▶ In the row of the connection, select **Rotate token**.
 - ↳ The **Rotate bot token** dialog opens.



The dialog box is titled "Rotate bot token". It contains four input fields: "Platform" with the value "Mattermost", "Workspace" with the value "muster-team", "Agent" with the value "briefing", and "New bot access token" with the placeholder text "leave blank to keep the current token". At the bottom right, there are two buttons: "Cancel" and "Rotate token".

Figure 186: Rotate bot token dialog

- ▶ Under **New bot access token**, enter the new token.
- An empty field keeps the current token.
- ▶ Click on the **Rotate token** button.
- Myra AI Workspace replaces the credentials of the connection. The bot continues to answer from the same workspace.

4.3.2.3.1 Entries in the Rotate bot token dialog

The **Rotate bot token** dialog changes the credentials of a connection that exists. The fields follow the platform of the connection.

The dialog gives the following entries as read-only:

Entry	Description
Platform	Chat platform of the connection.
Workspace	Identifier of the workspace on the platform.
Agent	Agent that answers.

You change the following entries in the dialog:

Entry	Effect of the change
New bot access token	Replaces the access token of the bot at Mattermost. The field appears for Mattermost.
New bot token	Replaces the bot token of the Slack app. The field appears for Slack.
New signing secret	Replaces the request-signing secret of the Slack app. The field appears for Slack.
New client secret	Replaces the secret of the Azure Bot registration. The field appears for Microsoft Teams.

The following options are available in the dialog:

Option	Description
Rotate token button	Takes over the credentials that you entered. While it stores, the button carries the Loading... label.
Cancel button	Closes the dialog without changing the credentials.



Every field carries the **leave blank to keep the current token** placeholder. A blank field leaves the related credential unchanged. If every field stays blank, the dialog changes nothing.



Change the credentials at the platform first and enter them here afterwards. Between the two steps, the bot does not answer.

4.3.2.4 Deleting a chat integration

Proceed as follows to delete a chat integration:




Figure 187: Chat Integrations view

- ▶ In the **Settings** area, open the **Connections** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Chat Integrations** tab.
- ▶ In the row of the connection, select **Delete**.
 - ↳ The **Delete this chat connection? The bot will stop answering and its access token is revoked.** query opens.




Figure 188: Delete chat connection query

- ▶ Confirm the query.
- Myra AI Workspace deletes the connection and revokes its access token. The platform does not reach the agent any more.

4.3.2.4.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Connection	Myra AI Workspace removes the connection from the Chat Integrations tab. It cannot be restored.
Credentials	Myra AI Workspace revokes the access token of the bot. The bot registration on the platform stays. Remove it there if the bot is no longer to appear.
Chat platform	The bot no longer answers in the workspace. The messages so far stay in the channel.
Agent	The agent that answers stays and continues to work unchanged in Myra AI Workspace.
	The usage already billed stays in the report.

Object	Effect
Gatewa	
y	

4.4 Gateway

A gateway is the named installation of a tenant. It determines the available providers and models, the rules for content, the budget, and the permitted API tokens. The following configurations set up the gateway and determine how it checks and forwards a request.

The following sections describe these configurations:

- [Managing gateways](#)
- [Configuring guardrails](#)
- [Managing routing rules](#)
- [Managing governance templates](#)

4.4.1 Managing gateways


A gateway is the named installation of a tenant. It sets the providers and models that are available, the rules for the content, the budget, and the API tokens permitted.

4.4.1.1 Creating a gateway

The following requirements are met:

- ☒ Your user account has the **Manage gateways and routing** permission.

Proceed as follows to create a gateway:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	-	0s	-	-	7/19/2026

Figure 189: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- Click on the **New Gateway** button.
 - ↳ The **New Gateway** dialog opens.


The image shows a 'New Gateway' dialog box with the following fields and options: 'Name' (text input with 'prod', note: 'Lowercase letters, numbers, and hyphens only.'), 'Gateway Budget (USD)' (text input with 'unlimited', note: 'Applies to this gateway only. The tenant's overall spend is capped separately by its own budget.'), 'Cache TTL (s)' (text input with '300', note: '0 = disabled'), 'Retry Count' (text input with '2', note: 'Default 2 - a failed upstream call is retried up to twice. Set 0 to disable retries.'), 'Timeout (ms)' (text input with '120000', note: 'How long to wait for the upstream before failing. The default suits most models; raise it only for very long generations.'), a checked checkbox for 'Require auth token', a 'Guardrails (0)' section with buttons for 'Regex / Pattern', 'Keyword', 'Jailbreak', 'Presidio (NLP)', 'Prompt Guard', 'Prompt Injection', 'PII Protector', and 'Custom PII Blacklist', and a note: 'No guardrails configured. Add one above to start scanning requests or responses.' At the bottom are 'Cancel' and 'Create Gateway' buttons.

Figure 190: New Gateway dialog

- ▶ In **Name**, enter the identifier of the gateway.
- The identifier contains lowercase letters, numbers, and hyphens.
- ▶ Set the budget, the cache, the retries, and the timeout.
- ▶ Click on the **Create Gateway** button.
- Myra AI Workspace creates the gateway.

4.4.1.1.1 Entries in the New Gateway dialog

You set the following entries when you create a gateway:

Entry	Description
Name *	Identifier of the gateway, for example prod . Only lowercase letters, numbers, and hyphens are permitted. The identifier is part of the endpoint URL. The entry is required.
Gateway	

Entry	Description
Budget (USD)	Spend cap of this gateway. If you leave the field empty, the unlimited entry applies. The budget applies to this gateway only. The tenant's overall spend is capped separately by its own budget.
Cache TTL (s)	Lifetime of the cache in seconds. The value 0 disables the cache.
Retry Count	Number of retries of a failed call to the provider. The default is 2 . The value 0 disables the retries.
Timeout (ms)	Time to wait for the provider before the request fails. The default suits most models. Raise it only for very long generations.
Require authentication	If the check box is set, every request to the gateway must carry an authentication token.
Guardrails ()	Detectors that examine the requests and the responses of this gateway. With the Regex / Pattern , Keyword , Jailbreak , Presidio (NLP) , Prompt Guard , Prompt Injection , PII Protector , and Custom PII Blacklist buttons you add detectors while you create the gateway. As long as no detector exists, the No guardrails configured. Add one above to start scanning requests or responses. entry appears. See the Adding a guardrail detector section.

The following options are available in the dialog:

Option	Description
Create Gateway button	Creates the gateway. While the gateway is created, the button shows the Creating... label.


Option	Description
Cancel button	Closes the dialog without creating the gateway.



You set all further settings – rate limit, circuit breaker, webhook, SIEM, web search, image generation, prompt caching, compaction, tools, and tracing – in the **Edit Gateway** dialog after you create the gateway. Afterwards you no longer change the guardrails in this dialog, but in the **Guardrails** section of the detail view of the gateway.

4.4.1.2 Enabling a model on a gateway

Proceed as follows to enable a model on a gateway:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	–	0s	–	–	7/19/2026

Figure 191: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- Open the gateway.
 - ↳ The detail view of the gateway opens.





Figure 192: Detail view of a gateway

- Click on the **Add Model** button.
 - ↳ The **Add Model** dialog opens.



The image shows a modal dialog box titled "Add Model". At the top, there are two tabs: "Use My Own Key" (which is selected and highlighted with a blue border) and "Use Platform Key". Below the tabs, a line of text reads "Enable a provider's models on this gateway." Underneath, there are two input fields. The first is labeled "Provider" and contains a dropdown menu with "anthropic" selected and a downward arrow. The second is labeled "Alias" and contains the text "default". Below these is a section for the "API Key *". It features a text input field with the placeholder text "Paste API key..." and a small eye icon on the right. Below the input field, a line of text states "Stored encrypted. Replaces existing key for this provider/alias." At the bottom right of the dialog, there are two buttons: "Cancel" and "Store Key".

Figure 193: Add Model dialog

- ▶ Select the method **Use My Own Key** or **Use Platform Key**.
- ▶ Select the **Provider** and, for your own key, enter the **API Key** and, optionally, give an **Alias**.
- ▶ Click on the **Enable** button at the model you want.
- Myra AI Workspace shows **Key stored and encrypted successfully**. The model is available for selection in the chat.




Figure 194: Detail view of a gateway



A model made available through the platform key carries the **Billed to your plan / trial budget** mark.

4.4.1.2.1 Entries in the Add Model dialog

You set the following entries when you enable a model:

Entry	Description
How to add a model	Sets the key the model runs with: Use My Own Key or Use Platform Key .

Entry	Description
Provider	Provider whose models the gateway is to reach, for example anthropic or openai .
Aliases	Label of the key, for example default . The alias distinguishes several keys of the same provider.
API Key	Key of the provider. Myra AI Workspace stores it encrypted and thereby replaces an existing key for the same provider and alias. Most providers expect the key as a single character string. AWS Bedrock expects the credentials separated by colons in the form ACCESS_KEY_ID:SECRET_ACCESS_KEY or ACCESS_KEY_ID:SECRET_ACCESS_KEY:SESSION_TOKEN . Vertex AI expects the JSON key of a Google Cloud service account. Some providers need no key. Then the dialog shows This provider does not require an API key .

The following options are available in the dialog:

Option	Description
Show button	Shows the API key you entered as plain text. When the key is shown, the button carries the Hide label.
Store Key button	Stores the key encrypted. While the key is stored, the button shows the Storing... label.
Enable button	Enables the model of the row on this gateway. For an enabled model, the button carries the Disable label.
Cancel button	Closes the dialog.



A model provided through the platform key is usable only once the workspace has a budget. If no models are available, the dialog shows **No Myra-provided models are available**.




Several keys of the same provider exist next to each other on one gateway under different aliases. If you store a key again for the same provider and alias, it replaces the previous key. Thus you exchange a key without an interruption. A request selects an alias other than **default** with the `x-aig-byok-alias` header. If the specified alias does not exist, the gateway rejects the request with a configuration error. It does not fall back to the **default** alias.

4.4.1.3 Deleting a provider key

A deleted provider key cannot be restored. Requests that resolve to this provider and alias fail until a new key is stored.

Proceed as follows to delete a provider key:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	–	0s	–	–	7/19/2026

Figure 195: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- Open the gateway.
 - ↳ The detail view of the gateway opens.




Figure 196: Detail view of a gateway

- In the **Provider Keys** area, click on the **Delete** row action of the key.
 - ↳ The **Delete key for /? Routing to this provider will fail.** query opens.




Figure 197: Delete key query

- Confirm the query.
- Myra AI Workspace removes the key from the **Provider Keys** area.




Figure 198: Detail view of a gateway

4.4.1.3.1 Effects of the deletion of a key

The deletion has the following effects:

Object	Effect
Key	Myra AI Workspace removes the key immediately. The deletion cannot be undone.
Requests	Requests that resolve to this provider and alias fail until a new key is stored. The gateway does not fall back to a different alias.
Other aliases	The other keys of the same provider remain. A request reaches them with the <code>x-aig-byok-alias</code> header.
Keys of the tenant	A key stored on the tenant does not step in. It serves the administration and the analysis of the usage data, not the answering of requests.




The deletion of the key with the **default** alias makes every request fail that specifies no alias.

The following options are available in the query:

Option	Description
Confirm button	Deletes the key.
Cancel button	Closes the query without deleting the key.

4.4.1.4 Editing a gateway

Proceed as follows to edit a gateway:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED	
produktion	required	€460.00	300s	60/60s	3	5/14/2026	 
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026	 
labor	open	–	0s	–	–	7/19/2026	 

Figure 199: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the gateway.
 - ↳ The detail view of the gateway opens.





Figure 200: Detail view of a gateway

- ▶ Click on the **Edit** button.
- ↳ The **Edit Gateway:** dialog opens.



The image shows a modal dialog titled "Edit Gateway: produktion". It contains several sections for configuring the gateway. The "BUDGET & LIMITS" section includes fields for Gateway Budget (USD) set to 250, Budget Period set to Month, Cache TTL (s) set to 300, Retry Count set to 2, Timeout (ms) set to 120000, Rate Limit (req) set to 600, and Rate Window (s) set to 60. There are two checked checkboxes: "Require auth token" and "Log request/response payloads". Below this are sections for "PROVIDER BASE URLS" and "CIRCUIT BREAKER", each with a text input field. At the bottom right are "Cancel" and "Save Changes" buttons.

BUDGET & LIMITS	
Gateway Budget (USD)	Budget Period
250	Month
Cache TTL (s)	Retry Count
300	2
Timeout (ms)	Rate Limit (req)
120000	600
Rate Window (s)	
60	
☒ Require auth token	
☒ Log request/response payloads	
PROVIDER BASE URLS	
CIRCUIT BREAKER	

Figure 201: Edit Gateway dialog

- Change the routing, the budget, the rules, or the URL overrides of the providers.
- Click on the **Save changes** button.
- Myra AI Workspace applies the change for the next request.




Figure 202: Detail view of a gateway

4.4.1.4.1 Entries in the Edit Gateway dialog

The dialog groups the entries in sections. You change the following entries in the **Budget & limits** section:

Entry	Effect of the change
Gateway Budget (USD)	Sets the spend cap of this gateway. When the cap is reached, the gateway rejects further requests.
Budget Period	Sets what the cap refers to: Monthly , Daily , or Lifetime .
Cache TTL (s)	Changes the lifetime of the cache. The value 0 disables the cache.
Retry Count	Changes the number of retries of a failed call to the provider.
Timeout (ms)	Changes the time to wait for the provider. A value that is too small aborts long generations.

Entry	Effect of the change
Rate Limit (req) and Rate Window (s)	Limit the requests per time window. The gateway rejects requests beyond that.
Require auth token	If the check box is not set, the gateway accepts requests without a token.
Log request/response payloads	If the check box is set, Myra AI Workspace stores the payloads of the requests and responses in the request log.

You change the following entries in the remaining sections:

Section	Effect of the change
Provider Base URLs	Overrides the upstream URL for a provider, for example for a self-hosted OpenAI-compatible endpoint. The Add URL override button adds a row, the Remove button deletes it.
Circuit Breaker	Stops the routing to a provider after repeated failures and probes again after a cooldown. The Failure threshold (default 5), Window (s) (default 60), and Cooldown (ms) (default 30000) entries set when the circuit breaker opens and when it probes again.
Webhook	Sends HTTP POST notifications on gateway events to the URL you enter. If you leave the URL empty, the webhook is disabled. A signing secret adds the X-AIG-Signature header.
SIEM Integration	Streams security events to an external SIEM and thereby overrides any SIEM configuration at tenant level. The Events to forward selection sets which events the gateway sends. The all (every request) entry sends every request.
Web Search	Enables the server-side web search for all requests on this gateway. The Search provider is Brave – US or Linkup – EU . Without a selection, the tenant default applies. If the gateway enforces EU data residency, Linkup – EU is required. Without an API key of the search provider, the web search cannot be enabled.
Image Generation	Lets the assistant create images on request through the generate_image tool. The image model is a self-hosted EU model. The default is flux.2-dev .

Section	Effect of the change
Anthropic Prompt Caching	Sets a cache breakpoint on the system prompt so that Anthropic reuses cached tokens across turns. The Cache TTL is 5 minutes (1.25× write cost) or 1 hour (2× write cost) . The setting applies to Anthropic models only.
Context Compaction (Anthropic)	Summarises the conversation history as soon as the input tokens exceed the Threshold (tokens) . The default is 200 000 tokens. The Keep last N turns verbatim entry sets how many turns stay uncompactd.
Compact error threshold (tokens)	Returns a context too long error to the client as soon as the estimated number of input tokens exceeds this value. If you leave the field empty, the response is disabled.
Tools & Access	Groups the IP allowlist , the code interpreter , the agentic web fetch , the malware scan of the uploads, and the maximum number of parallel tool calls. See the following table.
Request Tracing	Records a full step-by-step trace for every request. The Retention (hours) entry sets how long the traces are kept (1 to 720).

You change the following entries in the **Tools & Access** section:

Entry	Effect of the change
IP allowlist	Restricts the inference to the source IP ranges you enter in CIDR notation, for example 10.0.0.0/8 or 2001:db8::/32 . An empty list allows all source IP addresses. The list applies to the inference only, not to the admin API.
Enable code interpreter	Lets the assistant run code. With Persistent kernel (stateful) , the variables are kept across turns.

Entry	Effect of the change
Enable agentic web fetch	Lets the assistant fetch web pages on its own.
Scan uploads for malware (ClamAV)	Scans files that are uploaded to this gateway before they are processed. The gateway rejects infected or unscannable files. The scan requires a provisioned ClamAV service.
Max parallel tool calls	Limits the tool calls per round. If you leave the field empty, the default 4 applies.

The following options are available in the dialog:

Option	Description
Save Changes button	Applies the changes. While the changes are saved, the button shows the Saving... label.
Cancel button	Closes the dialog without applying the changes.



If an entry of the IP allowlist is not a valid CIDR notation, the dialog shows **Invalid CIDR:** . If the web search is enabled without an API key, the dialog shows **An API key for the selected search provider is required to enable web search.**


4.4.1.5 Creating an auth token

An auth token identifies an application to the gateway. Create one token per application so that you can revoke it individually.

The following requirements are met:

- ☒ Your user account has the **Administrator** or **Tenant Administrator** role.

Proceed as follows to create an auth token:



NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	–	0s	–	–	7/19/2026

Figure 203: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the gateway.
 - ↳ The detail view of the gateway opens.





Figure 204: Detail view of a gateway

- In the **Auth Tokens** section, click on the **Generate** button.
 - ↳ The **Create Auth Token** dialog opens.



The image shows a 'Create Auth Token' dialog box. It has a title bar and a main content area. The content area contains several input fields and buttons. The first field is 'Label (optional)' with the placeholder text 'e.g. ci-pipeline'. The second field is 'Expires At (optional)' with the placeholder text 'mm/dd/yyyy --:--'. Below this field is a note: 'Leave blank for a non-expiring token – recommended for service and scheduler tokens.' The third field is 'Spend cap (USD, optional)' with the placeholder text 'unlimited'. Below this field is a note: 'Block this token once cumulative cost exceeds this amount.' The fourth field is 'Rate limit (req. optional)' with the placeholder text 'unlimited'. The fifth field is 'Window (s)' with the placeholder text '60'. At the bottom of the dialog are two buttons: 'Cancel' and 'Generate Token'.

Create Auth Token

Label (optional)
e.g. ci-pipeline

Expires At (optional)
mm/dd/yyyy --:--

Leave blank for a non-expiring token – recommended for service and scheduler tokens.

Spend cap (USD, optional)
unlimited

Block this token once cumulative cost exceeds this amount.

Rate limit (req. optional) Window (s)
unlimited 60

Cancel Generate Token

Figure 205: Create Auth Token dialog

- ▶ Enter the label.
- ▶ If necessary, set the expiry, the spend cap, and the rate limit.
- ▶ Click on the **Generate Token** button.
- Myra AI Workspace creates the token and shows it once in the **Token Created** dialog.




Figure 206: Token Created dialog



A token always applies to exactly one gateway. The section carries the hint **An authentication token is a secret key an application sends with each request to prove it may use this gateway – like a password for software. Create one per app or integration so you can revoke it individually.**

4.4.1.5.1 Entries in the Create Auth Token dialog

You set the following entries at creation:

Entry	Description
Label (optional)	Label of the token, for example ci-pipeline . The label appears in the Label column of the section. Without a label, the section lists the token under the first 16 characters of its hash value.
Expires At (optional)	Time at which the token loses its validity. The field carries the hint Leave blank for a non-expiring token – recommended for service and scheduler tokens. If you set an expiry, the owner is warned (in-app and by email) 7 days before it lapses.

Entry	Description
Spend cap (USD, optional)	Amount above which Myra AI Workspace blocks the token. The field carries the hint Block this token once cumulative cost exceeds this amount . An empty field means: unlimited.
Rate limit (req, optional) and Window (s)	Limit the requests with this token per time window. An empty Rate limit (req, optional) field means: unlimited. The Window (s) field carries the value 60 .

The following options are available in the dialog:

Option	Description
Generate Token button	Creates the token and shows it once. During the creation, the button shows Generating...
Copy button	Copies the token shown to the clipboard. After the copy, the button shows Copied! . A click on the token itself does the same.
Done button	Closes the Token Created dialog.
Cancel button	Closes the dialog without creating a token.



Myra AI Workspace shows the token only immediately after the creation. The dialog carries the hint **Copy this token now. It will not be shown again**. Copy the token before you close the dialog.


4.4.1.6 Editing an auth token

You change the spend cap, the rate limit, and the expiry of an existing token without revoking the token. The token itself stays valid. An application that is already set up continues to work.

The following requirements are met:

- ☒ Your user account has the **Administrator** or **Tenant Administrator** role.
- ☒ At least one auth token exists for the gateway.

Proceed as follows to edit an auth token:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED	
produktion	required	€460.00	300s	60/60s	3	5/14/2026	 
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026	 
labor	open	–	0s	–	–	7/19/2026	 

Figure 207: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the gateway.
 - ↳ The detail view of the gateway opens.





Figure 208: Detail view of a gateway

- In the **Auth Tokens** section, click on the **Edit** row action at the token.
 - ↳ The **Edit Token** – dialog opens with the current entries.



The image shows a modal dialog titled "Edit Token — ci-pipeline". It contains several input fields and instructions. The "Expires At (optional)" field is set to "12/31/2026 23:59". Below it, a note says "Leave blank for a non-expiring token — recommended for service and scheduler tokens." The "Spend cap (USD, optional)" field is set to "50.00", with a note below it saying "Leave blank to remove the spend cap." There are two side-by-side fields: "Rate limit (req, optional)" set to "120" and "Window (s)" set to "60". At the bottom right are "Cancel" and "Save Changes" buttons.

Edit Token — ci-pipeline

Expires At (optional)

12/31/2026 23:59

Leave blank for a non-expiring token — recommended for service and scheduler tokens.

Spend cap (USD, optional)

50.00

Leave blank to remove the spend cap.

Rate limit (req, optional) **Window (s)**

120 60

Cancel **Save Changes**

Figure 209: Edit Token dialog

- Change the data.
- Click on the **Save Changes** button.
- Myra AI Workspace applies the entries. The **Auth Tokens** section shows the changed values.




Figure 210: Detail view of a gateway



Without a label, the dialog carries the first 16 characters of the hash value instead of the label, for example **Edit Token – a1b2c3d4e5f60718...**

4.4.1.6.1 Entries in the Edit Token dialog

You change the following entries:

Entry	Effect of the change
Expires At (optional)	Moves the time at which the token loses its validity. If you empty the field, the token no longer expires. The Expires column then shows never .
Spend cap (USD, optional)	Changes the amount above which Myra AI Workspace blocks the token. The field carries the hint Leave blank to remove the spend cap . If you empty the field, requests with this token consume without a limit.

Entry	Effect of the change
Rate limit (req, optional) and Window (s)	Change the permitted requests per time window. If you empty the Rate limit (req, optional) field, the rate limit of the token no longer applies.

The following entries cannot be changed in this dialog:

Entry	Description
Label	Label of the token. It is fixed at creation.
Hash (first 16)	Hash value of the token. It is fixed at creation.
Gateway	A token always applies to exactly the gateway that you created it for.

The following options are available in the dialog:

Option	Description
Save Changes button	Applies the changed entries. During the save, the button shows Saving....
Cancel button	Closes the dialog without changing the entries.

4.4.1.6.2 Effects of the changes


The change has the following effects:

Object	Effect
Token	The token itself stays valid without a change. An application that uses it continues to work unchanged.
Requests	The new limits apply from the next request. A lowered spend cap and a lowered rate limit take effect immediately.
Consumption so far	The recorded costs of the token remain. If the consumption so far already exceeds the new spend cap, Myra AI Workspace blocks the token.

4.4.1.7 Revoking an auth token

A revoked token cannot be restored. Applications that use it receive no answer any more.

Proceed as follows to revoke an auth token:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	-	0s	-	-	7/19/2026

Figure 211: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the gateway.
 - ↳ The detail view of the gateway opens.




Figure 212: Detail view of a gateway

- In the **Auth Tokens** section, click on the **Revoke** row action at the token.
 - ↳ The **Delete this token? Requests using it will fail.** query opens.




Figure 213: Delete token query

- Confirm the query.
- Myra AI Workspace removes the token from the **Auth Tokens** section.




Figure 214: Detail view of a gateway

4.4.1.7.1 Effects of the revocation

The revocation has the following effects:

Object	Effect
Token	Myra AI Workspace removes the token from the section. The revocation cannot be undone.
Requests	The gateway answers requests with this token with Missing or invalid gateway token .
Spend cap and rate limit	The entries of the token are removed with the token. A new token gets its own entries.
Gateway	The gateway remains. The other tokens of the gateway stay valid.

The following options are available in the query:

Option	Description
Confirm button	Revokes the token.
Cancel button	Closes the query without revoking the token.

4.4.1.8 Deleting a gateway

Proceed as follows to delete a gateway:




Figure 215: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Delete the gateway.
 - ↳ The **Delete gateway ""? This cannot be undone.** query opens.




Figure 216: Delete gateway query

- ▶ Confirm the query.
- Myra AI Workspace deletes the gateway. The tokens of this gateway lose their validity.

4.4.1.8.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Gateway	Myra AI Workspace deletes the gateway permanently. The endpoint URL is no longer reachable afterwards.
Auth tokens	The tokens of this gateway lose their validity. Requests with these tokens fail.
Provider keys	The keys stored on the gateway are removed with the gateway.
Routing rules and guardrails	The rules and the selection of the detectors are removed with the gateway.



The deletion cannot be undone. Create a gateway with the same name again if you need the identifier once more.

4.4.2 Configuring guardrails


A guardrail detector examines the requests and the responses of a gateway for content and blocks, scrubs, or flags them. A gateway can have any number of detectors. Their sequence controls the order in which they examine the data.

4.4.2.1 Adding a guardrail detector

The following requirements are met:

- ☑ Your user account has the **Manage gateways and routing** permission.
- ☑ The tenant has a gateway.

Proceed as follows to add a guardrail detector:



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	-	0s	-	-	7/19/2026

Figure 217: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- In the row of the gateway, click the **Open** → row action.
 - ↳ The detail view of the gateway opens.




Figure 218: Detail view of a gateway

- In the **Guardrails** area, click the button of the necessary detector type, for example **PII Protector**.
 - ↳ Myra AI Workspace adds the area of the detector to the end of the list.

4.4.2 Configuring guardrails




Figure 219: Guardrails area

- Expand the area of the detector.
- Change the data of the detector.
- Click the **Save Guardrails** button.
- Myra AI Workspace shows the **Saved** message.




Figure 220: Guardrails section



You specify the detectors of a new gateway in the **New gateway** dialog. The builder there is the same as in the detail view. See the **Creating a gateway** section.

4.4.2.1.1 Detector types

The following detector types are available in the toolbar:

Detect or type	Time	Default action and target	Description
Regex / Pattern	1	block, request	Examines the text against supplied pattern sets and against your own patterns. The detector operates fully in the gateway. The default name is pii-check .

Detect or type	Tier	Default action and target	Description
Keyword		flag, request	Examines the text for a list of keywords. The default name is keyword-check .
Jailbreak	1	flag, request	Examines the text for phrases with which users try to bypass the instructions of the model. The default name is jailbreak-check .
Presidio (NLP)	2	block, request	Finds personal data with a language model and flags or removes the data. The detector needs the Presidio service. The default name is presidio-pii .
Prompt Guard	2	block, request	Rates the request with the Llama Guard classifier by safety category. The default name is prompt-guard .
Prompt Injection	2	block, request	Rates the request with a classifier and blocks or flags the request when the injection probability found is at or above the threshold. The default name is prompt-injection .
PII Protector	2	reversible, request and response	Replaces the personal data found with opaque tokens before the model and puts the data back into the response. The default name is pii-protect .
Custom PII Blacklist	1	reversible, request and response	Replaces the terms that you specify with opaque tokens before the model and puts the terms back into the response. The detector operates fully in the gateway. The default name is custom-pii .



The tier controls the sequence of the examination. Myra AI Workspace examines all detectors of tier 1 first and the detectors of tier 2 subsequently. In a tier, the sequence of the areas applies.

4.4.2.1.2 Data common to each detector

You specify the following data for each detector type:

Data	Description
Name	Name of the detector, for example detector-name . The name is shown in the area, in the execution plan, in the Detectors column of the guardrail events, and in the request log.
Action	Treatment of a match: block rejects the request, scrub removes the text that matched, and flag lets the request continue and records the match. The PII Protector and Custom PII Blacklist detectors have no Action field. They always operate reversibly.
Target	Object of the examination: request examines the request, response the response, and both the request and the response. For the PII Protector and Custom PII Blacklist detectors, the target is permanently set to request and response.

4.4.2.1.3 Data for the Regex / Pattern detector

You also specify the following data:

Data	Description
Pattern sets	Supplied sets that group several patterns: pci_pan (credit card, CVV, expiry, IBAN, routing; less than 2 % false positives, safe for the block action), hipaa_structured (SSN, MRN, NPI, date of birth, phone, email, IP; up to 26 % false positives, use the scrub action), gdpr_structured (email, phone, IP, IBAN, national ID, passport; up to 26 % false positives, use the scrub action), credentials (API keys, JWTs; 0 % false positives, safe for the block action), and pii_basic (email, phone, SSN; up to 26 % false positives, use the scrub action).
Individual patterns	Patterns that you can select one by one: email , phone , ssn , dob , ip_address , cc , cvv , card_expiry , iban , routing_number , mrn , npi , national_id , passport_number , api_key , and jwt .
Custom pattern	Your own patterns in the Lua notation, for example %d%d%d%d%- %d%d%d%d . Anchor a pattern at word boundaries to prevent false

Data	Description
Patterns (Lua regex)	positives. Each pattern that you add is shown as a tag below the field, where you can also remove it.
Scrub placeholder	Text that Myra AI Workspace puts in the place of the match. The field is shown only for the scrub action. The default is [REDACTED] .

4.4.2.1.4 Data for the Keyword detector

You also specify the following data:

Data	Description
Keywords	Terms that the detector examines the text for, for example confidential . Each term that you add is shown as a tag below the field, where you can also remove it.
Block-safe example button	Sets the action to block , selects Whole words only , and enters unambiguous internal terms. These terms cause almost no false positives.
Flag-only example button	Sets the action to flag , selects Whole words only , and enters general sensitive terms. For the block action, these terms cause too many false positives.
Whole words only	If the check box is selected, a term matches only as a separate word. The setting is recommended and selected by default. Clear it only for substrings such as product codes.
Case-sensitive	If the check box is selected, a term matches only in exactly the notation that you entered.



For the **block** action, the detector shows the note that each match blocks the request fully. Use only unambiguous terms such as internal code names or product identifiers. General words cause high false-positive rates.

4.4.2.1.5 Data for the Jailbreak detector

You also specify the following data:

Data	Description
Default phrases	Supplied phrases that the detector finds. The list is read-only and shows the note that your own phrases replace it.
Active phrases	Your own phrases. As soon as the list has a phrase, the detector shows the Custom list active message and examines the text only for your own phrases. If you remove all phrases, the default phrases apply again.
Whole words only	If the check box is selected, a phrase matches only with exact word boundaries. The setting is cleared by default, so that the detector also finds inflected forms.
Case-sensitive	If the check box is selected, a phrase matches only in exactly the notation that you entered.

4.4.2.1.6 Data for the Presidio (NLP) detector

You also specify the following data:

Data	Description
Entity types	Types of personal data that the detector examines the text for. The Focused PII (0% FP) button selects the 14 entities without false positives, the All entities (~10% FP) button the full catalog. In the field for custom entities, you add more entity types, for example IN_PAN .
Score threshold	Minimum confidence between 0 and 1 at which a match counts. Usual values are between 0.7 and 0.85 . For the PERSON , LOCATION , and DATE_TIME entities, which cause many false positives, the gateway automatically sets 0.9 and for ORG the value 0.85 .
Match mode	exact reports only full matches, partial also partial matches.

Data	Description
Allow list	Comma-separated values that the detector never flags as personal data.
Fail open	If the check box is selected, the gateway lets the requests continue while the Presidio service is not available. If you clear it, the gateway rejects the requests during this time.



The **All entities** mode shows a warning. **PERSON**, **LOCATION**, and **DATE_TIME** cause approximately 10 to 15 % false positives on usual text, also with the increased threshold. A threshold above **0.85** conversely no longer reports the less certain matches. This data then goes to the model unmasked.

4.4.2.1.7 Data for the Prompt Guard detector

You also specify the following data:

Data	Description
Safety categories	Categories of the Llama Guard classifier from S1 to S14 . Each category shows its false-positive risk as low FP , medium FP , or high FP . The Recommended block (~1.7% FP) button selects the S1 , S3 , S4 , S9 , S11 , S12 , and S14 categories, the All categories (~18.6% FP) button all fourteen.
Context prompt (optional)	Text that Myra AI Workspace puts before each user message, to tell the classifier the purpose of the application. The context decreases the false positives, for example for the S2 category from 14.5 % to 7.2 %.
Timeout (ms)	Time to wait for the classifier. The default is 3000 milliseconds.
Fail open	If the check box is selected, the gateway lets the requests continue while the classifier is not available.

4.4.2.1.8 Data for the Prompt Injection detector

You also specify the following data:

Data	Description
Injecti on thresh old	Probability between 0 and 1 at which the detector blocks or flags. The default is 0.5 . A higher value decreases the false positives, a lower value also finds the weaker attempts.
Fail open	If the check box is selected, the gateway lets the requests continue while the classifier is not available. If you clear the check box, the gateway rejects the requests during this time. The default is cleared.



The platform administration specifies the endpoint of the classifier and its credentials for the full system. In the detector you specify only the threshold and the behaviour during a malfunction.

4.4.2.1.9 Data for the PII Protector detector

You also specify the following data:

Data	Description
Entity types	Types of personal data that the detector replaces. The selection is the same as for the Presidio (NLP) detector.
Score threshold	Minimum confidence between 0 and 1 at which a match counts. The default is 0.7 .
Language calibration	Auto finds the language and keeps the German threshold 0.6 for names and addresses as a fallback, German enforces the German model, and English / Other deactivates the German threshold.
Find German ID and tax numbers	If the check box is selected, the detector finds the German tax identification number, the health insurance number, and the VAT identification number by their check digit. The setting operates independently of the language calibration and is cleared by default.
Skip system	If the check box is selected, the detector examines only the messages of the users.

Data	Description
and assistant messages	
Allow list	Comma-separated values that the detector never replaces with a token.
Fail open	If the check box is selected, the gateway lets the requests continue while the Presidio service is not available.



The detector replaces each match in the request with an opaque token in the form **[PII:a3f9b2:1]** and puts the initial value back before it sends the response. Thus, the model never processes the personal data, but the users still get a usable response.

4.4.2.1.10 Data for the Custom PII Blacklist detector

You also specify the following data:

Data	Description
Sensitive terms	Your own terms that the detector masks, for example customer names, project code names, or employee names. Each term that you add is shown as a tag below the field.
Whole words only	If the check box is selected, the detector masks a term only as a separate word. The setting is cleared by default. The detector then also masks a term in longer strings.
Case-sensitive	If the check box is selected, the detector masks a term only in exactly the notation that you entered. Select the check box for names with umlauts or accents.

4.4.2.1.11 Execution plan

Below the detector areas, the **Execution plan** section shows the detectors in the sequence in which the gateway examines the data.

The following information is available:

Column	Description
Tier	Tier of the detector: 1 for the detectors in the gateway, 2 for the detectors with their own service.
Name	Name of the detector with a dot in the color of its type.
Phase	Object of the examination with an arrow: → request examines the request, ← response the response, and ↔ both the request and the response.
Mode	Treatment of a match: block , scrub , flag , or ↺ reversible for the two detectors that replace and restore data.


4.4.2.2 Changing the sequence of the guardrail detectors

The sequence of the areas controls the order in which the detectors of a tier examine the data.

The following requirements are met:

- ☑ Your user account has the **Manage gateways and routing** permission.
- ☑ The gateway has a minimum of two detectors.

Proceed as follows to change the sequence:



Guardrails

128 blocked 64 scrubbed 311 flagged avg 42 ms guardrail latency - last 24h

Guardrails (3) Regex / Pattern Keyword Jailbreak Presidio (NLP) Prompt Guard Prompt Injection PII Protector Custom PII Blacklist

Block - request X

Flag - request X

PII Protector pi-protect protect - both X

Name pi-protect Target request + response (required)

Score threshold 0.7

Global minimum confidence (0 - 1). High-FP entities (PERSON, LOCATION, DATE, TIME) automatically use 0.7 and ORG uses 0.85 - raising the global threshold above these values will also filter them. Values below 0.85 have little effect on high-FP entities.

Entity types

Focused PII (92% FP) All entities (~10% FP)

Email Address low FP Phone Number low FP SSN low FP

Credit Card low FP US Bank Account low FP IDIN low FP

US Passport low FP Passport (multilingual) low FP US Driver License low FP

ITIN low FP Crypto Address low FP IP Address low FP

Medical License low FP URL low FP Organisation medium FP

Person Name high FP Location high FP Date / Time high FP

Custom entity (e.g. NLP, URL, etc.) Add

Allow list e.g. Myra Security, John Doe Match mode exact

Comma-separated values that will never be tokenized as PII.

☒ Skip system & assistant messages (scan user turns only)

☒ Fail open (pass requests through if Presidio is unreachable)

German PII (DSGVO). Without an explicit language (Automatic), the German 0.6 confidence floor for names and addresses applies as a fail-safe; 'English' turns it off. 'German (ID recognizers)' works independently of the language choice. Off by default.

Language calibration Automatic - detect language (German DSDVO floor applies as a fail-safe)

German forces the German model for names and addresses. English keeps the English calibration (no German floor).

☐ German ID recognizers (Blaues-ID, KVV-Nr, USt-IdNr, ...)

Defects German tax and health insurance identifiers by checksum - independent of the language choice. Off by default.

EXECUTION PLAN

Tier	Name	Phase	Mode
1	pi-check	request	block
2	jailbreak-check	request	flag
3	pi-protect	both	reusable

Show recent activity

Figure 221: Guardrails area

- Open the detail view of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- In the area of the detector, click the **Move up** or **Move down** icon.
- Click the **Save Guardrails** button.
- Myra AI Workspace shows the **Saved** message.



The **Move up** icon is unavailable in the top area, the **Move down** icon in the bottom area. The tier is more important than the sequence of the areas. A detector of tier 2 examines the data subsequently also when its area is at the top.

4.4.2.3 Removing a guardrail detector

A detector that you removed no longer examines the requests of this gateway.

The following requirements are met:

- ☒ Your user account has the **Manage gateways and routing** permission.

Proceed as follows to remove a guardrail detector:




Figure 222: Guardrails area

- Open the detail view of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- In the area of the detector, click the icon to remove it.
 - ↳ The **Remove the <Name> detector from this guardrail?** confirmation prompt opens.




Figure 223: Confirmation prompt before the removal of a detector

- ▶ Confirm the prompt with the **Remove detector** button.
- ▶ Click the **Save Guardrails** button.
- Myra AI Workspace shows the **Saved** message.



For the **Presidio (NLP)**, **PII Protector**, and **Custom PII Blacklist** detectors, the prompt also warns that the removal deactivates the masking of personal data for this gateway and that the data subsequently goes to the model unmasked.

4.4.2.3.1 Effects of the removal

The removal has the following effects:

Object	Effect
Requests	The requests to this gateway no longer go through the detector. Requests that it blocked before now get to the provider.
Guardrail events	The events recorded for the detector stay. The area records no new events for it.

Object	Effect
Other gateways	The detectors of other gateways do not change. A detector always applies only to the gateway in which it is located.
Counter	The Detectors column of the Gateways view shows the remaining quantity.

4.4.3 Managing routing rules


For the requests that its conditions match, a routing rule overrides the default routing of the gateway. Thus, you can for example send all requests of one tenant to a less expensive model, or distribute the load of one model to several providers.

4.4.3.1 Creating a routing rule

The following requirements are met:

- ☒ Your user account has the **Manage gateways and routing** permission.
- ☒ The tenant has a gateway.

Proceed as follows to create a routing rule:



Gateways

Tenant: **muster-gmbh**

Search gateways by name... **New Gateway**

Select Tenant Filter tenants by name...

muster-gmbh **beispiel-ag** **probe-kg**

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	-	0s	-	-	7/19/2026

Figure 224: Gateways view

- In the **Settings** area, open the **Gateways** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- In the row of the gateway, click the **Open** → row action.
 - ↳ The detail view of the gateway opens.




Figure 225: Detail view of a gateway

- In the **Routing Rules** area, click the **New Rule** button.
- ↳ The **New Routing Rule** dialog opens.

New Routing Rule

Priority (higher = first)

100

☒ Enabled

Conditions (all must match) Add

model

eq

claude-sonnet-4-6

x

header

x-customer-tier

eq

premium

x

Action

Direct route

Load balance

Provider

anthropic

Model

claude-opus-4-6

Fallbacks (tried in order) Add

openai

gpt-4.1

x

Cancel

Create Rule

Figure 226: New Routing Rule dialog

- Specify the data of the rule.
- Click the **Create Rule** button.
- Myra AI Workspace creates the rule and shows it in the table of the area.




Figure 227: Detail view of a gateway

4.4.3.1.1 Data in the New Routing Rule dialog

You specify the following data:

Data	Description
Priority (higher = first)	Rank of the rule as an integer. The gateway examines the rules from the highest to the lowest priority and applies the first rule whose conditions match. The default is 0 .
Enabled	If the check box is selected, the gateway examines the rule. If you clear it, the rule stays but has no effect. The check box is selected by default.
Conditions (all must match)	Conditions in which the rule applies. See the Conditions section.
Action	Target of the routing: Direct route sends the request to exactly one model, Load balance to several targets. See the Action: Direct route and Action: Load balance sections.

4.4.3.1.2 Conditions

With the **Add** button, you add a condition. With the button with the cross, you remove it again. All conditions of a rule must match. While the rule has no condition, the dialog shows the **No conditions = matches all requests.** note.

Each condition has a field, an operator, and a value:

Data	Description
Field	Object of the examination: model examines the model requested, provider the provider requested, tenant_id the identifier of the tenant, header an HTTP header of the request, and meta a metadata key of the request.
Name	Name of the header or of the metadata key, for example x-customer-tier or region . The field is shown only for the header and meta fields.
Operator	Comparison: eq matches on an exact match, neq on each difference, and prefix if the value matches at the start.
Value	Comparison value. For the model and provider fields, you select it from the catalog with the eq and neq operators. With the prefix operator and for the tenant_id , header , and meta fields, you enter it freely.

4.4.3.1.3 Action: Direct route

You specify the following data for the **Direct route** action:

Data	Description
Provider	Provider to which the gateway sends the request. The – any provider – entry keeps the provider of the request.
Model	Model to which the gateway sends the request. The – any model – entry keeps the model of the request.
Fallbacks (tried in order)	More pairs of a provider and a model. If the first target does not answer a request, the gateway tries the fallbacks in the sequence shown. With the Add button, you add a fallback. With the button with the cross, you remove it.

4.4.3.1.4 Action: Load balance

You specify the following data for the **Load balance** action:

Data	Description
Strategy	Distribution of the requests: Weighted random distributes them by the weight of the targets, Round robin in turn.
Sticky field (optional)	Field whose value keeps a target, for example meta.user_id . The gateway sends the same value to the same target for the duration of the sticky TTL.
Sticky TTL (s)	Duration of the connection to a target in seconds. The field is shown only if the sticky field has a value. The default is 3600 seconds.
Targets	Pairs of a provider and a model with their weight. With the Add target button, you add a target. With the button with the cross, you remove it.
Weight (0 = disabled)	Share of a target as an integer between 0 and 100 . Adjacent to the field, the dialog shows the share calculated from it as a percentage. The weight 0 deactivates a target without removing it. The dialog then shows it as off .




The gateway tries the active targets that are not selected as fallbacks on a failure.

4.4.3.2 Editing a routing rule

The following requirements are met:

- ☒ Your user account has the **Manage gateways and routing** permission.
- ☒ The gateway has a minimum of one routing rule.

Proceed as follows to edit a routing rule:



Gateway: `produktion` Edit Delete

BUDGET

€342.00 / €460.00 74% used

This month
= converted at today's rate; the cap is stored and enforced in USD.

CACHE TTL
5 m

RETRIES
2

TIMEOUT
60 s

AUTH
required

RATE LIMIT
60/60 s

LOG PAYLOADS
yes

WEBHOOK
all events

TRACING
on

WEB SEARCH
on

IMAGE GEN
off

PROMPT CACHING
on

COMPACTION
on

COMPACT ERR
32k tok

Endpoint URL
anthropic
https://ai-gateway.beispiel.de/v1/ruster-gmbh/produktion/anthropic/chat/completions Copy

Provider Keys Add Model

PROVIDER	ALIAS	ADDED
anthropic	Editorial Anthropic	5/14/2026 🗑️
openai	Analysis OpenAI	6/2/2026 🗑️


Auth Tokens Generate

An authentication token is a secret key an application sends with each request to prove it may use this gateway – like a password for software. Create one per app or integration so you can revoke it individually.

LABEL	HASH (FIRST 16)	RATE LIMIT	SPEND CAP	EXPIRES	CREATED
Editorial system	a7b2c3d4e5f60718...	60/60 s	€50.00	12/31/2026, 11:59:00 PM	4/2/2026 🗑️
Analysis	9f8d76c5d4a3021...	10/60 s	–	never	6/11/2026 🗑️

Figure 228: Detail view of a gateway

- ▶ Open the detail view of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- ▶ In the **Routing Rules** area, click the **Edit** row action in the row of the rule.
 - ↳ The **Edit Rule** dialog opens.



The 'Edit Rule' dialog box contains the following fields and controls:

- Priority (higher = first):** A text input field with the value '50'.
- Enabled:** A checkbox that is checked.
- Conditions (all must match):** A section with a dropdown menu set to 'provider', a dropdown menu set to 'eq', and a text input field with the value 'anthropic'. There is an 'Add' button and a red 'x' icon.
- Action:** Two buttons: 'Direct route' and 'Load balance'.
- Strategy:** A dropdown menu set to 'Weighted random'.
- Sticky field (optional):** A text input field with the value 'meta.user_id'. Below it, a note says 'Routes same value to same target for TTL seconds'.
- Sticky TTL (s):** A text input field with the value '3600'.
- Targets:** A table with three rows. Each row has a dropdown menu, a text input field, a number input field, a percentage input field, and a red 'x' icon. The first row has 'anthropic', 'claude-sonnet-4-6', '70', and '70%'. The second row has 'openai', 'gpt-4.1', '30', and '30%'. The third row has 'mistral', 'mistral-large', '0', and 'off'.
- Buttons:** 'Cancel' and 'Save Rule'.

Figure 229: Edit Rule dialog

- Change the data of the rule.
- Click the **Save Rule** button.
- Myra AI Workspace applies the changes.



The **Edit Rule** dialog has the same data as the **New Routing Rule** dialog. See the **Data in the New Routing Rule dialog** section. The changes apply from the next request. Requests in progress keep the target that was found for them.


4.4.3.3 Changing the sequence of the routing rules

The table of the area shows the rules in the sequence in which the gateway examines them.

The following requirements are met:

- ☑ Your user account has the **Manage gateways and routing** permission.
- ☑ The gateway has a minimum of two routing rules.

Proceed as follows to change the sequence:



The screenshot shows the 'production' gateway configuration. It includes a budget bar at €342.00 / €460.00 (74% used), a 'This month' note, and various settings like CACHE TTL (5 m), RETRIES (2), TIMEOUT (60 s), AUTH (required), RATE LIMIT (60/60 s), LOG PAYLOADS (yes), WEBHOOK (all events), TRACING (on), WEB SEARCH (on), IMAGE GEN (off), PROMPT CACHING (on), COMPACTION (on), and COMPACT ERR (32k tok). The endpoint URL is 'anthropic' with a full URL provided. Below are 'Provider Keys' for 'anthropic' and 'openai', and 'Auth Tokens' for 'Editorial system' and 'Analysis'.

PROVIDER	ALIAS	ADDED
anthropic	Editorial Anthropic	5/14/2026
openai	Analysis OpenAI	6/2/2026

LABEL	HASH (FIRST 16)	RATE LIMIT	SPEND CAP	EXPIRES	CREATED
Editorial system	a7b2c3d4e5f6d7f8...	60/60 s	€50.00	12/31/2026, 11:59:00 PM	4/2/2026
Analysis	9f8c7d6c5b4a3d21...	10/60 s	-	never	6/11/2026

Figure 230: Detail view of a gateway

- ▶ Open the detail view of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- ▶ In the **Routing Rules** area, click the **Move up (evaluate earlier)** or **Move down (evaluate later)** row action in the row of the rule.
- Myra AI Workspace exchanges the rule with its adjacent rule and writes the priorities again.



The **Move up (evaluate earlier)** row action is unavailable in the top row, the **Move down (evaluate later)** row action in the bottom row. While Myra AI Workspace writes the new sequence, both row actions are unavailable in all rows.

4.4.3.4 Deleting a routing rule

The following requirements are met:

- ☑ Your user account has the **Manage gateways and routing** permission.

Proceed as follows to delete a routing rule:




Figure 231: Detail view of a gateway

- Open the detail view of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▷ Open the **Gateways** group.
 - ▷ Open the gateway with the **Open** → row action.
- In the **Routing Rules** area, click the **Del** row action in the row of the rule.
 - ↳ The **Delete this routing rule?** confirmation prompt opens.




Figure 232: Confirmation prompt before the deletion of a routing rule

- ▶ Confirm the prompt.
- Myra AI Workspace deletes the rule.

4.4.3.4.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Reques ts	The requests that the rule matched again follow the default routing of the gateway or the subsequent rule whose conditions match.
Other rules	The other rules keep their priority and their sequence.
Provide r keys	The keys of the gateway stay. A rule refers only to providers and models. It has no credentials of its own.
Reques t log	The requests recorded stay.



A rule that you only want to suspend temporarily is better cleared with the **Enabled** check box than deleted. The data then stays.

4.4.4 Managing governance templates

A governance template maps compliance obligations onto the settings of a gateway.

4.4.4.1 Creating a governance template

The following requirements are met:

- ☑ Your user account has the **Manage governance rules, prompts, and approvals** permission.

Proceed as follows to create a governance template:





Figure 233: Security view

- ▶ In the **Settings** area, open the **Security** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Click on the **New template** button.
 - ↳ The **New governance template** dialog opens.



The dialog box is titled "New governance template". It contains a disclaimer: "Templates created here are non-authoritative and are not a legal compliance guarantee. Authoritative content is signed off by the Data Protection Officer." Below this, there are input fields for "Name", "Template key", and "Version". The "Version" field is pre-filled with "1". Under the heading "TARGET CONTROLS", there is a descriptive text: "Each control sets a governable gateway leaf on or off. The residency floor may refuse a downgrade at apply time." Below this text are two dropdown menus: the first is labeled "eu_region_routing" and the second is labeled "on". There is also an "Obligation reference (optional)" field and a red "Remove" button. At the bottom, there is an "Add control" button and "Cancel" and "Save" buttons.

Figure 234: New governance template dialog

- ▶ Enter the **Name**, the **Template key**, and the **Version**.
- ▶ Click on the **Add control** button and set the **Control**, the **Value**, and, optionally, the **Obligation reference**.
- ▶ Save the template.
- Myra AI Workspace creates the template.

4.4.4.1.1 Entries in the New governance template dialog

You set the following entries when you create the template:

Entry	Description
Name	Name of the template in the list. The name holds a maximum of 255 characters. The entry is mandatory.
Template key	Technical key of the template, for example gdpr-basic . The key holds a maximum of 190 characters and appears in the Key column. The entry is mandatory.
Version	

Entry	Description
	Version of the template as an integer from 1 on. The version appears in the Key column after the key, for example v1 .
Target controls	List of the controls the template sets on a gateway. Each row carries the Control and Value drop-down lists and the Obligation reference input field.
Control	Governable option of the gateway the template sets.
Value	State of the option: on or off .
Obligation reference	Reference to the compliance obligation this control satisfies, for example the article of a regulation. The entry is optional.

The following options are available in the dialog:

Option	Description
Add control button	Adds a row for one more control. The button is disabled if no governable option is available.
Remove button	Removes the row from the target controls.
Save button	Creates the template. The button is disabled while Name or Template key is empty.
Cancel button	Closes the dialog without creating the template.



The dialog states: **Templates created here are non-authoritative and are not a legal compliance guarantee. Authoritative content is signed off by the Data Protection Officer.** A template created this way carries the **Non-authoritative** label in the list.

4.4.4.1.2 Entries in the Edit governance template dialog

You change the following entries in the dialog:

Entry	Effect of the change
Name	Changes the name in the list and in the title of the Preview & apply dialog.
Target controls	Changes which options the template sets. The change takes effect only when you apply the template to a gateway again. Controls that were applied already stay unchanged on the gateways.



Template key and **Version** cannot be changed after the creation. The dialog shows both fields only when you create the template. For an authoritative template the **Edit** and **Delete** row actions are disabled and carry the note **Authoritative templates are read-only**.

4.4.4.2 Applying a governance template

Proceed as follows to apply a governance template to a gateway:




Figure 235: Security view

- ▶ In the **Settings** area, open the **Security** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the template and select the gateway.
 - ↳ The preview shows which settings of the gateway change.
- ▶ Apply the template.

→ Myra AI Workspace puts the settings into the gateway.



Myra AI Workspace always enforces the residency minimum. A template does not lower it.

4.4.4.2.1 Entries in the Preview & apply dialog

The **Preview & apply** – dialog compares the template with a gateway before it applies the template.

You set the following entries in the dialog:

Entry	Description
Gateway	Gateway the template acts on. The drop-down list shows the gateways of the selected tenant. While no gateway is selected, the Preview and Apply buttons are disabled.

The preview shows the following information for each control:

Column	Description
Control	Governable option of the gateway.
Requested	Value the template wants to set: on or off .
Result	Value the option carries after the apply.
Status	Will change, Unchanged, Blocked by residency floor or Set on tenant .

The following options are available in the dialog:

Option	Description
Preview button	Compares the template with the selected gateway and shows the table of the differences. The gateway does not change.
Apply button	

Option	Description
	Writes the controls to the gateway and shows Template applied . The button is disabled if no control changes the gateway. It then carries the note Nothing to apply – no control changes the gateway .
	If the template changes nothing, the dialog shows This template changes nothing on the selected gateway . After the apply, the dialog shows Applied to the gateway .
	Myra AI Workspace always enforces the residency floor. A downgrade carries the Blocked by residency floor status with the note The tenant residency floor refuses this downgrade – not applied . PII masking carries the Set on tenant status with the note PII masking is set on the tenant, not on a gateway .

4.4.4.3 Deleting a governance template

Proceed as follows to delete a governance template:




Figure 236: Security view

- ▶ In the **Settings** area, open the **Security** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Delete the template.
 - ↳ The **Delete the template "?"** query opens.




Figure 237: Delete template query

- ▶ Confirm the query.
- Myra AI Workspace shows **Template deleted**. The settings already applied to the gateways stay as they are.

4.4.4.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Template	Myra AI Workspace removes the template from the list. It is no longer available for the preview and the apply.
Gateways	The controls that were applied already stay on the gateways. The deletion resets no gateway.
Obligation references	The references of the template are removed with it. The record of the mapping no longer exists.



An authoritative template cannot be deleted. The **Delete** row action is disabled and carries the note **Authoritative templates are read-only**.

4.5 User Management

User management determines who uses Myra AI Workspace and what they are allowed to do in it. A user account gives a person access and carries their role. Groups, roles, and tenants organize the accounts and their permissions.

The following sections describe these configurations:

- [Managing groups](#)
- [Managing user accounts](#)
- [Managing roles](#)
- [Managing tenants](#)

4.5.1 Managing groups

A group holds users together to give them the access to agents, projects, and connectors as one.

4.5.1.1 Creating a group

The following requirements are met:

- ☑ Your user account has the **Manage groups and memberships** permission.

Proceed as follows to create a group:





Figure 238: Groups view

- In the **User Management** area, open the **Groups** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- Click on the **New group** button.
 - ↳ The **New group** dialog opens.



The image shows a 'New group' dialog box. It has a title 'New group' at the top. Below the title, there are two input fields. The first is labeled 'Name' and contains the text 'e.g. Engineering'. The second is labeled 'Description' and contains the text 'Optional'. At the bottom right of the dialog box, there are two buttons: 'Cancel' and 'Create'.

Figure 239: New group dialog

- ▶ Enter the name of the group.
- ▶ Click on the **Create** button.
- Myra AI Workspace creates the group.

4.5.1.1.1 Data in the New group dialog

You set the following data at creation:

Item	Description
Name	Name of the group, for example Engineering . The name appears in the Name column of the list. The item is mandatory.
Description	Short description of the group. The description appears in the Description column of the list. The item is optional.

The following options are available in the dialog:


Option	Description
Create button	Creates the group. During the creation, the button shows Saving....
Cancel button	Closes the dialog without creating the group.



If the name field stays empty, the dialog shows **Name is required**. A new group holds neither members nor grants. You assign both in the dialog of the group.

4.5.1.2 Editing a group

Proceed as follows to edit a group:




Groups New group

Tenant: muster-gmbh

NAME	DESCRIPTION	MEMBERS	CREATED
Editorial	Handbook and website	8	2/18/2026
Sales	Quotes and tenders	12	3/2/2026
Engineering	Operations and development	15	1/9/2026

Figure 240: Groups view

- ▶ In the **User Management** area, open the **Groups** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the group.
 - ↳ The dialog with the name of the group opens.



Engineering

Name
Engineering

Description
Technical writing

Members 1

Alex Muster
alex.muster@beispiel.de

Add member (pick from the list or type an email)
user@example.com

Add member

Access grants 1

Quotes 2026 Editor

Grant access to a knowledge space Role

Select a pro... Viewer – ca... Add grant

Figure 241: Dialog of a group

- Click on the **Add member** button and select a user account from the list, or enter its email address.
- Click on the **Add grant** button and select the object to be shared.
- Myra AI Workspace shows the addition with **added to the group**.




Figure 242: Dialog of a group

4.5.1.2.1 Data in the dialog of the group

You change the following data in the dialog:

Item	Effect of the change
Name	Changes the name of the group in the list and in every grant.
Description	Changes the short description of the group in the list.
Members area	Takes users in or removes them. A user taken in receives every access that the group grants. A user removed loses it.
Access grants area	Grants or revokes the access to a project. The grant applies to all members of the group at once.

The following options are available in the dialog:

Option	Description
--------	-------------

Save button

Option	Description
	Applies the changed data. During the saving, the button shows Saving....
Cancel button	Closes the dialog without applying the changes.

4.5.1.3 Deleting a group

Proceed as follows to delete a group:




Figure 243: Groups view

- ▶ In the **User Management** area, open the **Groups** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Delete the group.
 - ↳ The **Delete the group ""? Its memberships and grants are removed.** query opens.




Figure 244: Delete group query

- ▶ Confirm the query.
- Myra AI Workspace deletes the group with its memberships and grants.

4.5.1.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Group	Myra AI Workspace removes the group from the list.
Memberships	Myra AI Workspace removes all memberships of the group. The user accounts themselves remain.
Access grants	Myra AI Workspace revokes every access that the group granted. The members lose it unless they hold it by another route.



If the group cannot be deleted, it remains in the list.

4.5.2 Managing user accounts


A user account gives a person access to Myra AI Workspace and carries their role.

4.5.2.1 Creating a user account

The following requirements are met:

- ☑ Your user account has the **Manage tenant users** permission.
- ☑ The email address of the person is available.

Proceed as follows to create a user account:



Users
3 users

Search users by name or email... ☐ Show deleted All tenants ▾ Export CSV Export XLSX

New User

EMAIL	NAME	ROLE	TENANT	LAST LOGIN	CREATED
alex.muster@beispiel.de	Alex Muster	Tenant Administrator	muster-gmbh	8/16/2026	1/12/2026
kim.probe@beispiel.de	Kim Probe	Member	muster-gmbh	8/15/2026	3/4/2026
sam.beispiel@beispiel.de Deleted	Sam Beispiel	Viewer	beispiel-ag	never	6/21/2026

Figure 245: Users view

- In the **User Management** area, open the **Users** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- Click on the **New User** button.
 - ↳ The **New User** dialog opens.


A screenshot of a 'New User' dialog box. The dialog has a title 'New User'. Below the title, there are four input fields: 'Tenant *' with a dropdown menu showing 'muster-gmbh', 'Email *' with a text input showing 'alice@example.com', 'Name' with a text input showing 'Alice', and 'Role' with a dropdown menu showing 'Member'. Below these fields is a line of text: 'We'll email them an invitation to set up their account.' At the bottom right of the dialog are two buttons: 'Cancel' and 'Create User'.

Figure 246: New User dialog

- ▶ Select the tenant.
 - ▶ Enter the email address and select the role.
 - ▶ Click on the **Create User** button.
- Myra AI Workspace creates the user account. The person signs in with a one-time code.

4.5.2.1.1 Data in the New User dialog

You set the following data at creation:

Item	Description
Tenant *	Tenant that the user account belongs to. The list holds the tenants of the work area. The item is mandatory.
Email *	

Item	Description
	Email address of the user account, for example alice@example.com . Myra AI Workspace sends an invitation to this address to set up the account. The item is mandatory.
Name	Name of the user. The item is optional.
Role	Role of the user account: Administrator, Tenant Administrator, AI Manager, Finance, Member, Viewer , or Demo User . The role sets which views and functions the account uses.

The following options are available in the dialog:


Option	Description
Create User button	Creates the user account and sends the invitation. During the creation, the button shows Creating....
Cancel button	Closes the dialog without creating the user account.



If the email address is already part of a workspace, the dialog shows **That email address is already part of a workspace**. If the plan holds no further seats, the dialog shows **You've reached your plan's seat limit. Remove a member, or change your plan to add more seats**.

4.5.2.2 Editing a user account

Proceed as follows to edit a user account:



Users
3 users [New User](#)

Search users by name or email... ☐ Show deleted All tenants ▾ [Export CSV](#) [Export XLSX](#)

EMAIL	NAME	ROLE	TENANT	LAST LOGIN	CREATED
alex.muster@beispiel.de	Alex Muster	Tenant Administrator	muster-gmbh	8/16/2026	1/12/2026
kim.probe@beispiel.de	Kim Probe	Member	muster-gmbh	8/15/2026	3/4/2026
sam.beispiel@beispiel.de Deleted	Sam Beispiel	Viewer	beispiel-ag	never	6/21/2026

Figure 247: Users view

- ▶ In the **User Management** area, open the **Users** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the user account.
 - ↳ The detail panel opens.





Figure 248: Detail panel of a user account

- ▶ Click on the **Edit** button.
- ↳ The **Edit:** dialog opens.



Edit: alex.muster@beispiel.de

Tenant *

muster-gmbh

Email *

alex.muster@beispiel.de

Name

Alex Muster

Role

Tenant Administrator

☐ Restrict processing (GDPR Art. 18)

While restricted, this account cannot send requests to any model. A reversible legal hold, separate from deletion.

Conversation retention (days)

Leave blank to disable

Permanently delete this user's conversations with no activity for this many days (30-3650). Leave blank to disable. A user first may only shorten retention versus the organisation policy, never lengthen it. Deletion is irreversible and also requires the operator to enable it deployment-wide.

Cancel **Save Changes**

Figure 249: Edit dialog

- Change the tenant, the name, the role, the processing restriction, or the conversation retention.
- Click on the **Save Changes** button.
- Myra AI Workspace applies the change at the next request of the user account.




Figure 250: Detail panel of a user account

4.5.2.2.1 Data in the Edit dialog

You change the following data in the dialog:

Item	Effect of the change
Tenant *	Assigns the user account to another tenant.
Email *	Changes the address that the person logs in with.
Name	Changes the name in the list of the user accounts and in the detail view.
Role	Changes the role of the user account: Administrator, Tenant Administrator, AI Manager, Finance, Member, Viewer, or Demo User . Myra AI Workspace applies the new role with the next request of the user account.
Restrict processing	While restricted, this account cannot send requests to any model. A reversible legal hold, separate from deletion.

Item	Effect of the change
g (GDPR Art. 18)	
Conversations retention (days)	Permanently deletes this user's conversations with no activity for this many days (30–3650). To disable the deletion, leave the field empty. A user period may only shorten retention versus the organization policy, never lengthen it.

The following options are available in the dialog:

Option	Description
Save Changes button	Applies the changes. During the saving, the button shows Saving...
Cancel button	Closes the dialog without applying the changes.



The deletion by the retention period is irreversible and also requires the operator to enable it deployment-wide.

4.5.2.3 Deleting a user account

Proceed as follows to delete a user account:




Figure 251: Users view

- ▶ In the **User Management** area, open the **Users** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the user account.
 - ↳ The detail panel opens.




Figure 252: Detail panel of a user account

- ▶ Click on the **Delete User** button.
 - ↳ The **Delete user ""? Their tokens will also be revoked.** query opens.




Figure 253: Delete user query

- ▶ Confirm the query.
- Myra AI Workspace deletes the user account and revokes its tokens. The **Show deleted** check box shows it again with the **Deleted** mark.

4.5.2.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
User account	Myra AI Workspace marks the account as Deleted . The person no longer logs in.
Tokens	Myra AI Workspace revokes all tokens of the user account. Requests with these tokens then fail.
List	The Show deleted check box shows the account again. The Restore row action restores it.



If the deletion fails, the view shows **Couldn't remove that user. Please try again.**

4.5.3 Managing roles


Your own role puts a set of permissions together for the tenant.

4.5.3.1 Creating a role

The following requirements are met:

- ☑ Your user account has the **Manage roles within the tenant** permission.
- ☑ Your user account has every permission that the new role is to get.

Proceed as follows to create a role:



ROLE	TYPE	PERMISSIONS	USERS	ACTIONS
Tenant Administrator	System	38	2	Edit
Member	System	12	18	Edit
Editorial	Custom	7	6	Edit

Figure 254: Roles & permissions view

- In the **User Management** area, open the **Roles & permissions** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- Click on the **New role** button.
 - ↳ The **New role** dialog opens.


The image shows a 'New role' dialog box. It has two input fields at the top: 'Name' with the placeholder 'e.g. Support agent' and 'Description (optional)' with the placeholder 'What this role is for'. Below these is a 'Permissions' section with a note: 'Pick permissions per module. A permission you do not hold yourself is disabled – you cannot grant what you lack.' There are two rows of permission modules. The first row includes 'Agents 1/3' (highlighted with a blue underline), 'Analytics 0/4', 'Gateways 0/3', 'Groups 0/1', and 'Connectors 0/2'. The second row includes 'Projects 0/1', 'Roles 0/2', and 'Users 0/2'. Under the 'Agents 1/3' module, there are three checkboxes: 'Create and manage own agents' (checked), 'Review, approve, and moderate agents' (unchecked), and 'Curate the tenant agent catalog' (unchecked). At the bottom right of the dialog are two buttons: 'Cancel' and 'Create role'.

Figure 255: New role dialog

- ▶ Enter the name of the role.
 - ▶ Select the permissions in the modules.
 - ▶ Click on the **Create role** button.
- Myra AI Workspace creates the role. It is available for selection in the **Users** view.



You do not give a permission that your own user account does not have. The editor shows it as not assignable.

4.5.3.1.1 Data in the New role dialog

You set the following data at creation:

Item	Description
Name	Name of the role, for example Support agent . The name appears in the Role column of the list. The item is mandatory.

Item	Description
Description (optional)	Describes what the role is for. The description appears below the name in the list.
Permissions	Permissions of the role. The editor lists them per module. One tab per module opens its permissions.

The following options are available in the dialog:


Option	Description
Create role button	Creates the role. During the creation, the button shows Saving...
Cancel button	Closes the dialog without creating the role.



If the name field stays empty, the dialog shows **A role name is required.** If the name is already taken, the dialog shows **A role with that name already exists.** If you grant a permission that you do not hold yourself, the dialog shows **You cannot grant a permission you do not hold yourself.**

4.5.3.2 Editing a role

Proceed as follows to edit a role:



ROLE	TYPE	PERMISSIONS	USERS	ACTIONS
Tenant Administrator	System	38	2	Edit
Member	System	12	18	Edit
Editorial	Custom	7	6	Edit

Figure 256: Roles & permissions view

- ▶ In the **User Management** area, open the **Roles & permissions** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the role.
 - ↳ The **Edit** dialog opens.




Figure 257: Edit role dialog

- ▶ Change the selection of the permissions.
- ▶ Click on the **Save changes** button.
- Myra AI Workspace applies the change for every user account with this role.

4.5.3.2.1 Data in the Edit NAME dialog

You change the following data in the dialog:

Item	Effect of the change
Name	Changes the name of the role in the list and in the assignment on the user account.
Description (optional)	Changes the description below the name.
Permissions	

Item	Effect of the change
	Changes the permissions of the role. The change applies to every user account that the role is assigned to. Effective permissions are the union of the platform role of the account and every assigned role.

The following options are available in the dialog:


Option	Description
Save changes button	Applies the changes. During the saving, the button shows Saving....
Cancel button	Closes the dialog without applying the changes.



If the role includes permissions that you do not hold yourself, the dialog shows **This role includes COUNT permission(s) you do not hold yourself. You can rename it, but changing its permissions requires an admin who holds them.**

4.5.3.3 Deleting a role

Proceed as follows to delete a role:



Roles & permissions					New role
ROLE	TYPE	PERMISSIONS	USERS	ACTIONS	
Tenant Administrator	System	38	2	Edit	
Member	System	12	18	Edit	
Editorial	Custom	7	6	Edit	

Figure 258: Roles & permissions view

- ▶ In the **User Management** area, open the **Roles & permissions** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Delete the role.

- ↳ The **Delete the role ? Its permissions and every user assignment are removed with it.** query opens.




Figure 259: Delete role query

- ▶ Confirm the query.
- Myra AI Workspace deletes the role. If the step fails, the view shows **Could not delete the role.**

4.5.3.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Role	Myra AI Workspace removes the role from the list.
Permissions	Myra AI Workspace removes the permissions of the role.
User assignments	Myra AI Workspace removes every assignment of the role. The user accounts concerned keep only the permissions of their platform role and of their remaining roles.



If the role cannot be deleted, the view shows **Could not delete the role.**



The permissions of the **Platform** module and the **Manage system (Myra-curated) roles**, **Manage users across all tenants**, **View cross-tenant cost**, **Create tenants**, and **Purge tenants** permissions are not available for selection in the editor. They stay reserved for the **Administrator** role.

4.5.4 Managing tenants

A tenant is the highest unit of Myra AI Workspace. It holds the gateways, the user accounts, and the cost of an organization.

4.5.4.1 Creating a tenant

The following requirements are met:

- ☑ Your user account has the **Create tenants** permission.

Proceed as follows to create a tenant:





Figure 260: Tenants view

- In the **User Management** area, open the **Tenants** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- Click on the **New Tenant** button.
 - ↳ The **New Tenant** dialog opens.



The 'New Tenant' dialog box contains the following fields and options:

- Name ***: Text input with 'my-company'. Below it: 'Lowercase letters, numbers, and hyphens only.'
- Plan**: Dropdown menu with 'Enterprise' selected.
- Tenant Budget (USD)**: Text input.
- Budget Period**: Dropdown menu with 'Monthly' selected.
- Trial ends (demo accounts)**: Text input with '9/30/2026'. Below it: 'Last day of the trial. After this the tenant is blocked from inference until upgraded. The trial budget is the Tenant Budget (USD) above – set the period to Total for a one-off trial cap. Leave empty for no trial.'
- Default user type**: Dropdown menu with 'Consumer' selected. Below it: 'Chat persona for /easy – picks chip labels, example prompts, and system prompts.'
- Assistant name**: Text input with 'Myra AI'. Below it: 'How the assistant identifies itself in chat (e.g. "Myra AI", "Bundestag Assistant"). Leave blank to use the platform default.'
- Buttons: 'Cancel' and 'Create Tenant'.

Figure 261: New Tenant dialog

- ▶ In **Name**, enter the identifier of the tenant.
- The identifier contains lowercase letters, numbers, and hyphens.
- ▶ Select the plan and set the tenant budget and the budget period.
- ▶ If necessary, set the end of the trial, the default user type, and the assistant name.
- ▶ Click on the **Create Tenant** button.
- Myra AI Workspace creates the tenant. Then you create its gateways and user accounts.



The identifier in **Name** cannot be changed after the creation. It is part of every address of a gateway of this tenant.

4.5.4.1.1 Entries in the New Tenant dialog

The dialog groups the entries into the **General**, **Branding**, **Data & Retention**, **Routing**, and **Identity & SSO** tabs. You define the following entries in the **General** tab:

Entry	Description
Name *	Identifier of the tenant, for example my-company . Only lowercase letters, numbers, and hyphens are permitted. The entry is required.
Plan	Plan of the tenant: Free , Standard , or Enterprise .
Tenant Budget (USD)	Spending limit across all gateways of this tenant. If you leave the field empty, then the unlimited entry applies.
Budget Period	Period of the limit: Monthly , Daily , or Total .
Trial ends (demo accounts)	Last day of the trial. After that the tenant is blocked from inference until it is upgraded. The trial budget is the tenant budget. For a one-off trial limit, set the budget period to Total . If you leave the field empty, then no trial applies.
Default user type	Chat persona of the tenant. It determines the labels, the example prompts, and the system prompts of the simplified chat.
Assistant name	Name the assistant uses for itself in the chat. If you leave the field empty, then the platform default applies.
Organization policy	Organization-wide instruction that Myra AI Workspace adds to every chat and agent interaction, for example on tone or language rules. Up to 4000 characters are permitted.
Allow sharing with the whole organization	If the check box is not selected, then no project can be shared org-wide, and existing org shares no longer grant access.
Enable Agents for this workspace	If the check box is not selected, then the Agents area is hidden, and scheduled agent runs are paused.

Entry	Description
Enable Workflows for this workspace	If the check box is not selected, then the Workflows area is hidden. The setting is off by default.
Enable Scheduled Tasks for this workspace	If the check box is not selected, then the Scheduled Tasks area is hidden, and existing scheduled tasks no longer run.
Code interpreter file size limit (MiB)	Largest single file that the code interpreter may deliver to this tenant. The default is 4 . Values from 4 to 25 are permitted. Only platform administrators change this entry.

The following options are available in the dialog:

Option	Description
Create Tenant button	Creates the tenant. During the creation, the button carries the Creating... label.
Cancel button	Closes the dialog without creating the tenant. If unsaved changes exist, then the Discard changes? prompt opens.



If the trial date is in the past, then the dialog reports **This date is in the past – the tenant is blocked from inference immediately**. If a trial is set without a budget, then the dialog reports **This trial has a time limit but no spend cap. Set a Tenant Budget (USD) above to cap trial spend**.

4.5.4.2 Editing a tenant

Proceed as follows to edit a tenant:




Figure 262: Tenants view

- ▶ In the **User Management** area, open the **Tenants** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the tenant.
 - ↳ The detail view of the tenant opens.





Figure 263: Detail view of a tenant

- ▶ Click on the **Edit** button.
 - ↳ The **Edit:** dialog opens.



Edit: muster-gmbh

General Branding Data & Retention Routing Identity & SSO

Plan
Enterprise

Tenant Budget (USD)
500.00
Caps total spend across all of this tenant's gateways.

Budget Period
Monthly

Trial ends (demo accounts)
9/30/2026
Last day of the trial. After this the tenant is blocked from inference until upgraded. The trial budget is the Tenant Budget (USD) above – set the period to Total for a one-off trial cap. Leave empty for no trial.

Default user type
Consumer
Chat persona for /easy – picks chip labels, example prompts, and system prompts.

Assistant name
Myra AI
How the assistant identifies itself in chat (e.g. "Myra AI", "Bundestag Assistant"). Leave blank to use the platform default.

Cancel Save Changes

Figure 264: Edit dialog

- Change the settings in the **General**, **Branding**, **Data & Retention**, **Routing**, or **Identity & SSO** tab.
- Click on the **Save Changes** button.
- Myra AI Workspace applies the change.




Figure 265: Detail view of a tenant

4.5.4.2.1 Entries in the Branding tab

You change the following entries in the **Branding** tab:

Entry	Effect of the change
Product name	Overrides the product name in the application – the browser tab title, the logo label, and the login page. Up to 80 characters are permitted.
Home greeting	Replaces the greeting on the home screen. Up to 200 characters are permitted.
Chat disclaimer	Replaces the disclaimer below the chat input. Up to 128 characters are permitted.
Custom sidebar links	Adds up to four custom navigation links. Every link needs a Label and a URL with the http or https scheme.
Brand color	Colors the headings and links in the PDF export, for example #1a6fb5 .
Brand font	Sets the font family of the body text in the PDF export, for example Helvetica Neue . Only letters, numbers, spaces, and hyphens are permitted.

Entry	Effect of the change
Application logo (dark theme)	Replaces the default logo in the sidebar, in the chat header, and on the login page. PNG and JPEG up to 512 KB and up to 4000×4000 pixels are permitted.
Application logo (light theme)	Replaces the logo in the light theme. Without an entry, the dark-theme logo applies in both themes.
Favicon (browser tab icon)	Replaces the icon in the browser tab. A small square PNG works best.



After the upload, the dialog reports **Logo updated.**, after the removal **Logo removed.** For a wrong format, the dialog reports **Only PNG or JPEG images are allowed.**, for a file that is too large **The image is too large (max 512 KB).**, and for dimensions that are too large **The image dimensions are too large (max 4000×4000 px).**

4.5.4.2.2 Entries in the Data & Retention tab

You change the following entries in the **Data & Retention** tab:

Entry	Effect of the change
Conversation retention (days)	Permanently deletes chat conversations without activity after this many days (30–3650). If you leave the field empty, then the deletion is disabled.
Request-log retention (days)	Permanently deletes request logs with prompts and responses after this many days (30–3650). The period is independent of the conversation retention.
Show the PII masking preview automatically during chat	If the check box is selected, then the masking preview opens before every message. If it is not selected, then the users open the preview themselves.
Disable request	If the check box is selected, then Myra AI Workspace stores no request logs for this tenant, and the Request Logs view stays

Entry	Effect of the change
logging for this tenant	empty. Billing and quotas are not affected. Only platform administrators change this entry.



The deletion by a retention period is irreversible and additionally requires the operator to enable it deployment-wide.

4.5.4.2.3 Entries in the Routing tab

You change the following entries in the **Routing** tab:

Entry	Effect of the change
Bedrock region (AWS)	Sets the EU region for Amazon Bedrock models of the gateways of this tenant. Every gateway without a region of its own inherits it. A gateway region takes precedence. The Deployment default (inherited) entry inherits the platform setting.
Vertex region (Google Cloud)	Sets the EU region for Google Vertex models of the gateways of this tenant. The same precedence rules apply as for the Bedrock region.
Web search provider	Sets the default web search backend for the gateways of this tenant. Every gateway without a provider of its own inherits it. A gateway provider takes precedence. Under an active EU data residency, Myra AI Workspace raises a provider outside the EU to the EU tenant default. The web search API key of the gateway must match the effective provider.

4.5.4.2.4 Entries in the Identity & SSO tab

You change the following entries in the **Identity & SSO** tab:

Entry	Effect of the change
Single sign-on (OIDC) enabled	Lets the users of this tenant sign in via your identity provider. The access applies to existing accounts only. The Issuer URL , Client ID , Client secret , and Subject claim entries describe the identity provider.

Entry	Effect of the change
Allowed email domains	Restricts the sign-in to the comma-separated domains entered. Your identity provider signs in users of these domains only.
Group attribute and Group mapping	Map an external attribute value to an internal group ID, for example {"Engineering": "<group-id>"}. Only the listed values grant a membership.
SCIM provisioning	Lets your identity provider create, update, and deactivate the user accounts of this tenant, and sync groups. You enter the SCIM base URL and the generated bearer token in your identity provider.
SAML 2.0 SSO	Sets up the sign-in via a SAML 2.0 identity provider. The IdP entity ID , IdP SSO URL , IdP signing certificate (PEM) , NameID format , and Allowed email domains entries describe the identity provider. You register the ACS URL there.

The following options are available in the tab:

Option	Description
Test & save SSO button	Checks the configuration and saves it. After the saving, the dialog reports SSO configuration saved .
Remove SSO button	Removes the OIDC configuration. After that, the dialog reports SSO configuration removed .
Generate token button	Generates a SCIM bearer token. The token appears once. Copy it immediately. For an existing token, the button carries the Rotate token label.
Remove SCIM button	Removes the SCIM credentials.
Save SAML config button	Saves the SAML configuration. After the saving, the dialog reports SAML configuration saved .
Remove SAML config button	Removes the SAML configuration.



If the tenant was changed by someone else in the meantime, then the dialog reports **This tenant was changed by someone else. Reload and try again.** The **Save Changes** button applies the changes of all tabs. During the saving it carries the **Saving...** label.

4.5.4.3 Deleting a tenant

The following requirements are met:

- ☒ Your user account has the **Purge tenants** permission.

Proceed as follows to delete a tenant:




Figure 266: Tenants view

- ▶ In the **User Management** area, open the **Tenants** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **User Management**.
- ▶ Open the tenant.
 - ↳ The detail view of the tenant opens.




Figure 267: Detail view of a tenant

- Click on the **Delete Tenant** button.
 - ↳ The **Delete tenant ""? This will also delete all their gateways, keys, and tokens.** query opens.




Figure 268: Delete tenant query

- ▶ Enter the identifier of the tenant in the field and confirm the query.
- Myra AI Workspace deletes the tenant with its gateways, user accounts, and data.



Myra AI Workspace does not undo the permanent deletion of a tenant. The query enables the acceptance only when the identifier is in the field exactly.

4.5.4.3.1 Effects of the deletion

The deletion has the following effects:

Object	Effect
Tenant	Myra AI Workspace deletes the tenant. During the deletion, the button carries the Deleting... label.
Gateways	All gateways of the tenant are removed. Their endpoint URLs are no longer reachable.
Provider keys	All stored keys of the tenant are removed.

Object	Effect
Authentication tokens	All tokens of the gateways lose their validity. Requests with these tokens fail.



Export the data of the tenant via the **Export data** button before you delete it.



These tasks need the **Create tenants** and **Purge tenants** permissions. They stay reserved for the **Administrator** role. A role of your own does not take them.

4.6 Billing

Billing determines how Myra AI Workspace charges the costs of a tenant. The following configurations maintain the billing settings, the plans, the prices of the models, and the vouchers.

The following sections describe these configurations:

- [Managing the billing](#)
- [Managing model prices](#)
- [Managing plans](#)
- [Managing vouchers](#)

4.6.1 Managing the billing

The billing settings set how Myra AI Workspace bills the cost of the tenant.

The following requirements are met:

- ☑ Your user account has the **Manage billing settings** permission.

Proceed as follows to edit the billing settings:




Figure 269: User Costs view

- ▶ In the **Settings** area, open the **Costs** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **User Costs** tab.
- ▶ In the **Billing settings** section, change the entries.
- ▶ Save the settings.
- Myra AI Workspace shows **Saved**.



If the save operation fails, the view shows **Could not save. Please try again.**

4.6.1.1 Entries in the billing settings

The following entry is changed in the section:

Entry	Effect of the change
Invoice recipient	Changes the email address Myra AI Workspace sends the invoices of the tenant to, for example billing@example.com . The field carries the hint The email address invoices are addressed to. Leave empty to clear. If you empty the field, Myra AI Workspace removes the recipient. The tenant then receives no invoice at a stored address any more.

The following entry cannot be changed in this view:

Entry	Description
Plan	Plan of the tenant. The plan determines the prices and the limits of the tenant. As long as no plan is assigned, the entry shows —. The operator assigns the plan in the Tenants view.

The following option is available:

Option	Description
Save button	Applies the invoice recipient. While saving, the button carries the label Saving... , afterwards the view reports Saved .

4.6.1.2 Effects of the changes

The change has the following effects:

Object	Effect
Invoices	Myra AI Workspace sends the following invoices to the new address. Invoices already sent do not change.
Costs	The change affects neither the recorded costs nor the plan. The Costs view reports the same amounts as before.
Tenant	The setting applies to the whole tenant, not to a single user account.



If the billing settings cannot be loaded, the section shows **Could not load billing settings**. If the address entered is not a valid email address, the view shows **Enter a valid email address**.

4.6.2 Managing model prices

A model price gives a model of a provider its cost. Myra AI Workspace uses it for the cost allocation.

4.6.2.1 Creating a model price

The following requirements are met:

- ☑ Your user account has the **Manage model prices and deprecations** permission.

Proceed as follows to create a model price:




The screenshot shows the 'Model Prices' section under the 'Provider Costs' tab. It displays a table with 3 entries for cost attribution. The table columns are: PROVIDER, MODEL, DISPLAY NAME, HOSTED IN, INPUT \$/1K TOKENS, OUTPUT \$/1K TOKENS, CACHE READ \$/1K, and UPDATED. The entries are for anthropic (Claude Sonnet 4.5), openai (GPT-4o mini), and mistral (Mistral Large). There are buttons for 'Sync from Providers' and 'New Price'.

PROVIDER	MODEL	DISPLAY NAME	HOSTED IN	INPUT \$/1K TOKENS	OUTPUT \$/1K TOKENS	CACHE READ \$/1K	UPDATED
anthropic	claude-sonnet-4-5	Claude Sonnet 4.5	EU	0.003	0.015	0.0003	8/12/2026
openai	gpt-4o-mini	GPT-4o mini	US	0.00015	0.0006	0.000075	8/12/2026
mistral	mistral-large	Mistral Large	EU	0.002	0.006	–	7/30/2026

Figure 270: Model Prices view

- In the **Settings** area, open the **Costs** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- Open the **Provider Costs** tab.
- Click on the **New Price** button.
 - ↳ The **New Model Price** dialog opens.



The image shows a 'New Model Price' dialog box. It contains the following fields: 'Provider' (a dropdown menu with 'anthropic' selected), 'Model *' (a text box with 'claude-sonnet-4-6'), 'Input \$/1K tokens *' (a text box with '0.003'), 'Output \$/1K tokens *' (a text box with '0.015'), 'Cache Write 5m \$/1K (optional)' (a text box with a dash), 'Cache Write 1h \$/1K (optional)' (a text box with a dash), and 'Cache Read \$/1K (optional)' (a text box with a dash). Below these fields is a note: 'Prices are per 1,000 tokens. Providers usually publish per 1,000,000 – divide their figure by 1,000.' At the bottom right are 'Cancel' and 'Add Price' buttons.

Figure 271: New Model Price dialog

- ▶ Under **Provider**, select the provider and, under **Model**, enter the identifier of the model.
- ▶ Enter the prices.
- The prices are per 1,000 tokens.
- Providers usually publish their prices per 1,000,000 tokens.
- ▶ Click on the **Add Price** button.
- Myra AI Workspace creates the entry and uses it for the cost allocation that follows.

4.6.2.1.1 Effects of the new model price

The new model price has the following effects:

Object	Effect
Cost attribution	Myra AI Workspace rates the following requests to this model with the prices entered. Costs already recorded do not change.

Costs view

Object	Effect
	The view reports the costs of this model per tenant, user account, and gateway.
Provider Costs tab	The entry appears in the list. The subtitle reports the count increased by the entry.



The prices are per 1,000 tokens. Providers usually publish their prices per 1,000,000 tokens. Divide their figure by 1,000. For the **linkup/web_search** row the price is per query. The **Input \$/1K tokens** field carries the price per 1,000 queries, and the **Output \$/1K tokens** field stays at **0**.

4.6.2.2 Editing a model price

Proceed as follows to edit a model price:




Figure 272: Model Prices view

- ▶ In the **Settings** area, open the **Costs** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Provider Costs** tab.
- ▶ In the row of the entry, select **Edit**.
 - ↳ The **Edit: /** dialog opens.

Edit: anthropic/claude-sonnet-4-5

Provider

anthropic

Model *

claude-sonnet-4-5

Input \$/1K tokens *

0.003

Output \$/1K tokens *

0.015

Cache Write 5m \$/1K (optional)

0.00375

Cache Write 1h \$/1K (optional)

0.006

Cache Read \$/1K (optional)

0.0003

Prices are per 1,000 tokens. Providers usually publish per 1,000,000 – divide their figure by 1,000.

Cancel

Save Changes

Figure 273: Edit dialog

- ▶ Change the prices.
- ▶ Click on the **Save Changes** button.
- Myra AI Workspace uses the new prices for the cost allocation that follows.

4.6.2.2.1 Effects of the changed prices

The change has the following effects:

Object	Effect
Cost attribution	Myra AI Workspace rates the following requests with the new prices. Costs already recorded keep the prices that applied at the time of the request.
Updated column	The column carries the time of the change.
Budget and limits	A higher price uses up the budget of a tenant faster, a lower price slower.



The **Provider** and **Model** entries are locked in the **Edit** dialog. To assign an entry to a different model, delete it and create it again.

4.6.2.3 Deleting a model price

Proceed as follows to delete a model price:



The screenshot displays the 'Model Prices' section under the 'Provider Costs' tab. It shows a table with 3 entries used for cost attribution. The table columns are: PROVIDER, MODEL, DISPLAY NAME, HOSTED IN, INPUT \$/1K TOKENS, OUTPUT \$/1K TOKENS, CACHE READ \$/1K, and UPDATED. Each row includes edit and delete icons.

PROVIDER	MODEL	DISPLAY NAME	HOSTED IN	INPUT \$/1K TOKENS	OUTPUT \$/1K TOKENS	CACHE READ \$/1K	UPDATED
anthropic	claude-sonnet-4-5	Claude Sonnet 4.5	EU	0.003	0.015	0.0003	8/12/2026
openai	gpt-4o-mini	GPT-4o mini	US	0.00015	0.0006	0.000075	8/12/2026
mistral	mistral-large	Mistral Large	EU	0.002	0.006	–	7/30/2026

Figure 274: Model Prices view

- ▶ In the **Settings** area, open the **Costs** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Provider Costs** tab.
- ▶ In the row of the entry, select **Delete**.
 - ↳ The **Delete pricing for /?** query opens.




Figure 275: Delete pricing query

- ▶ Confirm the query.
- Myra AI Workspace removes the entry.

4.6.2.3.1 Effects of deleting a model price

The deletion has the following effects:

Object	Effect
Model price	Myra AI Workspace removes the entry from the Provider Costs tab. No price is available for this model any more.
Cost attribution	The following requests to this model no longer get costs attributed. Costs already recorded remain.
Model	The model itself remains usable. The model price controls the cost attribution only.



The **Sync from Providers** button creates a deleted entry again as soon as the provider still lists the model.



These tasks need the **Manage model prices and deprecations** permission of the **Platform** module. It stays reserved for the **Administrator** role. A role of your own does not take it.

4.6.3 Managing plans

A self-serve plan sets the models, the budget, the caps, the degrade configuration, and the Stripe price ids of the self-registration. A plan can be neither created nor deleted.

The following requirements are met:

- ☑ Your user account has the **Manage plan and entitlement config** permission.

Proceed as follows to edit a plan:




Figure 276: Self-serve Plans view

- In the **Account** area, open the **Plans** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- In the area of the plan, select **Edit**.
 - ↳ The **Edit plan:** dialog opens.




Figure 277: Edit plan dialog

- Change the budget, the caps, the models, the degrade configuration, or the Stripe price ids.
- Click on the **Save Changes** button.
- Myra AI Workspace applies the change.



Myra AI Workspace permits only one plan edit at a time. If one is already running, the dialog shows **Another plan edit is in progress. Please retry in a moment.**

4.6.3.1 Effects of the plan change

The change has the following effects:

Object	Effect
Tenants on this plan	The change applies to all tenants that carry the plan, not to a single tenant.

Object	Effect
Models	A model removed is no longer available to the tenants on this plan. A model added is available to them in addition.
Budget	The new budget applies to the following billing period of the tenants on this plan.
Degrade configuration	If you empty the degrade multiplier, Myra AI Workspace no longer degrades above the budget. If you set it, the degrade provider with the degrade model answers the requests above the budget.
Self-service signup	The Stripe price ids determine which amount Stripe charges for the plan.
Last updated entry	The area of the plan carries the time of the change.



This task needs the **Manage plan and entitlement config** permission of the **Platform** module. It stays reserved for the **Administrator** role. A role of your own does not take it.

4.6.4 Managing vouchers

A voucher gives a new registration a one-time trial credit.

4.6.4.1 Creating a voucher

The following requirements are met:

- ☒ Your user account has the **Manage vouchers** permission.

Proceed as follows to create a voucher:



Vouchers						New voucher
CODE	TYPE	CREDIT	REDEEMED	EXPIRES	ACTIVE	ACTIONS
TEST-2026-....	Unique (single use)	\$25	1 / 1	12/31/2026	Yes	Deactivate
MESSE-....	Generic (campaign)	\$10	42 / 100	10/15/2026	Yes	Deactivate
ALT-....	Generic (campaign)	\$10	8 / 50	Never	No	Reactivate

Figure 278: Vouchers view

- In the **Account** area, open the **Vouchers** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- Click on the **New voucher** button.
 - ↳ The **New voucher** dialog opens.

New voucher

Code

MESSE-2026

Type

Generic (campaign) ▾

Credit (USD)

10

Max redemptions

100

Expires (optional)

Unique codes are single-use. Generic codes are capped by max redemptions. Credit is a one-time trial bonus applied to the new tenant first billing period.

Cancel

Create voucher

Figure 279: New voucher dialog

- ▶ Under **Code**, enter the code.
 - ▶ Under **Type**, select the type.
 - ▶ Under **Credit (USD)**, enter the credit and, if necessary, set an expiry date.
 - ▶ Click on the **Create voucher** button.
- Myra AI Workspace creates the voucher. The credit applies once, to the first billing period of the new tenant.

4.6.4.1.1 Effects of the new voucher

The new voucher has the following effects:

Object	Effect
Self-service signup	A new signup redeems the code and gets the credit. A unique code is redeemable once, a generic code up to the number under Max redemptions .
Billing	Myra AI Workspace applies the credit once to the first billing period of the new tenant.

Object	Effect
Vouchers view	The voucher appears in the list. The Redeemed column shows 0 and the Active column shows Yes .



An expiry date applies until the end of the day selected. If you leave the **Expires (optional)** field empty, the **Expires** column carries the value **Never**.

4.6.4.2 Deactivating a voucher

A voucher cannot be deleted. You deactivate it and reactivate it if necessary.

Proceed as follows to deactivate a voucher:

Vouchers							New voucher
CODE	TYPE	CREDIT	REDEEMED	EXPIRES	ACTIVE	ACTIONS	
TEST-2026-....	Unique (single use)	\$25	1 / 1	12/31/2026	Yes	Deactivate	
MESSE-....	Generic (campaign)	\$10	42 / 100	10/15/2026	Yes	Deactivate	
ALT-....	Generic (campaign)	\$10	8 / 50	Never	No	Reactivate	

Figure 280: Vouchers view

- In the **Account** area, open the **Vouchers** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- In the row of the voucher, select **Deactivate**.
 - ↳ The **Deactivate voucher ? It can no longer be redeemed.** query opens.




Figure 281: Deactivate voucher query

- ▶ Confirm the query.
- The voucher can no longer be redeemed. The **Reactivate** row action undoes the step.

4.6.4.2.1 Effects of the deactivation

The deactivation has the following effects:

Object	Effect
Voucher	The Active column shows No . The code can no longer be redeemed.
Redemptions so far	Credits already redeemed remain. The Redeemed column does not change.
Vouchers view	The voucher remains in the list. The Reactivate row action takes the place of the Deactivate row action.



These tasks need the **Manage vouchers** permission of the **Platform** module. It stays reserved for the **Administrator** role. A role of your own does not take it.

5. FAQ & Error messages

5.1 Sign-in

This section names the messages that appear during the sign-in, their cause, and their remedy.

5.1.1 Messages

The following sections give the cause and the remedy for these messages:

- That code is incorrect or has expired
- Too many attempts
- Could not reach the server

5.1.2 That code is incorrect or has expired

The sign-in shows **That code is incorrect or has expired. Please request a new one.**

5.1.2.1 Cause

The one-time code entered does not match the one sent, or its validity has expired.

5.1.2.2 Solution

- ▶ Click on the **Didn't get the code? Resend** link.
 - ↳ The link becomes available again 30 seconds after the last send. Until then it shows the remaining time.
- ▶ Enter the new code from the email and click on the **Sign in** button.
- Myra AI Workspace starts with the navigation bar of your role.

5.1.3 Too many attempts

The sign-in shows **Too many attempts. Please request a new code.**

5.1.3.1 Cause

Myra AI Workspace has used up the number of the attempts for this code and does not accept it any more.

5.1.3.2 Solution

- ▶ Click on the **Didn't get the code? Resend** link.
- ▶ Enter the new code.
- The sign-in accepts the new code.

5.1.4 Could not reach the server

The sign-in shows **Could not reach the server. Check your connection and try again.**

5.1.4.1 Cause

The request did not reach the server.

5.1.4.2 Solution

- ▶ Check your network connection.
- ▶ Reload the page and sign in again.
- The sign-in appears again.

5.1.5 Frequently asked questions

This section answers the questions most frequently asked about the sign-in.

5.1.5.1 How long is a one-time code valid?

A one-time code has six digits and is valid for 15 minutes. After that, the sign-in refuses it with the **That code is incorrect or has expired. Please request a new one.** message. Request a new code with the **Didn't get the code? Resend** link. See [Getting access](#).

5.1.5.2 How often can I enter a code?

Myra AI Workspace accepts five attempts for each code. After that, the sign-in shows **Too many attempts. Please request a new code.** and no longer accepts this code. A new code sets the count back. See [Getting access](#).

5.1.5.3 Why do I get no email with the code?

The sign-in always confirms the send with **If that email is registered, a code has been sent.** and thus does not disclose whether a user account for an address exists. Myra AI Workspace sends an email only to an address that has a user account. If the email does not arrive, check the spelling of the address and the spam folder. The **Didn't get the code? Resend** link becomes available again 30 seconds after the send. See [Getting access](#).

5.1.5.4 How long do I stay signed in?

A session is valid for 8 hours. The **Stay logged in for 30 days on this device** checkbox extends it to 30 days. The checkbox applies only to the device on which you select it. See [Getting access](#).

5.1.5.5 Can I sign in with a password?

No, Myra AI Workspace uses no password. The sign-in uses a one-time code sent by email. If your tenant has an identity provider, the **Continue with SSO** button also appears. See [Getting access](#).

5.2 Gateway

This section names the messages that the gateway shows for a request, their cause, and their remedy.

5.2.1 Messages

The following sections give the cause and the remedy for these messages:

- Rate limit exceeded
- The usage budget of this workspace is used up
- Request blocked by content policy
- The prompt is too long for this model
- The request took too long to complete
- The requested model is not available on this gateway
- This gateway requires EU model hosting
- The model provider returned an error

5.2.2 Rate limit exceeded

The gateway shows **Rate limit exceeded – please slow down.**

5.2.2.1 Cause

The number of the requests went over the rate limit of the token or of the gateway in the current time window.

5.2.2.2 Solution

- ▶ Wait for the end of the current time window and send the request again.
- ▶ Contact your administration if the limit is permanently too low.
- The gateway answers the requests again.

5.2.3 The usage budget of this workspace is used up

The gateway shows **You've used up this workspace's usage budget. It renews at the start of the next cycle.** The cost of the workspace reached the budget set.

5.2.3.1 Cause

Myra AI Workspace blocks new requests.

5.2.3.2 Solution

- ▶ Wait for the start of the next billing cycle.
- ▶ Contact your administration if the budget must be increased before that.
- After the renewal or the increase, the gateway accepts the requests again.

5.2.4 Request blocked by content policy

The gateway shows **Request blocked by content policy**. A guardrail of the gateway rejected the input or the output.

5.2.4.1 Cause

The request does not reach the provider.

5.2.4.2 Solution

- ▶ Write the request without the content objected to.
- ▶ Contact your administration if the rule rejects the content incorrectly.
- The gateway sends the request to the provider.

5.2.5 The prompt is too long for this model

The gateway shows **The prompt is too long for this model. Start a new chat or shorten the input.**

5.2.5.1 Cause

Together with the new input, the conversation goes over the context window of the selected model.

5.2.5.2 Solution

- ▶ Shorten the input or remove an attachment.
- ▶ Start a new conversation.
- ▶ Select a model with a larger context window.
- The model accepts the request.

5.2.6 The request took too long to complete

The gateway shows **The request took too long to complete. Try a simpler request, narrow the scope, or turn off web search.**

5.2.6.1 Cause

The model did not answer within the time limit.

5.2.6.2 Solution

- ▶ Narrow the scope of the request.
- ▶ Turn off the web search for this chat.
- ▶ Send the request again.
- The model answers within the time limit.

5.2.7 The requested model is not available on this gateway

The gateway shows **The requested model is not available on this gateway. Check the model id, or pick a model this gateway serves.**

5.2.7.1 Cause

The selected model is not enabled on the gateway, or the model ID is spelled incorrectly.

5.2.7.2 Solution




Figure 282: Chat

- ▶ Open the model selection in the chat.
- ▶ Select a model from the list.
- ▶ Contact your administration if the model you want is missing.
- The gateway sends the request to the provider.

5.2.8 This gateway requires EU model hosting

The gateway shows **This gateway requires EU model hosting. The selected model routes to a provider that is not EU-hosted.**

5.2.8.1 Cause

A residency requirement applies to the gateway that the selected model does not meet.

5.2.8.2 Solution




Figure 283: Chat

- ▶ Open the model selection in the chat.
- ▶ Select a model of an EU provider – Myra, Mistral, or an EU-region adapter.
- The gateway sends the request to the provider.

5.2.9 The model provider returned an error

The gateway shows **The model provider returned an error. Please try again.** The provider accepted the request and answered it with an error.

5.2.9.1 Cause

Related messages are **The model provider is rate-limiting requests.**, **The model provider is temporarily unavailable.**, and **All configured providers failed.**

5.2.9.2 Solution

- ▶ Wait a moment and send the request again.
- ▶ Select a model of a different provider.
- ▶ Contact your administration if the error stays.
- The provider answers the request.

5.2.10 Frequently asked questions

This section answers the questions most frequently asked about the gateway.

5.2.10.1 Which budget is used up?

The **The usage budget of this workspace is used up.** message does not name the scope. The notice banner above the **Dashboard** view names it:

- **Tenant**
- **Gateway**
- **Token**
- **Plan allowance**
- **Rate limit**

A budget renews at the start of the next billing cycle. See [Dashboard](#).

5.2.10.2 How high is my rate limit?

The rate limit applies to each token. The **My Account** view lists it in the **Tokens** section in the **Rate Limit** column as requests for each time window. A gateway carries a rate limit of its own in addition, and the two apply independently of each other. The **Rate limit exceeded – please slow down.** message does not name which of the two applied. See [Managing tokens](#).

5.2.10.3 Why are other models available to me than to other users?

The gateway through which your request goes decides the available models, not your user account. For a model that is not enabled on the gateway, the gateway answers with **The requested model is not available on this gateway.** Users who work through a different gateway therefore see a different selection. See [Managing gateways](#).

5.2.10.4 What does EU model hosting mean?

A gateway can require that requests only go to providers that host in the EU. For a model outside this requirement, the gateway answers with **This gateway requires EU model hosting. The selected model routes to a provider that is not EU-hosted.** The Myra and Mistral providers and the EU region adapters meet the requirement. See [Managing gateways](#).

5.3 Gateway token

This section names the messages that concern a gateway token, their cause, and their remedy.

5.3.1 Messages

The following sections give the cause and the remedy for these messages:

- Access denied for this gateway
- Missing or invalid gateway token


5.3.2 Access denied for this gateway

The gateway shows **Access denied for this gateway**.

5.3.2.1 Cause

Your user account or your token has no access to this gateway.

5.3.2.2 Solution



My Account
Your profile, sign-in, tokens, and data export.

Profile

Email: alex.muster@beispiel.de Role: member

Tokens New Token

LABEL	GATEWAY	EXPIRES	BUDGET	RATE LIMIT	CREATED	
Editorial script	redaktion	12/31/2026	\$12.40 / \$50.00 • \$37.60 left	60/60 s	4/2/2026	✖
Analysis	produktion	never	\$10.00 • usage unknown	10/60 s	6/11/2026	✖

Your data

Download a machine-readable copy of your account data (conversations, memories, agents, projects, usage statistics and more) under GDPR Article 15.

[Export my data](#)

Danger Zone

Deleting your account signs you out and locks you out of MYRA AI. Your data is retained per our privacy policy; only an administrator can restore the account.

[Delete Account](#)

Permanently erase your account and all associated data (GDPR Article 17, right to erasure). This is irreversible and cannot be undone by you or an administrator.

[Erase all my data](#)

Figure 284: My Account view

- ▶ In the **Tokens** section of the **My Account** view, check which gateway your token applies to.
 - ▶ In the same row, check the **Rate Limit** column.
 - ▶ Contact your administration if your user account must be assigned to the gateway.
- After the assignment, the gateway accepts your requests.

5.3.3 Missing or invalid gateway token

The gateway shows **Missing or invalid gateway token**.

5.3.3.1 Cause

The request carries no token, or the token used is revoked or expired.

5.3.3.2 Solution




Figure 285: My Account view

- In the **Account** area, open the **Profile** group.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Account**.
- In the **Tokens** section, check the **Expires** column.
- Create a new token with **New Token** and enter it in the calling application.
- The gateway accepts the requests again.

5.3.4 Frequently asked questions

This section answers the questions most frequently asked about gateway tokens.

5.3.4.1 Where do I find my tokens?

In the **Account** area in the **Profile** group. The **My Account** view lists them in the **Tokens** section with the **Label**, **Gateway**, **Expires**, **Budget**, **Rate Limit**, and **Created** columns. See [Managing tokens](#).

5.3.4.2 Does a token apply to more than one gateway?

No, each token belongs to exactly one gateway. For a second gateway, create a second token. For a request to a gateway to which the token does not apply, Myra AI Workspace answers with **Access denied for this gateway**. See [Managing tokens](#).

5.3.4.3 Does a token expire?

A token expires only if you make an **Expires At** entry in the **New Token** dialog. The field carries the **Leave blank = never expires** note. Without an entry, the **Expires** column shows **never**. See [Managing tokens](#).

5.3.4.4 What do I do if a token gets into the wrong hands?

Revoke it in its row and confirm the **Revoke this token?** query. After that, the gateway answers each request with this token with **Missing or invalid gateway token**. Then create a new token with **New Token** and enter it in the calling application. See [Managing tokens](#).

5.4 Your own objects

This section names the messages that appear when you edit your own agents, workflows, and MCP connectors, their cause, and their remedy.

5.4.1 Messages

The following sections give the cause and the remedy for these messages:

- The agent could not be saved
- The workflow could not be created
- Changed in another tab
- Duplicate step
- Save the connector first, then test
- No gateways available

5.4.2 The agent could not be saved

The **New Agent** or **Edit Agent** dialog shows **The agent could not be saved**.

5.4.2.1 Cause

A mandatory entry is missing, or the server rejected the change. The dialog shows a missing entry as **Name is required**, **Slug is required**, or **Model is required**.

5.4.2.2 Solution

- ▶ Check the **Name**, **Slug**, and **Model** entries.
- ▶ Check whether **Max tokens** contains a positive integer.
- ▶ Click on the **Save** button.
- Myra AI Workspace saves the agent.

5.4.3 The workflow could not be created

The **New workflow** dialog shows **The workflow could not be created**.

5.4.3.1 Cause

The name is missing, or the server rejected the request.

5.4.3.2 Solution




Figure 286: New workflow dialog

- ▶ Enter a name in the **Name** text field.
- ▶ Click on the **Create** button.
- Myra AI Workspace creates the workflow and opens the editor.



Figure 287: Editor of a workflow

5.4.4 Changed in another tab

The editor of a workflow shows **Changed in another tab**.

5.4.4.1 Cause

The same workflow was edited and saved in another browser tab or by another person.

5.4.4.2 Solution

- ▶ Click on the **Reload** button.
 - ↳ The editor loads the current state of the workflow.
- ▶ Repeat your change.
- The editor shows **Saved**.

5.4.5 Duplicate step

The editor of a workflow shows **Duplicate step**. Two steps carry the same key.

5.4.5.1 Cause

Myra AI Workspace does not save the workflow.

5.4.5.2 Solution

- ▶ Rename one of the two steps.
- The editor saves the workflow.

5.4.6 Save the connector first, then test

The **New MCP Connector** dialog shows **Save the connector first, then test**.

5.4.6.1 Cause

The test needs a saved connector.

5.4.6.2 Solution

- ▶ Save the connector.
- ▶ Start the test again.
- The dialog shows the result of the test.


5.4.7 No gateways available

The **New Token** dialog shows **No gateways available**.

5.4.7.1 Cause

A token always applies to one gateway. No gateway exists for your tenant.

5.4.7.2 Solution



The image shows a 'New Token' dialog box with the following fields and options:

- Gateway ***: A dropdown menu showing 'muster-gmbh/produktion'.
- Label**: A text input field containing 'dev laptop'.
- Expires At**: A date input field showing 'mm/dd/yyyy --:--'. Below it, a note says 'Leave blank = never expires'.
- Budget (USD)**: A text input field containing 'unlimited'.
- Rate limit (req/window)**: Two buttons, 'none' and '60s'. Below them, a note says 'requests / window_sec (blank = gateway default)'.
- At the bottom, there are two buttons: 'Cancel' and 'Create Token'.

Figure 288: New Token dialog

- ▶ Contact your administration so that it creates a gateway.
- ▶ Open the **New Token** dialog again.
- The gateway is available for selection.

5.4.8 Frequently asked questions

This section answers the questions most frequently asked about your own agents, workflows, and MCP connectors.

5.4.8.1 Why does the editor show Changed in another tab?

When it saves, the editor of a workflow checks whether the state is still the one that it loaded. If the same workflow was saved in another browser tab or by another person in the meantime, the editor does not save but asks for the current state with **Reload**. Thus neither of the two edits overwrites the other unnoticed. See [Managing workflows](#).

5.4.8.2 Why can an MCP connector only be tested after it is saved?

The test calls the connector through the server, not from the browser. While the connector is not saved, the server does not know its address and its credentials. The dialog shows this with **Save the connector first, then test**. See [Managing MCP connectors](#).

5.4.8.3 Who can edit my agents, workflows, and connectors?

You and the administration. A permission for your own objects applies to the objects of your user account only. A role of the administration goes beyond that. An agent that you share can be started by the recipients but not edited. See [Managing agents](#).

5.5 Administration

This section names the messages that appear in the administration, their cause, and their remedy.

5.5.1 Messages

The following sections give the cause and the remedy for these messages:

- [Internal gateway error](#)
- [Gateway configuration error](#)
- [No API key is configured for the selected provider](#)
- [The gateway couldn't authenticate with the model provider](#)
- [You do not have the permission](#)
- [Another plan edit is in progress](#)

5.5.2 Internal gateway error

The gateway shows **Internal gateway error – the request never reached the model.**

5.5.2.1 Cause

The request failed in the gateway before it went to a provider.

5.5.2.2 Solution




Figure 289: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group and check the provider configuration of the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Open the **Request Logs** group and check the trace of the request.
- ▶ Remove the cause reported and send the request again.
- The gateway sends the request to the provider.


5.5.3 Gateway configuration error

The gateway shows **Gateway configuration error – check the gateway's provider settings**.

5.5.3.1 Cause

The provider settings of the gateway are incomplete or contradictory.

5.5.3.2 Solution



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	–	0s	–	–	7/19/2026

Figure 290: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group and open the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- ▶ Check the providers, the enabled models, and the URL overrides.
- ▶ Click on the **Save changes** button.
- The gateway accepts the requests again.


5.5.4 No API key is configured for the selected provider

The gateway shows **No API key is configured for the selected provider on this gateway.**

5.5.4.1 Cause

No key for the provider that routes the selected model exists on the gateway.

5.5.4.2 Solution



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED	
produktion	required	€460.00	300s	60/60s	3	5/14/2026	 
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026	 
labor	open	–	0s	–	–	7/19/2026	 

Figure 291: Gateways view

- ▶ In the **Settings** area, open the **Gateways** group and open the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
 - ▶ Click on the **Add Model** button.
 - ▶ Select **Use My Own Key** and enter the API key, or select **Use Platform Key**.
 - ▶ Click on the **Enable** button at the model.
- Myra AI Workspace shows **Key stored and encrypted successfully**.




Figure 292: Detail view of a gateway


5.5.5 The gateway couldn't authenticate with the model provider

The gateway shows **The gateway couldn't authenticate with the model provider – a provider key is missing or invalid.**

5.5.5.1 Cause

The key stored is withdrawn, expired, or entered incorrectly.

5.5.5.2 Solution



Gateways

Tenant: muster-gmbh

Search gateways by name...

New Gateway

Select Tenant

Filter tenants by name...

muster-gmbh beispiel-ag probe-kg

NAME	AUTH	BUDGET	CACHE	RATE LIMIT	DETECTORS	CREATED
produktion	required	€460.00	300s	60/60s	3	5/14/2026
redaktion	required	€110.40	3,600s	30/60s	2	6/2/2026
labor	open	-	0s	-	-	7/19/2026

Figure 293: Gateways view

- In the **Settings** area, open the **Gateways** group and open the gateway.
 - ▷ On the start page, click on your name at the bottom of the navigation bar.
 - ▷ In the user menu, click on **Settings**.
- Enter the valid API key of the provider.
- In the **Providers** group, check the availability status of the provider.
- The gateway reaches the provider again.




Figure 294: Providers view

5.5.6 You do not have the permission

The editor of the **Roles & permissions** view shows a permission as not assignable.

5.5.6.1 Cause

A role gives only the permissions that the user account editing it has itself.

5.5.6.2 Solution

- ▶ Let a user account that has the permission give it.
- ▶ Click on the **Save changes** button.
- The role carries the permission.

5.5.7 Another plan edit is in progress

The **Edit plan** dialog shows **Another plan edit is in progress. Please retry in a moment.**

5.5.7.1 Cause

Myra AI Workspace permits only one plan edit at a time.

5.5.7.2 Solution

- ▶ Wait a moment.
- ▶ Click on the **Save Changes** button again.
- Myra AI Workspace applies the change.

5.5.8 Frequently asked questions

This section answers the questions most frequently asked about the administration.

5.5.8.1 What is a gateway configuration error?

The message appears when a gateway is available but its settings permit no request – for example because no provider is enabled, no model is assigned, or a routing rule points to a model that the provider does not carry. The **Gateways** view shows the installation concerned. The configuration is corrected there. See [Managing gateways](#).

5.5.8.2 Where do I store the API key of a provider?

You store the API key with the provider of the gateway. If it is missing, the gateway answers each request to this provider with **No API key is configured for the selected provider**. If the provider refuses the key stored, the gateway shows **The gateway couldn't authenticate with the model provider**. See [Providers](#).

5.5.8.3 Why does Myra AI Workspace show that I do not have a permission?

Each action of the administration is bound to a permission of the catalogue, and the roles carry different parts of it. The **Roles & permissions** view shows which permissions a role carries. The **Users** view shows the role of a user account. See [Roles & permissions](#).

5.5.8.4 Why can a plan not be edited at the moment?

Myra AI Workspace permits one edit for each plan at the same time and shows a second one with **Another plan edit is in progress**. Wait until the current edit is saved or cancelled, and do the procedure again. See [Managing plans](#).

5.6 Interface

This section names the message that the interface itself shows, its cause, and its remedy.

5.6.1 Something went wrong

The interface shows **Something went wrong** or **The app couldn't finish loading**.
Please reload to try again.

5.6.1.1 Cause

A part of the interface did not load.

5.6.1.2 Solution

- ▶ Click on the **Try again** button.
- ▶ If the message stays, reload the page.
- The interface appears completely.

5.6.2 Frequently asked questions

This section answers the questions most frequently asked about the interface.

5.6.2.1 The message stays after a reload – what now?

Sign out and sign in again. If the message stays after that too, speak to your administration and name the view in which it appears and the time. Myra AI Workspace logs the errors of the interface and uses that to find them. See [Home page and GUI layout](#).

5.6.2.2 How do I report an unsuitable response of the AI?

With the **Report** action on the response. The **Report this response** dialog asks for the type of the complaint in the **Reason** selection field and accepts a description in the **Notes (optional)** text field. After the **Submit report** button, Myra AI Workspace confirms with **Report submitted**. See [Report this response dialog](#).

5.6.2.3 Why do I not see a view that this manual describes?

The navigation bar shows only the views for which your role has the permission. The **Product description** section names the functions of your role. For a view that you need, speak to your administration. See [Product description](#).

6. Reference

6.1 Notational conventions

The pages of this documentation use the following notational conventions throughout.

Formatting	Description
Text in bold	Identifies screen elements of the interface, such as views, tabs, buttons, fields, and messages.
Links	Identifies references to other sections of this documentation and to external sources.
Text in monospace	Identifies commands, code, endpoints, addresses, and input.
<ABC>	Identifies a placeholder that the application fills with the respective value, such as the name of the selected object.
☑	Identifies a prerequisite that must be fulfilled before the first step.
▶	Identifies a step that you perform.
▷	Identifies a subordinate step.
↳	Identifies the interim result after one or more steps.
→	Identifies the result that is available after the last step.
■	Identifies an item of a list, such as the sections that an overview summarizes.
	Identifies information that supports the procedure. It appears as a separate box or, within a table cell, as a short form in the cell.
	Identifies information that, if not observed, leads to incorrect settings, additional cost, or data loss.