QDT QUANTUM DATA TECHNOLOGIES

QDT PROPRIETARY MARKETING INTELLIGENCE

Q.Mind Developer API

Build canonical marketing intelligence into products, workflows, and decision systems.

API v1

Q.Mind 0.22.33 · October 17, 2026

PATENT PENDING

01 · QUICK START

Quick start

One asynchronous interface for counsel, analysis, and files.

Production base URL

```
https://qmind.qdt.ai/api/v1
```

Create a key in Q.Mind !’ Account !’ Developer API. Send a question to POST /responses, then poll the Location URL until the response reaches a terminal status.

```
export QMIND_API_KEY="qmind_live_…"

curl https://qmind.qdt.ai/api/v1/responses \
  -H "Authorization: Bearer $QMIND_API_KEY" \
  -H "Content-Type: application/json" \
  -H "Idempotency-Key: northstar-growth-001" \
  -d '{"input":"For fictional Northstar Laundry, should growth prioritize penetration or
premiumization?"}'
```

ASYNC BY DESIGN

Quick counsel and long governed analysis use the same response resource. The status and progress fields tell the caller what Q.Mind is doing without imposing a short HTTP timeout.

02 · AUTHENTICATION

Authentication

Account-scoped keys, separate from Q.Mind operator credentials.

Send the key only as a Bearer token:

```
Authorization: Bearer qmind_live_…
```

- A key belongs to exactly one Q.Mind account.
- The secret is shown once; Q.Mind stores only its SHA-256 digest.
- Keys cannot reach internal admin, worker, raw-status, or memory routes.
- Revocation in the Account screen is immediate.
- Keep keys in a server-side secret manager. Never ship them in browser or mobile code.

03 · RESPONSE ENDPOINTS

Response endpoints

Create, retrieve, cancel, and download.

POST	`/responses`	Create a canonical Q.Mind turn. Returns 202 Accepted and a Location header.
GET	`/responses/{response_id}`	Retrieve status, progress, output, canonical integrity metadata, and artifacts.
POST	`/responses/{response_id}/cancel`	Cancel queued or running work.
GET	`/artifacts/{artifact_id}/content`	Download a response PDF, PowerPoint, creative set, or campaign calendar.
GET	`/models`	List the stable qmind-asimov model alias.
GET	`/capabilities`	List available capability preferences.

03A · CREATE A RESPONSE

Create a response

Request body, routing preference, and continuity fields.

```
{
  "input": "Diagnose fictional Northstar Laundry and create a PDF.",
  "processing_mode": "analysis",
  "capability": "growth-diagnosis",
  "project_id": "optional-uuid",
  "conversation_id": "optional-uuid"
}
```

FIELD	REQUIRED	PURPOSE
input	Yes	Question or work request, up to 12,000 characters.
project_id	No	Attach work to an existing account project.
conversation_id	No	Continue a prior API conversation.
processing_mode	No	auto, conversation, or analysis.
capability	No	Routing preference; defaults to cmo-counsel.

04 · POLLING AND OUTPUT

Polling and output

Observe progress until work becomes terminal.

Poll GET /responses/{id}. Terminal states are completed, failed, blocked, and cancelled; public jobs do not enter waiting_approval. Start near two seconds and back off to roughly ten seconds. Every poll counts toward the per-key rate limit.

```
{
  "id": "9a4d…",
  "object": "qmind.response",
  "status": "completed",
  "model": "qmind-asimov",
  "conversation_id": "b71e…",
  "output_text": "Use penetration as the working priority…",
  "artifacts": [],
  "metadata": {
    "canonical": true,
    "canonical_provider": "openclaw-a2a",
    "canonical_answer_sha256": "…",
    "canonical_integrity_verified": true
  }
}
```

EXACT WRITTEN ANSWER

Canonical Asimov generates the answer once. The bridge signs the exact answer bytes with SHA-256; Q.Mind independently verifies the receipt and fails closed if it is absent or mismatched.

Conversation continuity

Save conversation_id and include it on the next call. Q.Mind retrieves the bounded thread, project facts, prior answers, task ledger, and governed evidence. The client does not need to rebuild prompt context.

```
{
  "input": "What evidence would falsify your first recommendation?",
  "conversation_id": "b71e…"
}
```

05 · LONG ANALYSIS AND FILES

Long analysis and files

The same resource can outlive a three-minute request window.

Set processing_mode to analysis for deterministic governed work, or use auto for intent-based routing. Progress remains visible while Q.Measure, Q.Marketing, authoring, rendering, and delivery stages run.

```
{
  "input": "Diagnose fictional Northstar Pet Nutrition. Return the written answer, a PDF,
and an editable PowerPoint.",
  "capability": "growth-diagnosis",
  "processing_mode": "analysis"
}
```

Completed work lists protected artifact URLs:

```
{
  "id": "31f9…",
  "object": "qmind.artifact",
  "type": "qmind-report-pdf",
  "filename": "Pedigree_Growth_Diagnosis.pdf",
  "content_type": "application/pdf",
  "bytes": 248331,
  "download_url": "/api/v1/artifacts/31f9…/content"
}
```

Use the same Bearer key to download each artifact. File access remains inside the key's account and exposes only customer-facing output kinds referenced by the response.

Types: qmind-report-pdf, qmind-report-pptx, qmind-creative-set, and qmind-campaign-calendar. URLs are stable authenticated API paths, not expiring signed URLs.

06 · IDEMPOTENCY

Idempotency

Safe retries without duplicate analyses.

Send a unique Idempotency-Key on every POST /responses. Retrying the same payload with the same key returns the original response. Reusing it with a different payload returns 409 idempotency_conflict; use a new key for genuinely new work.

07 · ERRORS AND LIMITS

Errors and limits

Predictable envelopes and operational receipts.

```
{
  "error": {
    "type": "invalid_request_error",
    "code": "conversation_not_found",
    "message": "The conversation does not exist in this account.",
    "param": "conversation_id"
  }
}
```

STATUS	MEANING	CLIENT ACTION
400	Invalid request	Correct the named param.
401	Invalid or revoked key	Use an active Bearer key.
403	Missing scope	Use a key with the required scope.
404	Absent or foreign resource	Check the ID and account.
409	Context/state conflict	Correct the context or wait.
413	Request too large	Reduce JSON below 1 MB.
429	60 requests/minute/key	Honor Retry-After.
500	Unexpected failure	Retry with the same idempotency key.

- JSON request limit: 1 MB.
- Input limit: 12,000 characters.
- Every response includes X-Request-Id; retain it in support logs.
- Treat IDs as opaque strings and preserve the returned Location URL.
- Generated answers and files may contain confidential account data.

Keys use responses:write for create/cancel and responses:read for poll/download. Polling counts toward the 60 requests/minute limit. Completion webhooks and a published whole-response maximum runtime are not currently available; callers should set their own deadline and cancel when needed.

08 · REFERENCE LINKS

Reference links

Keep the machine and human references together.

Web documentation https://qmind.qdt.ai/developers/

OpenAPI 3.1 https://qmind.qdt.ai/developers/openapi.json

Q.Mind portal https://qmind.qdt.ai/app/