White paper 3 of 4

Grounded answers across 400+ apps: connectors, permissions and citations

How connectors, per-user access, retrieval and citations produce answers people can check — and what "400+" does and does not mean.

About 6 min read · 1,357 words Print-friendly

CONTENTS

Summary

How connectors, per-user access, retrieval and citations produce answers people can check — and what "400+" does and does not mean.

The grounding problem

A language model on its own knows nothing about your customers, your cases or your contracts. Ask it anyway and it may produce a confident answer with nothing behind it. Grounding is the practice of giving the model the right context from your own systems at the moment of the question — and showing where each part of the answer came from.

Doing that well depends on three things: reaching the systems where the facts live, respecting who is allowed to see what, and returning answers with citations a person can open. This paper walks through each on the AgenTorQ platform.

Four ways to reach your systems

AgenTorQ has one governed connector layer, and the same connectors run in the Cloud workspace and inside Salesforce. Every connector goes through the same governance: per-user OAuth or a vaulted credential, scoped access, approvals for writes and an immutable audit trail.

Route	How it works
SaaS apps	400+ apps behind one framework. Each user connects with their own OAuth 2.0 grant, and tokens are stored and refreshed for you.
REST and GraphQL	Any HTTP/JSON service becomes a connector with a typed schema; credentials live in the vault. In the Salesforce edition, Remote Site Settings are wired automatically.
Databases	SQL Server, Oracle, PostgreSQL, MySQL and Snowflake through the credential vault, with read-first, parameterized SQL. A DML guard holds back writes until they are approved.
MCP, webhooks and custom	Any Model Context Protocol (https://modelcontextprotocol.io/) server with OAuth 2.1 PKCE and automatic tool discovery, signed webhooks, or a custom connector with role-based actions.

What "400+" means

The number deserves a precise definition. The 400+ figure reflects apps reachable through the connector layer, plus REST, GraphQL, databases, MCP servers and webhooks. It is not a claim of 400 native, certified integrations. Flagship apps such as Salesforce, Microsoft 365, Google Workspace, Slack and GitHub are first-class; long-tail and enterprise systems are reachable through the connector layer, REST or MCP.

By category, that includes CRM (Salesforce, Dynamics 365, HubSpot), productivity (Microsoft 365, Google Workspace, Slack, Teams, Zoom), storage (Google Drive, OneDrive, Dropbox, SharePoint, Amazon S3) and developer tools (GitHub, GitLab, Bitbucket, Azure DevOps), with ERP, ITSM and HR systems reached through the connector layer, REST/OData or MCP.

Permissions: the user's access, not a service account's

Grounding must never widen access. With per-user OAuth (the authorization framework described in RFC 6749 (OAuth 2.0) (https://datatracker.ietf.org/doc/html/rfc6749)), each person connects under their own grant, so a Workmate reaches only the systems and records that specific person is entitled to. MCP servers use OAuth 2.1 with PKCE (RFC 7636 (PKCE) (https://datatracker.ietf.org/doc/html/rfc7636)). Databases and HTTP APIs use credentials in an encrypted vault, isolated per organization.

- Tokens are refreshed automatically and never shared or logged.
- Credentials are preserved on reconnect, never re-keyed to the wrong identity.
- In the Salesforce edition, reads pass a SOQL guard in USER_MODE, so field-level security and sharing hold.
- Results from search respect each source's access rules.

Retrieval: a knowledge base you control

Beyond live records, teams build a retrieval layer from their own material. Content is organized into folders, files and chunks with embeddings, and Workmates answer only from approved sources scoped to the user's access. Object-level include and exclude controls decide what may be grounded on, and personal data is masked before it is ever embedded or sent.

Documents in PDF, Word, Excel and CSV are read natively, and OCR handles scans and screenshots, so a model without native vision can still reason over them. Model-agnostic web search adds current, sourced results with dates when a question depends on public information.

Citations people can open

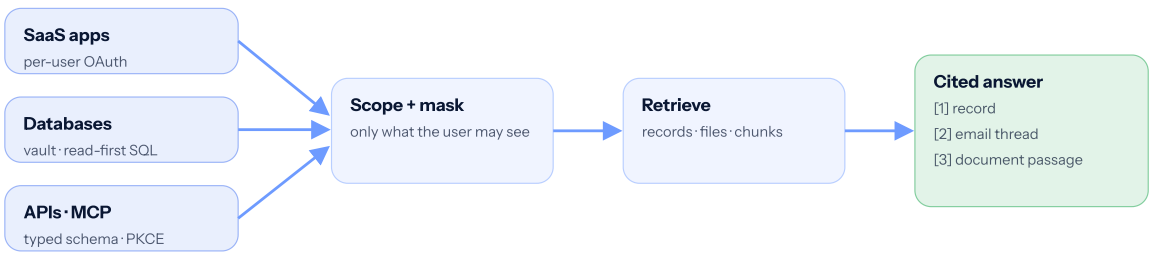

Figure 1. From connected systems to a cited answer.

Every grounded answer links back to the records, files and passages it used. That changes how people use the answer: instead of trusting or distrusting the assistant, they check the source. Teams can also see which sources shaped each answer and refine the knowledge base as content changes.

> **What to check in an evaluation:** ask a question whose answer spans a record and a document, open every citation, and confirm that a user without access to one of the sources receives a different, narrower answer.

Live data, not stale copies

The Data Bridge reads live records and connected apps, so answers reflect the current state of the business rather than a nightly copy. Beyond CRM it reaches Gmail, Calendar, Drive and REST services. Related email and calendar items are pulled via per-user OAuth, and databases answer with real numbers cited to the row, with every query logged.

From answer to action

Connectors are not read-only. They support governed actions — create and update records, run Salesforce Flows, send email, commit code and open pull requests, transition Jira issues and more. Write actions are shown as approval cards, destructive steps always confirm, and every action is written to the audit trail. The same approval model applies whether the action originates in the browser or on a Salesforce record.

Memory is not retrieval

Two mechanisms are easy to confuse. Retrieval brings in facts from your systems and knowledge base at the moment of a question, with citations. Memory is what a Workmate keeps about how a particular person works: concepts, entities, workstreams and notes, applied on future turns on top of cross-session semantic recall.

Keeping them separate matters for trust. Retrieved facts are shared, sourced and access-scoped; memory is personal, isolated per user and private by default, and people can view and manage it. An answer should cite retrieved sources for its facts, while memory only shapes how the answer is framed for the person asking.

Enterprise search across sources

Not every question needs a generated answer. Sometimes people need to find the right document or record. Enterprise search runs one query across connected apps, documents and knowledge, and the results respect each source's access rules. The Knowledge Bridge adds include and exclude scoping, so an administrator can decide which objects and folders are searchable and groundable at all.

Search and grounded answers share the same permission model. If a person cannot open a record in its source system, it does not appear in their search results and it cannot shape their answers.

Logging every connector call

Grounding creates a new question for security and operations teams: which systems did the AI touch, on whose behalf? Every connector call is logged, and each call is attributed to the acting user in the audit trail. For databases, every query is recorded; for writes, the approval that allowed the change is recorded with it.

Those logs are also the best diagnostic tool when an answer looks wrong. If a citation points to an unexpected record, the call log shows how it was retrieved and under which user's access.

Packaged connector or custom connector?

With the long tail reachable through several routes, teams sometimes hesitate over which to use. A simple rule works well. Use a packaged connector when one exists, because per-user OAuth and token refresh are handled for you. Use a REST or GraphQL connector with a typed schema for internal services and APIs that have no packaged connector. Use an MCP server

when a system already exposes tools that way, or when you want tool discovery. And use signed webhooks when a system should push events rather than be polled.

Whichever route you pick, the governance is the same: credentials in the vault, scoped access, approvals for writes and an audit trail — in the browser edition and inside Salesforce.

A grounding checklist

1. List the five systems people ask about most, and connect them first with per-user OAuth.
2. Decide what to exclude from grounding before indexing anything.
3. Build the knowledge base in folders that match how teams already organize content.
4. Test citations with users who have different access levels.
5. Reach the long tail through REST, GraphQL or MCP rather than waiting for a packaged connector.

Sources

Product statements in this paper restate the AgenTorQ website:

- agentorq.com/connectors
- agentorq.com/platform
- agentorq.com/standalone
- agentorq.com/security

General context links point to the public pages named in the text. This paper contains no market statistics.