CONTEXT SYSTEMS

Context Boundary Mapper

A practical worksheet for deciding what an AI system can see, remember, retrieve, and produce before implementation begins.

WHAT THIS TEMPLATE HELPS YOU DECIDE

Turn privacy and data-flow questions into an operating architecture view: approved sources, context permissions, memory limits, retrieval paths, review triggers, and evidence needs.

BEST FOR

- Teams connecting AI to documents, CRM records, inboxes, or internal knowledge
- Leaders who need data boundaries before adding automation
- Workflows where context quality and privacy both matter

OUTPUTS

- A clear system-of-record map
- Allowed, restricted, and no-retention context rules
- Architecture assessment inputs

STEP 1

Map the operating context

AI does not need access to everything. It needs the right context, from the right source, for the right decision, with a visible owner.

Workflow frame

Workflow or decision supported
Name the specific operating moment this AI system supports.

Primary human owner
Name the role accountable for quality, exceptions, and escalation.

Systems of record
List the authoritative sources. Do not include convenience copies.

Architecture layers

Model infrastructure	Which model or AI service will reason over the context?
Context system	Which sources can be retrieved, summarized, searched, or cached?
Workflow interface	Where will people see, approve, edit, or reject AI output?
Governance layer	Where are permissions, logs, retention rules, and review gates enforced?

STEP 2

Define context permissions

A context boundary is useful only when it separates what the system may use from what must remain restricted, redacted, or outside memory.

Allowed context

- [ ] Approved policy, product, service, or operating documentation
- [ ] Customer or account records needed for the specific workflow
- [ ] Prior decisions that are authorized for reuse
- [ ] Structured reference data from a named system of record

Restricted context

- [ ] Personal information that is not required for the task
- [ ] Legal, financial, HR, or medical details without human review
- [ ] Draft strategy documents not approved for operational use
- [ ] Third-party data with unclear licensing or consent

No-retention context

Inputs that must not be stored

Examples: raw customer messages, credentials, sensitive attachments.

Outputs that require deletion or expiry

Name the expiry rule, retention window, or deletion owner.

Treat memory as infrastructure. If a system remembers context, retrieves it, or uses it to shape future outputs, it needs ownership and an evidence trail.

STEP 3

Make traceability visible

The strongest AI-native systems do not merely generate answers. They show where the answer came from, what rule governed it, and when a person must intervene.

Traceability map

Input source	Which record, document, or event triggered the AI workflow?
Retrieval rule	What determines which context can be pulled into the session?
Output consumer	Who receives the output, and is it internal, customer-facing, or operational?
Evidence log	What should be logged so the decision remains reviewable later?
Review trigger	Which conditions require human approval before the output is used?