



CLIENT GUIDE

RAIN Intelligence documentation

Optional hosted reasoning, tools, knowledge retrieval, gated actions, chat, artifacts, telemetry, and vertical intelligence for RAIN OS.


Contents

1. BYOK, credits, billing, and budgets
2. Chat, models, and routing
3. Command Center and workspace
4. Comparable-property and deal intelligence
5. Data handling and retention
6. Files, charts, code, and record artifacts
7. Flow and background automation
8. Health, observability, and operations
9. Knowledge, schema, metadata, and retrieval
10. RAIN Intelligence overview, licensing, and prerequisites
11. Safety gates and actions
12. Salesforce connection and access
13. RAIN Intelligence troubleshooting
14. Voice and multimodal interaction

BYOK, credits, billing, and budgets

Choose organization credits, enabled personal credits, or provider keys deliberately, then monitor balances, fallback policy, usage cost, and warning-only budgets.

Billing modes



The model control exposes automatic routing and the provider groups available to the active workspace. Billing mode, key status, credits, and budgets are separate settings and are not shown in this view.

RAIN Intelligence supports three billing approaches when enabled for the environment:

- **Organization credits** use the organization's RAIN-managed balance synchronized through the supported Salesforce relationship.

- **Personal credits** use an individual hosted balance and can support purchase, subscription, and configured auto-reload behavior when the feature is enabled.
- **Bring your own key** uses the user's or organization's configured model-provider key for supported providers.

The selected mode must also support the selected model. Property, speech, embedding, reranking, and other third-party services can have separate usage and credentials.

Configure a mode

1. Review which modes the organization permits.
2. Select the default mode in the supported billing and model settings.
3. For organization credits, refresh and verify the Salesforce-managed balance.
4. For personal credits, confirm the feature is enabled before checkout or auto-reload.
5. For provider keys, add each key through the protected hosted settings UI and test one low-cost model request.
6. Set model preferences and disabled models.
7. Choose whether provider-key exhaustion stops the turn or permits an allowed fallback.
8. Review usage after the first representative task.

Never enter a provider key into chat, a Salesforce record, a workspace file, or a support message.

Provider-key scope and security

Keys entered through supported settings are protected and are not recoverable or displayed after submission. Depending on administration and user settings, a key can apply to an individual or an organization default. Use least-privilege provider scope, keep individual and shared ownership clear, and rotate or revoke keys when access or ownership changes.

Provider billing, retention, rate limits, and account security remain governed by that provider. RAIN can report observed usage and failures but cannot guarantee the provider account's total spend outside RAIN.

Exhaustion and fallback

When a balance or key is unavailable, policy can stop the request or move to another allowed billing path or model. A fallback can change price, latency, model behavior, and provider data handling.

Keep stop-on-exhaustion for tightly controlled or regulated use cases unless the alternative path has been reviewed.

Personal credit features can be disabled by environment or organization policy. An unavailable checkout is not necessarily a service outage.

Budgets

RAIN Intelligence records supported service costs and can display global or organization monthly budgets and alert thresholds. Budget alerts do not provide a guaranteed hard spending stop.

Use provider-side hard limits, controlled key scope, credit balances, disabled models, and operational monitoring when strict cost containment is required.

Reconciliation

Cost estimates depend on model catalog pricing, measured tokens or units, provider responses, and successful telemetry. Compare RAIN usage with provider invoices and Salesforce credit records. Investigate missing attribution, repeated charges, and mode changes.

Troubleshooting

Organization credits are unavailable: Refresh the Salesforce balance, check the connected org and credit permission, and confirm the feature is enabled.

Personal credits cannot be selected: The environment or organization may not have enabled them, or checkout configuration may be incomplete.

A provider key is reported missing: Confirm the selected model's provider, key scope, key validity, and whether the active user can use the organization default.


A turn stops on quota: Review exhaustion policy and provider balance before retrying. Do not loop through repeated chargeable failures.

Spend exceeds the displayed budget: Budgets are warning thresholds. Pause costly features, inspect usage by provider and operation, and apply enforceable limits at the credit or provider layer.

Chat, models, and routing

Use persistent hosted chat with an available model or automatic routing, and interpret fallback, partial, and provider-health signals before retrying.

Chat lifecycle




The empty workspace keeps the session state safe while showing the controls that frame every turn: chat history, model and gate mode, readiness, workspace tools, and the composer.

RAIN Intelligence stores chat sessions and messages so users can continue work over time. A turn can stream progress, tool activity, gated actions, generated render blocks, and a final response. The service records execution and usage information when the related stores are available.

Each turn is bounded. A complex request may return a verified partial answer when time, tool, or model limits are reached. Continue the same session with a narrower follow-up rather than starting

over and losing context.

Select a model



The open selector is scrolled to complete model rows so provider grouping and model descriptions remain readable. Availability still depends on organization configuration, billing, and provider health.

Users can select an available model or allow routing to choose one. The available list depends on runtime provider configuration, billing mode, organization preferences, disabled models, and service health. Catalog contents can change as providers and supported models evolve.

Automatic routing evaluates the task and available context, then selects one driver model for the run. It may fall back or escalate when a provider is unavailable or a response cannot be completed safely. Model selection does not change Salesforce permissions or action-gate policy.

Start a useful session

1. Confirm the active Salesforce org.
2. State the business goal, relevant record or object, output format, and time range.

3. Identify whether the request is read-only or may change data.
4. Attach only files required for the task.
5. Select a review-oriented gate mode for consequential work.
6. Review tool activity and source grounding as the turn runs.
7. Validate important conclusions against authoritative records.

Use a new session when the organization, audience, confidentiality boundary, or unrelated business objective changes.

Context and memory

RAIN Intelligence uses conversation history and retrieved context within model limits. Long sessions may be compressed or summarized to preserve the most relevant work. A session is not a substitute for a durable Salesforce business record.

When an exact decision, approval, or customer commitment matters, record it in the authoritative system through an approved action. Do not rely on the chat transcript as the only operational state.

Provider fallback

Fallback behavior depends on billing mode and preferences. A provider key can be missing, exhausted, rate-limited, or incompatible with the selected model. RAIN may switch to another available model only when policy allows it. A different provider can produce different style, latency, cost, and tool behavior.

Review the active model shown for the turn before comparing results. When repeatability matters, select a specific available model and keep the same data, prompt, tools, and gate mode.

Safety and output quality

Chat can use Salesforce, retrieval, file, code, chart, and development tools. Tool results improve grounding, but model output remains decision support. Inspect the underlying records, calculations, generated code, deployment diff, and proposed actions before relying on consequential output.

Troubleshooting

No models are available: Check provider or billing configuration, organization model preferences, and model health.

The selected model is replaced: The model may be disabled, unavailable, unhealthy, or unsupported by the active billing mode.

A response stops early: Continue with a narrower request in the same session and inspect whether a tool, provider, or runtime budget caused the partial result.


Chat loses a fact: Restate the authoritative record or attach the required evidence; long-session compression may omit incidental detail.

A provider repeatedly fails: Stop rapid retries, check provider health and quota, then choose an available alternative or wait for recovery.

Command Center and workspace

Use the organization-scoped chat and development workspace for analysis, files, reviewed version-control changes, and gated Salesforce deployment.

Workspace responsibilities



The workspace keeps chat, model and gate choices, files, development tools, and activity evidence in the active organization context.

RAIN Intelligence combines persistent chat with a per-user, per-organization working area. Users can create, read, search, edit, upload, and remove supported files, review proposed changes, use an approved version-control connection, and prepare gated Salesforce development or deployment work.

The workspace is an execution aid. It is not the organization's authoritative document-management, business-record, change-management, or backup system.

Organization and user context

Every workspace action should be associated with the active hosted user and Salesforce organization. Confirm that context before opening files or running a connected action. When switching Salesforce orgs, start in the corresponding working context; do not reuse files across customer or environment boundaries without an approved transfer.

Workspace access does not authorize a Salesforce deployment or a version-controlled change. An administrator must approve the relevant connection and scope, and the user must still satisfy the applicable permission, policy, review, and action-gate requirements.

Work with files

1. Start in the intended organization and session.
2. Create or upload only the files required for the task.
3. Keep secrets and unnecessary production data out of workspace files.
4. Review generated edits and diffs.
5. Run relevant checks and tests.
6. Submit approved changes through the organization's normal review process.
7. Remove temporary files when the task is complete.

Generated files can be incomplete, stale, or based on misunderstood requirements. Use the same review standard as human-authored work.

Connected change management

When version control is enabled, use only the organization-approved connection and target. Review the complete proposed change, follow branch protection and peer-review requirements, and resolve ambiguous conflicts through the organization's normal change-management process.

Never enter connection credentials in chat or workspace files. Administrators are responsible for approved connection scope, periodic access review, rotation, revocation, and reauthorization through supported settings.

Salesforce deployment

Development tools can inspect Salesforce metadata, prepare a reviewed change, run supported checks, and start a deployment. Deployments are consequential actions and should use a review-oriented gate.

Before deployment, verify the target org, approved change scope, tests, destructive changes, dependencies, and rollback plan. After deployment, inspect the Salesforce result rather than relying only on the chat summary.

Workspace boundary

The workspace supports chat, working files, change review, approved version control, and gated development operations. It is not a general support inbox, organization-wide work queue, or durable business approval system.

Use Salesforce and approved operational systems for durable assignments, approvals, customer conversations, and business records.

Troubleshooting

Workspace files are missing: Confirm the active user, organization, session, and whether the file was temporary or deleted.

A version-controlled change fails: Check the approved connection, user authorization, target, branch protection, service availability, and unresolved conflicts.

A deployment is blocked: Review safety gate, connected Salesforce permissions, target environment, metadata validation, and test results.

The wrong org is active: Stop all writes and deployments, switch context, then re-read the target before continuing.

A team needs shared durable work: Use the organization's authoritative change-management and work-management systems; the hosted workspace is not the team work queue.

Escalate persistent workspace isolation, missing-file, connected-action, or deployment failures after the active user, organization, permissions, connection status, and target are verified. Include the timestamp and time zone, affected action, user-visible error, and a sanitized session or activity reference; do not include credentials or customer data.

Comparable-property and deal intelligence

Generate transparent comparable-property and deal analysis from supported market data, then review source coverage, adjustments, confidence, and property condition before making a decision.

Analysis types

RAIN Intelligence can analyze retail comparable properties, as-is comparables, and available deal or buyer context. A typical request uses the subject address and known property facts, then searches supported data sources for candidate properties and returns selected comparables, adjustments, confidence, diagnostics, and a value-oriented summary.

The output is decision support. It is not an appraisal, title report, inspection, lending decision, or guarantee of price.

Required inputs


Provide the most accurate available subject address and, when known, location, beds, baths, living area, property type, year, features, and relevant condition. Optional radius, market, and filtering choices can narrow the search.

Missing or incorrect subject facts affect candidate discovery and every later comparison. Verify the subject before judging the output.

Run a comparable-property analysis

Comparable-property analysis: review before deciding

Each stage can expose missing coverage, weak matches, or assumptions that require human review.



DECISION-SUPPORT OUTPUT

Transparent ranges and diagnostics, not an appraisal or price guarantee

Refresh stale facts, document manual pins or removals, and compare the result with current local evidence.

Source basis: comp engine, as-is comp engine, normalizer, and evaluation paths listed in the article metadata.

This process diagram shows the review sequence. Available controls and provider evidence vary by workspace, market, license, and analysis type.

1. Confirm the subject property and intended valuation question.
2. Review the candidate map or preview where available.
3. Set a reasonable market and distance scope.
4. Run the analysis and inspect included and excluded candidates.
5. Review sale recency, distance, physical similarity, data source, and condition.
6. Inspect adjustments, outlier handling, confidence, and warnings.
7. Pin or remove candidates only with a documented market reason.
8. Compare the result with local expertise and current listing or transaction evidence.

The engine can use location, size, age, beds, features, recency, market movement, and local pricing context. Exact weighting and provider coverage evolve and should not be treated as a public fixed formula.

As-is and condition analysis

As-is work depends heavily on actual condition. Photos and descriptions can improve context but cannot replace an inspection. Distinguish current condition from after-repair assumptions and avoid mixing retail-ready comparables into an as-is conclusion without an explicit adjustment.

Image analysis can be incomplete or wrong when photos are old, staged, low quality, or not representative of the entire property.

Deal intelligence

When connected providers contain retail comps, buyer information, or related deal data, RAIN Intelligence can normalize it into a common decision view. Availability varies by market, provider, license, and record quality. Missing data should be reported as missing rather than inferred.

Cost and caching

Property providers can charge per search, detail, or image request. The service can cache supported property details to reduce repeated calls, so a repeated analysis may use a recently cached fact. Refresh when market timing or corrected property data makes freshness material.

Troubleshooting

No candidates are found: Verify address and coordinates, expand the scope carefully, check provider coverage, and relax only filters that are not essential to comparability.

Candidates are poor matches: Review property type, living area, location boundaries, condition, and subject facts before manually pinning records.

Value looks too high or low: Inspect the selected set, recency, outliers, adjustments, and whether retail and as-is assumptions were mixed.


Provider data conflicts: Prefer the most authoritative and current source, record the discrepancy, and avoid false precision.

Confidence is low: Treat the result as a research lead and obtain more local evidence or professional review.

Data handling and retention

Minimize hosted data, understand what is retained or sent to providers, and apply the documented retention and deletion process for every enabled data class.

Data handled by the service



The connection screen states the visible access categories and separates Production from Sandbox before OAuth begins. It does not prove field-level minimization, retention, or third-party processing terms.

RAIN Intelligence can process Salesforce records and metadata, user prompts, chat responses, uploaded files, generated assets, tool results, usage and failure telemetry, model preferences, provider connection information, Salesforce connections, workspace files, property data, and billing information.

Salesforce records are generally queried live for a task. The hosted service retains data needed for connection, chat, retrieval, artifacts, workspace, telemetry, and billing according to the applicable service terms. Record-level indexed content or samples should be enabled only through an intentional synchronization decision.

Third-party processing

The relevant content for a request can be sent to configured model, retrieval, result-ordering, speech, property, or other connected providers. Each provider has its own retention, training, region, and account settings.

Before enabling a provider, review the data classes it will receive and its contractual controls. A provider key supplied by the customer changes billing ownership but does not eliminate provider data processing.

Minimize data

- Query only the fields and records required for the task.
- Upload a redacted extract instead of an unrestricted document where possible.
- Keep credentials, tokens, private endpoints, and unrelated personal data out of prompts and files.
- Avoid broad raw-record indexing unless the use case, access, retention, and deletion process are approved.
- Use non-production or synthetic data for development and evaluation.
- Persist final business decisions in Salesforce or the approved system of record.

Retention scope

Retention periods are feature-specific across chat, assets, activity, executions, insights, operational telemetry, workspace data, and property information. Provider-held data and contracted backup timelines can follow separate terms.

Organizations must document required periods with the service owner and verify the applicable schedule for every enabled data class. A chat or file deletion confirms only the scope named by that control; it does not by itself confirm removal of provider-held data, contracted backups, indexed content, telemetry, or separately retained artifacts.

Deletion and offboarding

1. Identify the Salesforce org, hosted users, chat sessions, files, indexed content, connections, provider accounts, and billing data in scope.
2. Stop new hosted automation and synchronization.
3. Delete user-visible sessions and assets through supported controls.
4. Submit a support request for service-managed data that does not have a self-service deletion control.
5. Revoke Salesforce OAuth and external provider credentials.
6. Request removal or refresh of retrieval results that contain in-scope content.
7. Apply the provider-deletion process and the contracted backup-retention timeline.
8. Verify that access is removed and record the completed scope and exceptions.

Connection security responsibilities

Create Salesforce and provider connections only through supported settings. Customer administrators are responsible for least-privilege scope, approved account ownership, periodic access review, rotation, and revocation. Never place credentials or tokens in prompts, workspace files, Salesforce business records, or support messages. Use the applicable security and data-processing terms for assurance about service-managed protection.

Incident response

If data is exposed to the wrong user, org, workspace, provider, or log, stop the affected feature, preserve access-controlled evidence, revoke credentials as needed, identify every downstream copy, and follow the organization's notification and deletion process.

User questions

Does deleting a chat delete everything? It removes the scope named by the chat control, but related telemetry, provider processing, indexed content, contracted backups, or separately retained artifacts may have different retention and deletion procedures.


Is Salesforce data copied in full? Live query is the normal task path. Approved schema, aggregates, insights, knowledge, metadata, and explicitly enabled record samples can be retained for retrieval.

Can retention be assumed from provider defaults? No. Review and configure every provider used by the enabled feature.

Files, charts, code, and record artifacts

Attach and generate bounded work artifacts, review their source and sensitivity, and persist final business results in the authoritative system.

Artifact types



The safe empty state identifies the attachment control in the composer and the Code Editor and Activity Log controls in the header. No generated file, record, or deployment result is represented here.

RAIN Intelligence can work with uploaded files and produce structured outputs such as charts, tables, metrics, code, record summaries, and list views. Chat stores supported message and asset context so users can return to the work.

The web workspace also supports working files used for analysis, code, version control, and deployment preparation. A generated artifact is a hosted work product; it is not automatically a Salesforce record, an approved report, or a deployed change.

Upload files

1. Confirm the active organization and session.
2. Remove material that is unrelated to the task.
3. Check the file for personal, regulated, contractual, or secret data.
4. Upload through the supported chat or workspace control.
5. State how the file should be used and what output is expected.
6. Verify the assistant processed the intended file and version.
7. Remove the hosted asset when it is no longer needed under the organization's policy.

File type, size, provider, and model limits apply. Images and complex documents may use a multimodal provider and can have different cost and data-handling behavior from text.

Charts and tables

Charts and tables are generated from retrieved or supplied data. Review the source query, filters, grouping, units, null handling, and time zone before sharing them. A visually plausible chart can still be based on incomplete or stale records.

When the result will become an operational metric, reproduce it in an approved Salesforce report, analytics platform, or governed dataset rather than relying only on the chat rendering.

Code and analysis

RAIN Intelligence can generate and, in supported environments, run bounded analysis or development tooling. Inspect generated code, dependencies, inputs, output files, and error logs. Code execution does not grant authority to deploy or modify a connected system.

Use synthetic or minimized data for exploratory code. Do not embed credentials in files, prompts, scripts, command history, or generated artifacts.

Record artifacts and actions

Record and list-view artifacts summarize live or retrieved Salesforce data and can lead into a gated action. Confirm the org, object, record identity, fields, and freshness before acting. The Salesforce record remains authoritative after any write.

When an artifact contains a proposal, preserve the distinction between generated content, user approval, tool execution, and final Salesforce status.

Version control and deployment

Workspace development features can prepare changes, show differences, interact with version control, and start supported Salesforce deployment work. Before deployment:

- Confirm the target org and branch.
- Review every changed component and generated file.
- Run relevant static checks and tests.
- Use check-only or equivalent validation when available.
- Require an appropriate gate and release approval.
- Keep a rollback or forward-fix plan.

Troubleshooting

An upload is rejected: Check size, type, browser completion, session state, and provider capability.

The wrong file was analyzed: Stop the turn, remove the unintended asset if required, and start a clearly scoped session with the correct version.

A chart is empty or misleading: Inspect the source data, filters, types, and missing values before changing presentation.

Generated code fails: Review the runtime, dependencies, data shape, permissions, and safety restrictions. Do not bypass restrictions by moving secrets into the code.


An artifact disappeared: Check the original session and workspace owner. Hosted assets are not a substitute for the organization's durable document repository.

AUTOMATION

Flow and background automation

Invoke hosted reasoning from Salesforce with a scoped goal and gate mode, choosing synchronous execution for returned results and background execution only for queue-and-forget work.

Cross-product execution



The workspace exposes automation entry points while the header keeps Ask Before Changes visible. It does not show a completed Flow run or background execution result.

Salesforce invokes RAIN Intelligence through the RAIN OS **Run RAIN Intelligence** Flow action. RAIN OS validates and sends the request; the hosted service authenticates it, selects tools and a model, applies gates, and returns a hosted execution result when the mode supports one.

The minimum useful request is a clear goal. Optional context includes a Salesforce record or object, agent version, files, preferred model, timing, gate mode, result destination, timeout, and advanced structured options.

Choose synchronous or background

Use **synchronous** execution when the Flow must inspect the hosted summary, result, actions, or error before continuing. Keep the goal bounded enough to fit practical Salesforce callout and hosted execution limits.

Use **background** execution when the Flow only needs to submit work and can continue without the completed result. Background mode queues Salesforce work that calls the hosted service. It does not return a later completed result to the originating Flow or update the source record automatically.

Do not build a polling or approval process around an assumed callback that does not exist.

Configure the request

1. Define one outcome and the authoritative target record.
2. Choose a gate mode that matches the highest-risk possible action.
3. Limit file and record context to what the task requires.
4. Select synchronous mode if downstream Flow logic needs a result.
5. Map returned success, status, summary, result, actions, run reference, and error explicitly.
6. Add a Flow fault path 🟡 that prevents duplicate resubmission.
7. Test with standard-user access and a low-risk record.

Use update destinations only when the target field, output format, and gate policy are deliberate. Do not configure unsupported destination values in production Flows.

Gates in automation

Plan-oriented modes constrain tools to analysis and proposed work. Review-oriented modes can pause a chat interaction for confirmation, but unattended automation cannot depend on a human being present in the originating Flow transaction. Choose a mode and action set that can produce a deterministic completed, blocked, or failed outcome.

For customer-visible, financial, security, bulk, deletion, code, and deployment work, route the result to a separately supported review process.

Result handling

Persist only the result needed by the business process. A hosted execution can be completed, partial, blocked, queued, or failed. Preserve the run reference for diagnosis, but do not expose it as customer-facing content.

If the hosted service completes actions against Salesforce, re-read the authoritative record before making the next Flow decision. Validation rules, locks, and downstream automation can alter or reject the intended result.

Troubleshooting

Authentication fails: Check the RAIN OS service configuration, hosted org registration, and credential rotation state.

Synchronous execution times out: Narrow the task. Background mode is appropriate only if the result is not needed in the current process.

Background work has no visible result: Queue acceptance is the current contract. Add explicit supported persistence or redesign the workflow.

An action waits for approval: The selected gate and tool risk require review. Do not automatically retry with broader autonomy.

The Flow repeats the request: Correct entry criteria, prior-value logic, fault handling, and queue-reference persistence before resuming.

Health, observability, and operations

Monitor service availability, provider health, Salesforce connectivity, execution activity, failures, costs, and capability-specific degradation as separate operating signals.


Health has multiple layers

RAIN Intelligence exposes a [authorization omitted] check. An authenticated model-health view can check configured model providers. Neither proves that a Salesforce connection, hosted datastore, retrieval service, speech provider, property-data provider, billing path, or every tool is healthy for a particular user.

Operate the product as a set of dependencies:

- Hosted application service.
- Hosted datastore and retrieval service.
- Salesforce connection and limits.
- Model and retrieval providers.
- Speech and property providers.
- Billing and credit synchronization.
- Workspace, version-control, and deployment tools.

Operational signals



Ready confirms that this browser workspace reached its idle state, and Activity Log is the visible execution-evidence entry point. Neither control replaces service, model, datastore, retrieval, or Salesforce health checks.

The service records supported execution activity, token and cost estimates, tool metrics, gate outcomes, user feedback, provider failures, and recurring failure patterns. Administrators can use these signals to find elevated error rates, expensive operations, provider concentration, and common remediation needs.

Activity and cost signals can themselves be delayed or unavailable. A missing metric is not proof that no request occurred.

Daily operating review

1. Confirm the hosted application service is available.
2. Review model-provider health and recent fallback patterns.
3. Test a read-only Salesforce request in each production org class.
4. Check failed and long-running executions.
5. Review cost, credits, and budget warnings.

6. Check retrieval freshness and synchronization failures.
7. Review voice, property, or deployment health if those features are in active use.
8. Triage repeated negative feedback or failure clusters.

Use synthetic or non-sensitive canary requests and avoid chargeable high-cost tests when a focused probe is sufficient.

Degraded operation

Capabilities can fail independently. The application service may remain available while chat history, retrieval, activity signals, billing, or a provider-specific feature is degraded. Diagnose the affected capability rather than treating basic service availability as complete readiness.

When hosted datastore status is degraded, avoid writes or conversations that depend on durable history. When one model provider is degraded, use an approved fallback only if its billing and data handling are acceptable. When Salesforce is degraded, stop actions and deployments but preserve non-Salesforce workspace work as appropriate.

Cost and budget monitoring

Compare execution cost activity with provider and credit records. Budget alerts are informational unless the applicable service terms explicitly define enforcement; do not treat an alert threshold as a guaranteed hard spending stop. Alert ownership and response must be operationally defined.

Investigate abrupt changes in model mix, repeated retries, property-detail calls, transcription volume, and high-cost partial turns.

Incident evidence

Capture environment, service version, time, affected org and feature, sanitized run or session reference, selected model, billing mode, gate mode, dependency health, and safe error category. Do not put provider keys, Salesforce tokens, private endpoints, full prompts, customer records, or uploaded files into general incident channels.

Support escalation

Escalate when the hosted application, datastore, or retrieval service remains degraded; when multiple entitled users or orgs show the same provider or activity failure; or when health indicators

conflict with user-visible outcomes. Include the timestamp and time zone, environment type, affected capability and org, user-visible status, sanitized activity reference, customer-managed checks already completed, and whether writes or autonomous work were paused.


Recovery

Restore authentication and data isolation before convenience features. Wait for the relevant hosted-service status to recover or for support to confirm recovery before resuming writes. Re-run one read-only Salesforce canary, one gated low-risk action, and the affected provider path. Resume background or autonomous use gradually while monitoring duplicate and cost signals.

Knowledge, schema, metadata, and retrieval

Connect approved Salesforce structure and knowledge for grounded retrieval, and enable record-level memory only through an intentional access and retention decision.

Retrieval layers



This operating view distinguishes live Salesforce facts from indexed context so administrators can choose the correct freshness and access boundary.

RAIN Intelligence can ground work with several kinds of context:

- Live Salesforce queries for current records.
- Salesforce schema and object descriptions.
- Salesforce and package metadata indexes.
- Approved knowledge content.
- Cached aggregates, insights, or model information.
- Record samples only when record-level synchronization is intentionally enabled.

These layers have different freshness and access characteristics. Use a live query for a current restricted fact and retrieval for discovery, explanation, and context.

Initial synchronization

After an approved Salesforce connection is established, administrators can synchronize selected schema, metadata, and knowledge. Supported metadata can include code, objects, Flows, triggers, layouts, permission metadata, Lightning components, pages, and other available types.

Indexing may be partial when the connected identity cannot describe an object or retrieve a metadata type. A successful connection does not prove complete indexing.

Configure retrieval

1. Identify the content and metadata users should be able to discover.
2. Scope the connected Salesforce identity and RAIN source filters.
3. Confirm the hosted retrieval capability is enabled for the organization.
4. Enable optional content-processing or result-ordering providers only when their data terms are approved.
5. Run the supported schema, metadata, or knowledge synchronization.
6. Test exact-name, semantic, no-result, access-denied, and stale-content cases.
7. Record synchronization ownership and freshness expectations.

Do not index broad raw record populations by default. Record-level memory should have an explicit use case, field allowlist, user access model, retention period, and deletion process.

Search behavior

Retrieval searches approved indexed content within the organization, type, object, and source boundaries configured for the request. Optional result ordering may affect which candidates are presented first.

A high similarity result is not proof that content is current or authoritative. Results can be stale, duplicated, incomplete, or drawn from a metadata version that has since changed. The assistant should cite or identify the operative Salesforce source, and users should verify consequential facts.

Data and provider considerations

Configured content-processing and result-ordering providers receive the content required for their operation. Review their data handling before indexing personal, regulated, contractual, or proprietary material. Indexed metadata can also expose code and configuration details to authorized hosted users; source filtering and workspace access must reflect that sensitivity.

Refresh and deletion

Refresh after material schema, metadata, package, knowledge, or permission changes.

Administrators should monitor the last successful synchronization and investigate missed or stale refreshes rather than assuming content is current.

Deleting Salesforce content does not necessarily remove previously indexed content or prior chat references immediately. Follow the supported data-removal process and verify that the content is no longer returned.

Troubleshooting

Search returns nothing: Check retrieval-service availability, organization and type filters, source permissions, synchronization completion, and terminology.

Results are stale: Run the supported refresh, check sync timing, and confirm the index represents the active org and package version.

A metadata type is absent: Confirm it is supported, retrievable by the connected identity, and included in the approved synchronization scope.


Users discover content they should not: Disable affected retrieval, narrow connection and source access, request removal of the affected indexed content, and verify the corrected results before restoring access.

Optional result ordering fails: Retrieval may continue with a less useful ordering. Check the configured provider and compare results before declaring the complete retrieval capability unavailable.

RAIN Intelligence overview, licensing, and prerequisites

Add RAIN Intelligence to an initialized RAIN OS org when hosted chat, routing, retrieval, artifacts, gated tools, or vertical intelligence are required.

What RAIN Intelligence adds



The workspace combines persistent chat entry, model and action-gate controls, workspace tools, activity, attachments, and voice controls in one session surface.

RAIN Intelligence is an optional hosted add-on for RAIN OS. It provides persistent chat, multi-provider model routing, Salesforce and development tools, grounded retrieval, generated artifacts, activity and cost telemetry, voice and multimodal input, and real-estate intelligence.

Install and initialize a supported RAIN OS release first. An authorized administrator then connects the Salesforce org to RAIN Intelligence through the approved connection flow.

What remains in RAIN OS

RAIN OS continues to own Salesforce agents, communications, phone records, outreach, suggestions, insights, scoring, Flows, Lightning surfaces, reports, permissions, and protected Salesforce-side connection settings.

RAIN OS remains available when RAIN Intelligence is unlicensed, disconnected, or temporarily unavailable. During a RAIN Intelligence service interruption, pause only the workflows that depend on its chat, routing, retrieval, artifacts, automation, voice, or vertical intelligence capabilities.

Licensing and entitlement

Hosted access requires an active commercial entitlement for the organization and the appropriate RAIN OS custom permission for each Salesforce user. Model usage also requires an available billing mode and the providers needed by the selected model or feature.

Licensing, Salesforce permission, hosted authentication, model billing, and third-party provider accounts are separate prerequisites. Satisfying one does not automatically satisfy the others.

Service and administrator prerequisites



Connect to Salesforce

Authorize RAIN Intelligence to access your Salesforce org

Environment

Production
login.salesforce.com

Sandbox
test.salesforce.com

RAIN Intelligence will have access to:

- ✓ Read and query your data
- ✓ Access metadata and schema
- ✓ Create, update, and deploy (when requested)

Want RAIN inside Salesforce?

Install the native RAIN OS package after connecting, or install it first if you are the Salesforce admin.

Install RAIN OS Package

Connect with Salesforce →

You can revoke access anytime from your Salesforce settings

The connection step separates Production and Sandbox, describes requested access categories, and links package installation before OAuth. Entitlement and service availability must still be validated separately.

Before onboarding:

- Install and validate a supported RAIN OS version.
- Establish an approved Salesforce connection with least-privilege API access.
- Complete organization registration and user sign-in through the supported flow.
- Configure at least one available model provider or organization billing path.
- Confirm that the hosted service and each enabled capability report available.
- Review data handling, retention limits, third-party model terms, and regional requirements.
- Decide default safety gates, model choices, disabled models, and who may manage billing or connections.

Voice, property data, retrieval, result ordering, version control, and deployment tools require their applicable provider accounts, approved connections, permissions, and service entitlements.

Onboarding validation

1. Confirm standalone RAIN OS workflows still operate.
2. Connect a non-production or limited-scope Salesforce org first.
3. Sign in as an entitled standard user.
4. Run a read-only chat request against known Salesforce data.
5. Verify the connected org, source grounding, and model used.
6. Test a low-risk action in a review-oriented gate mode.
7. Upload a non-sensitive file and verify artifact handling if files are enabled.
8. Review activity, cost, and provider-health information.
9. Test access denial for a user who lacks the RAIN Intelligence permission.

Availability boundaries

Capabilities are conditional. A model is usable only when its provider, billing path, policy, and service configuration are available. Retrieval requires an enabled retrieval service and approved content sources. Salesforce tools require a healthy connection. Voice requires browser and provider support. Property intelligence requires market data from configured sources.

RAIN Intelligence provides a chat workspace and per-user working files. Licensing does not by itself add a general support inbox, an organization-wide work queue, or durable business-process approvals. Use the supported RAIN Intelligence surfaces and keep operational records in their authoritative systems.

Support escalation

Escalate when an entitled user remains unable to sign in, a connected org cannot complete a read-only request, or a required hosted capability remains unavailable after its customer-managed connection, provider, billing, and permission prerequisites are verified. Include the environment type, affected capability, timestamp and time zone, user-visible status, and a sanitized session or activity reference. Do not include credentials, tokens, customer records, full prompts, or uploaded files.

When to use standalone RAIN OS


Prefer standalone RAIN OS for deterministic Salesforce-native communication, phone, outreach, suggestion, scoring, and Flow processes that do not need hosted reasoning. Use RAIN Intelligence

when the job benefits from conversational analysis, dynamic tool selection, cross-source retrieval, generated artifacts, multimodal work, or hosted domain intelligence.

Safety gates and actions

Match the gate mode to business risk, inspect every proposed write or deployment, and treat chat approval as authorization only for the presented action and current interaction.

Gate modes



Ask Before Changes is visibly active in this idle session. This screenshot establishes the selected policy, not approval of any proposed action; no action is pending or executed.

RAIN Intelligence supports four user-facing safety postures:

- **Plan First** emphasizes analysis and a proposed plan before changes.
- **Ask Before Changes** requires confirmation for write activity and is the appropriate general default.
- **Balanced Autonomy** can proceed with lower-risk actions while escalating higher-risk work.

- **Full Autonomy** permits the broadest supported automatic execution and should be restricted to proven workflows.

The executor also applies deterministic policy based on tool, target, operation, bulk size, sensitive fields, deletion, and deployment risk. A model cannot waive a required deterministic gate through persuasive text.

Review an action

Before approval, inspect:

1. The active Salesforce org and connected system.
2. The exact target records or metadata.
3. Fields, values, recipients, files, or deployment components.
4. Whether the action is reversible.
5. Validation, automation, and downstream customer impact.
6. The evidence that supports the action.
7. Any warnings, partial tool results, or unresolved ambiguity.

Approve only the action presented. If the target or values change materially, require a new proposal.

Gate continuation

When a turn reaches a gate, RAIN Intelligence retains the pending action context and can continue the exact approved or rejected batch. This is a chat-turn continuation mechanism. It is not a durable organization-wide approval queue, and it should not be used as the sole record for a long-running business authorization.

If the process or browser restarts, confirm whether the action actually ran before approving or resubmitting anything.

Salesforce remains authoritative

Salesforce validation, sharing, field access, locks, and automation can still reject an approved action. Conversely, the stored integration credential may have broader Salesforce access than the interactive user. Gate policy and least-privilege Salesforce connection design must both be correct.

For customer communications, financial changes, security settings, bulk updates, deletion, code, and deployments, use Plan First or Ask Before Changes plus an independent business approval where policy requires one.

Autonomous execution

Enable broader autonomy only after measuring representative outcomes. Limit the available tools, records, and connected credentials, then monitor false positives, failure patterns, cost, and downstream automation.

Full Autonomy means the current tool gate can approve supported actions without a chat confirmation. It does not remove the need for least-privilege access, policy, monitoring, outcome review, and incident response.

Troubleshooting

A read-only request asks for approval: The plan may contain a write, deployment, or sensitive tool. Inspect the proposed action rather than approving based on the original wording.

An approved action fails: Check Salesforce access, current record state, validation, locks, provider status, and whether the target changed after the proposal.

The approval prompt returns after refresh: Determine whether the first continuation was submitted and inspect the authoritative target before repeating it.

Balanced Autonomy performs too much: Switch to Ask Before Changes, narrow tools and connection access, and review the action's risk classification.

A user cannot approve: Confirm session ownership, authentication, organization context, and any required RAIN or Salesforce permission.

Salesforce connection and access

Connect each Salesforce org through the supported OAuth or application flow and scope the connected identity to the records, metadata, and actions RAIN Intelligence should use.

Connection model

RAIN Intelligence connects to Salesforce through an approved organization and application context. Configure the OAuth application boundary according to Salesforce's connected-app guidance [🔗](#). The service stores the connection needed for API work and associates hosted users and sessions with an organization.

One hosted environment can support multiple Salesforce orgs. Keep sandbox, staging, and production connections visibly separate, and verify the active org before any action or deployment.

Choose the connected identity

Use a dedicated Salesforce integration identity where practical. Grant only the objects, fields, metadata operations, files, and API features required by enabled RAIN Intelligence tools. Review Salesforce's user-access guidance [🔗](#) for the permission layers involved, and verify record visibility against the organization's sharing model [🔗](#).

The hosted service can act through the stored connection rather than through the exact permissions of the person typing in chat. Salesforce still enforces the connected identity's access, but administrators must not assume that every hosted operation is automatically limited to the interactive user's Salesforce profile.

RAIN Intelligence also applies its own authentication, session ownership, source filters, and action gates. These controls supplement Salesforce permission; they do not replace least-privilege connection design.

Connect an org



Connect to Salesforce
Authorize RAIN Intelligence to access your
Salesforce org

Environment

Production
login.salesforce.com

Sandbox
test.salesforce.com

RAIN Intelligence will have access to:

✓ Read and query your data

✓ Access metadata and schema

✓ Create, update, and deploy (when requested)

Want RAIN inside Salesforce?
Install the native RAIN OS package after connecting,
or install it first if you are the Salesforce admin.

Install RAIN OS Package

Connect with Salesforce →

You can revoke access anytime from your Salesforce settings

Choose Production or Sandbox before starting OAuth. The screen lists the broad access categories and keeps native RAIN OS package installation distinct from connecting the hosted service.

1. Confirm the RAIN OS package and hosted user entitlement.
2. Start the supported Salesforce connection flow from the intended environment.
3. Authenticate as the approved integration or user identity.
4. Review and approve the requested Salesforce scopes.
5. Confirm the returned org name and environment.
6. Run a read-only query against a test object.
7. Test one expected access denial.
8. Record the connection owner, purpose, scope, and rotation process.

If multiple RAIN applications use the same hosted service, keep their connection contexts separate so one application cannot silently replace another application's Salesforce credential.

User and session access

Users sign in to the hosted experience and operate within an organization context. Chat sessions and working assets are associated with that context and their owner. Switching organizations

should change both Salesforce data access and hosted storage scope.

The RAIN Intelligence permission in Salesforce controls access to the packaged entry point. It does not create a hosted user, establish OAuth, or grant data access.

Connection lifecycle

Reauthorize when the Salesforce refresh token is revoked, the connected user is deactivated, required scopes change, or the organization changes identity policy. Rotate the hosted organization credential or registration proof separately when required.

Before deleting a connection, identify scheduled or background requests, active chat sessions, indexed metadata, and users that rely on it. Disconnecting Salesforce stops new live operations but does not automatically delete all hosted records derived from earlier use.

Troubleshooting

Authorization loops or returns to the wrong environment: Check the selected hosted environment, callback configuration, browser session, and Salesforce login domain.

The connection succeeds but queries fail: Confirm API access, object and field access, connected-user status, token validity, and the active org context.

A user sees the wrong org: Stop writes, switch to the intended org, and verify session and connection linking before continuing.


Metadata or deployment tools fail while queries work: The connected identity may have data API access without the metadata, development, or deployment permissions required by those tools.

Access is broader than expected: Reduce the connected identity's permission sets and review RAIN tool policy and user entitlement. Do not rely only on chat-role labels.

RAIN Intelligence troubleshooting

Isolate hosted failures by authentication, organization context, Salesforce access, model and billing availability, storage, retrieval, tools, and result persistence before retrying.

Classify the failing stage



A coherent idle workspace provides the browser-side baseline: Ready is visible, the active model and gate policy are named, and Activity Log is available. Dependency health still requires separate service evidence.

Start with the current org, user, session, task, and authoritative target. Then identify the first failed stage:

1. RAIN OS permission or hosted sign-in.
2. Organization registration or Salesforce connection.

3. Session or chat loading.
4. Billing mode or model selection.
5. Retrieval or file preparation.
6. Tool planning, safety gate, or approval continuation.
7. Salesforce, provider, code, or deployment execution.
8. Result rendering or persistence.

Do not broaden permissions or switch to Full Autonomy just to make an unclear failure disappear.

Authentication and org context

If sign-in fails, check entitlement, browser session, hosted environment, and Salesforce connection. If data is unexpected, stop all writes and confirm the active Salesforce org and connected identity.

An authenticated chat session can still have an expired or revoked Salesforce token. Reauthorize the connection rather than recreating users or changing source data.

Models and billing

If no model can run, inspect provider availability, model preferences, disabled models, billing mode, credits, provider-key scope, and quota. A requested model may be unavailable even when another provider is healthy.

Repeated model retries can add cost. Use one focused health check, select an approved fallback, or wait for recovery.

Retrieval and files

For missing context, compare a live Salesforce query with retrieval. Check synchronization time, org and source filters, embedding health, file upload completion, and whether record-level data was ever enabled.

Do not treat a no-result retrieval response as proof the Salesforce record does not exist.

Actions and gates

If an action is blocked, inspect the exact risk and proposal. If approval was submitted but the result is unknown, read the Salesforce record or connected system before continuing. Do not send a

second approval until the first action's state is known.

Salesforce validation, sharing, locks, and downstream automation can fail after hosted approval. Correct the authoritative issue rather than weakening the gate.

Storage and degraded mode

If chat history, files, telemetry, billing, or retrieval is missing while the service still responds, suspect a store-specific initialization or database problem. Pause work that requires durable state and preserve local user evidence.

The simple service health indicator may remain green during an optional capability failure.

Safe retry rules

Retry a transient provider timeout only after checking whether the first action completed. Do not automatically retry authentication, permission, validation, ambiguous write, deployment, communication, or chargeable property requests.

For partial analysis, continue the same session with a narrower goal and reuse completed evidence.

Support evidence

Provide environment, approximate time and time zone, feature, active org category, sanitized session or run reference, user role, model, billing mode, gate mode, error category, and whether a read-only reproduction succeeds. Include recent connection, credential, deployment, package, or provider changes.

Do not provide credentials, tokens, private endpoints, literal customer data, uploaded files, full transcripts, or unrestricted logs outside the approved secure process.


Escalation priorities

Escalate immediately for cross-org data exposure, incorrect writes, credential compromise, duplicate customer actions, uncontrolled spend, or deployment to the wrong org. Disable the narrow capability, revoke trust where necessary, preserve evidence, and validate isolation before restoring service.

Voice and multimodal interaction

Use browser transcription, spoken output, realtime conversation, and image or document analysis only when the required provider, billing, browser, and privacy controls are available.

Voice modes



The empty composer shows the attachment control on the left and separate dictation and realtime voice controls on the right. No microphone permission, recording, transcript, or file is active.

RAIN Intelligence supports distinct voice experiences:

- **Transcription** converts recorded browser audio into text for an ordinary chat turn.
- **Spoken output** converts supported text responses to audio.
- **Realtime conversation** streams audio directly through a supported realtime model session.

These hosted voice modes are separate from RAIN OS phone numbers and provider-managed voice-agent calls. A browser microphone session does not create a telephony call record unless a separate Salesforce workflow does so.

Prerequisites

- Use a supported browser over a secure connection.
- Grant microphone permission only to the approved hosted origin.
- Configure the required speech or realtime provider.
- Select a billing mode that can use the chosen model.
- Confirm organization and provider policies allow audio and transcript processing.
- Use headphones and an appropriate environment for sensitive work.

Some speech features depend on more than one provider. The UI can present an option that later reports unavailable when the server-side provider or billing prerequisite is missing.

Use transcription

1. Confirm the active org and session.
2. Start recording and speak a bounded request.
3. Stop recording and review the transcript before submission.
4. Correct names, numbers, record references, and action language.
5. Submit the text as a normal gated chat turn.

Transcription errors can change the target or intent of an action. Never approve a write based on audio alone; inspect the transcribed proposal.

Use realtime conversation

Realtime mode establishes a short-lived provider session for low-latency audio. Confirm the displayed model and voice, then begin with a read-only request. Live voice has tighter response and tool budgets than ordinary chat and may perform only a focused data check before answering.

For complex analysis, deployments, bulk work, or anything requiring detailed review, switch to text chat so the inputs, tool results, and action proposal remain easy to inspect.

Multimodal files

Supported models can analyze images and documents attached to a session. Results depend on image clarity, file parsing, model capability, and supplied context. Visual analysis can assist with property condition, document extraction, and explanation, but users should verify material facts against the original file and authoritative records.

Do not upload identity documents, recordings, private contracts, or regulated images unless the organization's approved use case and provider terms permit it.

Troubleshooting

The microphone is unavailable: Check browser permission, secure origin, device selection, and operating-system privacy settings.

Recording succeeds but transcription fails: Check file size, speech-provider health, network continuity, language, and audio quality.

Realtime mode cannot start: Verify provider key or billing availability, supported model, session authentication, and browser WebRTC support.

The wrong record or number is transcribed: Correct the text before sending and use a gate that exposes the exact action target.

An image result is uncertain: Provide a clearer source, specify the decision question, and require corroborating Salesforce or property data.