

UPFLUX NOUS TECHNICAL RULES OF USE

Smart Process S.A. (UpFlux)

Document information

Field	Value
Code	PS-SEC-NOUS-001
Version	1.0
Status	In force
Issue date	May 25, 2026
Owner	Information Security and Engineering Management
Approvers	Information Security Committee (CSI) and Executive Board
Classification	Public
Review cycle	Every six months, or upon each major Nous release, or upon a relevant regulatory change

1. Purpose

This document describes the technical operating rules, the functional limits and the security and privacy controls that apply to UpFlux Nous, the conversational artificial intelligence layer of the UpFlux platform.

It is intended for three audiences:

- IT and Information Security teams at customers and prospects that evaluate the solution, are accountable for its internal approval, or conduct due diligence.
- Business teams that design solutions on top of Nous and need to know in advance which constraints cannot be worked around through configuration.
- Compliance and privacy functions that need to understand how the assistant handles data and where the human decision points are.

This document is subordinate to the UpFlux Generative AI Governance Policy. In the event of conflict, the policy prevails.

Notice: Nous evolves continuously. This document reflects the behavior of the product on the issue date. The version in force is always the one published on the UpFlux website. Specific contractual clauses, where they exist, prevail over the content described here.

2. Scope

This document applies to the following UpFlux Nous modules and components:

- Ask Nous and Chats, the conversational interface.
- Commands, also referred to as playbooks, defined per process.
- Scheduled routines.
- Findings and Inbox.
- Views and TV Mode.
- Mobile access through a progressive web application.
- The execution engine, the tools available to the language model, and the data access layer.

The following are outside the scope of this document: the Process Intelligence layer itself, meaning ingestion, transformation and process mining, which is covered in the platform architecture documentation; non generative machine learning models, which are covered by a dedicated policy; and modules, agents and integrations licensed separately, which have their own documentation and rules once contracted.

3. Definitions

Term	Definition
Nous	The UpFlux Enterprise AI platform. A conversational artificial intelligence layer built on top of the Process Intelligence foundation.
Conversation target	The combination of organization and process, set before the first message and immutable thereafter.
Command, or playbook	A standardized analysis defined in the configuration of a given process and invoked by typing a slash in the message field or by selecting a category chip.
Routine	A recurring scheduled analysis, defined by a natural language instruction, a frequency and a delivery channel.
Finding	An analysis saved from the chat into the Findings library, with a title, a description and tags.
View	A live dashboard created from a playbook analysis, refreshed on demand or automatically, and displayable in TV Mode.
Tier	The complexity classification assigned to a question, which determines the tools the language model is allowed to invoke.
Tool	A function the model can invoke during execution, such as a data query, code execution or reading a context file.
RLS	Row Level Security. Restricts which records the engine is allowed to read when producing an answer.
Hallucination	The generation of factually incorrect or fabricated information presented with an appearance of confidence. This is the most critical and most language model specific risk.
Module gating	Feature enablement according to the capability contracted by the organization. A module that is not enabled does not expose its menu.

4. Reference architecture

Layer	Definition
Delivery model	Pure software as a service. No infrastructure is required in the customer environment in order to use the platform. Access is through a web browser.
Infrastructure	AWS public cloud, us-east-2 region (Ohio). Infrastructure managed as code, validated by a continuous integration pipeline.
AI layer	The conversational layer that interprets the user question, plans the analysis, queries the process data and synthesizes the answer. It always operates within the context of the active process and is restricted to the data the user is authorized to access.
Language models	Approved provider: Anthropic, Claude family, with model selection according to the capability required at each stage of execution. Orchestration uses a provider abstraction, with no direct coupling between the application and the model, which allows models to be replaced or combined without impact on configured flows.
Transactional data layer	Managed document database, with a dedicated database per customer and a cluster of at least three nodes. Encryption in transit and at rest. The customer controls partial and full deletions as well as the backup and recovery policy.

Layer	Definition
Analytical layer	Vectorized columnar analytical layer, accessed in read only mode by the assistant, complemented by a high performance search engine.
Assistant persistence	Conversations, artifacts, context configuration and audit trail stored in a document database segregated by organization.
Identity and access	Federation through OpenID Connect, with support for corporate directories, multi factor authentication and single sign on where the customer has the integration configured. Token validation on every request.
Notifications	Queue based orchestration service, delivering through in app notification, email and webhooks.
Data ingestion	Near real time event log ingestion, with rule based violation detection, throughput control and automatic reprocessing. Asynchronous transformation through configurable pipelines and a staging area for raw data. Extraction over REST, SQL and files, with connectors for TOTVS, SAP, ServiceNow, Tasy and MV.
Observability	Application performance monitoring and tracing of model executions, with quality and consumption metrics.

5. Context and session rules

ID	Rule
CTX-01	Every conversation has a target, composed of an organization and a process. The target is confirmed before the first message and displayed in the selector below the message field.
CTX-02	The conversation target is immutable once the first message has been sent. Changing the process requires opening a new conversation.
CTX-03	Every answer is produced within the context of the active process. There is no cross process answer within a single conversation.
CTX-04	The conversation breadcrumb follows Organization, Process and Title. The title can be edited by the user.
CTX-05	Conversational context is retained within a single conversation and is not shared across separate conversations. Each new conversation starts without prior history.
CTX-06	The interface displays a badge with the date of the most recent data load. The badge turns amber when the data is more than seven days old, signaling to the user that the analysis may not reflect current operations.
CTX-07	Supported languages: Brazilian Portuguese, English and Spanish.
CTX-08	Default entry point: on desktop the platform opens on Ask Nous, on mobile it opens on Chat.

6. Chat and artifact rules

ID	Rule
CHT-01	Natural language input. Enter submits the message and Shift plus Enter inserts a line break. No knowledge of SQL or programming is required.
CHT-02	The home screen displays the Process Pulse card, with Cases, Activities, Average lead time and Median, together with category chips grouped into six profiles: Process Analyst, Audit and Compliance, Buyer and Operations, Executive and CPO, Finance and Payments, and Procurement and Sourcing Manager.
CHT-03	The output consists of a narrative with cited indicators, tables, visualizations and interactive dashboards.
CHT-04	During execution the assistant displays the build process, showing the steps taken and the tools used. Transparency of tool calls is standard product behavior.

ID	Rule
CHT-05	Where context is insufficient, the assistant presents guided questions, which can be answered or skipped. A gap in context is not filled by inference that has not been disclosed to the user.
CHT-06	Every answer can be rated as Helpful or Not helpful. A negative rating opens a feedback form, which is used in the continuous improvement of the configuration.
CHT-07	Comments and mentions allow collaboration on analyses between users of the same organization.
CHT-08	Available artifact actions: Expand, Download as HTML, Export to Excel, Export to PDF, Share, and Present in Story Mode.
CHT-09	Conversions available from the chat: Save as Finding, Turn into a View, and Turn into a routine.
CHT-10	Conversations are bound to the authoring user and are not visible to other users of the organization by default. Users can delete conversations from their own history.
CHT-11	Exporting an artifact transfers content outside the platform. From that point onward, handling of the file is the responsibility of the customer organization, in accordance with its own information classification and handling policy.

7. Reasoning mode rules

ID	Rule
RAC-01	Nous 2.1: direct and fast answer, with no planning step presented to the user.
RAC-02	Nous Deep: the assistant builds an investigation plan, presents it to the user and waits for explicit confirmation before running the full analysis, with explicit criteria, group comparison and concentration analysis.
RAC-03	Availability of Nous Deep is determined by organization level enablement, not per user.
RAC-04	The Nous Deep confirmation step is an execution containment control and is not suppressed by routine configuration.
RAC-05	Precision rules applied to the engine: numerical presentation at an appropriate granularity, inheritance of context and period in follow up questions, charts with value labels, and stable rendering across executions.

8. Command and playbook rules

ID	Rule
CMD-01	Commands are standardized analyses defined in the configuration of each process. They are process specific and not global.
CMD-02	Invoked by typing a slash in the message field, with incremental filtering by name or description, arrow key navigation and selection by Enter or click. The command is inserted into the field and still requires the user to submit it.
CMD-03	Alternatively invoked through the category chips, which list the commands for the corresponding profile.
CMD-04	A process with no published commands returns an empty state. Commands from another process are not used as a fallback.
CMD-05	Only analyses originating from a command or playbook can be promoted to a live View. Ad hoc analyses cannot.

9. Scheduled routine rules

ID	Rule
ROT-01	Configuration fields: Name, Process, the analysis instruction in natural language, Schedule and Delivery.

ID	Rule
ROT-02	Daily, weekly or monthly frequency, with a day of week or day of month and a time, plus the shortcuts Weekdays, Weekends, Every day, Day 1, Day 15 and Day 28.
ROT-03	Delivery through in app notification and by email, accepting multiple comma separated addresses. At least one delivery channel is mandatory.
ROT-04	The run now action triggers the routine outside the schedule, with the result delivered by email.
ROT-05	Available actions: pause and resume, edit and delete. The configuration and history tabs record success or error, automatic or manual origin, and execution duration.
ROT-06	After three consecutive failures the routine is paused automatically, with an on screen notice.
ROT-07	Before scheduling from a conversation, the assistant pre analyzes the question and raises the clarifications required, so that a routine is not configured on an ambiguous premise.
ROT-08	Routine results also arrive as messages in the Inbox, with an execution history.
ROT-09	The routine instruction is a configuration artifact, with its change history preserved.

10. Findings, inbox and sharing rules

ID	Rule
ACH-01	Findings are the library of analyses saved from the chat, with a Title, an optional Description and Tags. The page provides search, filtering by tag, a side preview, open, download, share, a link back to the source conversation, comments, and action plan generation. Available formats: HTML, PDF and shareable links.
ACH-02	Findings are private by default.
ACH-03	Link sharing operates in three scopes: Private, restricted to the author; Organization, for authenticated members; and Public, accessible to anyone holding the link, with an expiry date.
ACH-04	The Public scope does not require authentication. Its use is a decision of the customer organization and should not be applied to analyses containing personal data, sensitive personal data or confidential information. Where it is used, setting an expiry date is always recommended.
ACH-05	The Inbox is a menu distinct from Findings. It receives mentions and routine results, with filters for unread, mentions and routines.

11. View and TV Mode rules

ID	Rule
VIS-01	Only playbook analyses can be promoted to a View. Ad hoc analyses cannot.
VIS-02	Created from the chat through the turn into a View action, by setting a name and confirming.
VIS-03	Refresh on demand, with the time of the last refresh displayed, or automatically at intervals of 15 seconds, 30 seconds, 1 minute, 5 minutes and 15 minutes.
VIS-04	TV Mode provides full screen display, theme selection, a dedicated display link, and exit through ESC.
VIS-05	A failed refresh does not invalidate the dashboard. The previous version is retained together with a notice to the user, so that a blank dashboard or partial data is never displayed.
VIS-06	Deleting archives the View and removes it from the list. Renaming is available from the dashboard itself.
VIS-07	The dedicated display link grants access to live data. Deciding where the dashboard is displayed is the responsibility of the customer organization, which must consider the classification of the information shown before exposing it in a common area.

12. Mobile access rules

ID	Rule
MOB-01	Mobile access is provided through a progressive web application, installed from a banner in the browser. On Android through the install action, on iOS through share and add to home screen. It runs in standalone mode and opens on Chat.
MOB-02	Bottom navigation provides access to Chat, Routines, Findings and the More menu, which contains settings, theme, language, usage, process and organization switching, and sign out. The items displayed vary according to the modules enabled for the organization.
MOB-03	Push notification is opt in and requires browser authorization. A blocked state is signaled to the user. Notifications include an on screen alert, sound, and web push with deep links by organization and process.
MOB-04	The application signals when a new version is available and when there is no connection.
MOB-05	There is no native application published in an app store. Access is always through a browser, which removes the need for package management across the customer device fleet.

13. Execution engine and model tool rules

ID	Rule
EXE-01	Execution follows a five stage pipeline: question classification, planning, execution, rendering and answer synthesis.
EXE-02	Classification assigns a complexity level among simple, medium, complex and investigative. Most classifications occur without a call to the language model, through deterministic rules.
EXE-03	The tools available to the model are released according to the complexity level. Simpler levels operate with data querying only. Code execution and enriched rendering are released only at the levels that require them.
EXE-04	Data querying operates in read only mode. It accepts select statements exclusively, blocks any modification or removal statement, blocks multiple statements in a single call, and validates referenced columns against the actual database schema.
EXE-05	Code execution takes place in an isolated environment, with no network access and no file system access, with system, subprocess and server resource libraries blocked, and an allow list for the analytical libraries in use.
EXE-06	Reading context files operates in read only mode, with the path restricted to the context configuration directory.
EXE-07	At the rendering stage the model produces a structured configuration consumed by a library of standardized components. The model does not generate free form markup for standard visualizations, which eliminates execution errors and ensures visual consistency.
EXE-08	Where enriched content is generated, the data is injected by code from the query result, not by the language model.
EXE-09	Domain context resolution is cascading, from the most specific to the most generic: customer and process specific configuration, process template, domain context, and finally execution with no configured context.
EXE-10	Execution with no configured context is a degraded condition. It is monitored and is not accepted as a permanent state in production.
EXE-11	Context caching is segmented by organization and process. Cache segments are not shared across distinct organizations.
EXE-12	A request rate limit per minute is enforced, applied per user and per organization.

14. Grounding and hallucination mitigation rules

Hallucination mitigation is the most critical control in any application based on language models. In an audit or financial process context, fabricated information can lead to an incorrect decision with direct impact. The rules below are cumulative.

ID	Rule
ANC-01	The system prompt contains an explicit rule prohibiting the presentation of any numerical value not derived from the data.
ANC-02	Every answer is grounded in data retrieved during execution. Where the data does not exist or falls outside the configured scope, the assistant states that it does not know and does not improvise.
ANC-03	Visualizations reference the data actually retrieved. No illustrative series or placeholder value is generated.
ANC-04	Where enriched content is produced, the data is injected by code, which prevents the model from introducing a value at the presentation stage.
ANC-05	The user can request the sample records supporting any figure presented.
ANC-06	No regulatory or financial decision is executed solely on the basis of model output, without automated or human validation. Assistant output constitutes a recommendation.
ANC-07	Where confidence is insufficient, the system explicitly signals that the information requires human verification. Uncertain content is not presented as fact.
ANC-08	Regulatory precedence: applicable sector regulation and legislation in force override any generative output of the model.
ANC-09	Context and routine configuration cannot reduce or remove the grounding requirement, the uncertainty signal, or the display of how the answer was built.

15. Isolation, access and data security rules

ID	Rule
ISO-01	Each organization has an isolated space, with a dedicated database. Under no circumstances is data from one organization shared with or made visible to another.
ISO-02	Logical isolation is enforced through an organization identifier, with token validation on every request and propagation of the organization, process and user role context throughout the execution chain.
ISO-03	Identity is federated through OpenID Connect, with support for corporate directories, multi factor authentication and single sign on where the customer has the integration configured.
ISO-04	Access control is role and group based, administered through a dedicated security center, with access policies and exportable audit records.
ISO-05	Row level security is enforced across the entire query path, restricting which records the engine may read when producing an answer. The control does not depend on how the question is phrased.
ISO-06	Access precedence rule: the contracted capability defines what exists, the group defines who can see it, and row level security defines which records the engine may read.
ISO-07	The Admin Center defines which processes and information each user can access and is available only to administrator profiles within the customer organization.
ISO-08	Encryption of data at rest with AES-256 and in transit with TLS 1.2 or above. Key management with automatic rotation and policy based access control.
ISO-09	Every change to a group, an access policy or a row level security rule is recorded in an auditable and exportable trail.

16. Privacy and personal data processing rules

ID	Rule
PRI-01	Customer data is not used to train or fine tune language models, whether proprietary or third party. Any exception depends on formal and specific customer authorization, with prior assessment by the Data Protection Officer.
PRI-02	The model provider application programming interface is configured for zero data retention and no use for training, with corresponding contractual provisions.
PRI-03	Prompts are not retained after the request has been processed, except for internal audit records, whose retention is controlled under the UpFlux log retention policy.
PRI-04	Each request is processed in an isolated context, with no sharing across organizations.
PRI-05	Before content is sent to the language model, context sanitization is applied, with removal or masking of personal data and health data, together with minimization of the context to what is strictly necessary for the task.
PRI-06	Cloud hosting in the us-east-2 region, in the United States, and the use of a language model provider domiciled abroad constitute an international transfer of personal data under articles 33 to 36 of the LGPD and under the International Data Transfer Regulation approved by Resolution No. 19/2024 of the Brazilian National Data Protection Authority. UpFlux states this condition expressly in its security questionnaires and in its Record of Processing Activities. The safeguards applied are detailed in section 16.1.
PRI-07	The list of sub-processors involved in the processing is made available contractually and kept up to date by the Data Protection Officer.
PRI-08	Data subjects have the right to review of automated decisions under article 20 of the LGPD. Assistant output constitutes a recommendation subject to validation, which preserves the human decision point.
PRI-09	Anonymization and pseudonymization are configurable capabilities, applied according to the design agreed with each customer, and can be performed at source before data is sent to the platform.
PRI-10	The customer controls partial and full deletions within its dedicated database, as well as the backup and recovery policy.
PRI-11	Any feature that processes personal data is subject to a Data Protection Impact Assessment approved by the Data Protection Officer before it is made available.
PRI-12	All assistant output passes through data loss prevention filters before presentation. Output that fails these filters is blocked and routed for review, with the block recorded.

16.1 Safeguards applied to international transfers

The safeguards below are applied cumulatively to every international transfer arising from the operation of Nous.

Safeguard	Description
Transfer instrument	Incorporation of the standard contractual clauses approved by the Brazilian National Data Protection Authority, in the form of Annex II to Resolution No. 19/2024, into the instruments executed with each sub-processor that receives data abroad.
Data processing agreement	Executed with each sub-processor, with obligations of confidentiality, purpose limitation, information security, incident notification, and a prohibition on further processing incompatible with the stated purpose.
Zero retention and no training	The model provider application programming interface is configured for zero data retention and no use of content for training, with corresponding contractual provisions.
Encryption across the full path	TLS 1.2 or above in transit and AES-256 at rest, including the leg of communication with the model provider.

Safeguard	Description
Minimization and sanitization	Only the context strictly necessary for the task is sent, with removal or masking of personal data and health data before it reaches the model.
Prior sub-processor assessment	Review of the provider data use policy, processing location, retention, subcontracting and security controls before engagement, with periodic reassessment.
Record keeping and transparency	The operation is recorded in the Record of Processing Activities, with purpose, destination country, data categories and legal basis. The full text of the transfer clauses is made available to any data subject who requests it, within 15 days, subject to trade and industrial secrets.
Data subject channel	Requests relating to international transfers are received through the Data Protection Officer channel and answered within statutory deadlines.

17. Logging and audit rules

ID	Rule
AUD-01	The security center provides access audit records that the customer organization can export.
AUD-02	Routines retain an execution history, with status, automatic or manual origin, and duration.
AUD-03	The application maintains structured audit logging and performance observability, with tracing of model executions and quality and consumption metrics.
AUD-04	Log retention and immutability follow the UpFlux log retention policy, with protection against improper alteration and availability for audits and investigations.
AUD-05	Audit records relating to authentication and data maintenance can be consumed through an application programming interface, supporting traceability within the customer environment.

18. Operational limits and environment requirements

18.1 Operational limits

ID	Rule
LIM-01	Usage is tied to the package contracted by the organization and monitored through the Usage menu.
LIM-02	A request rate limit per minute is enforced in the execution engine.
LIM-03	A routine is paused automatically after three consecutive failures.
LIM-04	Minimum automatic refresh interval for a View: 15 seconds.

18.2 Minimum requirements for the web console

Category	Minimum	Recommended
Desktop browser	Chrome 90 or above is recommended, for better performance when rendering diagrams and charts. Firefox 88 or above is supported, with lower graphics rendering performance. Edge 90 or above is supported. Safari 14 or above has limited support due to WebGL constraints. Internet Explorer is not supported.	Chrome 90 or above
Mobile browser	iOS Safari 14 or above, Chrome Android 90 or above and Samsung Internet 13 or above are supported. Firefox Mobile 88 or above has limited support.	Chrome Android or iOS Safari, current version

Category	Minimum	Recommended
Processor	i3 or Ryzen 3, with two cores or more	i5 or Ryzen 5, with four cores or more. For high data volume, i7 or Ryzen 7 with eight cores or more
Memory	4 GB	8 GB. For high data volume, 16 GB
Graphics	WebGL 1.0	WebGL 2.0 with 1 GB of video memory. For high data volume, a dedicated GPU with WebGL 2.0
Resolution	1024 by 768	1920 by 1080. For high data volume, 2560 by 1440
Connectivity	HTTPS access on port 443 to UpFlux endpoints. There is no VPN dependency for the end user.	A stable corporate connection

19. Change management and versioning

ID	Rule
CIC-01	System prompts and context retrieval configurations are versioned artifacts, with integrity validation in the pipeline before deployment. Every change goes through review and approval.
CIC-02	The assistant backend maintains an automated test suite, including tests specific to the security of the isolated code execution environment.
CIC-03	Each major release undergoes prompt injection testing with known payloads and a red team exercise before release.
CIC-04	Controlled deployment, with gradual activation, rollback capability and intensive monitoring in the subsequent period.
CIC-05	Any change that alters the model provider or model family, introduces a new tool available to the model, or introduces a new data delivery channel is submitted to the Information Security Committee and is reflected in an update to this document.
CIC-06	Decommissioning an integration with a model provider requires contractual confirmation that the data has been purged at the provider.

20. Consolidated hard constraints

The constraints below cannot be worked around through configuration and must be taken into account in any solution design built on Nous.

- The conversation target is immutable after the first message.
- Only playbook analyses can be promoted to a View.
- Commands exist only where they have been published in the process configuration.
- A routine is paused automatically after three consecutive failures.
- A routine requires at least one delivery channel.
- Conversational context is not shared across separate conversations.
- Data querying by the model is read only and accepts select statements only.
- Code execution by the model has no network access and no file system access.
- Internet Explorer is not supported.

21. Compliance, references and contact

Applicable legislation and regulation

The Brazilian General Data Protection Law (LGPD), Law No. 13,709/2018, with particular attention to article 20, on the review of automated decisions, and to articles 33 to 36, on international transfers. The

International Data Transfer Regulation, approved by Resolution No. 19/2024 of the Brazilian National Data Protection Authority (ANPD). Sector regulation applicable to each customer context, including that of the Brazilian National Health Agency (ANS) and the TISS and TUSS standards in the healthcare sector.

Methodological references

The frameworks below are used as references in the definition and assessment of the controls described in this document, and not as a statement of certified conformity.

- OWASP Top 10 for LLM Applications.
- ISO/IEC 42001:2023, artificial intelligence management system.
- ISO/IEC 27001, information security controls.
- NIST AI Risk Management Framework 1.0.
- MITRE ATLAS.

On certifications. As at the date of this version, UpFlux does not hold formal certification or an independent third party audit report against ISO/IEC 27001, ISO/IEC 42001 or SOC 2. The frameworks cited above are adopted as methodological references in the design of the controls. Independent verification of the security program is carried out through semiannual penetration tests performed by an external specialist firm, whose executive summary is made available to customers under a non-disclosure agreement.

The set of internal Information Security, privacy and artificial intelligence governance policies underpinning this document is made available to customers and prospects under a non-disclosure agreement.

Contact for privacy and data protection matters: dpo@upflux.net

22. Revision history

Version	Date	Type	Description
1.0	May 25, 2026	Issue	Initial publication of the technical rules of use for UpFlux Nous.