Version: November 2026

GEM CORTEX AI

POLICY-AS-CODE AND AI GATEWAY

Commercial and technical product specification

Policy-as-Code and AI Gateway form a shared control layer for AI agents and applications within the GEM Cortex AI offering. Policy-as-Code converts approved rules into a form that can be versioned, tested and evaluated at runtime. AI Gateway and connected control points enforce these decisions. Within the agreed scope, the organisation can therefore control which data AI may use, where it may send that data and which actions it may perform.

Shared rules have their own lifecycle, independent of individual agents. Changes to limits or permissions are managed in one place and propagated to connected applications. Each decision records the version used, the reason and a link to the actual outcome.

The offering combines platform components, integrations, rule formalisation and GEM services for implementation and operation. The specific technologies, connected systems, models, level of autonomy and service parameters are confirmed in the solution design and the contractual scope of delivery. The technical variant described in this specification is a reference design to be validated for the specific deployment.

1. *Executive summary and position within the GEM offering*

The solution is intended for organisations seeking to extend the use of AI in operational processes while retaining accountability for access, approvals, costs and outcomes. The shared control layer connects technical control with the owners of business rules.

The relevant process, security or compliance owner defines the rule. The technical team converts it into an executable form, tests it and connects it to control points. An agent may prepare supporting information and propose a course of action, but authorisation for a protected action is based on approved rules and the current verified context.

- ***Consistent rules across connected agents and applications, with a recorded owner, version and effective date.***
- ***Controlled model calls based on identity, purpose, data profile and budget.***
- ***Control of business actions before execution, with a clearly defined handover to a human wherever approval is required.***
- ***A traceable sequence from the request through the decision to the actual outcome.***
- ***Safe rule changes through review, testing, phased rollout and rollback to a previous version.***

GEM SYSTEM

GEM CORTEX AI
– POLICY-AS-CODE A AI GATEWAY

Position within the product family

GEM Cortex AI

Capability block	Capability description
1 | Cortex Core	*Orchestration of intent, runtime path, model and tenant context.*
2 | Cortex Access	*SSO, roles, workload identity, provider accounts, projects and lifecycle management.*
3 | Cortex Knowledge	*Connectors, ingestion, ACL-aware retrieval, citations and index lifecycle management.*
4 | Cortex Flow	*Workflows, tools, approvals, human tasks and process state.*
5 | Cortex Control	*Policies, AI registry, telemetry, evaluations, budgets and FinOps.*
6 | Cortex Workspace	*A user channel for approved chat, knowledge and work tasks.*

Model Execution forms a separate architectural layer for model adapters, local inference, routing and fallback.

The solution can connect to the client's existing agents, applications and integration infrastructure. This overview shows how it relates to the product family; a specific delivery includes only the selected capabilities. The block names do not imply separate licences or a requirement to implement the entire product family.

2. Client problem and opportunity qualification

In an initial AI pilot, rules may be embedded in code, prompts or local configuration. As the portfolio grows, changes require increasing effort: teams identify affected applications, interpret the same requirements in different ways and test each deployment separately. The organisation also finds it difficult to demonstrate the basis on which a particular agent acted.

A second problem is fragmented access to models and tools. There is no shared link between identity, purpose, data and costs. An agent may have broader permissions than its process requires, or a direct route may exist that bypasses control. Policy-as-Code determines whether a decision is permissible; integrations and access controls ensure that it is actually enforced.

A suitable client is an enterprise with recurring decisions, shared rules and available integration points. Typical environments include financial services, insurance, document-intensive organisations, public services and enterprise IT. What matters is the process impact and the scope of control, rather than the sector itself or whether an application is labelled an agent.

Capability block	Capability description	Measurable evidence
1 | CIO and CTO	*Reuse of controls and less effort to implement changes across applications.*	*Time required to change a rule and number of application modifications.*
2 | CISO and security	*Control of data paths and permissions in the connected environment.*	*Control coverage, negative tests and exception records.*
3 | Management and compliance	*Demonstrable accountability and visibility into AI use.*	*Rule owners and traceability of a specific decision.*
4 | Process owner	*Preserved decision-making authority and predictable approvals.*	*Correct escalations, processing time and accepted outcomes.*
5 | CFO and procurement	*Transparent costs for the platform and individual processes.*	*Unit cost and realised benefit against the baseline.*

The initial opportunity must have an owner, documented rules, representative cases and a point at which an action can be stopped. A simpler profile may be sufficient for creative assistance or a one-off experiment. If an existing workflow or an already licensed service meets the same need safely and cost-effectively, GEM will propose using it or supplementing it with shared controls.

3. Client problem and opportunity qualification

Policy-as-Code separates the question of what is permitted under given conditions from how an agent performs its work. An executable rule may assess, for example, roles, financial limits, approval status, time conditions, permitted models or data classification. The generative model prepares supporting information; a rule evaluates the precise limit or permission using structured inputs.

The same rule version applied to the same complete context should return the same decision. This property facilitates testing and case reconstruction. The accuracy of an input value and the correctness of the conversion from a directive remain subject to separate controls. Information extracted from a document may therefore be insufficient evidence to authorise a write operation.

Two agent control profiles

The reference design distinguishes between an assistance profile and a process profile according to autonomy and impact. These are alternative ways of integrating the control layer, and their suitability is validated for each process.

Profil	Capability description	Measurable evidence
1 | Assistance profile	*Summarisation, text drafting or knowledge work without autonomously modifying a system..*	*Every model call passes through AI Gateway, with identity, data and consumption checks; a human assesses the result.*
2 | Process profile	*The agent prepares multiple steps and may propose a protected action.*	*Plan constraints, checks on model calls and tools, a fresh decision before the action and human approval where required.*

The full process profile is selected according to impact and the need for demonstrable evidence. Timers, queues, waiting and workflow state remain within process orchestration. The policy layer evaluates the conditions supplied by the orchestrator and trusted sources.

Enforcement must take place on every relevant path to a model, data or a protected action within the agreed scope. A gateway for model calls alone will not prevent a write operation in an ERP system or the use of a public AI service outside the connected environment. The production design therefore also maps network paths, service accounts and keys that could be used to bypass control.

Changing a rule does not require modifying the agent if the existing integration provides all the necessary attributes and can enforce the decision. A new action type, a new input field or a different approval process may require an interface change. Defining this boundary is also part of the delivery.

GEM Cortex AI

GEM Cortex AI

Defining the specific delivery

The scope of the offering identifies the connected agents and applications, selected rules, model endpoints, business actions and control points. Integrations, approval workflows and operational services are confirmed separately. Further processes are introduced on the basis of pilot results and an approved scope extension.

PaC uses enterprise identity, context and the results of other controls. IAM, content guardrails, security scanners and orchestration each have their own responsibilities. Replacing all of them is outside the scope of the PaC principle described here.

4. Reference architecture

The technical model in Sections 4 to 6 describes a reference variant to be validated for the specific deployment. The architecture, selection of delivered functions and solution scope are confirmed during discovery and delivery design.

The reference architecture separates policy management from evaluation of an individual request. The control plane manages rule sources, approval, testing and publication. The operational plane handles requests, obtains context, makes decisions and enforces the result.

Role	Meaning	Responsibility
1 | PAP	*Policy administration*	*Authoring, metadata, review, testing, approval and version distribution.*
2 | PDP	*Policy decision point*	*Evaluate the approved rule set against structured context.*
3 | PEP	*Policy enforcement point*	*Intercept the request and actually permit, block or impose conditions on its execution.*
4 | PIP	*Context source*	*Provide trusted roles, limits, classifications and the current case state.*
5 | Audit layer	*Records and oversight*	*Link the decision, version used, fulfilment of conditions and actual outcome.*

AI Gateway acts as the PEP for model calls. For business actions, this role is performed by an API gateway, an integration service, a workflow, a pipeline or the target system itself. The enforcement point is selected according to where the operation can be reliably stopped.

The reference technical variant uses Open Policy Agent as the PDP and Rego for executable rules. Rules and tests are authored in Git, pass through CI and are distributed to the operational environment as approved, signed bundles. Selection of the specific technology and supported versions forms part of the solution design.

In the reference variant, Git and the artefact registry are not synchronous dependencies of every decision. The runtime decision point uses a verified version that has already been loaded. The operational design defines its validity and the behaviour when other components are unavailable.

Integrations connect to the enterprise directory, approval workflows, model endpoints, data catalogues and operational monitoring. The client's existing API gateway and infrastructure can be used. A REST API or a tool protocol such as MCP provides an interface; controls over identity, permissions and the impact of an action must be provided separately.

5. Rule management and safe changes

Rule management is based on business ownership, versioning, review, testing and safe rollout. The reference design develops this approach into the lifecycle below. The deployment design specifies the metadata, tools and publication method.

Each rule has an owner and a link to an approved source. Records connect the directive version, effective date, criticality, approvals and technical implementation so that the origin of a decision can be demonstrated.

Conditions that can be evaluated by a machine are separated from matters requiring professional judgement. AI may assist with text analysis or drafting a rule, but a human remains the accountable author. An ambiguous directive must undergo substantive clarification before it is converted into an executable policy.



Policy lifecycle

- *The owner confirms the source, interpretation, effective date and specific examples of the required behaviour.*
- *The author prepares the rule, metadata and tests for permitted, prohibited, boundary-case and incomplete inputs.*
- *The change undergoes business and technical review. Critical rules require approval by a person other than the author.*
- *CI validates the schema and the composed policy set. Known cases are replayed against the new version, and every change in the expected result is explained.*
- *The approved release is produced as a versioned artefact with a signature, hash and traceable provenance.*
- *The new version is validated in shadow mode or within a limited group. It begins enforcing the new conditions only after approval.*
- *Operations monitors the impact of the change and the active versions. The lifecycle includes recertification, expiry, rollback to the previous release and retirement of the rule.*

Shadow mode compares the decisions and impacts of the new version without enabling its blocking behaviour. Existing controls remain in force. Rolling back the rule version restores the previous configuration; it does not automatically reverse business actions that have already taken place.

GEM Cortex AI

GEM Cortex AI

Rule composition and selection

Shared organisational rules are composed with domain and process rules. The hierarchy explicitly defines namespaces and the method for aggregating results. Baseline safeguards must form part of every applicable decision, and an exception must not silently disable them.

In the reference variant, the registry and verified context within the policy layer determine the correct set. The agent does not select the rule. If a required policy is missing, the request is denied; the control is not silently skipped.

The reference design provides for verification of the artefact's signature and scope before activation. It includes protection of signing keys, tracking of active versions, recertification and time-limited exceptions. These mechanisms are validated for the selected deployment.

6. Reference approach to plan and action control

The process profile applies a check before the agent starts work and a fresh assessment before a protected effect occurs. Between these points, the case state, permissions, amount or validity of supporting information may change. Permission to start a process therefore does not authorise all subsequent actions.

PLAN defines the scope of work

Before work begins, the policy determines the permitted checks, tools, templates and limits. The agent receives a defined framework within which it can prepare supporting information. For example, the plan may specify which documents to assess, which fields to extract and when further work serves no purpose. This can limit unnecessary model operations before they take place.

Each subsequent model call undergoes its own check in AI Gateway. Tool use is subject to the tool's permissions and current rules. Orchestration is responsible for process execution, waiting and aggregation of results.

DECIDE assesses the specific action

Immediately before a protected action, the decision point evaluates the check results and current context. A technical PASS result for a document may be insufficient to permit a write operation: authorisation, a valid time window or approval may still be missing. Incomplete or indeterminate information must not be silently treated as a fulfilled condition.

The output is a permit, a denial or a permit subject to mandatory conditions. Such a condition, referred to as an obligation, may require human approval, a lower limit or the use of another permitted model. The PEP executes the operation only once the required condition has been fulfilled.

Decision content and audit links

The reference design uses a subject, action, object and context. The decision includes an identifier, reason, policy version and any obligations. The specific interface schema is defined in the integration design.

Contract element	Required content
1 | Subject	*Verified user or workload, delegation and relevant permissions.*
2 | Action and object	*Specific operation, parameters, target system and case or resource identifier.*
3 | Context	*Purpose, environment, data classification, limits, current state and relevant time.*
4 | Decision	*Result, reason code, explanation, decision ID, policy versions and mandatory conditions.*
5 | Execution	*Evidence that conditions have been fulfilled, operation status and the result identifier in the target system.*

GEM Cortex AI

Trusted roles, limits and approvals are obtained from enterprise sources. The decision made before an action evaluates the current context, which may have changed since the process began.

Where an obligation requires human approval, the design provides for the action to be stopped and linked to the decision in the audit trail. The approval workflow remains the responsibility of the process team. Tool integrations are subject to requirements for permissions, limits, previews, approval, idempotency and audit.

By default, the pilot covers read operations or the preparation of proposals. Irreversible steps and those with legal significance remain under human control. Extending autonomy requires a separate decision for the specific process and validated control points.

7. AI Gateway data and cost control

AI Gateway brings model traffic from connected applications under common control. It assesses each call in the context of identity, purpose, data profile, permitted endpoint and cost accountability. The model registry records the owner, region, permitted data classification, contractual arrangements and operational parameters.

Call authorisation and content checks serve different purposes. The policy determines whether a request may proceed. Content checks, often referred to as guardrails, assess inputs and outputs, for example for sensitive content or compliance with a permitted format. Routing selects an approved model, while consumption metering provides the basis for FinOps.

- ***A check before transmission verifies whether the data path is permitted. Classification, sensitive-data detection or redaction is added according to the requirements.***
- ***Routing and fallback use only endpoints permitted for the relevant identity, purpose and data profile.***
- ***Limits constrain request size, frequency, runtime, token consumption and budget allocation.***
- ***Metering attributes actual consumption to an application, process, project or cost centre and links it to the completed task.***
- ***A structured denial explains the reason and enables the application to supplement or defer the request, or hand it over to a human, in a controlled manner.***

A hard budget stop applies to consumption passing through the controlled gateway. The estimate before a call and the final charge may differ; cost records are therefore reconciled with provider data. Specific limits and their enforcement are validated in the design and pilot. Model or context optimisation is accepted only while maintaining the required quality.

Data protection includes checking access permissions for sources before assembling model context. Checking a response cannot retrospectively prevent an unauthorised input from being sent. Zero retention, region and non-use of data for training are confirmed for the specific service and contractual configuration.

Public chat services, desktop tools and SaaS assistants require a separate assessment of integration options and enterprise controls. For developer copilots, the output can be checked in a pull request and in CI; this does not demonstrate control over all communication between that tool and its provider.

8. Audit security and operational responsibility

The audit trail links the plan, model calls, decisions made before an action, fulfilment of conditions and actual outcome. A record of permission does not confirm successful execution. This must be evidenced by the target tool or enterprise system. Unambiguous correlation makes it possible to establish where a process was stopped and why.

An audit record links the time, identity, process, action, context relevant to the decision, decision ID, rule version, reason and enforcement status. The design determines which inputs are retained, whether replay is possible, the retention period and access permissions for the records, according to data classification and the purpose of the control.

Standard telemetry does not automatically include complete prompts or documents. Operational, audit and forensic access are separated by role. Records require protection against unauthorised modification, a defined export mechanism and completeness checks that include denials and errors. The records support internal audit and regulatory assessment; the client confirms the applicability of legal requirements and responsibility for their interpretation.

GEM Cortex AI

Security requirements

The design assesses PEP bypass, spoofed identity or context, compromised rule bundles, audit tampering and unavailability of the decision service. It addresses prompt injection and tool misuse risks separately. Access to keys, repositories and policy changes follows least-privilege principles and separation of duties.

For protected actions, the reference design provides for a safe stop when no decision is available and for use of the last verified policy within a defined operating mode. Outage behaviour, policy validity, any caches and recovery are assessed and tested for the specific operating environment.

Operating model and responsibilities

The following allocation is a reference model of responsibilities. Responsibilities and service commitments are confirmed for the selected delivery.

Area	Client responsibility	GEM responsibility within the agreed scope
1 | Rules and process	*Business interpretation, limits, ownership and approval of autonomy.*	*Formalisation, implementation and technical tests.*
2 | Data and identities	*Classification, purpose, permissions and access to sources.*	*Integration and correct transfer of context.*
3| Platform	*Selection of the operating model and provision of client-side dependencies.*	*Application components, configuration and agreed administration.*
4| Acceptance	*Confirmation of correct outcomes and acceptance of the process.*	*Test evidence, defect resolution and handover.*
5| Incidents and changes	*Process escalations and approval of business changes.*	*Technical diagnostics, recovery and version management.*

GEM CORTEX AI
– *POLICY-AS-CODE A AI GATEWAY*

The infrastructure and model endpoints are operated by the party designated in the contract. The solution may run in the client's cloud, as a GEM service, in a hybrid arrangement or in a restricted local environment. Each option has its own data flows, capacity requirements and responsibilities.

The service design defines availability, latency, capacity and the RTO and RPO recovery parameters according to process criticality. Within the agreed operational scope, it includes monitoring, backups, recovery, version rollback and incident procedures. Performance is validated against a real rule set and the expected workload.

Are your AI agents under control?

Which rule governed the agent's decision?

Where can company data be sent?

Can you substantiate the decision to an auditor?

Policy-as-Code and AI Gateway help enforce rules and substantiate decisions.

GEM Cortex AI

GEM System a. s. | Štětkova 1638/18, 140 00 Praha 4 – Nusle, Czech Republic | sales@gemsystem.cz | www.gemsystem.cz

9. Use cases and initial sales scenarios

The following scenarios illustrate applications of the shared control layer. The initial process, its integration and acceptance metrics are defined during qualification and in the Pilot Charter.

Document and financial processes

AI prepares structured data from documents and checks that supporting materials are present. The policy assesses precise conditions, such as document validity, a signature, a time threshold or a financial limit. An incomplete case is referred to a specialist. A write operation or financial action takes place in the controlled system after the necessary approval.

A document process with a limited scope is selected for the first pilot. Preparation time, handover to a human, rule-based decisions and the audit trail are validated. The benefit lies in faster preparation and consistent control, while retaining accountability for the decision.

AI-assisted development and deployment

A developer or agent prepares a change. The pull request and CI pipeline check tests, security scanner results, secrets in code and approved dependencies. The policy evaluates these results and the conditions for the next step. A human approves the change, and production deployment follows an authorised process.

The pilot can be limited to a selected repository and pipeline. It validates the blocking of changes in accordance with the rules and the traceability of approvals. The agent has access only to authorised repositories and environments. Policy-as-Code uses scanner results as inputs.

Vulnerabilities and remediation management

The agent triages findings and prepares proposed tasks. The policy takes account of asset criticality, finding severity, internal deadlines and escalation rules. A production change requires an approved procedure. A successful remediation recommendation is distinct from authorisation to implement it.

The pilot selects a limited area of vulnerability and remediation management. It validates decisions against deadlines and escalation rules, together with their audit linkage. Shared rules make it easier to apply a new internal methodology across connected teams.

Data protection and model access

An application requests a summary or analysis of internal content. The gateway assesses purpose, classification and the permitted endpoint. The knowledge layer verifies the user's permissions before passing on the content. If no suitable route exists, the request is stopped or referred for an approved assessment.

The pilot selects a connected application and data scope. It validates model-call coverage, permitted routing, fallback and cost metering. Where the scope includes content checks, their effectiveness is validated separately.

Customer processes and commercial authority

AI prepares a proposed response, discount or complaint resolution. Rules verify decision-making authority, limits and case status. A binding commitment or contract amendment is referred to an authorised person and carried out as a controlled step in the relevant system.

The pilot selects a specific decision-making situation and its approval process. It validates compliance with decision-making authority, limits and escalation requirements. The preparatory output and the final commercial decision remain separately traceable.

GEM Cortex AI

Public services and sensitive administrative processes

AI třídí podání a připravuje evidenční údaje nebo kontrolní seznam. Pravidla řídí přístupy, předání mezi rolemi a podmínky workflow. Odborné posouzení a právně významné rozhodnutí zůstávají odpovědnému pracovníkovi.

Pilot pracuje s jednou agendou a reprezentativní sadou případů. Ověřuje oprávnění, chybné zatřídění, předání pracovníkovi a auditní vazby. AI nečiní meritorní rozhodnutí.

10. Delivery implementation and acceptance

GEM starts with a specific process and a verifiable value hypothesis. Discovery maps rules, existing access arrangements, data paths, accounts and every point at which a protected effect may occur. It identifies which controls can be reused from existing systems and where integration is required.

Phase	Work content	Client deliverable
1 | Qualification and discovery	*Process, owner, baseline, rules, risks and integration options.*	*Definition of the opportunity and recommended scope.*
2 | Design and pilot	*Decision contract, priority rules, historical case set and shadow mode.*	*Pilot Charter and evidence of functionality and impact.*
3| Limited production	*Identities, enforcement, approvals, monitoring, security and recovery.*	*Acceptance and operational handover.*
4| Development and operation	*Further processes, recertification, FinOps and controlled changes.*	*A portfolio with measured results and reuse of components.*

The Pilot Charter specifies inputs and outputs, what is in scope and out of scope, owners, control points, the test set, costs and criteria for a go, adjust or stop decision. Representative cases also include boundary cases and prohibited situations. The tuning set and independent validation set are kept separate.

The pilot validates a selected rule from approval through to its use at connected control points, the correctness of decisions, the audit trail and operational impact. The process profile adds control of the specific action and any required human approval.

Acceptance within the agreed scope

Validation areas are selected according to the pilot profile and architecture. Specific scenarios, expected results and numerical thresholds are approved in the Pilot Charter.

GEM Cortex AI

Validation area	Evidence for acceptance
1 | Rule-based decisions	*Permitted, prohibited and boundary cases produce the expected result.*
2 | Policy change and rollback	*Review, tests, publication and traceability of the version used; validated rollback.*
3 | Model path	*Calls from the connected application pass through the designated gateway and permitted models.*
4 | Process actions and approvals	*The PEP enforces the decision, and the action waits for a human where approval is required.*
5 | Outages and recovery	*Behaviour conforms to the agreed safe operating mode and recovery arrangements.*
6 | Audit trail	*The identity, rule version, reason and outcome of a selected case can be traced.*
7 | Operations and economics	*Measured latency, consumption, effort and impact against the baseline.*

Negative tests verify that no unauthorised action occurs within the agreed test set. Passing the test set is neither a general guarantee of error-free future operation nor automatic compliance certification.

The handover package is defined according to the delivery phase and scope. It includes the solution definition, architecture and data flows, security and test evidence, operational responsibilities, economics and agreed exports for service termination.

11. *Economics commercial offering and benefit measurement*

The economics distinguish between the shared platform and the introduction of individual processes. The platform has initial and recurring costs, while additional processes can reuse established identities, rule distribution, the gateway and audit capabilities. Process-specific rules, their approval and validation nevertheless require their own effort.

The offering separates discovery, platform implementation, integrations and process onboarding from model consumption, infrastructure, support and ongoing rule management. Pricing depends primarily on the number of paths requiring control, documentation quality, availability of identities, approval scope and the required service level. A binding cost estimate is prepared after the scope has been qualified.

Phase	Work content	Client deliverable
1 | Rule changes	*Change frequency, effort and number of affected applications.*	*A shared rule actually replaces repeated work.*
2 | Testing and oversight	*Time spent on regression testing, corrections and operational checks.*	*The required quality and validation coverage are maintained.*
3| Audit preparation	*Time required to retrieve and compile supporting evidence.*	*Records are complete and fit for the intended purpose.*
4| Model consumption	*Calls, tokens and cost per accepted result.*	*Optimisation preserves quality and adoption.*
5| Process throughput	*Case processing time, rework and completed volume.*	*The organisation makes use of the capacity released.*

Time savings are recorded as capacity until their use or a reduction in actual costs has been demonstrated. Risk reduction is expressed in a separate scenario. A more appropriate plan, context or routing can reduce model consumption; introducing PaC alone does not guarantee a general percentage saving.

The business case compares benefits over several years with implementation, onboarding and recurring TCO. It uses the expected portfolio, adoption and realistic rollout dates. The pilot is intended to validate individual assumptions and refine the cost estimate. The economics of a proposal for a particular client do not automatically transfer to another deployment.

12. Sales communication and next step

Core message

GEM Cortex AI provides an organisation with shared rules for model use and control of AI agent actions. Rules have owners, versions and tests. Connected gateways and enterprise control points apply them before a specific operation, and the outcome can be substantiated. The client can develop further AI processes on top of the established control layer.

Brief introduction to the offering

When a company rule changes, its different implementations across AI applications create both work and risk. GEM connects your agents to shared rule management and a controlled gateway to models. Precise limits and permissions are evaluated before an action, changes undergo testing and high-risk steps are referred to a human. We start with one process and validate control, operational impact and economics.

Addressing common objections

An existing agent or workflow can be retained if it offers suitable integration points. Discovery identifies rules that need a shared lifecycle and maps the required context. Phased integration reduces the extent of change and enables the new decision-making approach to be compared with the current state.

Local configuration is not inherently a problem. What matters is ownership, testing, publication and enforcement of the rule. B2B value arises where unified management eliminates repeated work and provides demonstrable control across multiple applications.

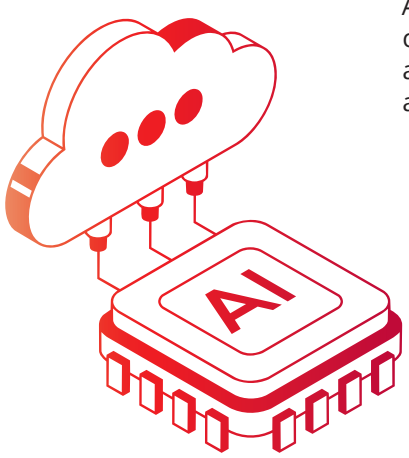

A policy engine evaluates rules. A production delivery adds context, control points, approvals, safe publication and operational accountability. This scope distinguishes a deployed platform service from a standalone library or demonstration prototype.

GEM Cortex AI

GEM SYSTEM

GEM CORTEX AI
– POLICY-AS-CODE A AI GATEWAY

Questions for the first meeting

- *Which process and decision should the solution improve, and who is accountable for them?*
- *Which rules are repeated across applications, and how much effort does changing them currently require?*
- *Which data, models and enterprise tools does the process use, and who manages access?*
- *Where can an actual action be stopped, and which steps require human approval?*
- *For a specific case, can the version used, the reason and the actual outcome be traced?*
- *What evidence must the pilot provide for the client to decide on production implementation?*

A GEM AI architect and the process owner prepare a limited pilot with specific rules, control points and measurement against the baseline. The output is an evidence-based decision on implementing and extending the solution.

One rule change applies to all

Rules separated from AI agent code

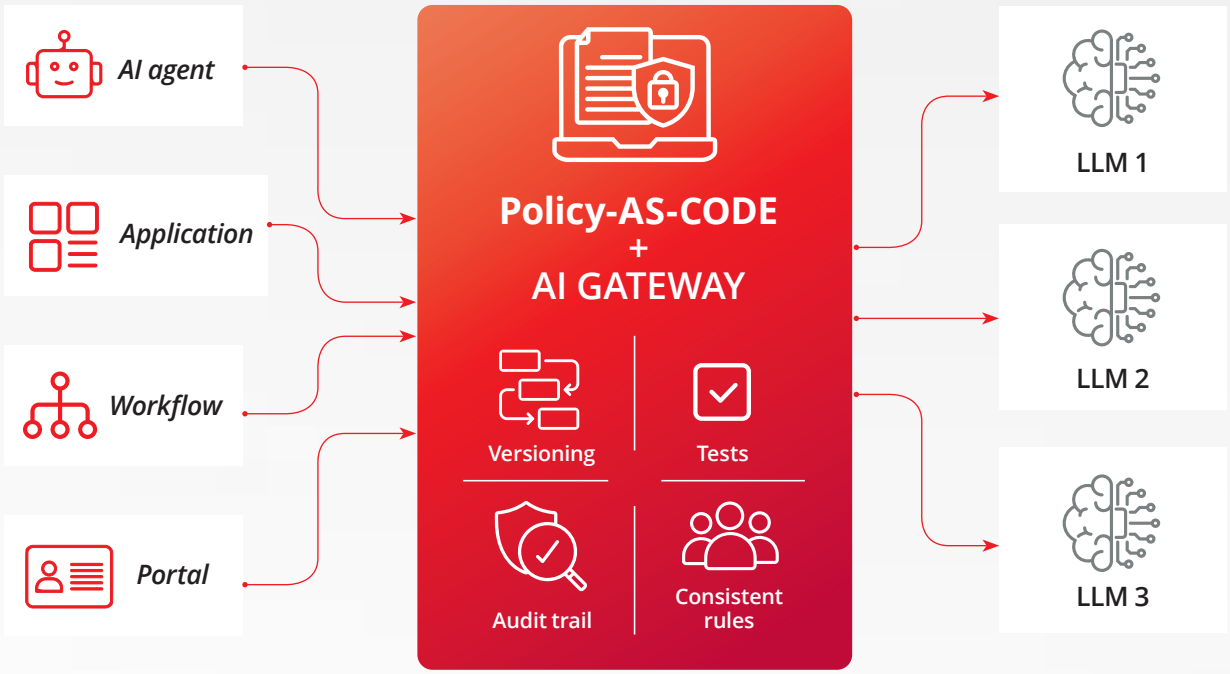

GEM Cortex AI

GEM System a. s. | Štětkova 1638/18, 140 00 Praha 4 – Nusle, Czech Republic | sales@gemsystem.cz | www.gemsystem.cz