

TRUST CENTRE

Responsible AI & AI-Assisted Features Policy

GOVERNANCE, CUSTOMER DATA, HUMAN OVERSIGHT, TRANSPARENCY AND SAFE USE

Effective Date: 6 October 2025

Classification: Public

Document ID: PCL-TRUST-AI-014

This document is intended for public distribution. Feature-specific AI architecture, provider configurations, prompts, model settings, security-sensitive implementation details and internal evaluation evidence are restricted and may be provided separately to eligible customers, auditors or regulators subject to appropriate access and confidentiality restrictions.

Contents

This policy should be read together with the PCL One Information Security Practices, Data Processing Addendum, Subprocessor & Third-Party Services Policy, Data Residency & Location Policy, Acceptable Use Policy, Secure Software Development Lifecycle Policy, applicable Service Description and the customer Ordering Document or Service Schedule.

01	Overview
02	Scope and applicability
03	Responsible AI principles
04	Human oversight and decision responsibility
05	AI-assisted feature use
06	Customer data and model training
07	AI providers and subprocessors
08	Inputs, outputs and confidentiality
09	Accuracy, limitations and verification
10	High-impact and sensitive use
11	Security, misuse and abuse prevention
12	Transparency, records and explainability
13	AI feature lifecycle and change management
14	Customer responsibilities
15	Assurance, exceptions and contractual precedence
16	Statement of changes

1. Overview

PCL One may provide AI-assisted features that help users analyze information, summarize content, identify patterns, generate suggestions, support workflows or improve productivity. This policy establishes the public baseline for how PCL One governs those capabilities and how Customers should use them.

Public baseline

AI-assisted features are intended to support authorized users, not replace accountable human judgment. Customer Data is not used by PCL One to train generalized or shared AI models for unrelated customers or public use unless the Customer expressly agrees to a different arrangement in writing.

2. Scope and applicability

This policy applies to AI-assisted capabilities provided by PCL One where this policy is incorporated into the applicable contract, Service Description or Trust Centre documentation. It applies whether a capability uses PCL One-developed logic, a cloud AI service, a third-party model or a combination of these components.

- Feature-specific behavior, provider use, supported data types and limitations may differ by service and are defined in the applicable Service Description, Order, Service Schedule or product documentation.
- This policy does not require a Customer to enable optional AI-assisted functionality unless the applicable service is expressly ordered or configured to use it.
- Where mandatory law or a negotiated contract imposes a stricter requirement, that requirement controls for the relevant Customer and use case.

3. Responsible AI principles

Principle	Public baseline
Human accountability	AI supports people and business processes; accountable users remain responsible for consequential decisions and approvals.
Purpose limitation	AI features are used for the documented service purpose and are not repurposed in a manner inconsistent with the applicable contract or privacy terms.
Data protection	Customer Data used by an AI feature remains subject to the applicable Agreement, DPA, security controls and data-location commitments.
Transparency	PCL One identifies AI-assisted functionality where reasonable and provides feature documentation appropriate to the use case.
Security and resilience	AI components are subject to risk-based security, access, monitoring and change-management controls appropriate to the service.
Quality and review	Outputs may require validation. PCL One does not represent AI-generated content as infallible, complete or a substitute for professional judgment.

4. Human oversight and decision responsibility

AI-assisted outputs may contain errors, omissions, outdated information, bias or context that is incomplete. Customers must apply appropriate human review before relying on outputs for material business, legal, financial, payroll, employment, eligibility, disciplinary, safety or regulatory decisions.

- Customers remain responsible for final decisions, approvals, communications and actions taken using AI-assisted output.
- Where a workflow presents recommendations, flags or classifications, authorized users should review the underlying record and available supporting information before taking consequential action.
- PCL One may design safeguards, confidence indicators, workflow approvals or other controls for selected features, but these controls do not eliminate the need for appropriate Customer governance.

5. AI-assisted feature use

AI-assisted functionality must be used only within the Customer's authorized service scope and for legitimate business or public-sector purposes. Customers should configure access so that only users with an appropriate role and business need can use AI features involving Customer Data.

- Do not intentionally use an AI feature to bypass access controls, approvals, segregation-of-duties rules, retention requirements or other service safeguards.
- Do not submit content that the applicable service, DPA, Order or feature documentation prohibits from being processed by the AI capability.
- AI features may be disabled, limited or made optional where supported by the applicable service configuration and contractual arrangement.

6. Customer data and model training

PCL One processes Customer Data used by AI-assisted features only to provide, secure, support and improve the contracted service as permitted by the applicable Agreement, DPA and service documentation.

No generalized training from Customer Data

PCL One does not use Customer Data to train generalized or shared AI models for unrelated customers or public use unless the Customer has expressly agreed to that use in writing. Feature telemetry or operational metadata may be handled separately where it does not constitute Customer Data or where its use is otherwise permitted by contract and law.

- Customer Data remains owned or controlled by the Customer in accordance with the applicable Agreement.
- Where an AI feature uses retrieval, prompts, context windows, embeddings or temporary processing, those data flows are governed by the applicable service architecture, DPA and subprocessor terms.
- Any retention of prompts, responses or related operational records depends on the feature and is subject to the applicable retention policy and service-specific configuration.

7. AI providers and subprocessors

Some AI-assisted capabilities may rely on approved third-party cloud or AI service providers. Where a provider processes Customer Data on PCL One's behalf, the provider is governed through the applicable subprocessor and data-protection framework.

- Provider applicability can vary by product, feature, deployment model and Customer configuration.
- The public Subprocessor & Third-Party Services Policy identifies approved providers and the general purpose for which they may be used.
- PCL One applies contractual and technical controls intended to limit provider processing to the authorized service purpose and applicable data-use settings.

8. Inputs, outputs and confidentiality

Customers should treat prompts, uploaded content, retrieved context and AI-generated outputs according to the sensitivity of the underlying information. Access to AI features and related records should follow least-privilege principles.

- Users should not place secrets, credentials or unnecessary sensitive information into an AI prompt or input when that information is not required for the task.
- AI-generated output may contain or infer information from Customer Data available to the authorized user; Customers should apply the same confidentiality controls to outputs as to the underlying source information.
- Outputs should not be assumed to be unique, confidential or free of third-party rights merely because they were generated by an AI capability.

9. Accuracy, limitations and verification

AI-assisted outputs are probabilistic or rules-assisted and may not always reflect the most current, complete or accurate information. PCL One may combine AI with deterministic rules, source citations, validation logic or business controls where appropriate, but output quality still depends on the feature, source data and user context.

Use pattern	Expected control
Summaries and drafting	User reviews the source material and verifies material facts before external or consequential use.

Use pattern	Expected control
Recommendations and anomaly flags	User evaluates the underlying records, business rules and context before approving action.
Compliance or regulatory assistance	Output is treated as decision support and is validated against authoritative requirements and Customer procedures.
Automated workflow assistance	Approvals, thresholds and exception handling remain subject to the configured business process and authorized human oversight.

10. High-impact and sensitive use

Unless expressly described and validated in the applicable service documentation, AI-assisted functionality is not intended to be the sole basis for decisions that materially affect a person's employment, compensation, benefits, legal rights, eligibility, access to essential services or other similarly consequential interests.

- Customers should apply heightened review, documented approval and appropriate domain expertise to high-impact use cases.
- Protected or sensitive characteristics should not be used to make unlawful or discriminatory decisions.
- Where a Customer configures a feature for a regulated or high-impact workflow, the Customer is responsible for confirming that the configuration, notices, approvals and human review satisfy its legal and policy obligations.

11. Security, misuse and abuse prevention

AI-assisted features are subject to the same security and acceptable-use expectations as other PCL One services. PCL One may apply technical safeguards intended to detect or limit malicious prompts, unauthorized access, excessive automated use, data exfiltration attempts or other activity that threatens service integrity.

- Users must not use AI features to develop, facilitate or materially assist malicious activity prohibited by the Acceptable Use Policy.
- Customers should protect API keys, service credentials and AI feature access paths in accordance with the applicable security controls.
- Suspected compromise, harmful output, data exposure or abuse involving an AI feature should be reported promptly through the designated support or security channel.

12. Transparency, records and explainability

PCL One aims to provide practical information about the purpose and operation of AI-assisted features at a level appropriate to the feature and Customer use case. Where a feature materially influences a workflow, PCL One may provide supporting context such as source references, reason codes, confidence information or audit records where technically supported.

- AI-related user actions and service events may be logged in accordance with the applicable service architecture, audit settings and retention rules.
- Feature documentation may identify important limitations, expected human review and known dependencies.
- Detailed model internals, proprietary prompts, security-sensitive configurations or third-party confidential information are not required to be publicly disclosed.

13. AI feature lifecycle and change management

AI models, providers, prompts, retrieval methods and safety controls can change over time. PCL One manages material AI feature changes through the applicable secure development, change-management, security and third-party review processes.

- Material feature changes are evaluated for security, privacy, data handling and expected user impact before production deployment.
- Provider or subprocessor changes are handled in accordance with the applicable Subprocessor & Third-Party Services Policy and contractual notice requirements.
- Feature-specific release timing, deprecation periods or migration requirements are governed by the applicable service documentation and Release, Change & Maintenance Policy.

14. Customer responsibilities

- Configure roles, permissions and approvals appropriate to the sensitivity and consequences of the AI-assisted use case.

- Provide lawful, accurate and appropriately minimized data to the feature and maintain required notices, consents or legal bases.
- Train users on feature limitations and require human verification where material decisions or external communications are involved.
- Validate Customer-created prompts, workflows, integrations and automation before production use and after material changes.
- Report suspected security, privacy, quality or safety issues so they can be investigated and addressed.

15. Assurance, exceptions and contractual precedence

AI-assisted features operate within PCL One's broader security, privacy, resilience, subprocessor, retention and support control environment. Assurance evidence may be provided subject to confidentiality, availability and the scope of the applicable audit or assessment.

- Customer-specific AI commitments, prohibited uses, data-location requirements or feature configurations may be documented in the applicable Order or Service Schedule.
- Exceptions to this policy require documented approval through the applicable PCL One governance process.
- If this policy conflicts with an executed Agreement, DPA, Order or Service Schedule, the governing contract controls according to its stated order of precedence.

16. Statement of changes

Effective date	Change
6 October 2025	Initial publication of this policy.