



Agentic Marketing Operations with *Responsible & Predictable AI*

How Pega Customer Engagement Studio
securely manages data, privacy, and AI.

— Pega Customer Engagement Studio™
Security & Privacy




Table of contents

01

What is Customer Engagement Studio (CES)?

04

Guardrails & Safety Controls

- Predictable Agent Execution
- External Agent Governance

07

Evaluation & Quality Assurance

02

CES Architecture

- Responsible AI

05

Monitoring & Observability

08

Human Control & Transparency

03

AI Governance

- How Pega Agentic works
- Data handling and flow

06

Access & Identity Management

09

Cloud Security



Pega Customer Engagement Studio™

Security & Privacy Summary

We built **Pega Customer Engagement Studio™ (CES)** to help marketers collaborate with AI while maintaining governance, transparency, and human control. CES is designed so organizations can accelerate marketing operations with agentic AI without compromising security, privacy, or compliance.



Governed AI from design through execution

- Every AI capability operates within Pega's ISO/IEC 42001-aligned AI governance framework and enterprise security standards.
- AI agents operate within defined permissions, business rules, approval workflows, and audit requirements.
- Human review and approval remain in control before changes can move into production.

Protected data and enterprise privacy

- CES agents work with configuration, policies, and performance analytics, not customer PII or individual customer records.
- Customer environments are isolated, preventing data sharing across tenants.
- Prompts, responses, and customer data are not used to train AI models.

Secure AI operations

- All AI requests are routed through the Pega AI Gateway Service, which provides authentication, security controls, content filtering, and observability.
- Prompt-injection detection, malicious-content protections, and provider-native guardrails help reduce AI risk.
- Every interaction, recommendation, and approval is auditable and traceable.

Enterprise-grade security

- Data is encrypted in transit and at rest using enterprise-grade security controls.
- Access is governed through role-based permissions, federated identity, and least-privilege principles.
- Continuous monitoring, compliance controls, disaster recovery, and operational safeguards.



Human approval remains in control



No AI training on customer data



Every action is governed and auditable




Enterprise-grade security built in

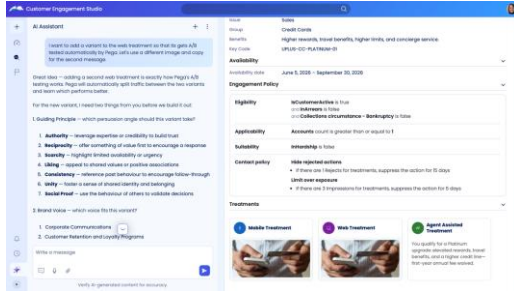
01

What is Customer Engagement Studio (CES)?

Pega's solution for agentic marketing operations



What is Customer Engagement Studio?



Customer Engagement Studio (CES) is Pega's new agentic user experience within Customer Decision Hub, where marketers and other users collaborate with an AI assistant to orchestrate the complete marketing operations lifecycle without sacrificing governance or human control.



Agents take-on repetitive tasks, marketers take-back strategy and the high-value decisions that truly require human judgment.

Receive actionable **insights and analyze performance** of campaigns to spot opportunities and resolve issues faster.



Governance foundation is built in from the ground up. Every capability is developed under strict security and ethical standards. Every output requires human validation. Every action is auditable.

Accelerate and de-risk your *agentic transformation* with Pega.



REALIZE YOUR AI POTENTIAL

Reimagine your customer engagement.

A free, browser-based AI-collaboration tool to brainstorm and visualize your 1:1 vision.

Pega Customer Engagement Blueprint™



MINIMIZE MARTECH WASTE

Orchestrate the full lifecycle of marketing.

Harness the power of agentic marketing with confidence and control.

Customer Engagement Studio
with Pega Customer Decision Hub™



RUN AGENTS PREDICTABLY

Orchestrate agents, people, & systems.


Drive predictable outcomes with predictable costs at scale.

Open Agentic Orchestration

All on a platform with security, trust, and governance built-in


Where does Customer Engagement Studio fit in?

Rapid design to jumpstart development




Customer Engagement Blueprint™

Business & Marketing collaboration tool to brainstorm and visualize on your 1:1 customer engagement vision.



Pega Infinity™ Pega Customer Engagement Studio

Collaborate with the AI assistant to create, curate and calibrate actions for the complete marketing operations lifecycle.



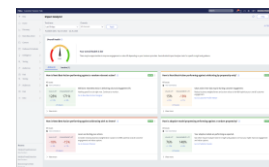
Customer Decision Hub™ 1:1 Operations Manager

Business and operations insights to manage change, monitor, troubleshoot, and optimize decisioning operations at scale.



Customer Decision Hub™ Next Best Action Designer

Provides a centralized interface to create, configure, and activate 1:1 Next Best Action programs, across channels.



Customer Decision Hub™ Simulation and Testing Capabilities

Simulate, test, and validate next-best-action strategies to reduce risk and improve business outcomes.

Reimagine



Build



Test & Deploy



Operate



Optimize

No PII *

No PII *

Potentially sensitive

Potentially sensitive

No PII *

Operated securely by Pega Cloud®

Private client deployment

Operated securely by Pega Cloud®

Private client deployment

Private client deployment

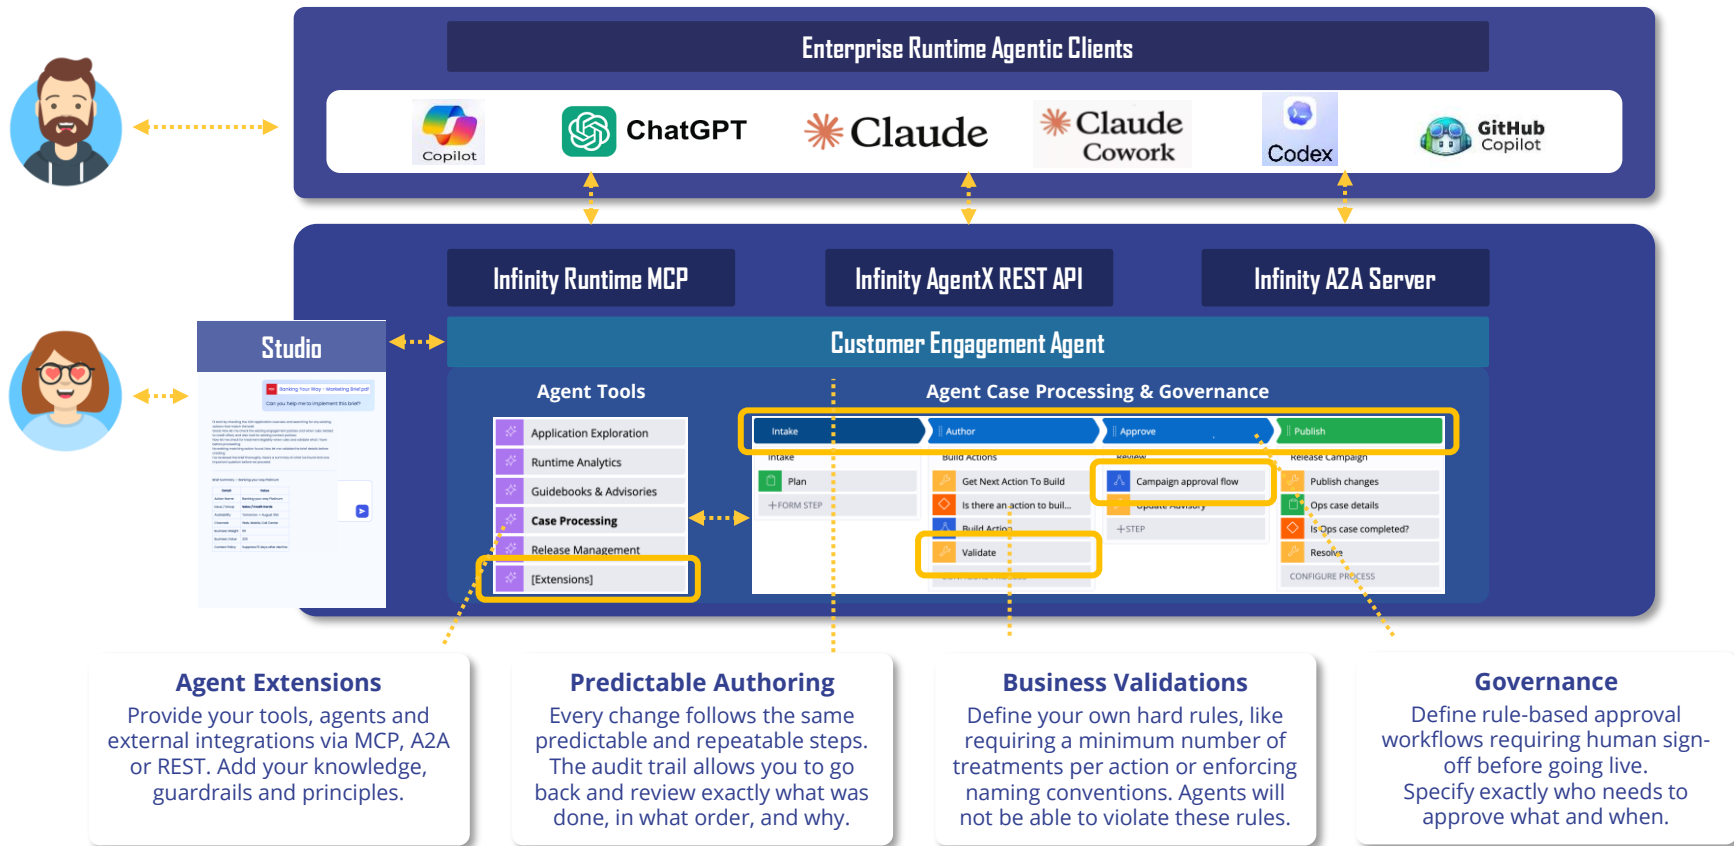
*Based on design time use cases, it is not advised to manage PII at these phases

02

Customer Engagement Studio Architecture




Customer Engagement Studio™ Architecture



Customer Engagement Studio™ Architecture

- CES Analyzes existing configurations, policies, analytics, and customer profile data through Pega 1:1 Operations Manager.
- It routes approved changes through existing change and release management processes before deployment to Production.
- The architecture separates customer engagement planning and orchestration from the Pega Customer Decision Hub™ production decisioning environment.



CES agents **cannot** access customer profile records or PII.


Agents can access:

- NBA configuration
- Customer profile configuration
- Analytics and interaction performance data

Agents reason over system configuration and performance, not individual customers.

Available at: docs.pega.com

Responsible AI at Scale



Who will win in the age of Agentic AI?


It's not organizations who deploy the most agents. It's the ones who govern them well enough to keep them running, earning trust with every decision.



Customer Engagement Studio was built with governance and trust at its core

- ✓ Guardrail
- ✓ Evaluation
- ✓ Human control
- ✓ Audit trail

These aren't boxes to check; they are the foundation of trustworthy AI.



You're not just getting AI capability

You're getting governance, transparency, and accountability built in.

You're getting a partner committed to your success and your responsibility to the customers and communities you serve.

03

AI Governance



AI Governance

At Pega

Every model, feature, and release is governed by strict security and compliance standards

- CES operates under **ISO/IEC 42001 -certified AI Management System**.
- Before any release, CES undergoes scenario-based evaluations and internal red-team exercises to identify robustness issues and potential risks.
- The product is validated across realistic, adversarial, and failure scenarios to ensure it behaves safely.
- AI Risk Register tracks and manages AI-related risks
- Annual internal and external audits validate compliance and governance controls

Pega's AI Governance Board

- Oversees all **AI utilization** across the platform
- AI Steering Committee oversees AI capabilities across Pega
- Comprised of experts from Product, Cloud Security, Cloud Operations, Legal, and Go-to-Market.
- Conducts continuous **AI security assessments**
 - Microsoft AI Red Team Methodology
 - OpenAI Safety Best Practices
 - OWASP Top 10 for LLM Applications
 - OWASP Cloud-Native Application Security Top 10

Pega ensures AI agents do not act outside their authorized scope

- **Tool Rules** define what agents are permitted to do.
- **Authentication** profiles secure agent connections.
- **1:1 Operations Manager** governance layer enforces planning, review, and approval before changes go live.

AI Governance at Pega

End-to-end oversight

Pega's AI Governance board is run by the Cloud Security team and oversees all AI utilization across Pega's products.

It brings together experts & owners from Product, Cloud Security, Cloud Operations, IT, Legal, & Go-to-market to ensure all utilization of AI in Pega is safe, responsible, secure.

Strategic Partnerships

In order to deliver on the unique needs of its enterprise clients, Pega has formed strategic relationships & overarching agreements with AWS, Google Cloud, and Microsoft to drive shared AI initiatives.

Pega & its cloud providers of LLM services meet regularly to review model options, performance, security, & issues.

Security-first

Pega's AI Governance board organizes & executes continuous security assessments of all AI powered capabilities.


Security assessments run include:

1. ISO 42001
2. Microsoft AI Red Team Methodology
3. OpenAI Safety Best Practices
4. Microsoft required mitigations
5. OWASP Top 10 for LLM Applications
6. OWASP Cloud-Native Application Security Top 10




How Pega Agentic Works

The agent connects to Pega's knowledge base, business rules, data, and case actions. And to third party tools and other AI systems. Everything flows through Pega LLM Gateway Service maintaining enterprise security and governance.



Data handling and flow



Data handling and retention

- Prompts and completions are not used for model training
- Pega does not retain prompt or completion content
- Model providers do not retain prompts or completions
- Human review services from model providers are disabled
- Models operate in a stateless manner for each request

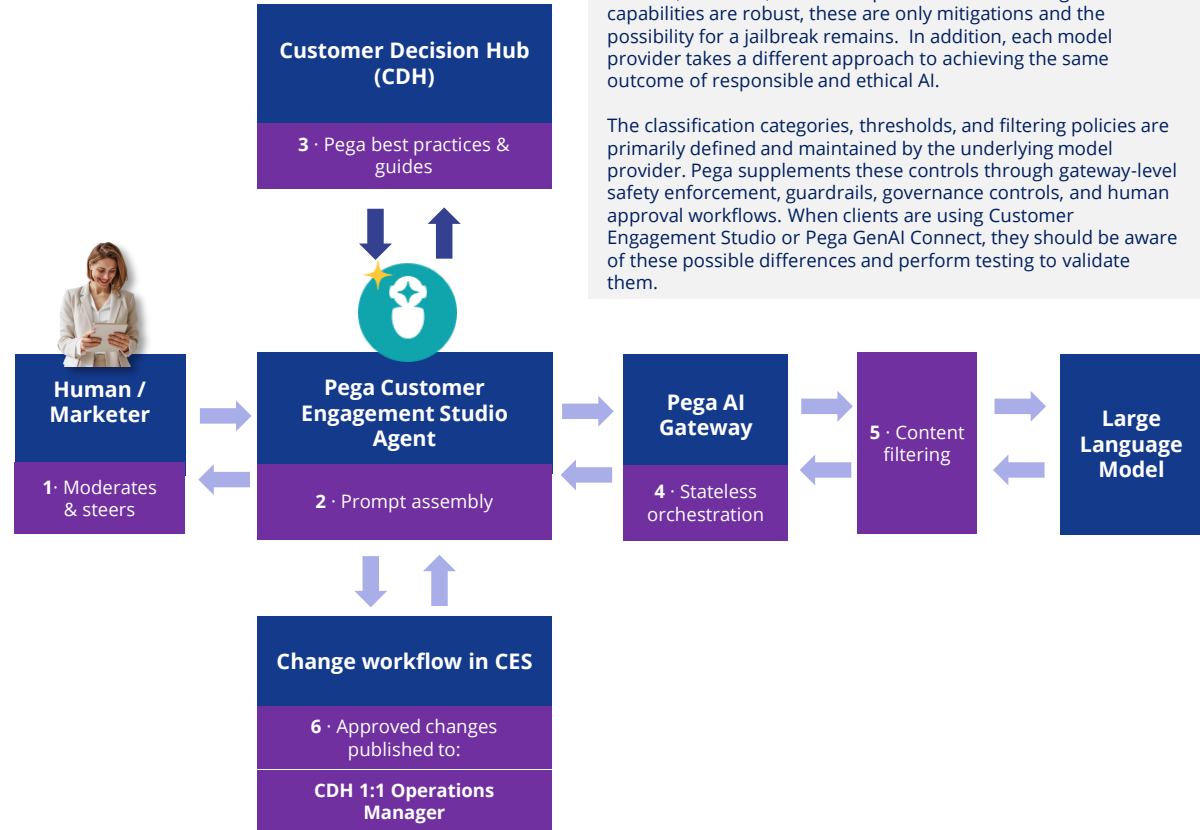


Data flow

- All LLM traffic is routed through the Pega AI Gateway Service
- Gateway performs authentication, routing, observability, and safety enforcement
- Prompt injection and malicious detection enabled by default
- Requests are blocked if safety controls cannot be executed ("fail closed")

Customer Engagement Studio AI Data Flow

- Pega **Customer Engagement Studio Agent** translates the user (marketer)'s intent into prompts
- The agent calls out to **Customer Decision Hub** when necessary to enrich the prompt with information about the application, Pega marketing best practices, and guides
- All **LLM** calls are brokered by the **Pega AI Gateway service** on Pega Cloud. This service provides a trusted layer of security and scalability for communication with large language model providers
- **Content filtering** applies to detect and prevent harmful content in both the prompt inputs and completion outputs
- Any **changes** to Customer Decision Hub artefacts are applied via a change workflow that ensures predictability and control. The changes are moved to 1:1 Operations Manager after they have been approved by a reviewer
- The **human** moderates the conversation, and steers the agent towards the goal, when required



Content filtering approach

Pega relies on the most proven Large Language Model Providers in its delivery of Pega capabilities. Within each model, are robust content filtering capabilities that will mitigate the possibilities of harmful, unethical, or toxic responses from occurring. While the capabilities are robust, these are only mitigations and the possibility for a jailbreak remains. In addition, each model provider takes a different approach to achieving the same outcome of responsible and ethical AI.

The classification categories, thresholds, and filtering policies are primarily defined and maintained by the underlying model provider. Pega supplements these controls through gateway-level safety enforcement, guardrails, governance controls, and human approval workflows. When clients are using Customer Engagement Studio or Pega GenAI Connect, they should be aware of these possible differences and perform testing to validate them.

04

Guardrails and Safety Controls



Guardrails and Safety Controls




Input Guardrails

- Every input prompt is subject to **pre-processing guardrails** before it reaches the AI model.
- CES validates and sanitizes inputs, detects prompt attacks, and redacts personal or sensitive data using **Amazon Bedrock Guardrails**.
- Jailbreak and data leakage scenarios have been performed by Pega.
- This multi-layer approach prevents bad data from triggering bad decisions.
- Malicious detection
- Provider-native safety controls
- Requests blocked when safety checks fail



Output Guardrails

- CES blocks harmful, sensitive, and irrelevant outputs through provider-side filtering combined with Pega's own **deterministic validation**.
- All outputs are constrained to defined purposes and cannot autonomously drive production change.
- Treatments, personas, and change requests all require human acceptance.
- **CES does not directly execute customer communications.**
- Execution and decisioning remain within CDH and 1:1 Operations Manager.



Operational Restrictions

- The **guidebook**, maintained by Pega, serves as the agentic integration layer between CES and CDH.
- Guidebook ensures that agents follow proven CDH best practices and structured guidance to deliver accurate, reliable, and governed outcomes. Agents don't rely on generic LLM knowledge.
- CES operates strictly within its approved authority. The system cannot exceed, modify, or self-expand its scope of permissions.
- If a model response is unsuitable, the associated tool call fails, and the workflow stops rather than proceeding unchecked.



Role-Based Access Controls

- Access to sensitive data is **explicitly authorized and limited** to approved use cases.
- CES enforces least-privilege principles, ensuring users can only access the data and tools required for their specific role.
- Access is regularly reviewed and revoked when no longer needed.

Predictable Agent Execution

Agents cannot violate business-defined rules

- **Structural validations** ensure that every change is technically sound. E.g., all required fields are present, all references exist, and the change is consistent with the system. An agent can never produce a structurally invalid change.
- **Business validations** allow your organization to define its own hard rules. CES enforces these automatically. Examples include naming conventions and treatment requirements. Agents cannot make changes that violate your business rules.

Configured approval workflows cannot be bypassed

- **Human approvals** give businesses control over who reviews what.
- Define rule-based approval workflows so that certain types of changes always require human sign-off before going live.
- Specify exactly who needs to approve what. Agents cannot skip or bypass configured approvals.

Changes follow a defined, repeatable process

- **Structured authoring** ensures that every change follows a defined sequence of steps based on the case, regardless of who or what initiates it.
- CES guides the agent through the authoring process step by step, making it easy for a human to follow and oversee what is happening.
- CES can make the same change ten times, it will follow the same steps every time.
- This consistency reduces cognitive load on the humans working alongside agents.

Every interaction is recorded and reviewable

- **Audit trail** every conversation between a user or external agent and the CES Agent is captured and recorded as part of the case.
- Backtrack and understand exactly what was done, in what order, and why.
- Client data remains isolated through dedicated **tenant boundaries**.
- CES prevents data sharing across customers and maintains environment-specific logging and audit records.

External Agent Governance (MCP)

CES supports MCP, A2A, and AgentX to allow external agents integration

- **External agents** can participate in workflows but cannot bypass configured governance controls
- **MCP (Model Context Protocol)** provides a standardized interface for AI agent integration while remaining an open-source, provider-agnostic protocol that allows agents built on different LLMs and frameworks to connect.
- MCP is not enabled out of the box.
- MCP requires the Pega MCP Server to be configured and accessible to external agents.

CES acts as the controlled gateway for all external MCP integrations

- Any MCP-compatible agent **connects through CES** rather than directly to CDH and 1:1 operations manager.
- MCP integrations inherit CES governance controls, including validations, approval workflows, and change management processes, regardless of the connecting agent.
- Users cannot bypass CES governance by using MCP.
- The guidebook is built on the MCP, which is an industry standard for connecting AI agents to enterprise systems.
- All agent interactions and decisions are recorded in auditable case histories.

Organizations control which MCP tool categories are exposed to external agents

CES MCP capabilities are split in 4 categories:

- Exploration Tools
- Guidebook Tools
- Change Management Tools
- Release Management Tools


For example, an organization may permit exploration access only while prohibiting release-management access. Enabling **least-privilege access** for external agents.

- Change management tools do not bypass configured validations or approvals.
- All AI-initiated changes remain subject to enterprise change controls and human approval processes.

External agents can automate significant portions of the lifecycle

- Pega's recommended approach keeps **humans in the loop** at key governance gates.
- Even highly automated workflows follow structured approval and review processes before changes go live in CDH.
- CES remains the governance and orchestration layer where humans can review, approve and override work.
- MCP guidebooks help reduce inconsistent behavior by providing governed, curated knowledge and standardized operating procedures.

Agentic Marketing Operations with Predictable AI Agents



05

Monitoring & Observability



Monitoring and Observability




Comprehensive Traceability

- CES captures **end-to-end visibility** into how every decision was formed. This includes request/response metadata, policy decisions, tool calls, latency metrics, failure rates, token usage, and interaction history.
- Supports investigations, audits, and continuous improvement.
- **Agent to agent (A2A)** interactions are auditable and compliant. Full history is retained. All changes remain subject to validation and approval processes which the agent cannot bypass.



Behavioural Monitoring


- CES continuously monitors for unusual behavior, deviations from baseline, anomalies, and performance trends.
- **Behavioral baselines** are established for each deployment, and any deviation triggers investigation.



Alerting & Response

- Threshold breaches, anomalous behavior, model drift, and policy violations trigger **automated alerts**.
- These feed into human review queues where compliance and quality teams can investigate and respond. No issue goes unnoticed.

Monitoring and Observability



Conversation retention

- **Conversation history** is retained within the customer's CES environment
- Retention follows customer-defined retention policies
- **Audit history** supports investigations, governance reviews, and compliance requirements



Drift Detection & Mitigation

- **Drift monitoring** ensures the system continues to perform as intended over time.
- CES monitors data inputs and model outputs against defined baselines and thresholds.
- Any performance degradation is identified, root cause is assessed, and corrective actions are implemented.

06

Access & Identity Management



Access and Identity Management




Unique System Identity

- CES operates under a unique, dedicated system identity with a clearly assigned owner responsible for access, operation, and entitlements throughout its lifecycle.
- Shared, generic, or embedded credentials are not permitted.
- The **system identity** is centrally managed, attributable, and auditable.



Secure Communications


- **TLS encryption** protects data in transit
- Authentication uses short-lived access tokens
- Environment identity is attached to requests for isolation and traceability



Federated Identity & Integration

- CES is integrated with enterprise identity providers and uses federated identity for authorization. This ensures all operations can be traced back to the individual or system responsible for them.
- CES supports protected agent-to-agent (A2A) communications using **OAuth 2.0 client registration** to securely authenticate external agents and establish a secure connection before interacting with CES.

Access and Identity Management



Access Provisioning & De-provisioning

- All access to data and systems must enforce least-privilege and require approval by the relevant service owner.
- Access is subject to periodic review and is revoked when no longer needed.
- **Identity lifecycle management** ensures rapid removal of access in the event of an incident.



Privileged Access Management

- Enterprise **Privileged Access Management (PAM) solution** - All privileged access is managed here.
- Privileged access is time-limited, requires multi-factor authentication, and is granted only following appropriate approvals and audit.

07

Evaluation & Quality Assurance



Evaluation & Quality Assurance




Scenario-Based Testing

- Before release, CES undergoes approximately 80 scenario-based evaluations.
- These combine deterministic ground-truth checks with behavioral metrics, covering **realistic, adversarial, and failure scenarios**.
- Every scenario is designed to **stress-test** the system and identify breaking points.




Quality & Accuracy Benchmarks

- CES must meet defined accuracy (factual correctness) and precision (relevance) **benchmarks**.
- Adherence to these benchmarks is continuously monitored to ensure the system detects and addresses toxicity, hallucination, and bias.



Human Evaluation & Scoring

- CES evaluations include human scoring for quality and accuracy.
- Subject matter experts assess outputs against business objectives and customer expectations.
- This **human-in-the-loop** evaluation ensures the system learns from real-world usage.



Failure Event Management

- Failure scenarios are explicitly defined and managed.
- The system captures **what went wrong, why**, and what safeguards prevented harm.
- When failures occur, they're treated as learning opportunities.
- This intelligence feeds back into continuous improvement.

08

Human Control & Transparency



Human Control & Transparency



Human-in-the-Loop by design

- No AI recommendation automatically becomes an action.
- All CES outputs are recommendations that require explicit human approval before they drive any change.
- Marketing teams review, evaluate, and decide. AI informs. **Humans decide.**



Feedback & Iteration

- Users provide feedback on CES recommendations. What worked, what didn't, what seemed off.
- This feedback is captured, interpreted, and used to continuously refine and improve the system.
- Feedback is not lost in a black box; it **drives measurable improvements.**



Transparency & Disclosure

- Organizations using CES are equipped with **clear guidance** on what decisions are being influenced by AI, where AI is being used, and what those limitations are.
- This transparency extends to customers where appropriate, ensuring trust is earned, not assumed.




Contestability & Control

- Individuals affected by CES recommendations can **challenge, contest, or request a review.**
- Organizations maintain clear processes for investigation and remediation. This isn't just compliance; it's about building trust.

Human-led trigger

What does that human triggered process look like?




Agent-led trigger

- Agent and human partnership drive resolution



HUMAN LED

PERFORMANCE AGENT IDENTIFIES GAP




AGENT LED

SELF DIAGNOSTICS

Performance agent determines:

- Too many decisions return nothing
- Filtering rules too narrow
- Not enough actions available



AGENT LED

EXAMINE PROBLEM

Business Guidance agent determines:

- Engagement policies of eligibility and suitability pass
- Applicability disqualifies too many customers
- All constraint rules pass



Engagement Policies

- ✓ Eligibility
- ✓ Applicability
- ✓ Suitability

Constraints

- ✓ Channel Limits
- ✓ Contact Policies
- ✓ Action Limits

AGENT LED

EXPLAIN GAP. SUGGEST SOLUTION.

Optimization Agent:

- Recommends relaxing applicability rule
- Suggests changes to criteria
- Sends for human approval



Outcome of loosening applicability policy:

15% more of the population will see these actions.

REVIEW GAP AND AGENT-SUGGESTED CHANGE. DECIDE NEXT STEP.

Human receives alert summary:

- Applicability issue explained
- Proposed solution(s) displayed
- Choose from available options



☒ Yes, run through Compliance Checking Agent to ensure policy adherence.

☐ No, do nothing.

☐ Suggest another change.

09

Cloud security





Visit the Pega Cloud Trust Center to learn more

Pega Customer Engagement Studio™ Cloud Security

Customer Engagement Studio runs on Pega's proven Pega Cloud® services, ensuring enterprise-grade security.

Transformation you can rely on.

- 24/7 operations monitoring, management, and support
- Secure-by-design architecture operations with strict access controls and operational safeguards-minimizing human touch through automation
- Enterprise-grade compliance, uptime, disaster recovery, and threat modeling

Operations

24/7 monitoring, environment support, & proactive response

[Details](#)

Access

Environment governed by automated operational controls and strict access protocols

[Details](#)

Compliance

Strict adherence to 20+ industry standards

[Details](#)

Disaster recovery

Comprehensive data & service back-up, failover, restoration

[Details](#)

Threat modeling

Follow red team methodology based on OWASP top 10

[Details](#)

Availability

Architecture leverages built-in high availability and disaster recovery to support near continuous uptime.

[Details](#)

The screenshot shows the Pega Cloud Trust Center website. At the top, there's a navigation bar with links for Home, Platform, Solutions, Customers, Learn, Services & Partners, Events, and About. Below this is a hero section with the title "Pega Trust Center" and a subtitle "Secure. Reliable. Compliant. Pega Cloud empowers the world's biggest brands to meet - and exceed - the challenges of today and tomorrow. Learn how." There's a link to "Learn more about Pega's security features." Below this is a section titled "Security" with a subtitle "Our security policies provide a framework for safeguarding against unauthorized access and preventing/mitigating attacks that compromise performance and availability." There are four sub-sections: "Authorization & access", "Network protection", "Secure system integration", and "Data encryption". Each sub-section has a brief description and a link to "View security features". Below this is a section titled "Supporting security documents" with a table of resources. The table has three columns: "Resources", "Last updated (YYYY-MM-DD)", and "Assessment scope". The table lists documents like "Digital Security Policies", "Veracode Statement", "Penetration Test Summaries", "Business Continuity Plan Summary", and "Disaster Recovery Test Results". Below this is a section titled "Privacy" with a subtitle "Use our services to enable you to implement your own privacy and compliance strategies. We continually evolve our platform to provide the features and security measures that you may use to support your security and privacy strategy." There's a link to "Read Pega's privacy notice". Below this is a section titled "Compliance certifications, attestations, and accessibility" with a subtitle "When evaluating the services listed under each compliance standard it should be noted that Pega relies on a common set of controls for the purposes of adherence. These common controls exist across the Pega Platform, the underlying infrastructure, and the operations, administration and management provided by Pega in Pega Cloud. Pega applications deployed within the Pega Platform inherit these controls which are attested to in the current scope." Below this is a section titled "Pega Cloud certifications" with a list of certifications: APRA, CS, CSA STAR, Cyber Essentials, Cyber Essentials Plus, CyberGIX, and Cybervalds. Each certification has a dropdown arrow next to it.

Resources	Last updated (YYYY-MM-DD)	Assessment scope
Digital Security Policies	2024-12-27	Pega Cloud AWS, GCP
Veracode Statement	N/A	Pega Cloud AWS, GCP
Penetration Test Summaries	2024-09-17	Pega Cloud AWS, GCP
Business Continuity Plan Summary	2024-05-03	Pega Cloud AWS, GCP
Disaster Recovery Test Results	2024-12-09	Pega Cloud AWS, GCP

Compliance certifications, attestations, and accessibility	
APRA	▼
CS	▼
CSA STAR	▼
Cyber Essentials	▼
Cyber Essentials Plus	▼
CyberGIX	▼
Cybervalds	▼

Learn more at:

- [CES FAQ \(Internal Only\)](#)
- [CES FAQ \(Client -friendly\)](#)



Pega is the leading Enterprise Transformation Company™ that helps organizations Build for Change® with enterprise AI decisioning and workflow automation. Many of the world's most influential businesses rely on our platform to solve their most pressing challenges, from personalizing engagement to automating service to streamlining operations. Since 1983, we've built our scalable and flexible architecture to help enterprises meet today's customer demands while continuously transforming for tomorrow. For more information on Pega (NASDAQ: PEGA), visit <http://www.pegas.com>