

AI Transparency Card Structure

Generative AI features are currently available as [Preview](#) & Beta features in Esri products; however, we are including transparency cards for AI features now for early awareness and transparency for our customers to make responsible AI decisions with our products. Esri welcomes your feedback via SoftwareSecurity@Esri.com as we expand our AI Transparency card coverage across our product features over time. Individual ArcGIS AI feature Transparency Cards may be found within the [ArcGIS Trust Center Trusted AI section](#).

This document provides detailed description information for each AI Transparency card field.

Section	Description
Product - Name	ArcGIS product name (links to doc) Example: ArcGIS Online , <i>Business Analyst</i> , ...
Product - Certification	Security and/or privacy certification status of the ArcGIS Product. Select One: <ul style="list-style-type: none"> - None - 3rd Party Tested (No certification in place, however a 3rd party has tested product) - In-progress (Provide timeline and certification name (e.g., ISO, SOC, FedRAMP)) - Certified (name of certification and scope)
Product - Deployment	Deployment model of the product Select One: <ul style="list-style-type: none"> - Software as a Service (SaaS) - Platform as a Service (PaaS) - Commercial off-the-shelf (COTS) Self-contained - Commercial off-the-shelf (COTS) External Service
All Sections from this point onward are specific to the AI Feature.	
Name	AI feature name in the product (links to doc) Example: Business Analyst Assistant , ...
Purpose*	Actions AI feature is expected to perform within the product. Includes use cases and scenarios if/when applicable.
Release Status	Release status of AI feature (See Product Lifecycle Support Policy) Select One: <ul style="list-style-type: none"> - Private Beta (Limited access, unsupported) - Beta (Incomplete, unsupported, limited privacy, limited security – See Beta FAQ) - Public Beta (Customers may opt-out) - Preview (No opt-out, limited security and privacy) - General Availability (List initial release Month/Year or version)
Certification	Certification status of AI feature or its subprocessors. Select One: <ul style="list-style-type: none"> - None - In-progress (timeline and certification name (e.g., ISO, SOC, FedRAMP)) - Certified (certification name and scope)

Section	Description
Deployment	<p>AI feature provided via what deployment model.</p> <p>Select One:</p> <ul style="list-style-type: none"> - Software as a Service (SaaS) - Platform as a Service (PaaS) - Commercial off-the-shelf (COTS) Self-contained - Commercial off-the-shelf (COTS) External Service
Management	<p>How AI feature can be enabled or disabled. Esri typically provides generative AI Features as an Opt-in option that can be controlled by your organization's administrator.</p> <p>Select One or More:</p> <ul style="list-style-type: none"> - Opt-in by Administrator (AI feature not made available to users unless the administrator enables it via the ArcGIS Online Allow AI Assistants interface) - Opt-in by Install (AI feature not installed by default) - No Opt-in or Opt-out (Rare)
Management – Feedback	<p>How user AI feedback be enabled or disabled.</p> <p>Select One:</p> <ul style="list-style-type: none"> - Opt-in by Administrator (Feedback not available as user option by default) - Opt-in by User (When feedback option is used by user, prompt/results may be stored) - Not utilized
Management - Telemetry	<p>How AI telemetry data collection can be enabled or disabled.</p> <p>Provide a Separate Answer for Each Telemetry Source Below:</p> <ul style="list-style-type: none"> - EUEI (list configuration choice from below) - Feature Specific (list configuration choice) - Subprocessors (list configuration choice) <p>Configuration Choices:</p> <ul style="list-style-type: none"> - Admin Config (Administrator can turn on/off for organization – ex. EUEI) - User Config (Each user can turn on/off) - Required (Telemetry data is collected) - None (No telemetry data collected) <p>Note: EUEI default for ArcGIS Online is opt-in for EU-based organizations and opt-out for US.</p>
Prompt Stored / Response Stored	<p>Are prompts or responses submitted to the AI stored?</p> <p>Select One:</p> <ul style="list-style-type: none"> - No (Not stored) - Not by default* (User can select an option such as feedback resulting in storage) - Yes, Required* (Stored) - Yes, Opt-Out Available* (List how to opt-out of storage) <p><i>*If other than No is selected, specify retention time and purpose of storage (options below):</i></p> <ul style="list-style-type: none"> - Retention time (# of months stored by Esri) - Storage Purpose (Select one or more): <ul style="list-style-type: none"> - User Recall (User can select previous prompts to modify and re-execute) - Specific Improvement: Used to improve responses for the specific customer - General Improvement: Used for AI experience across all customers
Personal Data	<p>Is personal data present in the training, testing, or validation datasets?</p> <p>Select One:</p> <ul style="list-style-type: none"> - No - Yes (Provide details on where/how personal data is utilized)

Section	Description
Processing Location	<p>Where data is processed by the product, feature, and LLM, including subprocessors.</p> <p>Note: Many SaaS-based capabilities from Esri are hosted within ArcGIS Online cloud-based infrastructure and may be offered only out of the US region, the location of such capability is listed as US AGO Infrastructure and as Global AGO Infrastructure if deployed across all ArcGIS Online regions around the globe.</p> <p>Provide a Separate Answer for Each Location Below:</p> <ul style="list-style-type: none"> - Product: (US or Global) Product processes data locally or leverages external infrastructure for processing. - Feature: (US or Global) AI feature processes data locally, on-premises, or via cloud-based services. Specify subprocessors involved in this processing - LLM: (US or Global) Large Language Model processes data (eg. locally, cloud, or hybrid). Detail any subprocessors involved in model operations (eg. used for hosting or executing model)
Intended Users	<p>Primary intended users of the AI feature</p> <p>Examples: <i>Administrator, GIS Analyst, Public, etc.</i></p>
Out of Scope Uses*	<p>Applications outside intended domain or data that deviates from training conditions (Such as supported languages or other limitations).</p>
Key Function	<p>Key capabilities and how the AI feature enhances workflows.</p> <p>Select One:</p> <ul style="list-style-type: none"> - Automate: AI automates tasks and repeats them at scale - Analyze: AI provides actionable insights into data - Augment: AI enhances creative exploration and problem-solving
Model Type & Technique*	<p>Specific architecture or algorithm that forms the foundation of the AI feature.</p> <p>Examples:</p> <ul style="list-style-type: none"> - Generative AI: Transformer, Variational Autoencoders (VAE). - Predictive Models: Decision Trees, Random Forests, Logistic Regression. - Classification Models: Feedforward Neural Networks, Support Vector Machines (SVMs). - Clustering Models: K-Means, Hierarchical Clustering
Model Used	<p>Specific AI model(s) used. (links to doc)</p> <p>Examples:</p> <ul style="list-style-type: none"> - <i>Azure OpenAI, GPT-4(x), Mistral 12B, custom-trained, Bring your own (list options), etc.</i>
Esri AI Framework Use	<p>Indicate if the AI feature utilizes the Esri AI Framework to process customer data.</p> <p>Select One: Yes, No</p> <p>Note: The framework has not completed Esri's secure SDLC and AI baseline review and should not be used with production data.</p>
Model License*	<p>License type of AI model backing the AI feature.</p> <p>Select One:</p> <ul style="list-style-type: none"> - Proprietary: Developed by Esri - Open Source: Developed by a third party and publicly available - Licensed: Licensed for use from a third party - Combination: Contains both Esri and third-party components

Section	Description
Training Data Sources*	<p>Esri data sources used for development of AI feature. Discuss customer-specific data usage and conditions.</p> <p>Select One or More:</p> <ul style="list-style-type: none"> - Open Source: Freely available data - Customer Content: Data provided by the customer - Synthetic Data: Generated data mimicking real data structure - Commercial: Data acquired under license, list the provider (ex. Azure AI Translator) - Mix: Multiple sources used (Specify the source types) - Esri Content: Esri owned data sources (such as documentation) - Customer Trained: Customer's proprietary data was used for training
Human-in-the-Loop	<p>Indicate if users can review or modify AI-generated outputs before final action, emphasizing user control.</p> <p>Select One:</p> <ul style="list-style-type: none"> - Yes (Include explanation how) - No (Explain why not)
Input/Output Formats	<p>Provide clear and concise information about the various input and output formats that AI feature can handle.</p> <p>Examples:</p> <ul style="list-style-type: none"> - Input format: Image (JPEG, PNG, TIFF), Structured Data (CSV, JSON, XML) - Output format: Text (plain text, HTML), Generated Images (e.g., object categories such as animals or items)
Bias/Ethical Mitigations*	<p>How biases are managed, especially in the data. Provide details if the data has been reviewed for bias before being integrated into the AI feature.</p> <p>Examples:</p> <p>Fairness Assessment, Ethical Impact Assessment, Bias Minimization Techniques, Data review</p>
Security Mechanisms	Detail the security mechanisms and validations employed beyond encryption to protect data.
Guardrails	<p>Technical measures, controls, or components are in place to protect AI and content.</p> <p>Provide a Separate Answer for Each Guardrail Below:</p> <ul style="list-style-type: none"> - LLM Integrated (list what aspects/checks enabled or not) - Adversarial Attack (Separate mechanism from LLM - eg. Azure Prompt Shield) - Objectionable Content (Separate mechanism from LLM - eg. Azure AI Content Safety) - Intellectual Property (Separate mechanism from LLM – eg. Azure Protected Material)
Encryption at Rest	<p>Is AES-256 or better used for at rest encryption?</p> <p>Select One: Yes, No, Configurable by customer</p>
Encryption in Transit	<p>Is TLS 1.2 or better utilized for end-to-end transit encryption in alignment with Zero Trust Architecture?</p> <p>Select One:</p> <ul style="list-style-type: none"> - Yes: End-to-end encrypted communication is in place. - Partially: Client-server communications encrypted, but not backend communications - No: Client-server communications are not encrypted. - Configurable by customer: Customer can disable HTTPS
Adversarial Testing	<p>Adversarial robustness tests / Protection against manipulation.</p> <p>Examples: Red Teaming, General holistic testing...</p>

Section	Description
Key Metrics*	Performance metrics relevant to ArcGIS, like data and AI feature accuracy or business insight reliability. Examples: accuracy, precision, recall
Evaluation/Testing*	Testing methodologies and how Esri evaluates performance. Examples: A/B testing, holdout datasets
Provenance	A data mapping or architecture diagram has been conducted that tracks data from source to output, covering training, validation, and real-time data. - Yes (Include mapping or diagram) - No (Explain why not)
Best Practices	Links to tutorials, and best practices for using the AI feature effectively in ArcGIS.
Continuous Improvement	Improvement plans Examples: audits, bias checks, transparency...

*Section aligns with Hugging Face and FedRAMP AI Model card templates