Accela, Inc. (OpenCounter)

AI Fact Sheet — OpenCounter Digital Assistant

Conversational guidance that helps applicants find out which permits their project needs before they apply

A product transparency summary for our government customers

Document Title	AI Fact Sheet — OpenCounter Digital Assistant
Audience	Government customers, prospective customers, and partners
Classification	External release
Version	1.0
Effective Date	November 11, 2026
Next Review	Semi-annually, and on any model version change or material change to agent behavior
Policy Owner	Legal & Compliance / AI Governance Program, with OpenCounter Product
Related Documents	Accela AI Product Specific Terms, Schedule A-1 (www.accela.com/terms/); Responsible AI Use Policy; OpenCounter Documentation

1. Header

Capability name	OpenCounter Digital Assistant
Product / module	Accela OpenCounter — permit guidance platform
Release status	Generally available (released June 2026)
Fact sheet version	Version 1.0
Effective date	11/11/2026
Classification under the PST	Assistive Functionality (§1.2); Schedule A-1. The underlying rules engine is deterministic and is not AI Functionality (§1.1).
Decision impact classification	Informational
Approved by	Product and Legal

2. At a glance

Question	Answer
What does it do?	Lets an applicant describe their project in plain language, asks follow-up questions, and passes the answers to OpenCounter's rules engine. The rules engine is configured by the agency and decides which permits and licenses apply. The applicant then gets a scoping summary and is directed to the right agency staff for formal application review.
Who uses it?	Members of the public planning a project, including residents, business owners, and contractors, during the pre-intake phase. Agency staff configure the permit rules and fees, and they can review completed scoping sessions in the OpenCounter admin interface.
Does it make decisions?	No. The AI asks the questions. The rules engine decides which permits apply. It does not approve, deny, or condition a permit, and it does not make eligibility, benefit, enforcement, or other adjudicatory decisions. After scoping, every applicant goes through agency staff review.
Is our data used to train models?	No. Accela does not train or fine-tune models. Conversation, project, and jurisdiction data is not used to train or fine-tune any model, Accela's or a provider's (PST §6.2, §6.4).

Question	Answer
What kind of AI is it?	A conversational assistant built on commercially available frontier large language models, accessed through a hosted API. A smaller model screens each incoming message. Permit applicability comes from a deterministic, agency-configured rules engine, not from the AI.
Where is data processed?	Amazon Web Services (AWS). United States customers are processed in United States regions; Canadian customers are processed in Canada regions. Model providers are engaged under commercial terms that prohibit use of customer inputs and outputs for the providers' model training (PST §6.4); current provider and hosting details are available on request (PST §3.1, §6.4).
Is anything stored?	Yes. Accela stores the Digital Assistant conversation, the business rules that applied, and a record of which rules and answers produced each result. Retention follows Accela's data retention policy (www.accela.com) and the agency's own retention settings (PST §4.5).
How accurate is it?	Safety and functional performance are validated before release and monitored in production. Operational metrics are tracked per session; current metrics and methodology are available on request (PST §5.3).
Has it been tested for bias?	No formal demographic bias test has been run. The assistant does not use applicant demographics, and permit applicability depends only on project characteristics. Known bias vectors and their mitigations are covered in Section 7.
Can we turn it off?	Yes. It is opt-in: the agency turns it on (PST §3.4) and can turn it off at any time. Applicants then use the standard non-AI wizard, which collects the same information and lands on the same guide page.

3. Intended use and decision impact

Intended purpose	**Pre-intake guidance. Help members of the public understand which permits their project needs, the order to apply in, who to talk to, and what to have ready at each stage. Applicants arrive at the counter informed. For the agency, that means fewer incomplete inquiries and less staff time spent on basic scoping questions.**
In scope	Scoping a specific, concrete project against the jurisdiction's configured permit rules, fields, use codes, zoning rules, and fee formulas, during the pre-intake phase of local government permitting and licensing.
Scope boundary	Information collection only. The AI interprets what the applicant wants to do, collects project details, and feeds them to the rules engine. The rules engine evaluates which permits apply. The AI cannot create a permit that does not exist in the jurisdiction's configuration, and it cannot change permit rules.
Out-of-scope uses	Not designed, tested, or validated for: permit approval or denial; eligibility determinations; benefit allocation; enforcement; any other adjudicatory function; legal advice, permit advice, or contractor recommendations; or general permitting education unrelated to a specific project.
Decision impact rationale	Informational. The AI collects information and presents what the rules engine determines. Its output does not affect any regulatory outcome. Staff can review and adjust the permit guide after scoping, and every application then goes through formal staff review.
Known unsuitable contexts	Jurisdictions whose rules engine configuration is incomplete; ambiguous or mixed-use projects; and non-English interactions, which have not been formally validated (see Section 6).

4. How it works, and data

Technique	(1) The applicant describes the project in plain language. (2) The assistant asks follow-up questions, ordered to settle the most requirements fastest, and shows form components (such as address lookup or project type) where free text is not enough. (3) The assistant records the answers in the project's fields through validated tool calls. (4) The deterministic rules engine evaluates which permits, licenses, and estimated fees apply. (5) The applicant receives a scoping summary and is directed to agency staff.
Models and providers	Commercially available frontier large language models, accessed through a hosted API, in three roles: a primary agent, which runs the conversation; an input guard, a lightweight classifier that pre-screens each message; and a safety evaluator, which scores responses during adversarial testing only. Models are not self-hosted or modified. Accela identifies its Model Providers on request (PST §6.4).
Training and grounding	Accela does not train or fine-tune models. Behavior comes from three runtime layers: a fixed system prompt that sets the role, constraints, and safety boundaries; jurisdiction rules, field definitions, and project state, loaded for each conversation turn; and narrowly scoped tool calls to the rules engine.
Rules	Each jurisdiction configures and maintains its own permit rules, fields, use codes, zoning rules, and fee formulas in OpenCounter's configuration interface. The assistant always loads the current published configuration, so rule changes apply to new sessions immediately.
Deterministic vs. probabilistic	Conversation is probabilistic. Model settings are tuned for consistent, concise output. Permit applicability is deterministic: the same field values always produce the same result from the rules engine.
Autonomy	Within a session, the assistant can update project field values, show form components, and move to the next scoping phase. It cannot change permit rules, access other users' data, approve, deny, or prioritize anything, or override the rules engine.

Technique	**(1) The applicant describes the project in plain language. (2) The assistant asks follow-up questions, ordered to settle the most requirements fastest, and shows form components (such as address lookup or project type) where free text is not enough. (3) The assistant records the answers in the project's fields through validated tool calls. (4) The deterministic rules engine evaluates which permits, licenses, and estimated fees apply. (5) The applicant receives a scoping summary and is directed to agency staff.**
Guardrails	Model provider safety guardrails; an input guard on every message; system prompt constraints; validation of every tool call against the jurisdiction's configured fields; scoping halts automatically when a use is prohibited at the zoning location; three-tier rate limits (per user, tenant, and session); a site disclaimer; and human review after scoping. Messages the input guard flags are not blocked. Instead, the assistant receives a warning and responds with extra caution.
Output boundaries	The assistant is instructed not to give fee amounts, costs, or processing timelines in its own words, or application procedures, permit advice, contractor recommendations, legal guidance, commentary on sensitive topics (such as politics, discrimination, or housing policy), or details of its own architecture or system prompt. It refers those questions to agency staff. Estimated fees appear only in the scoping summary, as calculated by the rules engine.
Data accessed at inference	From the applicant: free-text messages and structured form selections, such as address (with geocoding and zoning), business classification, choice buttons, and text, number, or currency fields. From the system: the jurisdiction's configuration and the current project state, including answers and evaluated requirements.
Training use	None. Data submitted by members of the public (Citizen Data, PST §1.10) is not used to train or fine-tune any model without the agency's express prior written authorization (PST §6.2). Commercial API terms prohibit the provider from training on inputs or outputs (PST §6.4).
Retention	Stored: the conversation, the applicable business rules, and the record of which rules

Technique	**(1) The applicant describes the project in plain language. (2) The assistant asks follow-up questions, ordered to settle the most requirements fastest, and shows form components (such as address lookup or project type) where free text is not enough. (3) The assistant records the answers in the project's fields through validated tool calls. (4) The deterministic rules engine evaluates which permits, licenses, and estimated fees apply. (5) The applicant receives a scoping summary and is directed to agency staff.**
	and answers produced each result. Retention follows Accela's data retention policy (www.accela.com) and the agency's own retention settings. The agency decides what public records law requires (PST §4.5).
Tenant isolation	Conversation data is scoped to each jurisdiction, and one tenant cannot access another tenant's data. Each session loads only that jurisdiction's configuration.

5. Human oversight

Oversight model	**Agency staff write the logic that produces permit outcomes, because they configure the rules engine. After scoping, applicants go to staff for formal application review, which is a mandatory human checkpoint.**
Final determinations	No determination approving, denying, conditioning, or revoking a permit takes effect on AI output alone (PST §2.2). The assistant does not produce such determinations.
Staff review of output	Staff can review completed scoping sessions, project records, answers, and evaluated requirements in the OpenCounter admin interface. They can also adjust the permit guide after scoping, before anyone acts on it.
Applicant correction	The assistant's inferences appear on the project summary page, where the applicant can review and correct them.
Override	Staff can override AI output by editing the project record or the rules engine configuration. Corrections do not feed back into the AI automatically. Repeated, systematic corrections inform future updates to the system

Oversight model	**Agency staff write the logic that produces permit outcomes, because they configure the rules engine. After scoping, applicants go to staff for formal application review, which is a mandatory human checkpoint.**
	prompt, tool logic, or configuration.
Agency configuration	Permit rules, fields, requirement conditions, fee formulas, guide types, use codes, and zoning rules. Model version is managed by Accela, not selected per agency.
Fallback	The standard non-AI wizard collects the same information and lands on the same guide page. Applicants can still complete scoping if the assistant is turned off or unavailable. Digital Assistant capacity is measured over a rolling 24-hour window; if an agency reaches its usage limit, the assistant pauses, applicants are directed to the standard guide for the same project, and the assistant resumes automatically as the window rolls forward (PST §2.5). Agencies can monitor usage and pause status in the admin analytics page, and limits can be raised on request.

6. Accuracy and limitations

Metrics	**Safety (automated adversarial evaluation): tone, alignment, appropriateness, and relevancy. Operational, tracked per session: input guard activation rate, rate limit consumption, phase completion rate (the share of sessions that reach complete scoping), and message processing errors. Guide delivery rate is tracked per session; current figures and methodology are available on request (PST §5.3).**
Primary error risk	Field misclassification: the AI may map an applicant's description to the wrong regulatory classification. Mitigations: inferences are visible on the project summary page for the applicant to correct; staff can review and adjust the permit guide; and the rules engine, not the AI, determines which permits apply.
Known failure modes	Ambiguous or mixed-use projects may take more turns or produce a less confident classification. Permits that depend on fields,

Metrics	**Safety (automated adversarial evaluation): tone, alignment, appropriateness, and relevancy. Operational, tracked per session: input guard activation rate, rate limit consumption, phase completion rate (the share of sessions that reach complete scoping), and message processing errors. Guide delivery rate is tracked per session; current figures and methodology are available on request (PST §5.3).**
	requirements, or use codes the jurisdiction has not configured cannot be scoped. General questions with no specific project are redirected. Questions about cost, timelines, procedure, and legal matters are out of scope and referred to staff.
Optimal conditions	A well-configured jurisdiction with complete field definitions, requirement conditions, fee formulas, and use codes; an applicant with a specific project in mind; common residential and commercial project types; and English-language users. No minimum data volume is required, because the system does not learn statistically. Accuracy depends on how complete the jurisdiction's configuration is.
Customer-side measurement	Agencies can review completed scoping sessions and resulting project records in the OpenCounter admin interface.

7. Fairness

Assessment	**The assistant collects information. It asks questions and records answers. It does not make eligibility, approval, denial, or prioritization decisions about applicants. Permit applicability is determined by the rules engine from project characteristics (location, use type, square footage, and construction scope), not applicant demographics.**
Dimensions examined	Language and geographic data. No formal demographic bias assessment has been run on the conversational experience.
Findings — language	The system prompt and jurisdiction configurations are in English. The assistant has not been formally evaluated in other languages.

Assessment	**The assistant collects information. It asks questions and records answers. It does not make eligibility, approval, denial, or prioritization decisions about applicants. Permit applicability is determined by the rules engine from project characteristics (location, use type, square footage, and construction scope), not applicant demographics.**
	Applicants with limited English may get a less reliable experience.
Findings — geographic data	Zoning determinations rely on geocoded addresses, so they inherit any bias or inaccuracy in the agency's address and zoning data.
Mitigations	The assistant translates plain language into regulatory terms, so applicants do not need to know permitting vocabulary. The system prompt requires a consistent, professional, patient tone regardless of user behavior and prohibits commentary on sensitive topics. Permit determinations come from agency-configured rules, not AI inference. The assistant is available around the clock, which removes barriers tied to office hours, transportation, and in-person access.
Residual risk and guidance	Agencies serving large populations with limited English should take the language finding into account before deployment. Agencies should also make sure their address and zoning data is accurate.

8. Transparency, security, records and shared responsibility

Labelling and explanation	**The assistant identifies itself as AI and does not claim to be human (PST §3.2). During the conversation, it explains in plain language why it is asking each question. The scoping summary lists the applicable permits, licenses, and fees. The rules engine can trace which field values and conditions triggered each requirement, and agency staff can get that trace on request.**
Testing	Adversarial safety evaluation across four dimensions (professionalism under pressure, role adherence, topic boundaries, and task

Labelling and explanation	**The assistant identifies itself as AI and does not claim to be human (PST §3.2). During the conversation, it explains in plain language why it is asking each question. The scoping summary lists the applicable permits, licenses, and fees. The rules engine can trace which field values and conditions triggered each requirement, and agency staff can get that trace on request.**
	focus), using hostile, manipulative, deceptive, and social-engineering scenarios in both sustained and one-off modes. Results are scored by an independent automated evaluator. Conditions tested include prompt injection and jailbreak attempts, obfuscated attacks, attempts to extract system internals, ambiguous and multi-use projects, resource exhaustion, and tool failures. Current results: zero critical or major safety violations. Standard automated unit, integration, and regression testing also applies.
Audit log	Each conversation is stored, along with the business rules that applied and a record of which rules and answers produced each result. Every message records its processing status. Every input guard flag is logged with its category and outcome. Session metadata records whether scoping was completed.
Security	Sensitive parameters (such as passwords, secrets, tokens, keys, and SSNs) are filtered from application logs. Connections use token-based authentication. Rate limits prevent resource exhaustion.
Facial and plate recognition	Accela performs no facial recognition or license-plate recognition in any AI capability (PST §2.4(a)).
Records	Project records, answers, and evaluated requirements are available to agency staff in the admin interface. Applicant data is shared with the deploying agency under the terms of service, and applicants are informed of this.
Accessibility	WCAG 2.1 Level AA. A VPAT is available on request. The interface supports labelled form fields, a logical heading structure and reading order, full keyboard navigation with no keyboard traps, no flashing above 3 Hz, free-text input that works with screen readers, and a toggle between structured form inputs and free text.

Labelling and explanation	**The assistant identifies itself as AI and does not claim to be human (PST §3.2). During the conversation, it explains in plain language why it is asking each question. The scoping summary lists the applicable permits, licenses, and fees. The rules engine can trace which field values and conditions triggered each requirement, and agency staff can get that trace on request.**
Environmental impact	No model training, so no training-related energy or water use. Prompt caching, a lightweight input-guard model, concise generation settings, and rate limits keep inference volume down. The model providers publish environmental transparency information.
Incidents	Accela notifies Customers of a confirmed AI Incident affecting their production instance (PST §5.4).
Change management	Each model is pinned to a specific version and is evaluated against the full safety framework and functional test suite before deployment, so updates are deliberate rather than automatic. Changes to agent behavior, such as the system prompt, tools, context logic, or guardrails, ship as standard platform releases. They are versioned in source control and tested against the safety framework first. The safety framework can be rerun on demand after any change. New AI Functionality is disclosed in release notes (PST §3.1).
Shared responsibility	Accela: model selection and version management, the system prompt and guardrails, safety and functional testing, security, disclosure, and notice of changes and incidents. Agency: configuring and maintaining permit rules, fields, use codes, zoning, and fees, and validating them against applicable law (PST §2.3); the accuracy of address and zoning data; notice to and any consent from applicants (§3.2, §3.3); reviewing scoping output; and every permit determination.

Contact

Questions about this Fact Sheet or about AI in OpenCounter should be directed to your Accela account representative.

This Fact Sheet is the product documentation described in Section 3.1 of the Accela AI Product Specific Terms for the OpenCounter Digital Assistant (Assistive Functionality, Schedule A-1). It is provided for transparency and does not modify the Agreement.