ProjectX

ASK · DATASHEET

Ask one question across every system, and see how it was answered.

Ask in plain English and get one answer drawn from footage, badge logs, sign-ins, servers and spend. Each answer shows how the question was read, which systems it asked, the queries it ran and what policy left out. The engine has no access of its own: it reads only what the person asking may read.

Why teams choose it

01
One question, every pillar
The question is split into a request to each system that holds part of the answer, and the parts are joined on people, resources, time and place.

02
An answer you can check
Every answer carries its plan: the queries that ran and the sources they read, so a conclusion can be traced rather than taken on trust.

03
It never sees more than you
Each part of the question is authorized as the person asking before it runs. The engine holds no privileges of its own to fall back on.

04
Withheld is not the same as none
When policy keeps part of an answer from you, the answer says what was left out and why, instead of reporting zero.

In numbers

7
worked cross-pillar questions from the specification, all answered correctly

8
business entities defined in the semantic layer

0
privileges held by the engine itself

What it does

Questions in plain language

A question is parsed into entities, constraints, a time window and a scope, resolved against the semantic layer, with synonyms for your own vocabulary.

The parser reports what it understood and what it did not use, so a dropped word is visible rather than silently ignored. Business terms map to entities through the semantic layer, and synonyms let “badge ins” and “door entries” mean the same thing.

Joins across pillars

Independent legs run in parallel against each pillar, then join on principal, resource, time and location, including “without” questions.

The engine picks the pillars that hold the answer and asks each one for a plan. “Without badging in” becomes an anti-join between footage and the access log, with the time tolerance stated in the answer.

Entitlements on every leg

Each leg is authorized at plan time as the asker. What was excluded, and the policy that excluded it, is part of the answer.

A leg the asker may not read is not run at all, and the answer names it. A deep-archive footage leg without a stated purpose is one example: the answer says purpose was required instead of reporting nobody.

A semantic layer across the platform

Entities such as customer, site, asset, employee, vehicle, camera, server and incident, their relationships, and mappings from tables, detections, IAM events and cloud telemetry.

The semantic layer is where “a person in the server room” becomes rows in a video table and an access table. One definition serves dashboards, alerts, questions and agents alike.

Governed, certified metrics

Metrics defined in governed SQL with versions, and marked certified, experimental or deprecated.

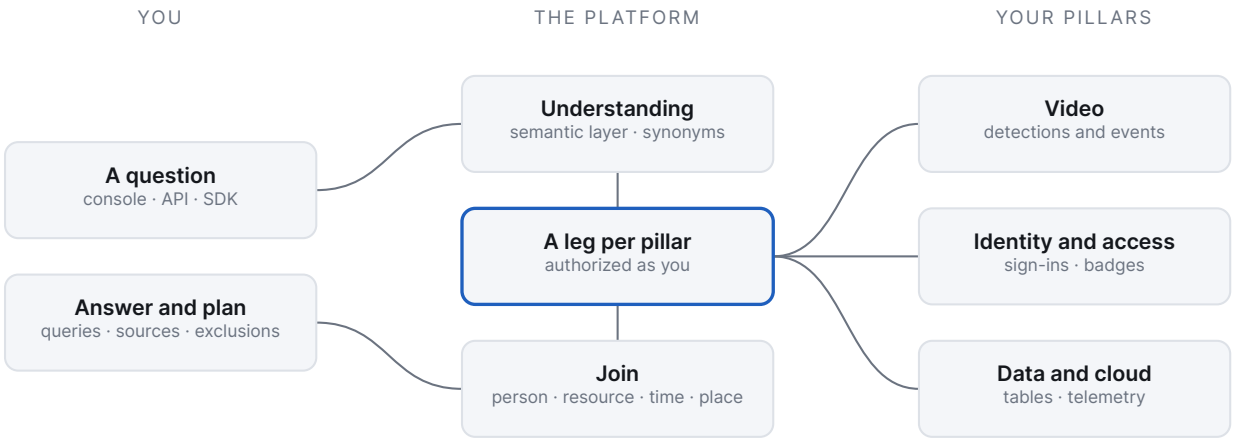
A metric such as badge entries or failed sign-ins is computed by one function wherever it is used, so an answer and a dashboard cannot disagree about it.

Saved, scheduled and structured

Save a question, share it and schedule it; or call the structured API with the same engine behind it.

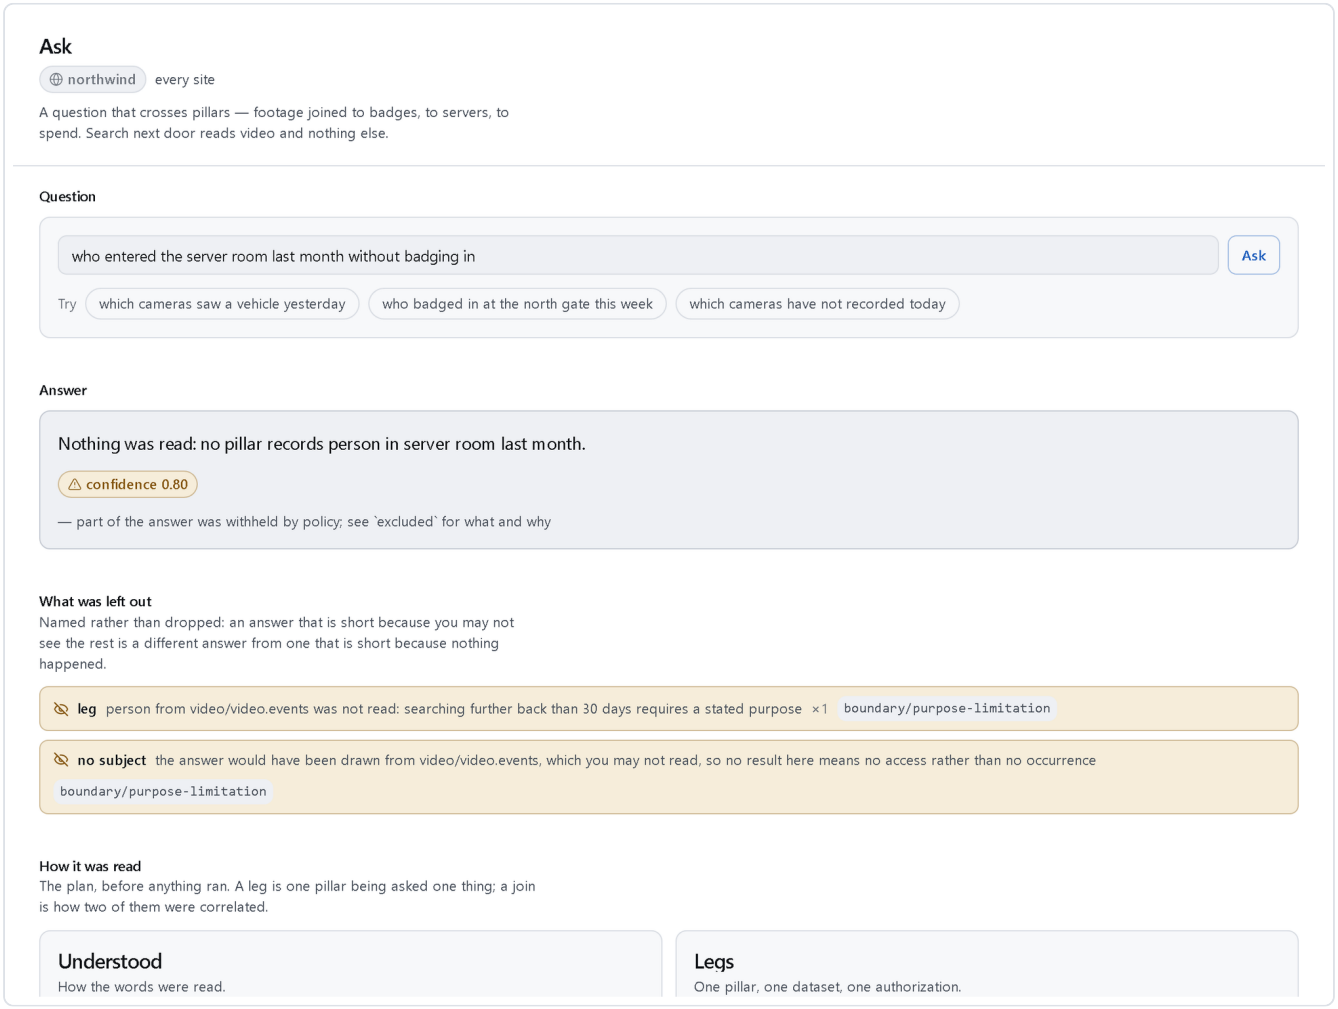
The deterministic parser answers first. An optional model can stand in when a question is outside it, constrained to the semantic layer, and can narrate an answer when you opt in.

How it fits together



A question is parsed and resolved against the semantic layer. The engine asks each pillar that holds part of the answer for a plan, authorized as the asker, runs the independent legs in parallel and joins them. The answer comes back with its plan and with anything policy excluded.

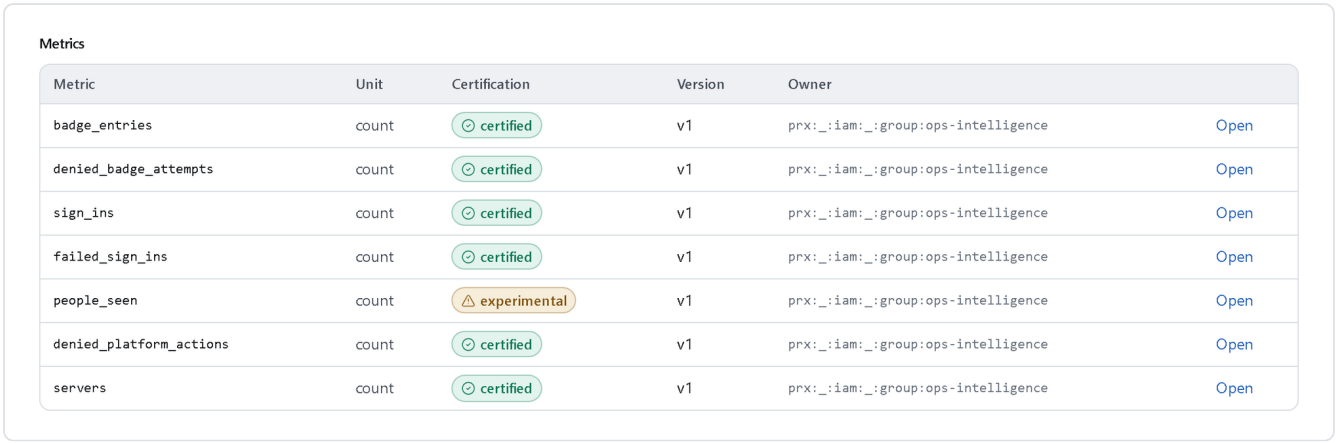
In the console



When policy withholds part of an answer, the answer says so. Screens shown with demo data.



The query that ran, not a summary of it. Screens shown with demo data.



Metrics

Metric	Unit	Certification	Version	Owner	
badge_entries	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open
denied_badge_attempts	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open
sign_ins	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open
failed_sign_ins	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open
people_seen	count	experimental	v1	prx:_:iam:_:group:ops-intelligence	Open
denied_platform_actions	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open
servers	count	certified	v1	prx:_:iam:_:group:ops-intelligence	Open

Metrics defined once, certified, and versioned. Screens shown with demo data.

Key ideas

Leg

One pillar being asked one thing as part of a question. Each leg has its own plan and its own authorization.

Semantic layer

The map from business words (a site, an employee, a vehicle) to the tables, detections and events that hold them.

Certified metric

A metric whose definition has been reviewed and marked certified, so it can be relied on wherever it appears.

Exclusion

A part of the answer policy kept from you, named in the answer with the reason and the policy.

In depth

No privileges of its own

The engine authorizes every leg as the person asking and has no account to fall back on, so a question can never return more than its asker could read directly.

Deterministic first

Questions are parsed without a model. A model can stand in only for questions outside the parser, constrained to the semantic layer, and narration is opt-in and degrades to the plain answer.

One metric everywhere

A metric is consumed the same way by dashboards, alerts, questions and agents, so a number does not change with where it is read.

How it works

01 **Ask** · Console
Type the question the way you would ask a colleague.

02 **Read how it was answered** · Console
See what was understood, which pillars were asked, the queries they ran and what was left out.

03 **Give a purpose when policy asks** · API
For deep-archive footage, state why you are asking and the footage leg runs.

04 **Save it** · Console
Name the question, share it with the team and put it on a schedule.

Who it is for

Security leads

Questions like “who entered the server room without badging in” that span cameras and access logs.

One answer joining footage to the badge record, with the queries behind it and the purpose it was asked for.

Operations managers

Numbers that mean the same thing on every dashboard and in every answer.

Certified, versioned metrics computed by one definition wherever they appear.

Audit and compliance

To know how an answer was reached and whether anyone saw more than they should.

Plans, sources and exclusions recorded with every answer, under the asker’s own permissions.

Works with

Investigator console · API · prx CLI · TypeScript SDK · Python SDK

Questions buyers ask

Which systems can a question reach?
Video detections, identity and access events, lakehouse tables and cloud telemetry, mapped into one semantic layer so a question can join them.

Does it depend on a language model?
No. A deterministic parser answers first. A model is optional: it can stand in for questions outside the parser, constrained to the semantic layer, and narrate an answer if you opt in.

What happens when I may not see part of the answer?
That part is not read, and the answer says what was excluded and which policy excluded it.

Can our own tools call it?
Yes. A structured API sits alongside the natural-language one, and both are in the CLI and SDKs.

How is it priced?
Ask is billed per 1,000 questions at a published rate in Platform Credits. The rate is on the pricing page.

Pricing

Current published rates, in your currency: https://adaydev.com/pricing/products/ask