



SALES & PARTNER PLAYBOOK

CosmicIntersection

CosmicIntersection gives research teams one governed home for their data, pipelines and results — so every finding traces back to the observation.

Data catalogue

Observation pipelines

Reproducible notebooks

Citable releases

Governed sharing

PREPARED BY

Burdenoff Consultancy Services Pvt.
Ltd.

PRODUCT

cosmicintersection.com

BOOK A CALL

calendly.com/vignesh-burdenoff/1-1



00 **How to use this playbook**

Read the opportunity signals and the one-question test first — they will tell you in a single conversation whether a prospect is worth pursuing. Use the personas section to pick the right opening line for whoever is in the room, and keep the objections pages to hand, because research buyers ask hard questions about maturity and lock-in. When you find something real, send it to us using the handoff checklist rather than trying to close it yourself.

IN A HURRY

Read sections 01–03 and 18. That is enough to recognise an opportunity and open a conversation.

BEFORE A MEETING

Add 06–08, 10 and 16 — who buys, what they buy, what it costs, and what they will push back on.

HANDING OVER A LEAD

Sections 19–22 tell you exactly what to send, to whom, and what you earn for it.

CONTENTS

01	Product Overview	13	Customer Proof & Credibility
02	One-Line Pitch	14	Demo Information
03	30-Second Elevator Pitch	15	Sales Talking Points
04	Problem We Solve	16	Objection Handling
05	Value Proposition	17	Discovery / Qualification Questions
06	Ideal Customer Profile	18	How to Identify a Good Opportunity
07	Buyer Personas	19	How to Introduce Us
08	Top Use Cases	20	Lead Handoff Process
09	Product Capabilities	21	Partner / Referral Commercials
10	Pricing & Plans	22	Contact Information
11	Competitive Landscape	23	Useful Links & Sales Assets
12	Why Us / Differentiators	24	Product Mockups — Complete Visual Appendix

01 Product Overview

COSMICINTERSECTION.COM

CosmicIntersection is a data platform for space and science research. It catalogues everything a team holds — telescope frames, satellite downlinks, ground-sensor readings, archive extracts, simulation output — as described, queryable datasets with their schema and provenance attached. Versioned pipelines turn raw observations into analysis-ready data and record what each step read and produced. Notebooks bind to the datasets they use and capture the exact runtime and package versions behind a result, and a finished notebook plus its datasets can be published as a versioned, citable release. Everything belongs to a workspace, so permissions, audit and multi-tenancy are there from the first day rather than added later.

CATEGORY

Science data platform — catalogue, processing pipelines, reproducible analysis and open-science publishing in one product, for research groups, observatories, agencies and data-intensive labs.

WHERE IT STANDS TODAY

Early. The product is live and open for sign-up, and the catalogue, pipeline, research, release and access-grant surfaces are all built and in use. It has no public reference customers, no published usage numbers and no third-party security certifications yet. Several headline capabilities are honestly partial: the platform records and validates pipeline runs while the compute still runs on the customer's own infrastructure, the persistent identifier is supplied at publish time rather than minted by a registrar, and the AI layer records attribution without yet running the detection model.

HOW CUSTOMERS GET IT

Cloud service — sign up and start; there is nothing to install. Regional pricing applies by billing country. Dedicated tenancy, data residency and self-hosting are Enterprise options rather than defaults. Onboarding for a lab is typically self-serve; institutions usually take a guided implementation covering source connections, pipeline migration and role design.

Web app — the complete product: catalogue, discovery, sources, collections, pipelines, runs, notebooks, analyses, releases and data access

GraphQL API — one typed API across the whole lifecycle, for notebooks, dashboards and partner tools

In-app assistant — available on every screen and aware of the record you have open

Agent tool surface — a published set of tools an external AI assistant can use to work the catalogue under your permissions

Documentation site — product guides and API reference, still being completed

Command-line client and Node/Python SDKs — specified and sequenced, not yet published

02 One-Line Pitch

SAY EXACTLY THIS

"CosmicIntersection gives research teams one governed home for their data, pipelines and results — so every finding traces back to the observation."

SHORTER, FOR A CHAT MESSAGE

One governed home for research data, pipelines and citable results.

THE CATEGORY SENTENCE

CosmicIntersection is a space and science data platform for research groups, observatories, agencies and labs that run the whole loop from raw observation to published result.

03 30-Second Elevator Pitch

SPOKEN, ROUGHLY 30 SECONDS

"Most research teams keep their data in one place, their processing scripts in another, and their analysis on somebody's laptop. Six months later nobody can reproduce the figure. CosmicIntersection puts all of it in one governed workspace: every telescope feed, sensor and archive catalogued, every processing pipeline versioned, every notebook pinned to the exact environment that produced it. When the work is ready, one step turns it into a citable release. The result is research that a reviewer, a funder or a collaborator can actually check."

IF THEY ARE TECHNICAL

CosmicIntersection models the whole research lifecycle as one connected graph — sources, datasets, schema fields, collections, pipelines, runs, steps, notebooks, analyses, releases and access grants — behind a single GraphQL API. Field types are science-native: coordinate, spectrum, geometry, image reference and timestamp, with source types for telescope, satellite, ground sensor, stream, simulation, archive and API. Access control, audit and usage quotas are enforced at the API edge on every operation, so a team can drive the platform from notebooks, dashboards or an AI agent without rebuilding the plumbing underneath.

IF THEY ARE A BUSINESS BUYER

Research groups today pay for archive access, notebook servers, orchestration tools, a repository, a file-sharing account and a home-grown governance layer — then pay again in the weeks their people lose stitching those together. CosmicIntersection replaces that sprawl with one platform, one bill, one identity and one audit trail, and it scales from a single graduate student on the free tier to an institution-wide deployment without re-platforming. For a funded lab, the argument is usually reproducibility and data-management compliance rather than licence cost.

04 Problem We Solve

These are the complaints you will hear almost word for word from research teams. Each one is a reason CosmicIntersection exists, and each one has a specific mechanism behind the answer.

The data is spread across a dozen portals

Who feels it: Research scientists, postdocs and graduate students

What it costs them: Every project starts with weeks of portal-hopping, each archive with its own query language, credentials and download workflow, before any science happens.

How CosmicIntersection answers it: Each upstream feed is registered once as a source with its connector configuration and a reference to the credential — never the secret itself. What lands is catalogued as a dataset with coverage, instrument metadata, format and per-column schema, so one search across the whole holding replaces a tour of the portals.

Nobody can reproduce a result six months later

Who feels it: Principal investigators facing a reviewer or a funder

What it costs them: The processing script is lost, the library versions have moved on, and a figure cannot be defended. Papers get delayed and data-management commitments go unmet.

How CosmicIntersection answers it: Notebooks bind to the datasets they used and capture the runtime and pinned package versions alongside the result. Pipeline runs record ordered steps with the inputs, parameters and outputs of each, so a chain runs from the raw observation through every step to the published figure.

Calibration lives in a folder of scripts on one laptop

Who feels it: Data engineers and research software engineers

What it costs them: The pipeline is one person's knowledge. When they leave, the team inherits something it cannot change safely, and downstream consumers have no idea how a table was produced.

How CosmicIntersection answers it: Processing becomes a versioned definition — an ordered step graph of ingest, transform, calibrate, analyse, export and notify — with a lifecycle and a trigger. Runs carry numbers, status, timing, resolved parameters and aggregated metrics, so the exact processing behind any product is on the record.

Sharing restricted data means emailing files

Who feels it: Research leads working with external or commercial partners

What it costs them: Two collaborators end up with different copies, there is no audit trail, and nothing can be revoked. With embargoed or proprietary data that is a genuine compliance exposure.

How CosmicIntersection answers it: Each dataset carries an access level — public, workspace, restricted or embargoed — and a collaborator is given a scoped grant with an optional embargo or expiry. Grants can be revoked at any point, every one carries a status and reason, and the whole trail is auditable.

Publishing a dataset is a project, not a button

Who feels it: Observatory directors and facility data leads with release obligations

What it costs them: Documentation gets written twice, identifiers are added late or never, users cite URLs that later break, and support email piles up after every release.

How CosmicIntersection answers it: A working notebook and the datasets behind it become a release with a version, abstract, authors and licence, moving through draft, review, published and retracted. The documentation is drawn from metadata already captured rather than written again, and the API lets people help themselves afterwards.

Governance gets bolted on after something goes wrong

Who feels it: Heads of research computing and institutional IT

What it costs them: Each research group ends up with its own authentication plumbing and access rules, which is fragile, expensive to support and impossible to report on.

How CosmicIntersection answers it: Roles, audit, tenancy and usage quotas come from the surrounding platform and apply to every operation, checked at the API edge rather than reimplemented per product. Six research roles — from principal investigator to read-only reproducibility auditor — are seeded into every workspace automatically.

Onboarding a new student burns a term

Who feels it: Principal investigators and lab managers

What it costs them: A new graduate researcher spends two or three months learning how to reach and process data before contributing anything, and the group's most senior people spend their time teaching plumbing.

How CosmicIntersection answers it: The free tier gives an individual the full product surface with no card required, and a new joiner sees the group's datasets, pipelines, notebooks and lineage in context on their first day. Guided walkthroughs and an in-app assistant that knows which record is open shorten the rest.

05 Value Proposition

Five arguments that hold up under scrutiny. Each one is grounded in something built rather than something planned — where a capability is partial, this playbook says so elsewhere.

One model for the whole lifecycle, not five tools bolted together

A source becomes a catalogued dataset, becomes a pipeline product, becomes a reproducible analysis, becomes a citable release — inside one connected model. Nothing is re-keyed between systems and no metadata is lost at a handover. That connection is also the reason a result can be walked backwards to the observation that produced it.

Evidence: Sources, datasets, schema fields, collections, pipelines, runs, steps, notebooks, analyses, releases and access grants are all first-class records in the same API, each with its own screen in the product.

Science-native, so observations are not forced into generic tables

Columns can be typed as coordinates, spectra, geometry, image references and timestamps, with units and a provenance note on each. Sources are typed by what they actually are — telescope, satellite, ground sensor, live stream, simulation, archive, API or file drop. General data platforms make researchers remodel observational data before they can store it.

Evidence: The typed field and source vocabularies are shipped and drive the schema-and-provenance preview that lets a researcher judge a dataset at a glance.

Reproducibility that is structural rather than heroic

Environment pinning, dataset binding and step-level lineage are properties of the platform, not disciplines a team has to maintain by hand. A collaborator can pick up a notebook and get the same numbers, and a reviewer can see which inputs and parameters produced a figure without asking anyone.

Evidence: Notebooks capture runtime and pinned package versions; runs record ordered steps with what each read and produced; the lineage view joins source, dataset, run, notebook, analysis and release into one chain you can click through.

Governance that was in place before any of it went live

Every record belongs to a workspace. Permissions are checked at the API edge on every read and write, writes are scoped to the tenant so an identifier guessed from another organisation cannot be touched, and every action lands in a searchable, exportable activity trail. Institutions do not have to build or buy this layer separately.

Evidence: Access control directives on every operation, six seeded research roles from principal investigator to read-only auditor, dataset access levels including embargo, and scoped grants with expiry and revocation.

Open science as the easy path, and restricted data on the same platform

A team can publish a versioned, citable release with authors, licence and a persistent identifier from the same workspace where its embargoed mission data sits behind scoped grants. Most tools force a choice between an open repository and a private file share; here both live under one policy and one audit trail.

Evidence: The release lifecycle and identifier field are built; access levels of public, workspace, restricted and embargoed apply per dataset, with time-bounded grants for named collaborators.

06 Ideal Customer Profile

A funded research group, observatory data centre or data-intensive lab of roughly five to fifty people that ingests data from more than one archive or instrument, has a reproducibility or public-release obligation, and currently holds the whole thing together with scripts and file shares.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
University research groups and labs	5–50 people, one or more grants	They feel every pain at once: portal-hopping, lost scripts, students onboarding for months, and funders now asking for data-management plans with versioning and provenance. Team seats, shared catalogue and embargo controls are the natural fit.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
Observatories, survey teams and facility data centres	20–200 people across science, engineering and operations	They have a public data-release mandate and a manual process to meet it. Automated documentation from captured metadata, persistent identifiers and self-service access for external researchers cut both the effort and the support load.
Space agencies and national facilities	Institution-wide, many missions and teams	They need consolidated governance, audit and usage visibility across many groups under one tenant, with enforced single sign-on, data residency and dedicated tenancy. This is the Enterprise conversation.
Commercial space and satellite operators	20–500 people	Proprietary mission data has to be shared with partners without handing over a copy. Access levels, scoped time-bounded grants, revocation and a full audit trail fit export-aware and contractual workflows.
Data-intensive science outside astronomy	Any team on the ingest-to-publish loop	Physics and instrumentation, Earth observation and remote sensing, signal and time-series work all run the same loop. The typed data model and lineage apply directly, and the platform does not assume telescopes.

GEOGRAPHIES

Sold globally with regional pricing bands by billing country — a standard global band covering the US, EU, UK, Japan, Korea, Australia, Canada and Singapore, and reduced bands for the Middle East and Eastern Europe, Latin America and South East Asia, and India and comparable markets. Early demand is strongest where research funders have started mandating data-management plans. Multi-lingual interface and locale-aware formatting are built in, which matters for international collaborations.

NOT A FIT – DO NOT CHASE

- A team that just needs somewhere to park finished files and get an identifier — a free repository does that well and CosmicIntersection would be over-specified
- Anyone shopping for raw compute at planetary scale, or a full machine-learning platform with model registries and experiment tracking; the compute story here is deliberately modest today
- Groups that want the authoritative mission archive itself — those archives are sources to connect to, not something this replaces
- Public-engagement and citizen-science projects wanting volunteer classification at scale; there is no volunteer crowd here
- Organisations that require an open-source, self-operated deployment as a precondition; self-hosting exists only as an Enterprise arrangement, and the software is not open source

07 Buyer Personas

Five people you will meet in a research organisation. The economic buyer is usually the principal investigator or the facility director, but the person who makes it real is almost always an engineer or a postdoc.

Principal investigator or research lead

WHAT THEY WANT

Publish work that stands up, win the next grant by demonstrating credible data management, and get their people spending more time on science than on wrangling.

Open with: If a reviewer asked you today to reproduce a figure from a paper you published two years ago, how long would that take — and would you get the same numbers?

They will ask: What actually happens to my data if we stop paying — can I get everything out, and does my published work survive?

WHAT FRUSTRATES THEM

Cannot reproduce older analyses on demand, shares data with international partners over email and cloud drives with no way to revoke, and watches each new student lose a term to technical setup.

Research software engineer or data engineer

WHAT THEY WANT

Stop writing the same archive integration for the twelfth time, cut the support burden, and build something reusable across the groups they serve.

Open with: How many one-off ingest scripts are you maintaining right now, and what happens to them the week you go on leave?

They will ask: Is the API good enough to build on, and can I run my existing processing without rewriting it into your engine?

WHAT FRUSTRATES THEM

A pile of bespoke scrapers, no shared conventions between groups, hours lost configuring notebook-server authentication, and constant requests to explain why last month's analysis no longer matches.

Observatory or facility data-centre director

WHAT THEY WANT

Meet the public release mandate on schedule, prove compliant governance to funders, and control cost while data volume grows faster than budget.

Open with: How much of your team's year goes into preparing public data releases that a self-service catalogue could serve directly?

They will ask: Can you show me the audit trail a funder would want, and what does an institution-wide deployment actually look like?

WHAT FRUSTRATES THEM

A manual, error-prone release process, information silos across fifty people in science, engineering and operations, and a support team spending a third of its time answering access questions.

Postdoc or staff scientist

WHAT THEY WANT

Get papers out on a short clock, build a portfolio that reads as technically serious, and work with distant collaborators without losing days to time zones.

Open with: What fraction of your last project was actually science, and what fraction was getting the data into a state where you could look at it?

They will ask: Does this fit the notebooks and libraries I already use, or is it another environment to learn?

WHAT FRUSTRATES THEM

Weeks of data wrangling per project, quiet anxiety about being asked to reproduce something, and a new reduction pipeline to learn for every instrument.

Head of research computing or institutional IT

WHAT THEY WANT

Standardise how research groups handle data, remove bespoke authentication plumbing, and be able to answer security and access questions across the institution.


Open with: If someone asked you today which external collaborators can reach which research datasets across your groups, could you answer it?

They will ask: How does single sign-on work, where does the data live, and what certifications do you hold today?

WHAT FRUSTRATES THEM

Every group runs a different stack, access reviews are impossible, and governance is retrofitted onto research tools after an incident rather than designed in.

Twelve situations the product was built for. Each one is a scenario illustrating what the platform does — not a record of a real customer, mission or result.



RESEARCH SCIENTIST

Find the right dataset in minutes instead of days


THE SITUATION

Answering one question means visiting three or four archive portals in turn, each with a different search interface and its own credentials, then guessing from filenames whether the files are the ones needed.

WHAT HAPPENS WITH COSMICINTERSECTION

Everything the workspace holds is searchable in one place. A keyword search combines with filters for access level, format, source, collection and status, with counts beside each filter drawn from the same live result set. Opening a dataset shows its columns, their units and a provenance note for each.

The working subset is identified and understood before a single file is moved.



DATA ENGINEER

Connect every feed once and watch it stay healthy


THE SITUATION

Ingest is a set of scripts that fail quietly. Nobody notices a feed has stopped until a researcher asks why the latest data is missing.

WHAT HAPPENS WITH COSMICINTERSECTION

Each upstream feed is registered as a typed source — telescope, satellite, ground sensor, stream, simulation, archive, API or file drop — with its endpoint, connector configuration and a reference to the stored credential rather than the secret itself. Each source carries a health status that moves through configuring, connected, syncing, degraded, error and disabled.

One screen tells you which data is flowing and which needs attention.



DATA STEWARD OR CATALOGUE OWNER

Describe a dataset so colleagues can trust it at a glance


THE SITUATION

A dataset arrives with no documentation. Anyone downstream has to reverse-engineer what the columns mean, what units they are in, and whether the numbers were measured or derived.

WHAT HAPPENS WITH COSMICINTERSECTION

Each column is declared as a schema field with a logical type — coordinate, spectrum, geometry, image reference, timestamp or a scalar — plus a unit, a nullability flag, a position and a provenance note saying where it came from. The dataset itself carries lifecycle status, access level, coverage, instrument metadata, format, size and version.

A researcher can judge whether a dataset fits their question without asking the person who made it.



RESEARCH SOFTWARE ENGINEER

Turn calibration scripts into a pipeline the team owns


THE SITUATION

Reduction and calibration live as scripts in one person's home directory. There is no version, no schedule, and no record of which version produced which product.

WHAT HAPPENS WITH COSMICINTERSECTION

Processing is composed as an ordered step graph — ingest, transform, calibrate, analyse, export and notify — with a lifecycle of draft, active, paused and archived, a version, and a trigger that can be manual, scheduled, on ingest or by webhook. Pausing or archiving a definition never loses its run history.

The processing behind every product is a versioned asset instead of one person's knowledge.



DATA ENGINEER

See exactly what a run did and where it went wrong


THE SITUATION

A fifty-step processing job fails somewhere in the middle. Debugging means reading log files on a cluster and guessing which inputs the failing stage actually consumed.

WHAT HAPPENS WITH COSMICINTERSECTION

Each run carries a number, status, timing, the parameters it resolved and aggregated metrics such as rows in and out, anomalies flagged and warnings. Every step records what it read and produced, plus its own logs and error messages, and runs can be advanced or cancelled through a validated set of transitions.

Debugging starts at the failing step with its inputs already named.



PRINCIPAL INVESTIGATOR OR PEER REVIEWER

Walk a published figure back to the raw observation


THE SITUATION

A reviewer asks where a number came from. The honest answer involves three people, two laptops and a folder nobody has opened in a year.

WHAT HAPPENS WITH COSMICINTERSECTION

Source, dataset, run, notebook, analysis and release join into a single chain that can be followed in either direction, assembled from the real links between the team's own records. A flag on each dataset records whether downstream processing has wired its provenance, so trust is explicit rather than assumed.

Provenance is a screen you open, not an archaeology project.



POSTDOC OR STAFF SCIENTIST

Keep an analysis re-runnable long after the session ends


THE SITUATION

A notebook produced the right figure in March. In September the libraries have moved, the notebook no longer runs, and the numbers cannot be defended.

WHAT HAPPENS WITH COSMICINTERSECTION

A notebook binds to the datasets it uses and captures its runtime and a pinned environment — package versions and a container identifier — alongside a share link and status. The environment travels with the notebook, so a collaborator opens exactly what the author had.

The environment behind a result is recorded rather than remembered.



RESEARCH LEAD

Record which findings a machine helped produce


THE SITUATION

Somebody used an automated detection method for part of an analysis. Six months on, nothing distinguishes that result from one derived by hand, which is exactly the question a reviewer will ask.

WHAT HAPPENS WITH COSMICINTERSECTION

An analysis records its method, its structured results — metrics, detected anomalies, summary statistics — a reference to the rendered visualisation, and an explicit mark saying whether a machine assisted. The mark is badged wherever the analysis appears, on desktop tables and mobile cards alike.

Machine assistance stays on the record instead of quietly disappearing.



RESEARCH LEAD OR FACILITY DATA MANAGER

Publish a dataset others can find, reuse and cite


THE SITUATION

Releasing data publicly means writing documentation for a second time, formatting files by hand, and putting them on a web page where the citation is a URL that will eventually break.

WHAT HAPPENS WITH COSMICINTERSECTION

A notebook and the datasets behind it become a release carrying a version, abstract, authors, licence and a persistent identifier, moving through draft, under review, published and retracted. Coverage and provenance already captured feed the release metadata rather than being written again.

Publication becomes a step in the work rather than a separate project.



PRINCIPAL INVESTIGATOR WORKING WITH EXTERNAL PARTNERS

Share restricted data without handing over a copy


THE SITUATION

Proprietary or embargoed data has to reach a collaborator abroad. Today that means an encrypted transfer, two divergent copies, no audit trail and no way to take access back.

WHAT HAPPENS WITH COSMICINTERSECTION

The dataset carries an access level of restricted or embargoed, and the collaborator receives a grant with explicit scopes and an optional embargo or expiry. Grants move through pending, active, revoked and expired, each with a reason, and can be revoked at any point.

Access opens and closes exactly when it should, and every decision is auditable.



RESEARCH LEAD OR SURVEY SCIENTIST

Group a mission or survey so it reads as one body of work


THE SITUATION

The datasets from one survey are scattered across instruments and epochs. Newcomers cannot tell which ones belong together, and the group has no way to expose the set publicly as a coherent whole.

WHAT HAPPENS WITH COSMICINTERSECTION

Related datasets are gathered into a collection representing a mission, a survey or a project, which can be marked public for open sharing while its members keep their own access levels. Collections sit alongside the catalogue rather than duplicating anything in it.

A body of work is discoverable as a unit, inside the team and outside it.



RESEARCH SCIENTIST OR GRADUATE RESEARCHER

Ask about the record already on your screen

THE SITUATION

Understanding an unfamiliar dataset, run or release means reading through several screens and asking a colleague who is in another time zone.

WHAT HAPPENS WITH COSMICINTERSECTION

An assistant is available on every screen and is told which record you have open before your question is sent, so a question about a dataset, source, collection, pipeline, run, notebook, analysis or release needs no identifier. It holds exactly your permissions and its activity lands in the same audit trail as anyone else's.

Context comes to the question instead of the person hunting for it.

09 Product Capabilities

What is built, what is in progress and what is planned. Anything not marked shipped is written here as it stands today, and the renderer flags it — a partner who repeats these labels accurately will never be caught out.

Data catalogue

Everything a team holds, described well enough to be found and trusted again.

Typed sources

Register telescope, satellite, ground-sensor, stream, simulation, archive, API and file-drop feeds with connector configuration and a reference to the stored credential rather than the secret.

Source health and sync

Each source reports configuring, connected, syncing, degraded, error or disabled, with a last-synced timestamp you confirm as fresh data lands.

Dataset lifecycle

Datasets carry draft, ingesting, available, deprecated and archived status, plus coverage, instrument metadata, format, row count, size and version.

Schema fields with provenance

Describe every column with a science-native type, unit, nullability and a note on where the value came from — instrument, derived formula or source column.

Collections

Group related datasets into a mission, survey or project, optionally public, without duplicating anything.

Faceted discovery

Keyword search combined with access-level, format, source, collection and status filters, with live counts that agree with the results.

Interactive querying with live previews planned

Slicing indexed columnar storage in the browser and carrying the result straight into a notebook is planned, not available today.

Lifecycle-aware tiered storage planned

Moving cold datasets to cheaper storage without them leaving the catalogue is on the roadmap.

Observation pipelines and lineage

Processing as a versioned, observable definition, with the record of what each step did.

Declarative step graphs

Compose ingest, transform, calibrate, analyse, export and notify steps into a reusable definition with its own version and lifecycle.

Triggers

Run a pipeline manually, on a schedule, when new data lands, or from a webhook.

Runs with metrics

Each execution records a run number, status, timing, resolved parameters and aggregated metrics including rows in and out, anomalies flagged and warnings.

Step-level lineage recording

Ordered steps record what they consumed and produced, with per-step logs and errors, and runs advance through a validated set of transitions.

Managed pipeline compute in progress

The platform records, validates and tracks runs today; the compute itself still runs on the customer's own infrastructure and reports back. An execution engine that runs step code in-platform is being worked on.

Compute co-located with large archives planned

Running pipelines and notebooks next to petabyte-scale archives, instead of over downloads, is planned.

Backfill and parallel-run validation planned

Tooling to migrate legacy processing with lineage reconstruction and side-by-side comparison is planned.

Reproducible research

Notebooks and analyses that can be re-run, checked and built upon.

Notebooks bound to datasets

A notebook records which datasets it uses, its runtime and its status, and can be shared by link so a collaborator opens exactly what you see.

Pinned environments

Package versions and a container identifier are captured with the notebook so the same libraries can be brought back later.

Analyses with structured results

Record the method, the resulting metrics, detected anomalies and summary statistics, plus a reference to the rendered visualisation; charts are drawn from the real numbers or the page says there is nothing to plot.

Machine-assistance attribution

Mark an analysis as machine-assisted and name the method; the mark stays on the record and is badged wherever the analysis is listed.

Hosted notebook execution planned

Running a pinned-environment notebook inside the platform is planned; today the platform records and pins the environment while execution happens on your own infrastructure.

Scientific visualisation surface planned

Interactive sky maps, light curves, spectra and three-dimensional scenes, shareable by link, are planned. Analysis results render as charts today.

Open science and publishing

Turning finished work into something other people can find, reuse and credit.

Citable releases

Bundle a notebook and its datasets into a versioned release with abstract, authors and licence, moving through draft, under review, published and retracted.

Persistent identifier on a release

A release carries an identifier field so the work can be cited like a paper.

Metadata aligned to FAIR practice

Coverage, provenance and licensing metadata travel with the release so others can judge what it is and how it may be used.

Registrar-minted identifiers in progress

Automatic registration with an identifier authority is being worked on; today the identifier is supplied at publish time rather than minted by the registrar.

Export and mirroring to external repositories planned

Standards-compliant export so institutions can also preserve a release in their own repository network is planned.

Governance, security and sharing

The controls an institution asks about, applied to every operation rather than added on top.

Access levels per dataset

Public, workspace, restricted and embargoed, enforced on reads as well as writes.

Scoped, time-bounded access grants

Give a named collaborator specific scopes with an optional embargo or expiry, and revoke at any time; each grant carries a status and reason.

Role-based access on every operation

Permissions are checked at the API edge for each read and write, with six research roles seeded into every workspace from principal investigator down to a read-only reproducibility auditor.

Tenant-scoped writes

Every change is filtered by workspace, so an identifier guessed from another organisation cannot be modified.

Audit and activity trail

A searchable, exportable timeline of every action with actor, target and result, including reads of sensitive catalogue records.

Secrets kept out of records

Sources reference a credential by name, and configuration that looks like a password, token or key is rejected at save time.

Compliance evidence export planned

Audit-driven reporting and streaming to a security information system is planned.

Independent certification planned

No third-party certifications are held. SOC 2, ISO 27001 and ISO 27701 are on a published roadmap with target dates in 2027; privacy-regulation readiness work is under way.

AI and agents

Deliberately labelled. The record-keeping is built; the model that does the science is not.

Assistant on every screen

Available across the catalogue, discovery, pipelines, runs, notebooks, analyses, releases, lineage and data access, and it knows which of eight kinds of detail page you have open.

Agent tool surface

A published set of tools lets an external AI assistant register sources, catalogue and publish datasets, declare fields, manage collections and pipelines, advance runs, write notebooks, run analyses, cut releases and manage access — holding your permissions and landing in your audit trail.

Guided agent workflows

Two walkthroughs ship with the tool surface: onboarding a dataset end to end, and taking research through to a release, so an agent follows a known order rather than guessing.

Confirmation before an agent deletes in progress

A shell-level gate refuses risky calls without an explicit confirmation; annotating each archive and revoke operation so the gate catches all of them is still being finished.

The model behind an assisted analysis in progress

Starting a machine-assisted analysis opens a record that waits for whatever model or worker performs the detection to report back. Choosing and running that model in-platform is open work.

Natural-language catalogue search planned

Asking for observations of a region, in a band, during an epoch and getting ranked datasets back is the target. Matching today is on words and filters only.

Anomaly and transient detection planned

Flagging the unexpected across streaming and archived observations for human review is planned, and needs a model provider and an evaluation set before any quality claim is made.

Field inference, run triage and release drafting planned

Three proposal-style capabilities — drafting a dataset's fields from the file, triaging a failed run against similar ones, and assembling a draft release from lineage — are designed and planned, not built.

Inherited platform services

Capabilities most standalone science tools have to build, buy or do without.

Single sign-on and identity

Enterprise sign-on, device-code login for terminal clients and keys for automation, all from the shared identity layer.

Billing, plans and usage quotas

Subscriptions with regional pricing, add-on capacity and quotas metered on datasets, notebooks and pipeline runs, enforced as the action happens rather than after the fact.

Files, tags and conversations

Durable storage for artefacts, notebook documents and rendered outputs referenced rather than copied into the catalogue, shared tagging, and inline discussion on datasets, pipelines, notebooks and releases.

Notifications and announcements

In-app, email and webhook delivery with user preferences, plus workspace-level announcements for releases and advisories.

Multi-lingual interface and accessibility

Language switching per user, locale-aware formatting, keyboard shortcuts with a command palette, screen-reader semantics and high-contrast theming.

Guided onboarding

In-product walkthroughs and contextual hints for people arriving at the product for the first time.

Command-line client and language SDKs planned

A terminal client and Node and Python SDKs are specified and sequenced, but not yet published to their registries.

10 Pricing & Plans

Priced by the size of the research data under management — datasets and connected sources — rather than by how many people log in, so growing a team inside a tier does not change the bill. Free, Pro, Growth and Business are flat monthly plans; Enterprise moves to an annual per-seat contract because that is how institutional procurement works. Regional bands apply by verified billing country and are the normal price for that market rather than a discount. Machine-assisted features run on the customer's own AI provider key, so model usage is billed by that provider and not included here.

PLAN	PRICE	WHO IT IS FOR	WHAT IS INCLUDED
Free	\$0 No card required, no expiry	An individual researcher or student trying the platform end to end	15 datasets under management; 1 live data source; 2 admin or steward seats; 1 workspace; 1 GB of shared storage; Community support plus one ticket a month; 30 days of analytics history; Add-ons cannot be purchased on Free
Pro	\$99/mo or \$79/mo annual India example: ₹3,999/mo or ₹3,199/mo annual	Independent researchers, graduate students and single-investigator projects	50 datasets under management; 3 live data sources; 5 admin or steward seats; 1 workspace; 25 GB of shared storage; 10 integrations; 90 days of analytics history; Email support under 48 hours
Growth	\$399/mo or \$319/mo annual India example: ₹14,999/mo or ₹11,999/mo annual; 14-day Business trial available	Scaling groups running data ingestion at volume across several projects	250 datasets under management; 15 live data sources; 15 admin or steward seats; 3 workspaces; 100 GB of shared storage; 50 integrations; 365 days of analytics history; Priority support under 24 hours, 99.5% uptime commitment
Business	\$999/mo or \$799/mo annual India example: ₹39,999/mo or ₹31,999/mo annual	Established organisations needing full control and a service commitment	1,000 datasets under management; Unlimited live data sources under fair use; 50 admin or steward seats; 10 workspaces; 1 TB shared storage pool; 250 integrations; 730 days of analytics history; Priority support under 8 hours, 99.9% uptime commitment
Enterprise	From \$12/seat/mo list or \$9/seat/mo annual at 250+ Annual contract; custom regional quote outside the standard band	Institutions, observatories, agencies and multi-lab organisations	Custom quotas across datasets, sources, storage and pipeline runs; Enforced single sign-on and directory provisioning; Audit streaming to a security information system; Data residency options; Multi-entity structures for many labs under one tenant; Dedicated tenancy or self-hosting; Compliance support and premium service levels

Add-ons

ADD-ON	PRICE
50-dataset capacity pack	\$39/mo
250-dataset capacity pack	\$149/mo
1,000-dataset capacity pack	\$499/mo
5 additional live data sources	\$29/mo
25 additional live data sources	\$99/mo
5,000 additional pipeline runs a month	\$29/mo

ADD-ON	PRICE
100 additional notebooks	\$49/mo
250 GB additional shared storage	\$49/mo
5 additional admin or steward seats	\$49/mo
Support retainer — 12 hours a month	\$4,000/mo

WHAT TO TELL A PROSPECT ABOUT PRICE

- Quote the plan prices exactly as printed here and let us confirm the final figure — regional band, taxes and any approved discount all move the number on the actual quote.
- Storage and API allowances are shared across the customer's whole account, so a customer already using another Burdenoff product brings a larger pool with them.
- Verified education and non-profit organisations can receive up to 50% off, and companies under three years old with fewer than 50 seats up to 40% off for twelve months; discounts generally do not stack.
- Pro and Business include a 30-day money-back guarantee on a first purchase, monthly plans cancel at the end of the paid period, and annual plans run to the end of their term.
- The dataset and template marketplace described in some older material is not available and must never appear in a quote, an order form or a checkout flow.
- Machine-assisted features require the customer's own AI provider key; that provider bills them separately for model usage.

11 Competitive Landscape

Seven alternatives you will actually hear named. Every one of them genuinely beats CosmicIntersection somewhere, and saying so out loud is what makes the rest of the comparison credible. Several of them are better positioned as sources or publication targets than as rivals.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
MAST and comparable mission archives	The authoritative, publicly funded archives for space-telescope missions — Hubble, Webb, TESS, Kepler and others — free to use and trusted worldwide.	They hold the authoritative mission data and the mission-specific search and preview tooling built around it. Nothing here replaces that, and CosmicIntersection has no equivalent mission query interface today.	An archive is read-oriented and single-organisation. CosmicIntersection is the workspace that consumes archive data alongside private instruments, adds versioned processing with lineage, reproducible notebooks, team permissions and citable output. The right pitch is complementary: keep the archive, add the workspace.
SciServer	A university-run science platform offering free virtual machines, pre-installed notebooks, a SQL query interface and access to more than two petabytes of public datasets.	Real co-located compute and a very large hosted public archive, free to academics. CosmicIntersection offers neither free academic machines nor a hosted archive at that scale.	It is not workspace-scoped: no built-in multi-tenant roles, audit, billing or enterprise sign-on, no versioned pipeline orchestration and no citable release workflow. CosmicIntersection adds the governance, lineage and publishing layers around the same kind of work.
Zenodo	A free, open repository launched by CERN and OpenAIRE that gives any researcher a persistent identifier for datasets, software and papers.	Free, unlimited deposits, real registrar-minted identifiers and a preservation mandate behind it. CosmicIntersection supplies the identifier at publish time rather than minting it, and has no comparable preservation network.	Zenodo is where finished files go. It does not catalogue live sources, run pipelines, hold notebooks or capture lineage — the active work happens elsewhere and the deposit arrives at the end. CosmicIntersection covers that active phase and can export to Zenodo for preservation.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Dataverse	An open-source institutional data repository from Harvard, widely deployed by universities for sharing, preserving and citing research data.	Open source, self-hostable, part of an established institutional repository network, and free apart from hosting. For an institution that requires open-source software it operates itself, that is decisive.	It is a repository rather than a live data platform: no observation pipelines, no real-time ingestion, no co-located compute, and no science-native model for telescope feeds and mission data. CosmicIntersection produces the datasets a repository archives, and can export to a Dataverse installation.
Zooniverse	The largest citizen-science platform, connecting researchers with millions of volunteers who classify images, transcribe text and mark features.	The volunteer community itself, free crowd classification at a scale nobody can replicate, and a proven public-engagement model. CosmicIntersection has no crowd and no plans for one.	Zooniverse stops at classification and is public by default with limited institutional controls. Its output can be brought in as a source and taken onward through processing, reproducible analysis, governed sharing and citable release.
Databricks	The enterprise lakehouse for data engineering and machine learning, with notebooks, SQL warehouses, model tooling and catalogue governance.	Far greater compute scale, a mature machine-learning platform and a large partner ecosystem. Anyone with heavy petabyte-scale processing or a full model lifecycle will get more from Databricks today.	It is generic: no understanding of instruments, coordinates or spectra, no citable release workflow, and consumption pricing that research budgets find hard to forecast. CosmicIntersection is science-native with predictable flat pricing and a publication step at the end.
Google Earth Engine	A planetary-scale platform pairing a multi-petabyte Earth-observation catalogue with large-scale geospatial analysis, free for academic use.	An enormous hosted Earth-observation archive and genuine planetary-scale raster compute. Neither exists here, and for large Earth-observation workloads Earth Engine is the better tool.	It is Earth-only and tied to one catalogue and runtime, with no citable-release layer and no workspace-scoped governance. CosmicIntersection spans astronomy, planetary, mission and instrumentation data, takes your own sources, and does not lock the data to one vendor.

THE POSITIONING LINE TO HOLD

CosmicIntersection does not try to out-archive the mission archives or out-publish the repositories — it is the governed workspace in the middle that turns their data into reproducible, citable research.

12 Why Us / Differentiators

If a prospect only remembers a handful of things about this product, make it these. The edge is integration and honesty rather than a longer feature list than anyone else.

The lifecycle is one model, not five integrations

A source becomes a dataset, becomes a pipeline product, becomes a notebook and an analysis, becomes a release — with the relationships between them recorded natively rather than inferred from filenames. Every competitor covers part of that chain well; none covers all of it, which is exactly why teams stitch tools together today.

Observations are modelled as observations

Coordinates, spectra, geometry, image references and timestamps are real field types with units, and sources are typed by what they are. General data platforms make a researcher flatten all of that into generic tables before storing anything, and the meaning goes with it.

Reproducibility is a property of the platform

Environment pinning, dataset binding and step-level lineage happen because that is how the records work, not because a team maintains a discipline. That is a different proposition from a repository that stores whatever documentation somebody wrote at the end.

Restricted and open data on one platform

Embargoed mission data behind scoped, expiring grants sits in the same workspace as a public citable release, under one policy and one audit trail. Most teams currently run an open repository and a private file share side by side and reconcile them by hand.

Enterprise controls arrived before the features did

Roles, audit, tenancy and quotas were in place before anything else was layered on, including the AI surfaces, so they apply everywhere rather than to whichever features remembered to ask. That is the opposite of the usual pattern of retro-fitting governance after an incident.

A typed API an agent can actually use

Because the catalogue is structured records rather than a folder of files, both the in-app assistant and an external AI agent can work with real datasets, fields, runs and lineage — under the caller's own permissions and inside the same audit trail. An assistant over a file share can only summarise whatever prose someone wrote.

The gaps are written down

Managed pipeline compute, hosted notebook execution, registrar-minted identifiers and the detection models are all labelled as in progress or planned on the public site and in this document. In a market of research buyers who read carefully, that is a competitive asset rather than a weakness.

13 Customer Proof & Credibility

There are no customer references to lean on, so credibility has to come from the depth of what is built, the governance posture, and the fact that limitations are stated rather than hidden. Show the product, show the roles, show the lineage chain, and be first to name what is not finished.

DEPTH OF THE SHIPPED DATA MODEL

Eleven connected record types across catalogue, pipelines and research — sources, datasets, schema fields, collections, pipelines, runs, steps, notebooks, analyses, releases and access grants — each with its own screen and its own operations in the API.

ACCESS CONTROL ON EVERY OPERATION

Permissions are declared on each read and write and enforced at the API edge rather than reimplemented per feature, with writes scoped to the workspace so identifiers from another organisation cannot be touched.

SIX RESEARCH ROLES SEEDED AUTOMATICALLY

Principal investigator, data steward, pipeline engineer, research scientist, research contributor and a read-only reproducibility auditor are created in every workspace, mapped to how research groups actually divide responsibility.

SECURITY POSTURE WRITTEN DOWN

Encryption in transit, multi-factor authentication, tenant isolation, per-environment secret management and a published subprocessor list, with security and availability controls designed against recognised criteria.

A PUBLISHED CERTIFICATION ROADMAP

SOC 2, ISO 27001 and ISO 27701 are named with target dates in 2027 and their prerequisites, alongside accessibility conformance and privacy-regulation readiness. Nothing is claimed as held.

PUBLIC HONESTY ABOUT STATUS

The product website labels each capability shipped, in progress or roadmap and states plainly that its scenarios are illustrative rather than records of real customers, missions or results.

Be straight about this. CosmicIntersection is a new product. It is live and open for sign-up, but it has no public reference customers, no published usage figures, and no third-party security certifications — SOC 2 and ISO 27001 are planned for 2027, not held today. Several headline capabilities are deliberately partial: the platform records and validates pipeline runs while the compute still runs on the customer's own infrastructure, the persistent identifier on a release is supplied at publish time rather than minted by a registrar, and the AI layer records attribution and holds the result without yet running the detection model. Say all of that before a prospect finds it. What they should weigh instead is the depth of the shipped catalogue and research model, the governance that was in place before anything was built on top of it, and the fact that every gap above is written down in public rather than glossed over.

14 Demo Information

Twenty minutes is enough. The demonstration works best told as one continuous story — a feed arrives, becomes a trusted dataset, gets processed, gets analysed, gets published — because the continuity is the product. Resist showing screens in isolation.

HOW TO GET A DEMO

Anyone can sign up on the free tier and follow along in their own workspace, which is usually more persuasive than watching. For a group or institution, we run a guided walkthrough shaped around the archives and pipelines the prospect actually has. Book it through the call link, or send us the opportunity and we will arrange it.

Runs about 20–30 minutes for a first walkthrough; 45–60 minutes for a technical audience who will want the API and the roles.

WHERE TO LOOK

- **Product website** — <https://cosmicintersection.com>
- **Sign in or sign up** — <https://app.cosmicintersection.com>
- **Where we win and where we do not, by competitor** — <https://cosmicintersection.com/vs>
- **Book a guided walkthrough** — <https://calendly.com/vignesh-burdenoff/1-1>

The walkthrough, in order

1 The dashboard

This is the front door — datasets, sources, pipelines, runs, notebooks and releases in one place, with recent activity. The point is that a researcher's day starts here instead of across six tabs.

2 Discovery with filters applied

One search across everything the team holds, narrowed by instrument type, access level, format and mission. The counts beside each filter are the number of datasets you would actually get.

3 A dataset with its schema fields

Every column has a type, a unit and a note on where the value came from. This is how somebody decides whether to trust a dataset without emailing the person who made it.

4 The sources list with health statuses

Each feed is registered once, with its credential referenced rather than copied. You can see at a glance which sources are flowing and which need attention.

5 A pipeline definition and its steps

Calibration and reduction as a versioned definition with a trigger, not scripts on a laptop. Worth pausing here: today the platform records and validates the run while the compute runs on your infrastructure.

6 A run, then one of its steps

The run carries its parameters and metrics; each step records what it read and produced, with its own logs. When something fails, debugging starts at the step rather than in a log file.

7 A notebook and its pinned environment

The notebook is bound to its datasets and records the exact runtime and package versions. This is the answer to the reviewer who asks you to reproduce a figure from two years ago.

8 The lineage chain, then a published release

Follow it back from the release to the observation it came from. Then show the release itself — version, abstract, authors, licence, identifier — and say plainly that the identifier is supplied at publish today, with registrar minting in progress.

15 Sales Talking Points

Eight lines that land with research buyers. Each one is a claim you can defend if pressed, which matters more here than in most markets — this audience checks things.

The lifecycle is one model, so a result can be walked back to the observation

Reproducibility is the pain that closes deals in research, and this is the mechanism behind it rather than a promise about it.

Observations are typed as observations, not flattened into generic tables

Any scientist who has remodelled spectra into a spreadsheet will recognise the difference immediately, and no general data platform can match it.

Governance came first, so it applies everywhere including the AI surfaces

Institutional buyers have been burned by tools that added access control after an incident. This inverts the usual order and is easy to demonstrate.

Embargoed data and public releases live on the same platform under one audit trail

Most teams currently run a private share and an open repository side by side, and reconcile them by hand. This collapses two systems into one.

It replaces a stack — archive access, notebook servers, orchestration, a repository, file sharing and home-grown governance

The financial argument is total cost and staff time, not licence price. Most research groups have never counted what the sprawl costs them.

The free tier is the whole product at small quotas, not a crippled preview

It removes the evaluation barrier entirely for a graduate student or a postdoc, and those people are usually the ones who bring the group across.

One typed API drives everything, so existing notebooks and dashboards keep working

Research software engineers are the gatekeepers, and this answers their real question: can I build on it without rewriting what I have.

We label what is not finished, in public

In a market that reads carefully, being the vendor who names its own gaps first is worth more than a longer feature list.

The objections below are close to what prospects actually say. The honest answers work better than clever ones here, and every response leans on something that can be checked.

"You have no customers I can call. Why would I bet a research programme on you?"

You would not, yet — and we would not ask you to. The realistic first step is one group, one archive and one pipeline on a paid tier, with a defined thing to prove. Nothing you put in comes out as a proprietary format, and a release you publish stands on its own. Judge the product on what you can see working rather than on somebody else's logo.

PROOF · Anyone can sign up on the free tier today and put real data through the whole loop before any money changes hands.

"Our archive is free. Why would we pay for this?"

The archive stays. It costs you nothing and it is authoritative — connect it as a source. What you are paying for is everything that happens after the download: the catalogue across all your holdings, the versioned processing, the pinned environments, the access grants and the citable output. The archive was never trying to do those things.

PROOF · The competitive material positions mission archives as sources and publication targets, not as rivals; the platform is designed to consume them.

"We could build this ourselves. We have engineers."

You could build parts of it, and many groups have. What tends not to get built is the boring half — roles, audit, tenancy, quotas, notifications, files, tagging, multi-lingual interface, accessibility — and that half is what turns a lab tool into something an institution will approve. The other question worth asking is what happens when the person who built it leaves.

PROOF · A whole platform layer is inherited here rather than reimplemented, which is why the governance is uniform across every product surface.

"Your pipelines do not actually run the compute. That is a big gap."

It is, and we say so in public rather than waiting for you to find out. Today the platform holds the definition, validates the run, records every step's inputs and outputs, and gives you the lineage — while the compute runs on your own infrastructure and reports back. Managed execution is being worked on. For most teams the lineage and versioning were the missing piece anyway; the compute they already have.

PROOF · The public site labels managed pipeline compute as in progress, alongside hosted notebook execution and co-located compute as planned.

"If the identifier is not minted by a registrar, is a release really citable?"

The workflow is complete — version, abstract, authors, licence, lifecycle and the identifier field — and the identifier is supplied at publish time rather than minted automatically. Automatic registration is being built. For teams that need registrar-minted identifiers today, the practical answer is to publish here and mirror to a repository for the identifier, which the export path supports.

PROOF · The release lifecycle is built and shipped; registrar registration is labelled in progress on the public site.

"What certifications do you hold? Our procurement will ask."

None yet, and that is the direct answer. SOC 2, ISO 27001 and ISO 27701 are on a published roadmap with target dates in 2027 and their prerequisites named. What exists today is encryption in transit, multi-factor authentication, tenant isolation, per-environment secret management, permissions enforced at the API edge and a full audit trail — plus a data-processing agreement on request. If a certification is a hard gate for your procurement, tell us now rather than after a pilot.

PROOF · The certification roadmap names each target, its owner and the work it depends on; the security page describes the current controls.

"What happens to our data if we leave, or if you disappear?"

Data above a reduced limit becomes read-only rather than deleted, and records you are legally required to keep are always retained and exportable. Releases are designed as standards-compliant packages so they can be preserved elsewhere, and export to an external repository is on the roadmap for exactly this reason. It is a fair question to put in writing before you sign, and we would rather answer it in the contract than in a conversation.

PROOF · Reference-rather-than-hoard architecture means artefacts sit in shared file storage under one policy, and the pricing terms commit to retained, exportable records on downgrade.

"You are a small company. What is your support like when something breaks at two in the morning?"

Support is tiered and published: community with one ticket a month on Free, email under 48 hours on Pro, priority under 24 hours with a 99.5% uptime commitment on Growth, and priority under 8 hours with 99.9% on Business. Enterprise adds premium service levels. Support-hour packs and a monthly retainer are available where a team wants named help. Small also means you will talk to the people who build it.

PROOF · Response targets and uptime commitments are stated per plan in the published pricing, along with support-hour add-ons.

"Our data is under embargo and some of it is export-controlled. Can we even use a cloud service?"

Access levels are enforced per dataset — public, workspace, restricted or embargoed — and a collaborator gets a scoped grant with an expiry that you can revoke at any point, with the whole trail auditable. Sources reference credentials rather than storing them. For genuinely sensitive programmes, data residency, dedicated tenancy and self-hosting are Enterprise options; that conversation should start early rather than late.

PROOF · Grants, embargoes and revocation are built and shipped; residency, dedicated tenancy and self-hosting are named Enterprise options.

"Is the AI just marketing? Everyone claims it now."

Half of it is deliberately labelled as not built, which should tell you something. What works today: an assistant on every screen that knows which record you have open, a published tool surface an external agent can drive under your own permissions, and an explicit mark on any analysis a machine helped produce. What does not: natural-language catalogue search, anomaly detection, and the automated proposals. The detection model that would produce results is open work.

PROOF · The public AI page labels every item shipped, in progress or roadmap, and states that permissions, audit and the assistance mark were all in place before any model was wired in.

17 Discovery / Qualification Questions

Ten questions that separate a real opportunity from a polite conversation. Ask about reproducibility and release obligations early — those are the triggers that move budget.

ASK THIS

LISTEN FOR

How many different archives or instruments does your team pull data from in a typical project?

Three or more is a strong signal. One archive and nothing else usually means the pain is not big enough yet.

If a reviewer asked you to reproduce a figure from a paper two years old, what would that take?

Laughter, a pause, or an honest 'we could not' is the opening. A confident answer means they have already solved it and you should move on.

Where do your calibration and reduction scripts live, and who else can run them?

One person's laptop or home directory is the classic bind. Listen for anxiety about that person leaving.

Does your current funding require a data-management plan with versioning and provenance?

A yes with a deadline is the single most reliable budget trigger in this market.

How do you share data with collaborators outside your organisation today?

Email, cloud drives or a physical disk. Follow up with: could you revoke that access if you needed to?

Do you have any public data-release obligation, and how much effort does a release take?

Facilities and publicly funded observatories usually have a mandate and a manual process they dread. Quantify the weeks.

How long does a new student or postdoc take before they produce their first result?

Two to three months of technical onboarding is common and is a cost a principal investigator feels personally.

Who maintains the plumbing — the ingest scripts, the notebook servers, the access rules?

A named research software engineer means you have a technical champion. Nobody means the pain is spread thin and the sale is slower.

What would have to be true for you to move this off its current setup?

A concrete condition — a certification, a compute story, a migration path — tells you exactly what to hand us. Vagueness means there is no real trigger yet.

If we did this, who else would need to approve it — IT, security, procurement?

Map the approval path before it becomes a surprise. Institutional buyers almost always have a security review waiting.

18 How to Identify a Good Opportunity

Score the conversation as you go. Green signals are worth acting on the same week; red ones are worth walking away from cheerfully, because a bad fit here costs everybody months.

STRONG SIGNAL – PURSUE

- They pull data from more than one archive or instrument and complain about it unprompted
- A funder or journal has asked them for a data-management plan, provenance or a reproducibility commitment
- They have a public data-release obligation and a manual process to meet it
- There is a named research software engineer or research computing lead who maintains shared plumbing
- They share restricted or embargoed data with external partners and cannot currently revoke access
- Somebody has recently left and taken the only working knowledge of a pipeline with them

Introduce us straight away.

WORTH A CONVERSATION

- Real pain but no budget line yet — often means waiting for the next grant cycle, so register it and stay in touch
- Enthusiasm from a postdoc or student with no route to the principal investigator who signs
- A hard requirement for a certification we do not hold yet, where the timeline might still work
- Heavy existing investment in an orchestration or lakehouse stack that a team is reluctant to disturb

Qualify further before an introduction.

POOR FIT – DO NOT CHASE

- They want free crowd classification or public volunteer engagement — that is a different kind of platform
- They need planetary-scale compute or a full machine-learning lifecycle as the primary requirement
- Open-source, self-operated software is a stated precondition rather than a preference
- They only want somewhere to deposit finished files and get an identifier; a free repository serves them better

Politely park it; suggest another product if one fits.

THE ONE-QUESTION TEST

Ask: 'If someone asked you today to reproduce a result from two years ago, could you?' A team that cannot — and knows it matters, because a funder or a reviewer has made it matter — is a real opportunity. A team that can, or does not care, is not.

19 How to Introduce Us

Keep the introduction short and specific to what you heard. Research people are sceptical of vendor language and respond to a named problem. Do not attempt to demonstrate the product or quote a price — pass it to us early and stay in the loop.

IN CONVERSATION

"Say something like: 'You mentioned you cannot reproduce that analysis any more. There is a platform built for exactly that — it catalogues everything you hold, versions the processing, pins the environment behind every notebook, and turns finished work into something citable. I know the founder. Want me to introduce you? It takes twenty minutes to see whether it fits.'"

EMAIL — SUBJECT: REPRODUCIBLE DATA PLATFORM FOR YOUR GROUP — WORTH 20 MINUTES?

Hi [Name],

After our conversation about [the archive sprawl / the reproducibility problem / the data-release deadline], I wanted to point you at something.

CosmicIntersection is a data platform built for research teams. It catalogues everything you hold across archives and instruments, turns calibration into versioned pipelines that record what each step did, pins the environment behind every notebook, and lets you publish a finished analysis as a citable release. Restricted and embargoed data is shared through scoped grants you can revoke.

It is early — no reference customers yet, and they are open about what is still being built. Worth a look for that reason as much as any.

Happy to introduce you to Vignesh, who leads it: vignesh@burdenoff.com. Or book directly: <https://calendly.com/vignesh-burdenoff/1-1>

Best,
[Your name]

SHORT MESSAGE / LINKEDIN

Saw your post about [reproducibility / the data release / archive sprawl]. There is a platform for research teams that puts the catalogue, the processing pipelines, the notebooks and the citable release in one governed place — CosmicIntersection, at cosmicintersection.com. Early product, honest about what is not finished yet. Happy to introduce you to the founder if it looks relevant.

WARM THREE-WAY INTRODUCTION

Introduce by problem, not product: 'You two should talk. [Name] runs a group that cannot reproduce its own analyses and has a funder asking for a data-management plan. Vignesh (vignesh@burdenoff.com, +91-7358445777) builds a platform for exactly that.' Copy us both in and let us take the first meeting from there.

Please do not: Do not promise a capability that is not shipped — managed pipeline compute, hosted notebook execution, registrar-minted identifiers, anomaly detection and natural-language search are all in progress or planned; Do not quote or negotiate a price; regional band, taxes and any approved discount all change the number, and every quote has to be approved by us; Do not mention a dataset or template marketplace — it does not exist and must not appear in any quote or order form; Do not imply certifications, reference customers or usage figures we do not have; Do not offer to run the demonstration yourself; the walkthrough works because it is shaped around the prospect's own archives and pipelines

20 Lead Handoff Process

Register the opportunity before you introduce anyone. Registration is what makes a commission payable, and it protects you from a second partner claiming the same account.

1 Register the opportunity first

Email the details to Vignesh T.V. at vignesh@burdenoff.com before procurement or any active sales engagement starts. There is no retroactive registration, and no commission if we already had an active opportunity on that account unless you bring a new buying centre or materially expand the scope.

2 Send the qualification summary

The organisation, the named contact and role, the trigger you heard, the archives and instruments involved, rough team size, and any deadline. Two paragraphs beat a form.

3 Make the introduction

Email or a warm introduction with us copied in, or send them straight to <https://calendly.com/vignesh-burdenoff/1-1>. Say what you already told them so we do not repeat it back.

4 Stay on the first call if you can

Your presence keeps the trust you built in the room. You do not need to present anything — we will run it and you can correct the record where you know the customer better than we do.

5 Track it to a decision

An approved registration is protected for 90 days, with one further 90-day extension possible where there is active progress. Check in with us monthly so it does not lapse quietly.

SEND US THIS — IT IS ALL WE NEED

- Organisation name, country and rough size of the team that would use it
- The named contact, their role, and whether they can sign or will need someone who can
- The trigger you heard — a reproducibility failure, a grant requirement, a release deadline, a departure, a collaboration that has outgrown email
- Which archives, instruments or data sources they work with, and roughly how much data
- Any hard constraint: a required certification, data residency, self-hosting, procurement process or a fixed timeline
- What you have already promised or implied, so we do not contradict you

WHAT HAPPENS NEXT, AND WHEN

We acknowledge a registration within one working day and confirm whether it is approved, already active with us, or needs more detail. First contact with the prospect happens within three working days of approval, and we tell you what came of it either way.

Deal registration: Registration is required before commission eligibility begins and cannot be applied retroactively. An approved registration is protected for 90 days, extendable once by a further 90 days where there is active progress. Where two partners claim the same opportunity, the first approved registration with real contribution takes priority, and any split needs written approval. We may ask the customer to confirm your involvement, and we control final pricing, discounting, contract terms and commission eligibility.

21 Partner / Referral Commercials

Partners are paid on collected net subscription revenue — what the customer actually pays us, after taxes, refunds, credits and pass-through costs are excluded — rather than on list price or contract value. Regional pricing does not reduce your eligibility; it simply means the collected amount is lower in that market. Choose the programme that matches how much you actually do: an introduction, an introduction plus evaluation support, or an introduction plus delivery.

PROGRAMME	YOUR ROLE	FIRST-YEAR	RENEWAL / ADD-ONS
Referral Partner	Registered introduction, no sales ownership	10% of first-year collected net subscription	0% unless separately approved · 3% of first-year collected net add-ons, capped at \$10,000
Data-Platform Consulting Partner	Introduction plus evaluation support	12% of first-year collected net subscription	0–3% if active in the success plan · 3% of first-year collected net add-ons, capped at \$15,000
Implementation Partner	Introduction plus implementation delivery	15% of first-year collected net subscription	3–5% if active in the renewal success plan · 5% of first-year collected net add-ons, capped at \$25,000
Authorised Reseller	Approved resale with first-line commercial ownership	Up to 15% reseller margin on the approved quote	Up to 8% if the account stays active and supported · Up to 5% on dedicated tenancy or managed services, only where approved
Managed service or strategic partner	Bundled or managed offering under the CosmicIntersection name	Custom	Custom · Custom

THE RULES THAT MATTER

- Register before procurement or active sales engagement — there is no retroactive registration, and an approved registration is protected for 90 days with one possible extension.
- Commission is calculated on collected net subscription revenue only; taxes, refunds, chargebacks, credits, payment fees, customer-owned cloud costs and AI provider costs are all excluded.
- Larger deals carry an extra incentive on first-year net subscription: 1% capped at \$2,500 from \$100,000, 2% capped at \$7,500 from \$250,000, and 3% capped at \$25,000 from \$500,000 — each requiring prior registration and an approved quote.
- Dedicated tenancy, self-hosted, air-gapped, compliance and managed-service packages carry lower partner economics and must be approved in the quote.
- Partner commission does not stack with affiliate rewards; a customer is attributed to one programme unless we approve an exception in writing.
- Never promise a feature, a compliance status, a discount or a timeline without our written approval, and never make side agreements, rebates or off-contract commitments.
- Resellers must be approved in writing before quoting, and every quote needs our approval before it reaches the customer.

GETTING PAID

Paid after revenue is realised, normally 15 days after successful realisation once settlement, refund, cancellation, fraud, tax and payout-document checks clear. Multi-year deals are paid as customer cash is collected rather than on total contract value. Refunds, chargebacks and customer credits reduce or reverse the commission. Partners handle their own tax forms and local compliance; we withhold and report where the law requires it. Currency follows the partner agreement.

These are the standard partner programme terms as published. They are indicative and not an offer. The terms that actually govern any payment are those in your signed partner agreement and in the quote we issue to the customer. We control final software pricing, discounting, contract terms and commission eligibility in every case.

22 Contact Information

YOUR CONTACT

Vignesh T.V.
vignesh@burdenoff.com
+91-7358445777

BOOK A 1-1 CALL

<https://calendly.com/vignesh-burdenoff/1-1>

Use this for anything — qualification, pricing, a joint call with your prospect, or a live demo.

EVERY BURDENOFF PRODUCT

<https://burdenoff.com/products>

If CosmicIntersection is not the right fit, the answer may well be on that page.

WHICH CHANNEL FOR WHAT

- **A live opportunity** — email with the handoff details in section 20, and copy the phone number in if it is time-sensitive.
- **A question you cannot answer** — book the 1-1 call; it is usually faster than email.
- **A prospect who wants to meet us** — send them the calendly link directly, or ask for a joint call.

23 Useful Links & Sales Assets

ASSET	WHERE	WHAT IT IS FOR
Product website	https://cosmicintersection.com	The public positioning, features and honest capability labels — start here
Sign in or sign up	https://app.cosmicintersection.com	The product itself; the free tier needs no card
Documentation	https://docs.cosmicintersection.com	Product guides and API reference, still being completed
Use cases	https://cosmicintersection.com/use-cases	Scenario stories by persona — useful for finding the right angle for a prospect
Comparisons	https://cosmicintersection.com/vs	Head-to-head pages with the honest gaps written out; read before a competitive conversation
AI capabilities	https://cosmicintersection.com/ai	Every AI item labelled shipped, in progress or roadmap
Security	https://cosmicintersection.com/security	Current controls and posture — send this when procurement asks
Book a 1-1 call with Vignesh T.V.	https://calendly.com/vignesh-burdenoff/1-1	Direct line for a walkthrough or a partner conversation — vignesh@burdenoff.com, +91-7358445777
All Burdenoff products	https://burdenoff.com/products	The wider portfolio, in case a prospect's need sits somewhere else

WHAT YOU CAN SHARE, AND WITH WHOM


- Share freely and publicly: anything on cosmicintersection.com, including the comparison pages, the use-case stories, the security page and the AI capability labels. It is written to be quoted.
- For prospects only, not public posts: this playbook, the plan prices and quotas as printed here, and the objection responses. Share them in a conversation, not on a website or a social post.
- Needs our written sign-off before use: any quote or price you put in writing, any use of the CosmicIntersection or Burdenoff name in your own marketing, any co-branded material, and any public statement that names a customer.
- Never share and never repeat: security findings, roadmap dates, internal architecture detail, and anything about the dataset or template marketplace, which does not exist.

24 Product Mockups — Complete Visual Appendix


32 SCREENS

Every screen in the CosmicIntersection design set — 32 in total, none omitted. Use these to show a prospect what the product looks like before a live demo.


PRODUCT SCREENS • 31 SCREENS




001 Cosmic 01 Dashboard




002 Cosmic 02 Datasets List




003 Cosmic 03 Dataset Detail




004 Cosmic 04 Dataset Form




005 Cosmic 05 Discovery Search




006 Cosmic 06 Sources List




007 Cosmic 07 Source Detail




008 Cosmic 08 Collections




009 Cosmic 09 Collection Detail




010 Cosmic 10 Pipelines List




011 Cosmic 11 Pipeline Detail




012 Cosmic 12 Pipeline Builder




013 Cosmic 13 Runs List




014 Cosmic 14 Run Detail




015 Cosmic 15 Notebooks List




016 Cosmic 16 Notebook Detail




017 Cosmic 17 Analyses




018 Cosmic 18 Releases Publish




019 Cosmic 19 Access Grants




020 Cosmic 20 Lineage Graph




021 Cosmic 21 Landing




022 Cosmic 22 Command Palette




023 Cosmic 23 Notifications




024 Cosmic 24 Mobile




025 Cosmic 25 Login




026 Cosmic 26 AI Assistant Dataset Context



027 Cosmic 27 AI Assisted Analysis Run




028 Cosmic 28 MCP Agent Catalog Session




029 Cosmic 29 AI Schema Proposal Review



030 Cosmic 30 AI Run Triage Annotation



031 Cosmic 31 AI Release Draft Review



032 Satellite Calibration Reprocessing Tracker 1 Reprocessing Campaign



Found someone who needs CosmicIntersection?

You do not need to be an astronomer to spot this opportunity. Listen for a team that cannot reproduce its own work, or a facility with a public data-release obligation it dreads, and make the introduction. We will take it from there, keep you informed at every stage, and pay on what the customer actually pays us.

EMAIL OR CALL

Vignesh T.V.

vignesh@burdenoff.com

+91-7358445777

BOOK A 1-1 WITH VIGNESH

<https://calendly.com/vignesh-burdenoff/1-1>

ALL BURDENOFF PRODUCTS

<https://burdenoff.com/products>



Burdenoff Consultancy Services Pvt. Ltd. · CosmicIntersection is available at cosmicintersection.com.

This playbook is confidential and intended for Burdenoff partners. Figures, plans and capabilities are current at the date of issue and may change; final commercial terms are those set out in your partner agreement and in the quote issued to the customer.