



SALES & PARTNER PLAYBOOK

FluidGrids

FluidGrids lets an operations team draw a process on a canvas and have that drawing run, reliably, in the background.

Visual canvas

Embeddable SDK

Run forensics

AI agents

Workspace RBAC

PREPARED BY
Burdenoff Consultancy Services Pvt. Ltd.

PRODUCT
www.fluidgrids.ai

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 within a minute whether a conversation is worth pursuing. Use the personas and objections sections live, in the room, rather than trying to memorise them. Everything in the capabilities table is marked shipped, in progress or planned, and those labels are accurate, so quote them as written and never soften a limitation into a promise.

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

WWW.FLUIDGRIDS.AI

FluidGrids is a visual workflow automation platform. You drag steps onto a canvas, wire them together, and FluidGrids turns that picture into a versioned workflow that runs in the background on its own worker pool. Steps come from a catalogue of connector nodes for services such as Slack, Stripe, GitHub, Salesforce, HubSpot and Google Workspace, plus built-in nodes for branching, looping, merging, filtering and running your own code. Every run keeps a record of what each step received, returned and how long it took, so a failure is diagnosed rather than guessed at. Uniquely, the whole builder is also published as a component library, so a software company can put FluidGrids automation inside its own product, where it keeps the FluidGrids name.

CATEGORY

Visual workflow automation and orchestration, for operations and automation teams, digital agencies, and SaaS product teams that want to offer automation to their own customers.

WHERE IT STANDS TODAY

FluidGrids is live and open for sign-up with a free tier, and the core loop of build, trigger, run and observe works end to end today. It is a young product: there are no public reference customers, no published usage figures and no third-party security attestations yet. Self-serve card checkout for paid plans is still being finalised, so paid subscriptions are arranged with us directly at the moment.

HOW CUSTOMERS GET IT

Fully managed in the cloud. There is no self-hosted edition today, and that is a genuine gap for buyers with a hard data-sovereignty requirement. A customer signs up, creates a workspace, and is building on the canvas within minutes; there is no installation and nothing to operate. Data residency defaults to India, with other regions available as an Enterprise arrangement. Enterprise customers can also take a dedicated tenant with isolated compute and data.

Web application at app.fluidgrids.ai: canvas, runs, connections, templates, analytics, billing and administration

Embeddable React package that puts the builder, the runs view and the connections UI inside a third-party product

Full API, so anything the interface does a script or another system can do

Hosted forms and webhook endpoints that let outside systems and people start a run

An in-app assistant that drafts workflows and scaffolds new connector nodes from a plain-English description

Marketing site and documentation at www.fluidgrids.ai and docs.fluidgrids.ai

Mobile builds of the application shell, in progress and not yet in the app stores

Command-line tool, language SDKs and a Model Context Protocol server, planned and not yet published

02 One-Line Pitch

SAY EXACTLY THIS

"FluidGrids lets an operations team draw a process on a canvas and have that drawing run, reliably, in the background."

SHORTER, FOR A CHAT MESSAGE

Draw the process on a canvas; FluidGrids runs it and shows you every step.

THE CATEGORY SENTENCE

FluidGrids is a visual workflow automation platform for operations teams, agencies and SaaS product teams that need automation their own customers can see.

03 30-Second Elevator Pitch

SPOKEN, ROUGHLY 30 SECONDS

"Most teams run on glue work: moving data between tools, chasing approvals, syncing systems that were never meant to talk. FluidGrids is visual workflow automation. You draw the process on a canvas, and that drawing is what runs. Every save becomes a version, every run records what each step did, and when one step fails you retry just that step. The same builder also ships as a component you can drop into your own product. Draw it once, and trust it to run."

IF THEY ARE TECHNICAL

FluidGrids stores a workflow as an immutable, versioned graph and executes it on a pool of background workers, so runs survive a closed browser and scale by adding replicas. Each node records its own state, inputs, outputs and a typed error class, so recovery is a single-node retry rather than a full replay. The entire builder is published as a React package, so a product team can mount the canvas, the runs view and the connections UI under its own routes and sign-in, following its light or dark theme.

IF THEY ARE A BUSINESS BUYER

FluidGrids replaces the spreadsheets, scripts and inbox chasing that hold a process together with one diagram that actually executes, and gives the team proof of what ran and what failed. Pricing is by usage rather than by feature, so the whole builder and the whole connector catalogue are available from the free tier. For a software company, the same product also becomes a way to offer automation to its own customers without building an engine.

04 Problem We Solve

These are the conversations that turn into FluidGrids opportunities. You are listening for a process that matters, spread across systems that do not talk, held together by a person.

The process document cannot execute

Who feels it: Operations leads, process owners, COOs

What it costs them: There is a wiki page describing how orders are handled or leads are routed, and a reality that drifts further from it every quarter. Nobody can say with confidence what actually happens today, so onboarding is slow and every audit begins with archaeology.

How FluidGrids answers it: FluidGrids makes the diagram the implementation. The canvas you draw is stored as the executable workflow, so there is no second artefact to keep in sync. Change the picture and you have changed the behaviour, and the version history shows exactly when it changed and who changed it.

Every integration is a bespoke script somebody wrote once

Who feels it: IT managers, platform engineers, small internal tooling teams

What it costs them: The scripts live on somebody's laptop or in a repository nobody reviews. When the person who wrote them leaves, or an API changes, the automation breaks quietly and the manual workaround comes back.

How FluidGrids answers it: Connector nodes are versioned packages that describe their own operations and render their own configuration form. Choosing an operation and filling in a form replaces reading API documentation, and the node pins its version so an upgrade elsewhere does not silently change an existing workflow.

A failed automation gives you nothing to work with

Who feels it: Operations engineers, support leads, anyone on call

What it costs them: The alert says the job failed. Finding out which step, on what input, and whether it is safe to re-run takes hours, and the usual answer is to run the whole thing again and hope no side effect is duplicated.

How FluidGrids answers it: FluidGrids records state per step, not per run: the input each step received, what it returned, its duration and a typed error class such as authentication, validation, timeout or rate limit. Recovery is a retry of the one failed step with its original input, so a pipeline that died at step four does not restart at step one.

Credentials are pasted into a dozen places

Who feels it: Security leads, operations admins, agency delivery teams

What it costs them: The same API key is duplicated across tools and configuration files, nobody is sure what breaks if it rotates, and rotation therefore never happens. That is how an automation project becomes a security finding.

How FluidGrids answers it: Connections are stored once per workspace and referenced by name. A saved workflow contains a reference, never the secret itself, and the secret is resolved at the moment the step runs. Rotating a credential is one edit in one place, and every workflow that uses it picks it up on its next run.

Editing the live automation is a leap of faith

Who feels it: Automation leads, finance and compliance owners

What it costs them: Someone changes the workflow that moves money, and there is no clean way to answer what changed, what ran before the change, or how to get back. Incident reviews become guesswork and change control becomes a spreadsheet.

How FluidGrids answers it: Every save mints an immutable version under a stable workflow key, and a run pins to the exact version it started with. Editing an active workflow creates a draft beside it while the active version keeps serving triggers. Rolling back is one operation that brings a previous version forward as the newest one.

Automation costs climb faster than the work does

Who feels it: Finance, RevOps, anyone who signs the renewal

What it costs them: Per-task and credit-based pricing means a workflow with more steps costs more every time it runs, so teams shorten workflows to save money and end up with worse automation. The bill at renewal is a surprise.

How FluidGrids answers it: FluidGrids meters a completed run as one run, whatever the step count, and gates plans by scale rather than by feature. The builder, the connector catalogue, triggers, versioning and the run history are all available on the free tier.

Customers keep asking for automation inside the product

Who feels it: Heads of product and engineering at software companies

What it costs them: Building a canvas, an execution engine, a connector catalogue and the governance around them is several quarters of platform engineering that does not differentiate the product. The alternative is sending customers to a third-party tool and losing the workflow.

How FluidGrids answers it: The builder that powers the FluidGrids application is published as a React package with several integration depths, from a single drop-in canvas to a fully routed automations module. It renders under the host product's own routes and sign-in and follows its light or dark theme, so the automation feature sits inside the product rather than in a separate tool.

Dashboards are fed by somebody exporting a spreadsheet

Who feels it: Data analysts, revenue operations, leadership reporting

What it costs them: The weekly numbers depend on a person remembering to export, reshape and paste. When they are on holiday, the dashboard is stale, and nobody trusts a figure they cannot trace.

How FluidGrids answers it: A single data sink node at the end of a scheduled workflow pushes the output into a live dashboard feed in BigConsole, our companion analytics console. The workflow becomes the pipeline: scheduled, observable, retryable, with the dashboard repainting each time the run lands fresh data.

05 Value Proposition

Five things carry a FluidGrids conversation. Lead with whichever one matches the pain the prospect just described, and leave the rest for the follow-up.

The drawing is the software

There is no gap between the diagram of a process and the implementation of it, because they are the same object. An operations manager can read the canvas and an engineer can extend it with a code step, without either of them translating for the other. When the process changes, the picture changes, and the change is versioned.

Evidence: Saving the canvas persists the whole graph as a versioned workflow definition; reload and the identical graph, positions and configuration come back.

Failures are diagnosed, not re-run

Every step keeps its own record: what it received, what it returned, how long it took and which class of error it hit. The run list is a triage queue and the run detail pinpoints the step that broke. Recovery is surgical, so a long pipeline that fails near the end does not have to repeat the work that already succeeded, and no side effect is duplicated.

Evidence: Cancel, pause, resume and single-node retry are ordinary operations in the product, and a live log streams while the run is still going, including anything a customer's own code step prints.

Change control that a compliance team recognises

Workflows behave like production software rather than like documents. Each save mints an immutable version, activation is a deliberate release that also arms the schedules, and deactivation disables them in the same act. A run permanently records which version it executed, so a failure attributes to a specific release rather than to the workflow in general.

Evidence: Draft, active and archived are explicit states; restore brings any historical version forward as a new draft; the permission to edit a draft is separate from the permission to activate one.

Automation you can ship inside your own product

This is the differentiator no competitor leads with. The builder is not a separately hosted tenant or an embedded frame; it is the same set of React components the FluidGrids application itself is built from, published as a package. A product team mounts it under its own routes, hides the header or sidebar where its own layout already has one, and mints the tokens, so tenancy and permissions stay under the host's control. FluidGrids is offered under its own name — there is no white-label or rebranded edition.

Evidence: Several integration depths are supported, from one drop-in builder to a fully routed automations module, with view-only mode, chrome controls, trusted-origin protection and callbacks for save, run and navigation.

Enterprise plumbing on day one

Multi-tenancy, role-based access control, audit trails, usage metering, billing, notifications and single sign-on are inherited from the wider Burdenoff platform rather than rebuilt for this product. That is unusual for a young automation tool, and it is why an access review does not turn into a project.

Evidence: Six roles are seeded into every workspace automatically, from Automation Owner down to a read-only Automation Auditor, and permission checks apply to every workflow and run operation, including calls made through the embedded builder.

Cost that tracks the work

One completed run counts as one run, however many steps it contains, and plans differ by quota rather than by which features are unlocked. Overages are blocked rather than silently billed, so a runaway workflow produces a stop, not an invoice. That predictability is the most common reason a team leaves per-task pricing.

Evidence: The full builder, the connector catalogue, all five trigger types, versioning and the run history are available from the free tier; paid tiers raise workflow, run, concurrency and retention limits.

06 Ideal Customer Profile

A 100 to 1,000 person company where one team owns a business-critical process that crosses four or five systems, has already outgrown a simple task-automation tool, and has somebody accountable when it breaks.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
Mid-market operations and shared-services teams	100 to 1,000 employees	They have real process complexity and audit expectations, but no capacity to build automation infrastructure. They are usually already paying for a task-automation tool and feeling the cost curve and the lack of branching, versioning and error detail.

SEGMENT	TYPICAL SIZE	WHY THEY BUY
SaaS and technology product companies	50 to 500 employees	Their roadmap has grown an automation line item because customers keep asking. The embed path turns several quarters of platform engineering into a sprint of integration, and the buying decision moves quickly because the buyer is technical.
Digital agencies, consultancies and managed service providers	10 to 200 staff, many client accounts	They sell the same automation repeatedly. One workspace per client, a template library, and per-client connections make the thirtieth delivery cheap, and read-only embedded views let clients watch their own automations without editing them.
Enterprise process owners in finance, HR, IT and customer service	1,000 employees and above, buying for one function	They need approvals inside automated runs, separation of duties, and evidence for an access review. Inherited role-based access and audit make the governance conversation short, though formal attestations are still pending.
Data and analytics teams	Any size with a reporting burden	They want a scheduled, observable pipeline that feeds a live dashboard without standing up a separate data tool. The data sink node makes the last step of a workflow the dashboard feed.

GEOGRAPHIES

Sold worldwide with pricing set in four regional bands, so the same plan costs less in India, South and South-East Asia, Latin America, Africa and the Middle East than in North America and Western Europe. India is the home market and the default data-residency region; other regions are available on Enterprise terms.

NOT A FIT – DO NOT CHASE

- A hard requirement to run the platform inside the customer's own infrastructure or an air-gapped network. There is no self-hosted edition, and open-source alternatives will win that deal honestly.
- The value of the deal rests on connector breadth across a very long tail of niche applications. The palette holds around fifty steps today; the large incumbents count theirs in the thousands.
- A procurement gate that requires a current SOC 2 Type II or ISO 27001 certificate before signature. Both are targeted, neither is held, and pretending otherwise wastes everyone's quarter.
- Desktop or legacy screen-scraping automation, or anything that needs a robot to drive a Windows application. That is a different category of product entirely.
- A single one-step notification that a free tier of a consumer automation app already handles. There is no problem there worth a sales conversation.
- Workflows that must survive for days and resume from the exact instruction after a crash. Durable queues with per-step state are not event-sourced replay, and a code-first orchestration engine is the honest answer.

07 Buyer Personas

The same product is bought for four quite different reasons. Open with the line that matches the job title in front of you rather than the feature you find most interesting.

Head of Operations or automation lead

WHAT THEY WANT

Get a process off email and spreadsheets, make it visible to the team, and stop being the person everyone messages when it breaks. Prove the improvement to their own leadership.

Open with: Which of your processes breaks most often, and how do you find out that it broke?

They will ask: How long before my team can build something themselves without an engineer?

WHAT FRUSTRATES THEM

Manual handoffs, no visibility into what ran, a task-automation bill that rises with volume, and no way to change anything without risking the live process.

Head of Product or Engineering at a software company

WHAT THEY WANT

Ship an automation feature customers are asking for without spending two quarters on a workflow engine, and keep it inside their own product rather than sending users away.

Open with: Are your customers asking to connect your product to the rest of their stack, and what have you told them so far?

They will ask: What does the integration actually look like, and who owns authentication and tenancy?

WHAT FRUSTRATES THEM

A build-versus-buy decision with no good buy option: embedding usually means an iframe or a separate tenant that sits outside their own navigation and sign-in.

Platform or integration engineer

WHAT THEY WANT

Retire a pile of one-off scripts, put the remaining glue somewhere observable, and stop being the only person who can fix it.

Open with: Where do your integration scripts live today, and what happens when one fails at three in the morning?

They will ask: Can I drop into code when the visual step is not enough, and can I write my own connector?

WHAT FRUSTRATES THEM

Scripts with no logs, no retries and no owner. Requests for automation queueing behind product work. Custom connectors that nobody has time to write.

Agency or consultancy technical director

WHAT THEY WANT

Deliver the same automation to many clients profitably, keep each client's data and credentials separate, and hand over something the client can watch without breaking.

Open with: How many of your clients are running some version of the same automation, and how long does the next one take to deliver?

They will ask: How is one client's workspace isolated from another's, and what does it cost me per client?

WHAT FRUSTRATES THEM

Every client engagement rebuilds the same thing. Credentials get mixed up. There is no clean way to show a client their automation without giving them edit rights.

Finance, compliance or risk owner

WHAT THEY WANT

Put approvals inside automated processes, prove separation of duties, and produce evidence for an access review without a manual trawl.

Open with: When an automated process is questioned six months later, what evidence can you actually produce today?

They will ask: Which certifications do you hold, and where is our data stored?

WHAT FRUSTRATES THEM

Automation that changes silently, approvals living in inboxes, and no defensible answer to who changed what and which version ran.

Data or analytics lead

WHAT THEY WANT

Replace the manual export ritual with a scheduled pipeline that keeps a dashboard current and can be traced when a number looks wrong.


WHAT FRUSTRATES THEM

Reporting that depends on a person remembering. No lineage from the figure on the dashboard back to the source system.

Open with: Who exports the spreadsheet that feeds your leadership dashboard, and what happens the week they are away?

They will ask: How do the numbers get from our systems into the dashboard, and can I see the run that produced them?

Fourteen situations FluidGrids is genuinely good at. Read the situation lines, not the titles, and use whichever one matches what the prospect has just described.



OPERATIONS OR AUTOMATION LEAD

Turn the process document into the process


THE SITUATION

The team has a wiki page describing how orders, leads or invoices are handled, and a reality that has drifted away from it. Nobody is confident describing what actually happens, and every new joiner learns it by asking somebody.

WHAT HAPPENS WITH FLUIDGRIDS

They open the canvas, drag on a trigger, then the steps: shape the data, branch on a condition, notify a channel. Wiring is drag-to-connect, and the validator refuses connections that make no structural sense before the workflow is ever saved. Dropping a new step onto an existing connection splices it in, and deleting one heals the graph.

One artefact instead of two: a diagram everyone can read that is also the thing that runs.



WORKFLOW DESIGNER OR BUSINESS ANALYST

Configure an integration by filling in a form


THE SITUATION

Connecting to a service means reading API documentation, working out which endpoint, which fields are required and what format a timestamp wants. That research is where most automation projects stall.

WHAT HAPPENS WITH FLUIDGRIDS

They search the palette for the service, drop the node on the canvas and pick an operation such as send message or create row. The configuration panel is generated from the node's own description of itself: required fields enforced, defaults filled in, types respected. Around fifty steps are available on the canvas today, and each connector is versioned so an existing workflow keeps the shape it was built against.

The integration knowledge lives in the node, so configuring one takes minutes rather than an afternoon of reading.



OPERATIONS ADMIN OR SECURITY LEAD

Rotate a credential once, everywhere


THE SITUATION

The same token is pasted into several tools and configuration files. Nobody knows the full list, so nobody rotates it, and the automation project has quietly become an audit finding.

WHAT HAPPENS WITH FLUIDGRIDS

An admin creates a connection once per workspace, either through a consent flow or a schema-driven form for a key. In the builder, a step's configuration offers only connections matching what that step needs, so designers pick from what admins approved and never handle a secret. The saved workflow holds a reference, and the secret is resolved when the step actually runs.

Rotation becomes one edit in one place, and one team's credentials are structurally invisible to another workspace.



DATA ANALYST OR REVENUE OPERATIONS

Retire the weekly spreadsheet export


THE SITUATION

A leadership dashboard is fed by a person who exports a file, reshapes it and pastes it somewhere visible. When they are away the numbers go stale, and nobody can trace a figure back to its source.

WHAT HAPPENS WITH FLUIDGRIDS

A scheduled workflow pulls from the source systems, shapes the data with built-in transform and aggregate steps, and finishes with a data sink node that pushes the result into a live dashboard feed. Leave the data field empty and the node captures the previous step's output automatically. Each delivery reports how many destinations updated and the payload size.

The dashboard updates itself on a schedule, and every number traces back to a run you can open.



OPERATIONS ENGINEER

Close the laptop and the run keeps going


THE SITUATION

Automation that only runs while someone watches is not automation. The team needs long processes to survive a closed browser and to keep up when volume spikes.

WHAT HAPPENS WITH FLUIDGRIDS

Every triggered run is queued and picked up by whichever background worker is free, entirely detached from the browser. Workers share the load without ever executing the same job twice, so adding capacity is adding replicas. Parallel branches genuinely run at the same time and reconverge at a merge step, with five different merge behaviours to choose from.

Throughput scales with capacity rather than with attention, and no run depends on a session staying open.



OPERATIONS ENGINEER OR SUPPORT LEAD

Fix the step that broke, not the whole pipeline


THE SITUATION

A thirty-minute pipeline fails at step four of six. The only recovery available in most tools is to run the whole thing again, which repeats the work that succeeded and risks duplicating side effects.

WHAT HAPPENS WITH FLUIDGRIDS

The failure is classified against a fourteen-type taxonomy covering authentication, network, validation, timeout, rate limit and more, and stored on that step alongside its input and the attempt count. The run list filters to failures; the run detail points at the step. Retry runs that one step with its original input and the run continues from there.

Diagnosis in a minute instead of an hour, and recovery without repeating anything that already worked.



PROCESS OWNER IN FINANCE, PROCUREMENT OR CUSTOMER SERVICE

Put a human decision inside an automated run


THE SITUATION

A refund over a threshold, a contract exception, a supplier onboarding. Full automation is the wrong answer, but so is a fully manual process where the request sits in somebody's inbox.

WHAT HAPPENS WITH FLUIDGRIDS

The run pauses at the point where judgement is needed, holding every completed step's output intact for as long as the decision takes. When the approver acts, resuming the run carries their payload back in, and a branch routes the approved path to fulfilment and the rejected path to a notification. Pause, resume and cancel are each permission-guarded, so you control who may release a gate.

Approvals happen inside the process with a trail that reconstructs itself, instead of in an inbox nobody can audit.



PRODUCT OR ENGINEERING LEAD AT A SOFTWARE COMPANY

Ship an automations tab in your own product


THE SITUATION

Customers keep asking to connect the product to the rest of their stack. Building a canvas, an engine, a connector catalogue and the governance around them is several quarters that do not differentiate the product.

WHAT HAPPENS WITH FLUIDGRIDS

They install the package and mount the builder inside their own application. A spike on day one proves it works with a test workspace. The real engineering is the middle of the sprint: mapping each customer account to a workspace and minting the token the builder forwards, which is the security boundary. Then the full module mounts under a route such as automations, inheriting their navigation, sign-in and light or dark theme.

A roadmap item measured in quarters becomes a sprint of integration, and the automations tab sits inside their own navigation.



AUTOMATION LEAD OR CHANGE MANAGER

Roll back a bad release in one operation


THE SITUATION

Someone edits the workflow that moves money and it starts failing. There is no clean answer to what changed, what ran before the change, or how to get back to the version that worked.

WHAT HAPPENS WITH FLUIDGRIDS

Every save mints a new version under a stable key, so version four never overwrites version three. Editing an active workflow creates a draft beside it while the active version keeps serving its triggers. Activation is the release act and arms the schedules; deactivation disables them in the same operation. Restore brings a previous version forward as the newest one, and schedules return to known-good behaviour in seconds.

Rollback in seconds, with in-flight runs untouched because each one pinned the version it started on.



TECHNICAL BUILDER OR AUTOMATION ENGINEER

Assemble an AI agent out of blocks


THE SITUATION

The team wants an AI step inside a real process, but every path they have tried means writing and maintaining agent code that nobody else on the team can read or change.

WHAT HAPPENS WITH FLUIDGRIDS

An agent node on the canvas takes sub-nodes for its model, its memory and its tools, and one of those tools can be another workflow. Model options include several commercial and local providers, and each tool is bound from the connector's own description, so the agent calls governed operations rather than holding keys. The agent is versioned, permissioned and observable exactly like any other step.

An AI experiment becomes a production step the whole team can read, change and roll back.



AUTOMATION LEAD WHO KNOWS THE PROCESS BUT NOT THE TOOL

Describe the automation and get a working draft


THE SITUATION

The blank canvas is the real barrier. Knowing which steps exist, what their operations are called and what fields they need is a day of learning before the first useful workflow.

WHAT HAPPENS WITH FLUIDGRIDS

An assistant sits on every screen and knows which workflow, run or connection is on screen, so instructions can be conversational. It reads the workspace's live catalogue before writing anything, mirrors the shape of a workflow known to run, and creates an ordinary versioned draft. It never invents a step that does not exist, credentials never travel with it, and nothing is activated on the user's behalf.

A first pass grounded in the customer's real catalogue, arriving as a normal draft they review, bind and activate themselves.



AGENCY OR CONSULTANCY TECHNICAL DIRECTOR

Deliver the same automation to thirty clients


THE SITUATION

The agency sells lead intake, enrichment and CRM routing repeatedly. The economics only work if the thirtieth delivery costs a fraction of the first, and today every engagement rebuilds it.

WHAT HAPPENS WITH FLUIDGRIDS

The master workflow is authored once in the agency's own workspace, with proper error branches and sensible naming, and keeps improving under its own version history. Each client gets a separate workspace, and onboarding is importing the template, which records which template it came from. Per-client work shrinks to binding that client's own connections and adjusting the parameters that genuinely differ.

Isolated copies with recorded lineage, so when the master improves you know precisely which clients are running the old shape.



ACCOUNTS PAYABLE OR FINANCE OPERATIONS LEAD

Clear the routine invoices and escalate the rest


THE SITUATION

The payables inbox is a stream of PDFs, every supplier using a different layout, and a person keys each one into the accounting system by hand. The mistakes are the expensive kind: a transposed total, a bill paid twice.

WHAT HAPPENS WITH FLUIDGRIDS

A mail trigger starts a run when an invoice arrives. An extraction step pulls the text, an AI step reads it and returns the fields with a confidence figure, and a parser pins the answer to a strict shape. A validation step redoes the arithmetic, checks required fields and guards against duplicates. Only invoices that pass every check and clear the confidence threshold post automatically; anything questionable goes to a review channel with the reason attached.

The routine majority clears itself, exceptions arrive with their reasoning, and every decision is inspectable in the run.



INTEGRATION ENGINEER OR TECHNICAL BUILDER

Build the connector that does not exist yet

THE SITUATION

Sooner or later the catalogue does not cover the internal API, the niche application or the bespoke protocol, and on most platforms that is where automation stops and a support ticket begins.

WHAT HAPPENS WITH FLUIDGRIDS

The team either writes the connector package themselves or describes it in plain English and lets the assistant scaffold a first pass, complete with the adapter that makes it callable by an agent. Either way it enters a governed lifecycle: register, validate the structure and description, run automated tests, then review and approve before it can be installed into the workspace catalogue. Once installed it behaves exactly like a built-in step.

The catalogue becomes a surface the customer extends. Note that the deployment side of this pipeline is still being re-platformed, so treat it as in progress.

09 Product Capabilities

What the product does today, grouped by the job it serves. The status against each line is accurate and is what you should quote; anything not marked shipped should be described as coming, never as available.

Designing a workflow

The canvas, and everything that makes a large workflow stay readable and correct.

Drag-and-drop canvas

A free-form canvas with a searchable step palette, drag-to-connect wiring, splice-a-step-onto-a-connection insertion, and layout that keeps a growing workflow readable.

Structural validation before save

Connections that make no structural sense and dangling branches are rejected in the builder rather than discovered at run time.

Self-describing configuration forms

Each connector renders its own configuration form from its published description, with required fields enforced and defaults filled in, so no step needs a hand-built form.

Data mapping between steps

Values are mapped from one step's output to the next step's input with expressions evaluated in a locked-down sandbox.

View-only canvas

A read-only rendering of any workflow, for dashboards, documentation and for showing a client their automation without giving them edit rights.

Templates and duplication

A gallery of the workspace's own proven workflows plus importable templates, with duplication creating a fully independent copy and imports recording where they came from.

Publishing your own templates for others in progress

The submission and review pipeline exists; the seller-facing flow is still being productised.

Connecting to other systems

The catalogue of steps, and how credentials are stored and used.

Connector catalogue

Around thirty connector steps are available on the canvas today, covering messaging, email, code hosting, CRM, payments, accounting, spreadsheets and documents, databases and AI providers, alongside twenty-one built-in flow-control steps.

Built-in flow control

Branch, switch, loop, merge, wait, set, filter, aggregate, transform, validate, limit, sort, de-duplicate, split, compare, log and call another workflow, all executed natively without a network hop.

Code steps

JavaScript and Python steps for the cases the visual palette does not cover, screened against dangerous patterns before they execute.

Stored connections

Credentials are created once per workspace through a consent flow or a form, referenced by name from any step, and resolved at run time so no secret is ever written into a workflow.

Broader catalogue in progress

A further set of connector packages is built but not yet published into the palette. Quote the roughly fifty steps available today, and treat the rest as coming.

Running and recovering

Background execution, and the record that makes a failure diagnosable.

Background execution

Runs are queued and executed by a pool of workers entirely detached from the browser, sharing load without ever executing the same job twice.

Per-step record

Every step keeps its state, input, output, duration, logs and error, live during the run and preserved afterwards.

Typed error classes

Failures are classified across fourteen types, so an expired credential, a bad payload and a downstream outage are told apart automatically.

Single-step retry

Retry one failed step with its original input and the run continues from there, instead of replaying everything that already succeeded.

Run control

Cancel, pause, resume with a payload, and terminate, each guarded by its own permission.

Live streaming run log

A run streams its log while it executes, in the way a build job does, including anything a customer's own code step prints.

Search across workflows and runs

Full-text search over workflows and run history for triage and audit.

Starting a run

The five ways a workflow begins, and how deliveries are audited.

Manual and API trigger

Start a run from the interface or from any client with a workspace token, synchronously or in the background.

Webhook trigger

A unique address per registration; every inbound delivery is validated, logged and individually replayable.

Schedule trigger

Cron-style schedules that are armed by activating a workflow version and disabled the moment it is deactivated.

Hosted form trigger

A hosted form whose submission starts a run, for intake and request processes that begin with a person.

Error trigger

One workflow starts when another fails, so alerting and recovery are themselves automations.

Event filters

Filter expressions on a registration so only matching events start a run.

Triggering from another Burdenoff product's own records

planned

Proposed, not built. Polling and per-application event triggers are deliberately not on the roadmap; point a webhook at the supplier instead.

AI on the canvas

Agents as ordinary steps, and an assistant that drafts workflows and connectors.

Agent step with pluggable parts

An agent step composes a chat model, memory and tools as sub-steps, including another workflow as a tool, so agentic behaviour is built visually rather than in code.

Connectors as governed agent tools

Each connector operation is already callable by an agent through its published description, so the agent uses approved operations and never holds the keys itself.

Assistant-drafted workflows

Describe the automation and the in-app assistant reads the workspace's live catalogue, mirrors a workflow known to run, and creates an ordinary versioned draft with its connections left open for the user to bind.

Assistant-scaffolded connectors

A second assistant mode produces a complete connector package from a description, including the adapter that makes it agent-callable. No competitor on our list generates a connector from a prompt.

Retrieval over the customer's own documents in progress

Grounding an agent step in the customer's own document set is being built.

Assisted repair of a failed run planned

Proposing a fix for a failed run as a change the user approves is on the roadmap, as is overnight triage of failures.

Exposing workflows to outside AI agents planned

Publishing a workspace's workflows as tools an external agent can call is planned, not built.

Embedding in another product

The part of FluidGrids that a software company buys.

Drop-in builder component

One React component that carries its own data layer: pass the endpoint, the workspace and a token, and the full canvas renders inside the host application.

Fully routed automations module

A single root component mounts the whole module, list, builder, runs and settings, under any route in the host's own router.

Headless core for deep integration

A core builder plus a provider, for hosts that want the builder to share their own data layer rather than run its own.

Embed chrome and theme controls

Hide the header, sidebar or log panel, follow the host's light or dark theme, switch to view-only, restrict trusted origins, and seed the canvas with a definition so it renders without a fetch.

Host callbacks

Hooks for save, run, step selection, navigation and back, so the host product can react in its own analytics and its own upgrade flows.

Composable surrounding modules

The runs list, run detail, templates gallery, connection selector and provider picker are all exported separately, so a host embeds exactly as much as it needs.

Governance and administration

The controls an enterprise buyer will ask about in the first meeting.

Workspace isolation

Every workflow, run, connection and credential belongs to a workspace, so one team's or one client's data is structurally invisible to another.

Six seeded roles

Automation Owner, Automation Lead, Workflow Designer, Integration Engineer, Run Operator and a read-only Automation Auditor are created in every workspace automatically, with permission checks on every workflow and run operation.

Separation of duties

The permission to edit a draft is separate from the permission to activate one, so releasing a workflow is a deliberate, attributable act.

Audit and activity trail

Workflow edits, activations and runs are recorded with the actor and the target, giving a who-changed-what record without a manual trawl.

Single sign-on and provisioning

Single sign-on with SAML and directory provisioning are included on Enterprise plans.

Usage metering, quotas and notifications

Runs are metered, quotas are enforced by blocking rather than by silent billing, and failures or quota events raise in-app, email and webhook alerts.

Formal attestations planned

SOC 2 Type II and ISO 27001 are targeted with an audit programme under way. Neither is held today.

Security of what runs

How untrusted code and secrets are handled.

Per-step sandbox isolation

External and custom connector code runs in an isolated sandbox rather than inside the shared worker process, with each run taking its own lease.

Screened code steps

Data-mapping expressions evaluate in a locked-down context, and customer code steps are screened against dangerous identifiers before they execute.

Secrets resolved at run time

Saved workflows carry a reference to a connection, never the secret. The secret is fetched when the step runs and is never written into a definition or an export.

Audited webhook deliveries

Every inbound delivery is validated and logged, and replaying one is an explicit action rather than an accident.

Encryption of stored run payloads in progress

Run inputs and outputs are stored but not yet encrypted at rest as a separate layer. This is a known item on the security backlog and should be disclosed to a security-led buyer.

10 Pricing & Plans

Pricing is by usage rather than by feature. The whole builder, the whole step catalogue, all five trigger types, versioning and the run history are available on the free tier; paid plans raise quotas for workflows, runs, users, concurrency, retention and storage. One completed run counts as one run, whatever the step count. Prices are set in four regional bands, so the same plan costs less in India, South and South-East Asia, Latin America, Africa and the Middle East than in North America or Western Europe. Overages are blocked rather than silently billed. Embedding in another company's product is a separate commercial arrangement, not a tier upgrade.

PLAN	PRICE	WHO IT IS FOR	WHAT IS INCLUDED
Free	\$0 forever No card required. Regional bands do not apply.	Makers, students and anyone evaluating the product	1 user; 10 workflows, 5 of them active; 1,000 runs per month; 3 concurrent runs, 20 steps per workflow; 5 minute maximum run time; 30 days of run history, 7 days of audit retention; 1 GB storage, 3 custom connectors; Community support
Starter	\$19/user/mo \$15/user/mo on annual billing. Band B \$14, Band C \$10, Band D \$6 or the local equivalent; India is Rs 999 monthly, Rs 799 annual.	Solo operators and small teams automating their core processes	1 paid user plus up to 5 guests; 50 workflows, 25 of them active; 10,000 runs per month; 10 concurrent runs, 50 steps per workflow; 15 minute maximum run time; 90 days of run history, 30 days of audit retention; 10 GB storage, 5 custom connectors; Email support within 48 hours
Pro	\$49/user/mo \$39/user/mo on annual billing. Band B \$35, Band C \$25, Band D \$15 or the local equivalent; India is Rs 2,499 monthly, Rs 1,999 annual.	Growing teams running automation the business depends on	1 paid user plus up to 25 guests; Unlimited workflows and runs under fair use; 50 concurrent runs, unlimited steps per workflow; 60 minute maximum run time; 1 year of run history, 90 days of audit retention; 100 GB storage, unlimited custom connectors; Priority support within 4 hours, 99.9% uptime commitment
Enterprise	Custom Annual contract. Quoted per customer.	Regulated teams, large deployments, and companies embedding the builder in their own product	Unlimited users, workflows and custom connectors; Custom run volume, concurrency and maximum run time; Custom run history and audit retention; Single sign-on with SAML and directory provisioning; Dedicated tenant and choice of data region; Negotiated uptime and support commitments; SDK embedding terms

Add-ons

ADD-ON	PRICE
Extra runs, 10,000 per month	\$10
Extra concurrency, 10 concurrent runs	\$15
Extra storage, 10 GB per month	\$5
Extra assistant queries, 100 per month	\$10
Extra API calls, 10,000 per month	\$10
Extra custom connectors, 5	\$5

ADD-ON	PRICE
Premium support, per workspace	\$199/month
Extended audit retention, 1 year	\$50/month
Embedding licence, per deployed host	Custom

WHAT TO TELL A PROSPECT ABOUT PRICE

- Every figure above is copied from our own pricing material and is indicative pre-launch. The written quote we issue to the customer is what governs.
- Self-serve card checkout for paid plans is still being completed, so a paid subscription is set up with us directly today. Do not promise a customer they can upgrade with a card this afternoon.
- Add-ons raise one dimension without forcing a tier upgrade, are pooled across the workspace, and are removed at the end of the period if the plan is downgraded.
- Unlimited on the Pro plan is subject to fair use with concurrency capped at the worker layer; a genuinely heavy workload should be quoted rather than assumed.
- Embedding the builder in another company's product is priced separately from subscriptions and always involves us.
- Never convert a band price into another currency yourself. Ask us for the number in the customer's own currency.

11 Competitive Landscape

Every one of these is a credible product and some of them will be the right answer. Concede the point they genuinely win before you make yours; a table where the competitor never wins is the fastest way to lose a partner's credibility.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Zapier	The best-known task automation service, built around linear step lists connecting a very large catalogue of applications, sold largely to business users.	Catalogue breadth is not close. They connect thousands of applications to our roughly fifty steps, and for a simple, low-volume, one-branch automation they are quicker to set up and cheaper to start. Their generation assistant is also more mature than ours.	Branching, looping and parallel paths are first-class on a canvas rather than bolted onto a list. They bill per task, so a five-step workflow running 10,000 times is 40,000-plus tasks against our 10,000 runs. And we version every save, pin runs to their version, and retry a single failed step.
Make	A visual automation platform with a strong canvas, a large catalogue and an affordable entry price, popular with agencies and technical operators.	Their canvas is the visual benchmark and their catalogue is far larger than ours. Their agent features are ahead of our agent step today, and entry-level pricing is close to ours.	They have no built-in version history, so restoring last week's scenario is a manual export and import. We mint an immutable version on every save and roll back in one operation. Their weighted credits vary by the type of step, so code and AI steps consume unevenly; we count a run as a run. And their embedded offering is a separate hosted tenant rather than components inside your product.
n8n	A developer-friendly automation platform with a free self-hosted community edition and a large, fast-moving library of AI steps.	If a team has the capacity to run it themselves, the licence cost is unbeatable and the data stays entirely in their own infrastructure. Their AI and agent step depth and their community ecosystem are both ahead of ours.	We are managed, so there is nothing to run, patch or scale. Their embedded offering is their whole editor on an instance you operate; ours is managed React components under the host's own routes and sign-in. And they are not natively multi-tenant, so an integrator embedding n8n builds tenant separation themselves, where we inherit it.

ALTERNATIVE	WHAT IT IS	WHERE THEY ARE STRONG	WHERE WE WIN
Pipedream	A developer-first automation platform where workflows are largely written in code, with a large connector catalogue and an embedded authentication component.	For an entirely engineering-led team writing short, code-heavy workflows, they are faster to get moving and their catalogue is far larger. If the requirement is genuinely just embedded authentication, their component is a simpler fit.	Their embedded component gives your users a connection widget, not a builder; we embed the whole builder, runs view and connections. Their steps are capped at a few hundred seconds each; our background workers have no hard per-step cap. And an operations team can read and fix our canvas without an engineer.
Workato	An enterprise integration platform with deep connectors into large business systems, mature governance and an embedded offering sold to software companies.	Their compliance attestations are real and mature, ours are pending. Their enterprise connector coverage across finance, HR and CRM systems is far deeper, and their environment and diff tooling is well established.	Their embedded product is a separate platform tenant; ours is components inside the host's own application. Our pricing is usage-based and published rather than a six-figure annual negotiation, and time to a first working workflow is measured in minutes rather than in an onboarding programme.
Microsoft Power Automate	Microsoft's automation platform, strongest at automating work inside the Microsoft estate, with desktop automation and a very large connector list.	If the work sits inside SharePoint, Teams, Outlook and their data platform, this is honestly the right tool. They also do desktop screen automation, which we do not do at all, and their compliance position is far ahead of ours.	Their embedding options require Microsoft identities in the path; ours are components that use the host's own sign-in. Branching, loops and parallel paths are visible on one canvas rather than nested in a sequential designer. And they meter across users, bots, flows and add-ons where we meter runs.
Activepieces	An open-source automation platform that can be self-hosted freely, with a growing catalogue and an embedded builder on its enterprise tier.	Free, open-source and self-hostable is a genuine advantage for a data-sovereign buyer, there is no vendor lock-in, and their catalogue has grown quickly and is now well ahead of ours.	We are managed, so nobody has to operate or secure it. Their embed provisions their own hosted builder; ours renders as components inside your codebase at several depths. Immutable versioning, single-step retry and inherited multi-tenant governance are ours rather than left to the integrator.
Appmixer	A white-label workflow builder sold specifically to software companies that want to offer automation inside their own product, deployable on the customer's own infrastructure.	They are a real, no-frame white-label product with a self-hosting story we do not have, they hold compliance attestations we do not, and their catalogue is larger than ours. In a deal where the automation layer must run in the customer's own infrastructure, they win.	They give you a builder to mount; we give you React components to compose at several depths under your own router and data layer. Immutable versioning with version-pinned runs, single-step retry and a live streaming run log are things they do not document. And extending their catalogue means hand-writing a connector, where our assistant scaffolds one from a description.

THE POSITIONING LINE TO HOLD

Other platforms give you automation you use; FluidGrids gives you automation you can also ship inside your own product, with a run record you can defend to an auditor.

12 Why Us / Differentiators

If a partner remembers only a few sentences about FluidGrids, make it these. Each one is checkable in the product within a couple of minutes.

The builder is a component, not a tenant

Competitors that embed give you a hosted surface: a frame, a separate tenant, or a prebuilt widget. FluidGrids publishes the same React components its own application is built from, so the canvas renders inside the host's tree, under the host's router and sign-in, following its light or dark theme, and the host mints the tokens. That is the difference between automation inside your product and a link out to a separate tool.

Every save is a version, and every run remembers which one it used

Nothing is ever overwritten. Editing an active workflow creates a draft beside it, activation is a deliberate release that also arms the schedules, and a run permanently records the version it executed. That combination means a failure attributes to a specific release, rollback does not disturb anything in flight, and the version an auditor is shown is provably the version that ran.

Recovery is one step, not the whole run

Failures are classified across fourteen types and stored on the individual step with its original input. Retry re-runs that step alone and the run continues from there, so a thirty-minute pipeline that dies at step four does not repeat steps one to three or duplicate their side effects. Most tools in this category can only replay the whole thing.

A workflow can end in a dashboard

One data sink step at the end of a scheduled workflow pushes its output into a live dashboard feed in BigConsole, our companion analytics console. Leave the data field empty and the step captures the previous step's output automatically. That turns any workflow into a reporting pipeline without a separate data tool, and no competitor on our list has a native equivalent.

Governance was there before the first workflow

Multi-tenancy, role-based access, audit, metering, billing and single sign-on come from the wider platform rather than being added later under enterprise pressure. Six roles are seeded into every workspace automatically, and the same permission checks apply whether a request comes from our application, from an embedded builder or from a script.

The assistant writes connectors, not just workflows

Plenty of tools now draft a workflow from a prompt, and several do it better than we do. What none of them do is generate an entire connector package from a description, complete with the adapter that makes it callable by an agent. When a prospect's blocker is a missing integration, that changes the answer from wait for the roadmap to build it this afternoon.

One run is one run

Metering counts a completed run regardless of how many steps it contains, and plans gate scale rather than features. Teams on per-task or credit pricing routinely shorten workflows to save money and end up with worse automation. Here the incentive points the other way.

The limitations are written down

Our own competitive material says plainly where rivals are ahead: catalogue breadth, self-hosting, compliance attestations and agent depth. A prospect who is told the gaps up front and finds them exactly as described trusts everything else on the page. That is a commercial asset, not a weakness.

13 Customer Proof & Credibility

There are no customer logos to show, so credibility comes from things a prospect can verify in a free account inside an hour: the depth of the run record, the way credentials are handled, the roles that appear in a new workspace without anyone creating them, and the fact that our own product is built on the package we ask partners to embed.

AN EXECUTION RECORD THAT SURVIVES SCRUTINY

Per-step state, inputs, outputs, duration, a fourteen-type error classification and a log that streams while the run is happening. Single-step retry, pause, resume and cancel are ordinary operations, not enterprise add-ons.

GOVERNANCE SEEDING AUTOMATICALLY

Six roles, from Automation Owner to a read-only Automation Auditor, are created in every workspace without anyone configuring them, and permission checks apply to every workflow and run operation including calls from an embedded builder.

SECRETS NEVER ENTER A WORKFLOW

Connections are stored once per workspace and referenced by name; the saved workflow holds a reference and the secret is resolved at the moment the step runs. Inspect an exported definition and there are no keys in it.

UNTRUSTED CODE RUNS IN ISOLATION

External and custom connector code executes in an isolated sandbox rather than inside the shared worker process, with each run taking its own lease, and customer code steps are screened before they execute.

WE EMBED OUR OWN PRODUCT

The FluidGrids application is built on the same package a partner's customer would embed. The embed path is exercised by our own team every day rather than maintained as a side door for other people.

REAL DEPTH, HONESTLY COUNTED

Around fifty steps on the canvas today, five trigger types, versioning, templates with recorded lineage, a per-workspace connector store, agent steps with pluggable models, memory and tools, and a full API behind everything the interface can do.

A STATED COMPLIANCE POSITION

SOC 2 Type II and ISO 27001 are targeted with an audit programme scheduled, and neither is held today. Data residency defaults to India, with other regions available on Enterprise terms and a dedicated tenant available.

Be straight about this. FluidGrids is early. It is live and open for sign-up with a working free tier, but there are no public reference customers, no published usage or revenue figures, and no third-party security attestations yet. Self-serve card checkout for paid plans is still being completed, so a paid subscription is arranged with us directly. There is no self-hosted edition, the connector catalogue is far smaller than the large incumbents', and run payloads are not yet separately encrypted at rest. What a prospect should weigh instead is what they can see for themselves in a free account within an hour: whether the canvas holds their real process, whether the run record tells them what they need when a step fails, and whether the embedded builder behaves inside their own application. If those three answers are yes, the missing logos are a timing question, not a product question.

14 Demo Information

Nobody needs a slide deck for this product. The canvas, a failed run and the embed are more persuasive in ten minutes than any amount of description, so aim to get the prospect looking at the product in the first meeting.

HOW TO GET A DEMO

Anyone can create a free account at app.fluidgrids.ai with no card and build immediately. For a guided walkthrough, or for anything involving embedding, pricing or security questions, book Vignesh on the 1-1 link and let us drive. Partners are welcome to demo the basics themselves from their own free account.

Runs about Ten minutes for the core loop, thirty for a full walkthrough including embedding and governance.

WHERE TO LOOK

- **Product website** — <https://www.fluidgrids.ai>
- **Application sign-up** — <https://app.fluidgrids.ai>
- **Use case library** — <https://www.fluidgrids.ai/use-cases>
- **Comparison pages against each alternative** — <https://www.fluidgrids.ai/vs>
- **Book a guided walkthrough with Vignesh** — <https://calendly.com/vignesh-burdenoff/1-1>

The walkthrough, in order

- 1 A blank canvas, then three steps dragged on and wired together**

This is the whole idea. What you are drawing is not a diagram of the process, it is the process. When you save this, that graph is what runs.
- 2 Clicking a connector step and filling in its configuration form**

Nobody hand-built this form. The connector describes its own operations and fields, so every integration is configured the same way and a new one arrives already usable.
- 3 The connection picker inside the step, filtered to matching providers**

The designer picks a credential your admin approved and never sees the secret. The saved workflow holds a reference, not a key, so rotating it later is one edit in one place.
- 4 Executing the workflow and watching the log stream step by step**

The run is on background workers now, not in this browser. Close the tab and it finishes anyway. The log streams the way a build job does, including anything your own code step prints.
- 5 A failed run, the error class on the step, then retrying that one step**

This is the part people remember. It tells you which step, on what input, and what class of failure. Then you retry that step alone rather than replaying everything that already worked.
- 6 The version history, and restoring a previous version**

Every save is a version and every run records which one it used. Rolling back is one operation, and anything already in flight keeps the version it started on.
- 7 The same builder rendered inside a different application's own layout**

That is the same builder you have been looking at, mounted inside somebody else's product under their routes and their sign-in. This is why software companies buy it.
- 8 The workspace roles list**

Nobody created these. They appear in every workspace, and the person who can edit a draft is not automatically the person who can put it live.
- 9 The assistant drafting a workflow from a sentence**

It reads your actual catalogue before it writes anything, so it cannot invent a step you do not have. What comes back is an ordinary draft you review and bind before it goes anywhere near live.

15 Sales Talking Points

Eight lines that land. Each one is short enough to say in a corridor and true enough to survive the follow-up question.

The diagram is the implementation. There is no second artefact to keep in sync.

This is the emotional hook for anyone who maintains a process document. It reframes automation from a technical project into finally making the picture true.

When something fails, you find out which step, on what input, and retry only that step.

Everyone in operations has re-run a whole pipeline to fix one broken step. This is the line that gets a knowing laugh and opens the real conversation.

Every save is a version and every run remembers which version it used.

It moves the conversation from convenience to change control, which is where finance, compliance and anyone with an audit obligation starts paying attention.

Credentials never live inside a workflow. Rotate once and every workflow follows.

Security leads assume automation tools are credential sprawl waiting to happen. Removing that objection early keeps them from blocking the deal later.

The whole builder is a component you can put inside your own product.

For a software company this is not a feature, it is a different purchase. It turns a tooling conversation into a roadmap conversation with a much larger budget.

One run is one run, whatever the step count.

Anyone on per-task or credit pricing has felt the cliff. It is the most concrete cost argument available and it needs no spreadsheet.

Roles, audit and tenant isolation were there before the first workflow, not added under pressure.

It explains how a young product can pass an access review, and it is verifiable in a free workspace in under a minute.

If the connector you need does not exist, the assistant scaffolds one from a description.

Missing integration is the most common reason an automation evaluation stops. This turns a dead end into a next step.

We are managed only, our catalogue is smaller than the incumbents', and our certifications are pending.

Saying the three weaknesses before the prospect finds them is what makes everything else believable, and it saves both sides a wasted cycle when one of them is a hard requirement.

Start with the process that broke last month, not the whole automation strategy.

It reduces the first commitment to something a single team can do this week, which is how these deals actually begin.

The objection column is what the prospect actually says. Answer in your own words, and when the honest answer is that we do not do that, say so and move to whether it is a hard requirement.

"Zapier already connects everything we use. Why would we change?"

For a simple, low-volume automation you probably would not, and I would say so. The reason teams move is branching and loops that a step list handles badly, a bill that grows with every step in every run, and the inability to see which step failed or retry just that one. If your workflows are one line and low volume, stay where you are.

PROOF · They bill per task, so a five-step workflow running 10,000 times is over 40,000 tasks; the same workload here is 10,000 runs. Their catalogue genuinely is far larger than ours, and I am not going to pretend otherwise.

"Can we run this in our own infrastructure?"

Not today. FluidGrids is managed only, and if running it yourself is a hard requirement then an open-source alternative is the honest recommendation. What we can offer is a dedicated tenant with isolated compute and data, and a choice of data region on Enterprise terms. India is the default region.

PROOF · There is no self-hosted edition and no timeline I would ask you to plan around. Our own comparison material lists this as a gap against the open-source alternatives.

"You do not have SOC 2 or ISO 27001. Our procurement will stop there."

That is a fair stop, and if the certificate is a gate before signature we should wait rather than waste your quarter. Both are targeted with an audit programme scheduled. In the meantime what we can put in front of your security team is the actual control set: workspace isolation, roles seeded automatically, permission checks on every operation, an audit trail, credentials resolved at run time and never stored in a workflow, and isolated execution of untrusted code.

PROOF · SOC 2 Type II and ISO 27001 are both on a dated roadmap and neither is held. One known gap worth disclosing: run inputs and outputs are stored but not yet separately encrypted at rest.

"You are a small company. What happens if you disappear?"

A reasonable question and the honest position is that we are early. Two things reduce the exposure. Your workflows are ordinary data you can export, and your credentials live in the connected services, not here. And this is one product on a shared platform rather than a single-product company, so the plumbing underneath it is used by a lot more than FluidGrids.

PROOF · The workflow definition, the run history and the connection list are all reachable through the API. Nothing is locked in a format only we can read.

"Our engineers would rather just build this ourselves."

They could, and for one workflow they should. What they would be signing up for is a canvas, a durable queue, a scheduler, a retry model, a connector catalogue and the governance around all of it, then maintaining it forever while it differentiates nothing. Ask them what else that quarter could ship.

PROOF · The embed path exists precisely because product teams reach that conclusion. Our own application is built on the same package, so the components are maintained as a product rather than as somebody's side project.

"You only have about fifty steps. We need integrations you do not have."

Tell me which ones and I will get you a straight answer, because some are already built and waiting to be published. Beyond that there are two real paths: a generic web request step covers any service with an API, and the assistant scaffolds a whole connector package from a description, which then goes through validation and review before it is installed. What we will not do is put a connector on a roadmap slide to win the meeting.

PROOF · The catalogue on the canvas today is roughly thirty connectors plus twenty-one flow-control steps. A further set of packages is built but not yet published, and I can check any specific one for you.

"This will end up as another tool nobody owns."

That risk is real and it is mostly a governance question, which is where the roles help. There is an owner for the automation programme, a lead who is the only person able to activate a version, designers who can build but not release, an operator who works the failure queue, and a read-only auditor. Those exist in a new workspace automatically, so the ownership question is answered on day one rather than after the first incident.

PROOF · Six roles are seeded into every workspace with no configuration, and the permission to edit a draft is separate from the permission to activate one.

"How do we know it will not break silently at three in the morning?"

You will know it broke, and you will know why. Every step keeps its state, input, output and a typed error class. Failures raise notifications in-app, by email or by webhook, and a failing workflow can start another workflow, so escalation is itself an automation. Recovery is a single-step retry rather than a full replay.

PROOF · Fourteen error types, a run list that filters to failures, a live streaming log, and an error trigger that starts a recovery workflow when another one fails.

"Embedding sounds like a big integration project."

The install and the first render are an afternoon. The genuine engineering is mapping your customer accounts to workspaces and minting the token the builder forwards, because that mapping is the security boundary. After that you mount the module under one of your own routes and it inherits your navigation and your light or dark theme.

PROOF · Several integration depths are supported, from a single drop-in canvas to a fully routed module, with view-only mode, chrome controls and callbacks for save, run and navigation. Our own application uses the same package.

"Your pricing is not on the website in the form you just described."

Correct, and I would rather flag that than let you find it. Self-serve checkout is still being completed, so what I have given you is our own pricing material and the written quote is what governs. If price is the deciding factor, let us get you a quote in writing before you go any further.

PROOF · The published site currently shows a simplified set of plans. The figures in this document come from our internal pricing reference, which is the source we quote from.

"We already have a workflow engine our engineers use."

Then keep it. Those tools are for engineers writing durable, long-running orchestration in code, and we do not claim parity on that. The question is who else in the business needs to change a process: if the answer is only engineers, you are fine. If it is operations, finance or support, they cannot read code, and that is the gap we fill.

PROOF · Our own comparison material says plainly that code-first orchestration engines are stronger on durability and are often a complement rather than something to replace.

"Our team is not technical enough to build workflows."

That is the design assumption. Steps render their own configuration forms, so configuring an integration is filling in fields rather than reading API documentation, and the assistant will draft a first workflow from a plain description grounded in the catalogue you actually have. Where something genuinely needs code, a code step is there for the one person who can write it.

PROOF · The template gallery, the self-describing forms and the assistant all exist to remove the blank canvas. A free account with no card is the cheapest way to test this claim.

17 Discovery / Qualification Questions

Ten questions that separate a real opportunity from a polite conversation. You are listening for a named process, a named owner, and evidence that the current situation already hurts.

ASK THIS

Walk me through a process your team runs every week that touches more than two systems. Who does what, in what order?

What automation do you have today, and what does it cost you a month?

When one of those automations fails, how do you find out and what do you do next?

If somebody changed that automation last month, could you tell me what changed and get back to the old behaviour?

Where do the credentials for these integrations live right now?

Does anything in this process need a person to approve it, and where does that approval happen today?

LISTEN FOR

Named systems, named people and manual handoffs. If they cannot describe it without saying then someone, you have found the workflow.

An existing tool and a rising bill is the strongest signal there is. A number they wince at means the cost argument will land.

Someone notices something did not happen, and the fix is to run it all again. That is the single-step retry pitch handed to you.

Hesitation, or a spreadsheet-based change log. Anyone with an audit obligation gets uncomfortable at this question, which is exactly the opening.

In the script, in a shared document, or nobody is quite sure. This gets a security lead engaged and gives you a second stakeholder.

Approvals in inboxes and chat. Human decisions inside an automated run is a capability most alternatives handle badly.

ASK THIS

LISTEN FOR

Are your customers asking to connect your product to the rest of their stack?

Only relevant if they build software, but if the answer is yes this is a far larger opportunity than internal automation and a much faster decision.

Who would build these workflows, and are they comfortable in a tool or do they need an engineer?

If only engineering can build it, automation is queued behind product work forever. That queue is the pain you are selling against.

Is there anything about where your data lives or which certifications a supplier holds that would rule us out?

Ask this early and take the answer at face value. A hard self-hosting or certification requirement means walk away now, kindly.

How does the number on your leadership dashboard get there?

A person exports a file. That is a scheduled pipeline and a live dashboard feed, and it is often the quickest first win.

If one process were automated properly by the end of the quarter, which one would you pick?

A specific answer means there is a champion and a scope. A vague answer means there is no urgency yet, so keep it warm rather than pushing.

Who else would need to be in the room before anything got signed?

Names and roles. If security, finance or procurement appear here, get them the honest limitations early rather than late.

18 How to Identify a Good Opportunity

Use this to decide where to spend your time. Green means make the introduction this week; amber means keep it warm and gather one more fact; red means say so honestly and keep the relationship.

STRONG SIGNAL – PURSUE

- They can name a specific process that crosses four or five systems and name the person who currently holds it together.
 - They already pay for an automation tool and the bill is rising faster than the work is.
 - They build software and their customers are asking for automation or integrations inside the product.
 - Someone recently could not explain why an automated process failed, or had to re-run a whole pipeline to fix one step.
 - They are an agency or consultancy delivering the same automation to client after client.
 - A compliance or audit event is forcing them to prove who changed a process and which version ran.
 - A leadership dashboard depends on a person exporting a spreadsheet every week.
- Introduce us straight away.

WORTH A CONVERSATION

- Interest is real but no single process has an owner. Get them to name one before you take up our time.
- They want it, but everything must go through an engineering team that has no capacity this quarter.
- Automation sits inside a broader digital transformation programme with no date attached.
- They are on the free tier of a consumer automation tool and have not yet hit its limits.
- The champion is enthusiastic but cannot describe how a decision gets made or who else must approve it.
- They ask about certifications early but cannot say whether a certificate is a hard gate or a preference.

Qualify further before an introduction.

POOR FIT – DO NOT CHASE

- The platform must run inside their own infrastructure or an air-gapped network. Recommend an open-source alternative and keep the relationship.
- A current SOC 2 Type II or ISO 27001 certificate is required before signature, with no exception process.
- The requirement is desktop or screen-driven robotic automation of a legacy application.
- The deal depends on connectors across a very long tail of niche applications we do not cover.
- They want long-running orchestration that resumes from the exact instruction after a crash, which is a different category of engine.
- There is no budget and no process owner, only curiosity about the technology.

Politely park it; suggest another product if one fits.

THE ONE-QUESTION TEST

Ask: which process would you fix first, and what does it cost you when it goes wrong? If they answer both halves with something specific, it is a real opportunity. If they can only answer the first, or neither, it is a conversation rather than a deal.

19 How to Introduce Us

Keep the introduction short and about their problem, not about our features. Your job is to get the right two people talking; ours is to run the evaluation. Name the process they described back to them, and let the product do the persuading.

IN CONVERSATION

"Say something like: you know that thing you described where the order has to go through five systems and someone chases it by email? There is a product called FluidGrids where you draw that on a canvas and the drawing is what actually runs, and when a step fails it tells you which one and lets you retry just that step. I know Vignesh who runs it. Want me to put you two together?"

EMAIL – SUBJECT: THE PROCESS YOU DESCRIBED, AND A TOOL THAT MIGHT FIT

Hi [Name],

You mentioned that [process] runs across [systems] and that most of the chasing still happens by hand. I came across something that fits that shape well and thought of you.

FluidGrids is visual workflow automation. You draw the process on a canvas and that drawing is what runs in the background. When a step fails it tells you which step, on what input, and lets you retry just that one rather than the whole thing. Every change is versioned, so rolling back is one action.

It is early, so no big customer list to point at, but there is a free tier and you can test it against a real process in an hour: <https://www.fluidgrids.ai>

If it looks worth a conversation, Vignesh who runs it is easy to talk to: <https://calendly.com/vignesh-burdenoff/1-1>

Happy to introduce you directly if you would rather.

[Your name]

SHORT MESSAGE / LINKEDIN

Short version for a message: you mentioned [process] still runs on exports and chasing. FluidGrids lets you draw that process on a canvas and the drawing is what runs, with per-step logs and a retry that only re-runs the step that broke. Free tier at www.fluidgrids.ai if you want to poke at it. Happy to introduce you to Vignesh who runs it.

WARM THREE-WAY INTRODUCTION

The strongest form is a three-line email with both parties on it: one line on what the prospect told you, one line on what FluidGrids does, and Vignesh's details. Copy vignesh@burdenoff.com and let us take it from there. Tell us privately what you already know about their timing, budget and who else must approve, so the first call does not repeat your discovery.

Please do not: Do not claim customers, usage figures or named references. We have none yet and one check on a call ends the credibility of everything else you said.; Do not promise a connector that is not already on the canvas. Ask us, get a straight answer, and go back with it.; Do not quote a price as final. Every figure in this document is indicative and the written quote governs.; Do not imply we hold SOC 2 or ISO certification, or that a self-hosted edition is available or coming soon.; Do not push a prospect whose requirement is a genuine red flag. Walking away honestly is worth more to us than a wasted evaluation.

20 Lead Handoff Process

Once the prospect is interested, hand over early rather than trying to run the evaluation yourself. The technical questions get deep quickly, and the sooner we are in, the sooner you get a decision.

1 Register the opportunity before the introduction

Email vignesh@burdenoff.com with the company name, the contact and one line on the process they want to fix. Registering first is what protects your commission, and it takes a minute.

2 Make the introduction

A three-line email with both parties on it, or book them straight onto <https://calendly.com/vignesh-burdenoff/1-1>. Stay on the thread so you can see what happens and step in if it stalls.

3 Brief us privately

Send anything you know that will not be said on the first call: budget signals, the real decision maker, a competing tool already in place, a compliance gate, or a date they are working towards.

4 Join the first call if you can

Optional but useful. You hold the relationship and the history; we hold the product detail. If you cannot attend, we will send you a short note afterwards.

5 Let us run the evaluation

We handle the guided walkthrough, the security questionnaire, any embedding conversation, and the quote. You stay informed and are copied on the commercial exchanges.

6 Stay in for the close and the handover

For consulting and implementation partners there is usually delivery work after signature. Tell us early if you want it, so it is scoped into the proposal rather than added afterwards.

SEND US THIS – IT IS ALL WE NEED

- Company name, website and roughly how many people work there
- The contact's name, role and email, and whether they can sign or need an approver
- The specific process they want to automate, in their own words
- Which systems it touches, and roughly how often it runs
- What they use today, if anything, and what it costs them
- Any hard constraint: a certification gate, a data-residency requirement, a self-hosting requirement, or a deadline
- Whether this is internal automation or embedding automation in their own product
- Your reading of urgency and who else must be involved

WHAT HAPPENS NEXT, AND WHEN

We acknowledge a registered opportunity within one working day and aim to be in front of the prospect within three. If a first call needs to happen sooner than that, call or message Vignesh on +91-7358445777 and we will make it work.

Deal registration: Registration is by email to vignesh@burdenoff.com before the introduction, and it is what establishes your claim to the deal. A registration holds for ninety days on a named company and contact and is renewed by any real activity on the account. If two partners register the same company, the earlier registration stands unless the later one is clearly the active relationship. Registering costs nothing and takes a minute; not registering is how a commission gets disputed.

21 Partner / Referral Commercials

FluidGrids has no product-specific partner scheme, so the standard Burdenoff partner programme applies. Commission is paid on collected net subscription revenue, and the tier you sit in depends on how much of the work you do: an introduction, an introduction plus evaluation support, an introduction plus delivery, or full commercial ownership as a reseller. Register the opportunity before you introduce it, and the rest is straightforward.

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%, capped at \$10,000
Consulting Partner	Introduction plus evaluation support	12% of first-year collected net subscription	0-3% if active in the success plan · 3%, 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%, capped at \$25,000
Authorised Reseller	Approved resale, 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 tenant or managed services, if approved
MSP or Strategic Partner	Bundled or managed offering	Custom	Custom · Custom

THE RULES THAT MATTER

- Register the opportunity in writing before the introduction. An unregistered deal has no protected claim, and this is the most common reason a commission is disputed.
- Commission is calculated on collected net subscription revenue, after discounts, taxes and any refunds, not on the list price or the invoice value.
- Add-on commission is paid at the rate and cap for your tier. Professional services and dedicated infrastructure are quoted separately and are not commissionable unless the agreement says so.
- Embedding the builder in another company's product is always quoted by us and commission on it is agreed deal by deal.
- Discounts beyond the standard programme need our approval in writing before you offer them to a customer.
- A registration holds for ninety days on a named company and contact, renewed by genuine activity on the account.
- Reseller terms require an approved agreement in place before any quote is issued in your name, and a resold or bundled subscription is sold under the FluidGrids name.

GETTING PAID

Commission is calculated once the customer's payment has cleared and is paid in the following monthly cycle against your invoice, in the currency named in your partner agreement. Renewal commission, where your tier earns it, is paid on the same basis at each renewal for as long as you remain active on the account.

These are the standard Burdenoff partner programme terms, published here so you can plan. They are indicative and not an offer. The commercial terms that actually apply are those in your signed partner agreement and in the quote issued to the customer, and those documents govern in any conflict with this page.

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 FluidGrids 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
FluidGrids website	https://www.fluidgrids.ai	Positioning, features, pricing and the comparison pages against each alternative. Safe to share with anyone.
Sign up and build	https://app.fluidgrids.ai	Free tier, no card. The fastest way for a prospect to test whether the canvas holds their real process.
Documentation	https://docs.fluidgrids.ai	Product and integration documentation. Still being expanded, so set expectations before sending an engineer straight to it.
Use case library	https://www.fluidgrids.ai/use-cases	Around forty worked scenarios by role and category. The quickest way to find one that mirrors what a prospect just described.
Pricing	https://www.fluidgrids.ai/pricing	The published plans. Note that the site shows a simplified set; use the plan table in this document and always confirm with a written quote.
Comparisons against alternatives	https://www.fluidgrids.ai/vs	Honest side-by-side pages against each named competitor, including where they are ahead. Useful to send after a competitive call.
AI capabilities	https://www.fluidgrids.ai/ai	What the assistant and the agent steps actually do, with each capability marked shipped, in progress or planned.
Security	https://www.fluidgrids.ai/security	The security posture in customer-facing language, for the moment a security lead joins the conversation.
Book a call with Vignesh	https://calendly.com/vignesh-burdenoff/1-1	Guided walkthrough, embedding conversations, pricing and security questions. Vignesh T.V., vignesh@burdenoff.com, +91-7358445777.
All Burdenoff products	https://burdenoff.com/products	The wider portfolio. Worth a look if your contact's problem turns out to sit next door to workflow automation.

ASSET	WHERE	WHAT IT IS FOR
General enquiries	mailto:contact@fluidgrids.com	Product and commercial enquiries. Security disclosures go to security@fluidgrids.com .

WHAT YOU CAN SHARE, AND WITH WHOM


- Share freely in public: anything on www.fluidgrids.ai, including the feature, pricing, comparison, AI and security pages, plus the use case library and the free sign-up link. It is written for a public audience and is already checked.
- Share with a named prospect, not publicly: this playbook, the cheat sheet, the plan and quota table, and the product screens in it. Send it to a person, do not post it, and do not publish it on your own website or a social feed.
- Get our sign-off first: any quoted price in writing, anything describing embedding terms, any security or compliance questionnaire response, and any use of the FluidGrids name or logo in your own marketing, on your website, or in a public announcement.
- Never share at all: our competitive battle cards as a document, anything implying we have customers, usage figures or certifications we do not hold, and any internal roadmap detail beyond the shipped, in progress and planned labels in this pack.
- When in doubt, ask. A one-line message to vignesh@burdenoff.com is faster than an correction after the fact, and we would far rather answer the question than unpick a claim.

24


59 SCREENS

Every screen in the FluidGrids design set — 59 in total, none omitted. Use these to show a prospect what the product looks like before a live demo.


PRODUCT SCREENS • 54 SCREENS




001 Fg 01 Builder Canvas




002 Fg 02 Node Config Ndv




003 Fg 03 Node Catalog




004 Fg 04 Builder Empty New




005 Fg 05 Data Mapper




006 Fg 06 Code Node




007 Fq 07 Builder Test Run




008 Fq 09 Runs List




009 Fg 10 Run Detail Success




010 Fg 11 Run Detail Failed




011 Fg 12 Trigger Schedule Builder




012 Fg 13 Webhook Deliveries




013 Fg 14 Connections List




014 Fg 15 Chat Trigger



015 Fg 16 Run Control Pause Resume




016 Fg 17 Flow Control Merge




017 Fg 18 Add Connection



018 Fg 19 Templates Gallery




019 Fg 20 Template Import




020 Fg 21 Datasink Config




021 Fg 22 Datasink Dashboard




022 Fg 23 Version History




023 Fg 24 Publish Activate




024 Fg 25 Workflows List




025 Fg 26 AI Agent Subnodes




026 Fg 27 AI Studio NI




027 Fg 28 AI Assistant Debug




028 Fg 29 Embed Host App




029 Fg 30 Embed Readonly Viewer



030 Fg 31 Node Store



031 Fg 32 Custom Node Lifecycle



032 Fg 33 RBAC Roles Matrix




033 Fg 35 Fleet Admin




034 Fg 36 Analytics Dashboard




035 Fg 37 Error Analytics




036 Fg 38 Billing Usage




037 Fg 39 Pricing Plans




038 Fg 40 Notifications Center



039 Fg 41 Comments Collab




040 Fg 42 Audit Log




041 Fg 43 Dev Portal API




042 Fg 44 Overview Dashboard




043 Fg 45 Global Search



044 Fg 46 Workspace Settings



045 Fg 47 Rule Builder




046 Fg 48 Scheduler Dashboard




047 Fg 49 Mobile Shells




048 Fg 50 AI Assistant Authored Workflow




049 Fg 51 AI Node Builder Mode




050 Fg 52 Agent Tool Run Trace




051 Fg 53 Vector Retrieval Node Config



052 Fg 54 AI Run Repair Proposal




053 Fg 55 Autonomous Triage Digest




054 Fg 56 MCP Tool Surface

BANKING FINANCE • 1 SCREENS




055 KYC Periodic Review Orchestration 1 Periodic Review Workflow

HEALTHCARE • 1 SCREENS




056 Hospital Bed Turnover And Evs Dispatch 1 Bed Turnover Run

SECURITY DEFENSE • 2 SCREENS



057 Soc Containment Playbook Approvals 1 Paused Containment Run



058 Soc Containment Playbook Approvals 2 Approval On A Phone

SOFTWARE SAAS • 1 SCREENS



059 SaaS Tenant Automation Fleet View 1 Tenant Automation Fleet



Found someone who needs FluidGrids?

You do not need to be able to build a workflow to sell this well. You need to recognise the moment someone describes a process held together by exports, reminders and goodwill, and make the introduction while that frustration is still fresh. Send it to us and we will take it from there.

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. • FluidGrids is available at www.fluidgrids.ai.

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.