

— FOR HEALTHCARE BUILDERS

You came to build healthcare's future. You're stuck cleaning its ~~past~~.

Lab PDFs. Wearable streams. Device readings. HL7v2. FHIR. CSV. 20-year-old feeds. Faxed forms. Scanned notes. Your engineers were hired to build the product. **JondaX cleans, codes, and structures all of it**, so they can.

SPEAKER

Suhina Singh

Founder & CEO · Jonda Health



PDF - LAB REPORT

Hb	13.2 g/dL
WBC	7.4 ×10 ⁹ /L
Glucose	11.8 mmol/L ↑
HbA1c	7.9 %
Cholesterol	5.6 mmol/L
LDL	3.4 mmol/L
Triglycerides	2.1 mmol/L

HL7V2 - MESSAGE

```
MSH|^~\&|EPTC|HOSP|LAB|20260415
PID|1|00482|SMITH^JOHN^A
OBR|1|CBC|F|20260415
```

WEARABLE - HEALTHKIT

HEART RATE **72** bpm

STEPS TODAY **8,421**

SLEEP **7h 12m**

SPO₂ **98%**

CSV - BULK EXPORT

ID	DATE	TEST	VAL
001	04-15	HGB	13.2
001	04-15	WBC	7.4
002	04-15	GLU	11.8
002	04-15	HbA1c	7.9
003	04-16	CHOL	5.6

FHIR - JSON

```
{
  "resourceType": "Observation",
  "status": "final",
  "code": {
    "text": "Glucose"
  },
  "valueQty": 11.8
}
```

EMAIL - ATTACHMENT

dr.tan@clinic.sg
Patient referral - attached labs

DEVICE - CONNECTED STREAM

08:30	BP cuff	142/88 mmHg ↑
12:15	Glucometer	8.2 mmol/L
14:00	Pulse oximeter	97 %
16:30	Smart scale	78.4 kg
22:00	Thermometer	36.8 °C

— THE SILENT TAX ON HEALTHCARE

And yes it's costing you more than you think.

Every roadmap delay, every late launch, every “we’ll integrate that next quarter” traces back to the same place. Fragmented data your team has to fix before they can build.

ENGINEERING TIME



60%

of data engineering time spent on cleanup.

Not building features. Not shipping product. Cleaning, mapping, normalising.

INTEGRATION SPEED



6–9 months

to onboard a single new data source.

Custom parsers. Field-by-field mapping. Edge cases that surface only in production.

OUTPUT QUALITY



1 in 5

manually mapped records contains an error.

Wrong codes. Wrong units. Wrong patient. Each one a downstream risk to clinical decisions.

● Industry estimates from peer-reviewed research and customer interviews. Your numbers will vary.

— THE DOWNSTREAM COST

When data doesn't flow, care doesn't either.

Every hour your engineers spend cleaning data, a clinician somewhere is making a decision without it. Behind every broken integration is a patient who waits.

PATIENT OUTCOMES



1 in 3

Diagnostic decisions made with incomplete records.

Lab results in a PDF the EHR can't read. Wearable readings in a portal nobody opens. Clinicians act on what they can see, not on what's been collected.

• WHO PAYS · The patient

CLINICIAN BURDEN



2+ hrs/day

Lost to data hunting and re-entry.

Looking for the latest result. Re-keying numbers from a faxed form. Reconciling units between systems. **Time that should belong to the patient.**

• WHO PAYS · The clinician

POPULATION HEALTH



80%

Of clinical data never enters analysis.

Locked in PDFs, scans, and unstructured notes. **The signal that could prevent the next outbreak, spot the next chronic disease cohort, fund the next health policy. Sits unread.**

• WHO PAYS · The system

THE THROUGHLINE | This isn't a healthcare IT problem. **It's a healthcare problem.** When your data flows, every part of the system upstream and downstream gets better.

— WHY THIS HASN'T BEEN SOLVED

The real barrier isn't AI. It's the data.

CONCEPT · WHITE BLOOD CELL COUNT

Take one clinical concept. Now multiply the variants at every layer between the source and the answer.

LAYER 01

Format

How the data arrives

- HL7 v2
- FHIR
- JSON
- PDF
- XLSX
- CSV
- Image
- API
- and more*

7+

INPUTS

LAYER 02

Language

How the source labels it

- English
- Spanish
- German
- French
- Indonesian
- Dutch
- and more*

33+

LANGUAGES

LAYER 03

Name

What the source calls it

- White Blood Cell Count
- WBC
- WCC
- White Cells
- and more*

28+

NAME VARIANTS

LAYER 04

Unit

How the value is measured

- /cumm
- /mm3
- 10x12/L
- K/CMM
- thou/uL
- THOUSAND/MCL
- and more*

58+

UNIT STRINGS

LAYER 05

Coding

Which medical coding standard

- International standards
- Regional terminologies
- and more*

5+

STANDARDS

— WHY THIS HASN'T BEEN SOLVED

The complexity isn't additive. It's **compounded**.

One clinical concept. Four layers of variation. Multiply, don't add.

$$\begin{array}{ccccccc} 33 & \times & 15 & \times & 58 & \times & 5 \\ \text{language} & & \text{name} & & \text{unit} & & \text{coding} \end{array}$$

$$= 143,550+$$

semantic variants of one clinical concept

× 7+ data formats each variant can arrive in

AND THEN

You have **thousands** of clinical concepts. That's **billions** of unique data variations.

— AND DOCUMENT AI ISN'T THE ANSWER

Looks structured. Loses what matters.

Today's document AI tools claim to harmonise health data. But when you look at the output, the clinical context that drives care, complete units, reliable reference ranges, consistent terminology and standard codes, doesn't consistently carry through. **Looks organised. Still incomplete.**

ORIGINAL LAB REPORT Source

LabCorp		Patient Report		
PID: ... · DOB: ... · Date: 01/16/2018				
TESTS	RESULT	FLAG	UNITS	RANGE
eGFR NonAfrican	128		mL/min/1.73	>59
eGFR African	149		mL/min/1.73	>59
BUN/Creatinine	18			9 - 20
Sodium, Serum	141		mmol/L	134-144
Potassium, Serum	4.4		mmol/L	3.5-5.2
A/G Ratio	1.1	Low		1.2-2.2
Alk Phosphatase	242	High	IU/L	39-117
ALT (SGPT)	62	High	IU/L	0-44

DOCUMENT AI OUTPUT Incomplete

eGFR NonAfrican Truncated unit

128 · mL/min/1. · range 73 >59

eGFR African Range lost

149 · mL/min/1. · range partially obscured

A/G Ratio Code missing

1.1 · range 1.2-2.2 · no terminology code

Alk Phosphatase Terminology unmapped

242 · IU/L · 39-117 · name not normalised

JONDAX OUTPUT Clinically complete

eGFR NonAfrican Unit complete

128 · mL/min/1.73 m² · range >59

eGFR African Range restored

149 · mL/min/1.73 m² · range >59 reconstructed

A/G Ratio Code assigned

1.1 · range 1.2-2.2 · mapped to standard code

Alk Phosphatase Terminology mapped





242 · IU/L · 39-117 · canonical name resolved

THE GAP | Even document AI vendors acknowledge that key clinical context, complete units, reliable reference ranges, terminology and codes, may not consistently carry through without additional harmonisation. **JondaX is that harmonisation, built in.**

— WHY CLIENTS CHOOSE US

When data drives care, QA isn't just nice to have.

Every transformation is validated, scored, reviewed, and traced. **Clinical-grade output requires clinical-grade quality assurance.**

<p>STEP 01</p>  <h3>Validate</h3> <p>Against the reference graph</p> <p>Every output is checked against our knowledge graph of synonyms, codes, units, and reference ranges.</p> <hr/> <p>3M+ MAPPED REFERENCES</p>	<p>STEP 02</p>  <h3>Score</h3> <p>Confidence on every field</p> <p>Each transformation gets a confidence score. Low-confidence outputs are flagged before they leave our system.</p> <hr/> <p>Per field AND PER FILE</p>	<p>STEP 03</p>  <h3>Review</h3> <p>Human-in-the-loop</p> <p>Edge cases route to medical data quality controllers. Your data quality bar stays high.</p> <hr/> <p>Reviewed BY HUMANS</p>	<p>STEP 04</p>  <h3>Audit</h3> <p>Every step traceable</p> <p>Every transformation is versioned, traceable, and inspectable. Full provenance, on demand.</p> <hr/> <p>End-to-end TRACEABILITY</p>
---	--	---	---

99% OUTPUT ACCURACY

Backed by humans. Auditable end-to-end. Built so your teams can trust the output.

— INTRODUCING



Finally. One engine for every health data problem.

One API. Any source. Any format. JondaX ingests fragmented clinical data, cleans it, maps it to medical coding standards, and returns it as **structured, regulation-ready output** your engineers can build on.



INGEST

Any source. Any format.

PDFs, scans, HL7v2, FHIR, CSV, XLSX, wearable streams, connected device data, faxed forms. If your data team has touched it, JondaX can read it.



TRANSFORM

Clean, code, structure.

Custom AI trained on clinical data. Maps to medical coding standards. Validates units, ranges, and context. Surfaces what a human reviewer would catch.



OUTPUT

Built for your stack.

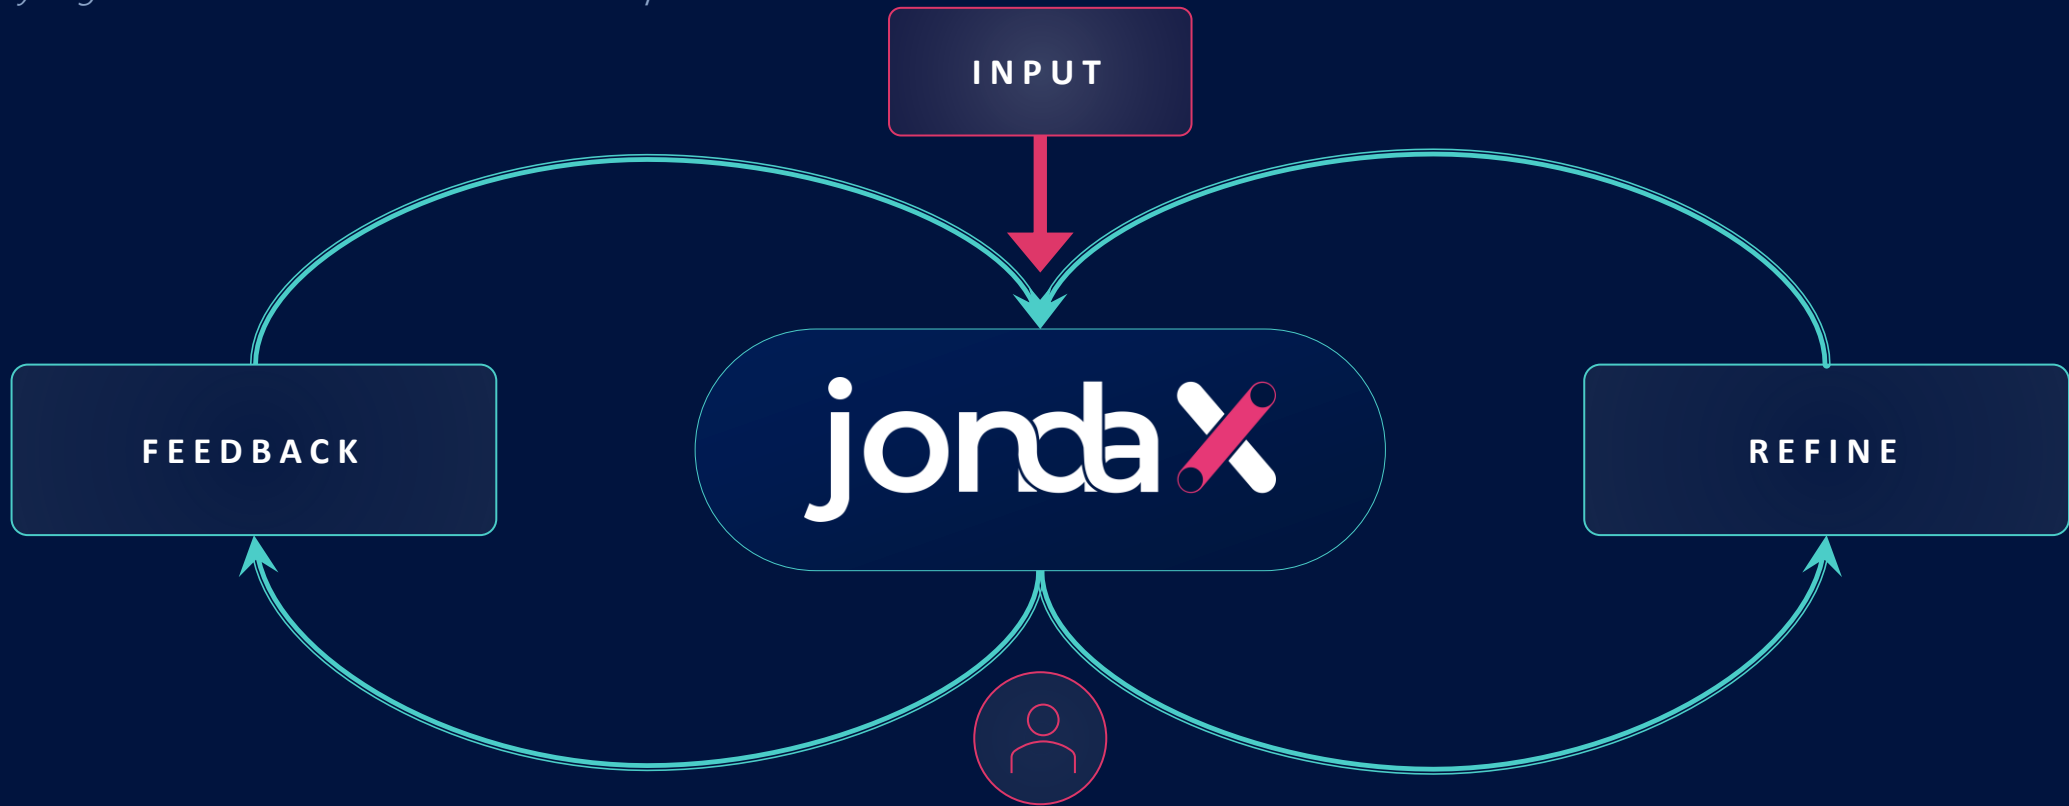
FHIR, JSON, structured exports. Connects to EHRs, LIMS, CMS platforms, research databases, and digital health apps. Plug it in and ship.

— THE SOLUTION

A Living Clinical Intelligence Layer

Continuously maintained and updated clinical knowledge and AI trained on real-world healthcare data across multiple languages and regions

Powered by insights from over 3 million real-world clinical data points



Human in the loop
Medical Data Quality Controller

— MODULAR BY DESIGN

Four modules. One engine.

Use what you need today. Add what you need tomorrow. **Pathology alone, combined with medical device data, or all-in.** Same API, same contracts, same clean output.

Pathology

More than blood. JondaX handles quantitative and qualitative results across blood, urine, stool, saliva, CSF, sputum, microbiology and more. Core for structured data, pro for PDFs and image files.

INPUTS

HL7 v2 FHIR JSON PDF Image

Medical Devices

Glucometers, BP monitors, pulse oximeters, thermometers, smart scales, and more. Photo of the device readout in. Structured data out.

INPUTS

Image

Wearables

300+ devices through our partner network. Heart rate, sleep, activity, HRV. One schema, no per-device maintenance.

INPUTS

API

Point-of-Care

Rapid diagnostic tests, lateral flow assays, pregnancy and antigen tests. Photo of the strip, structured result returned.

INPUTS

Image

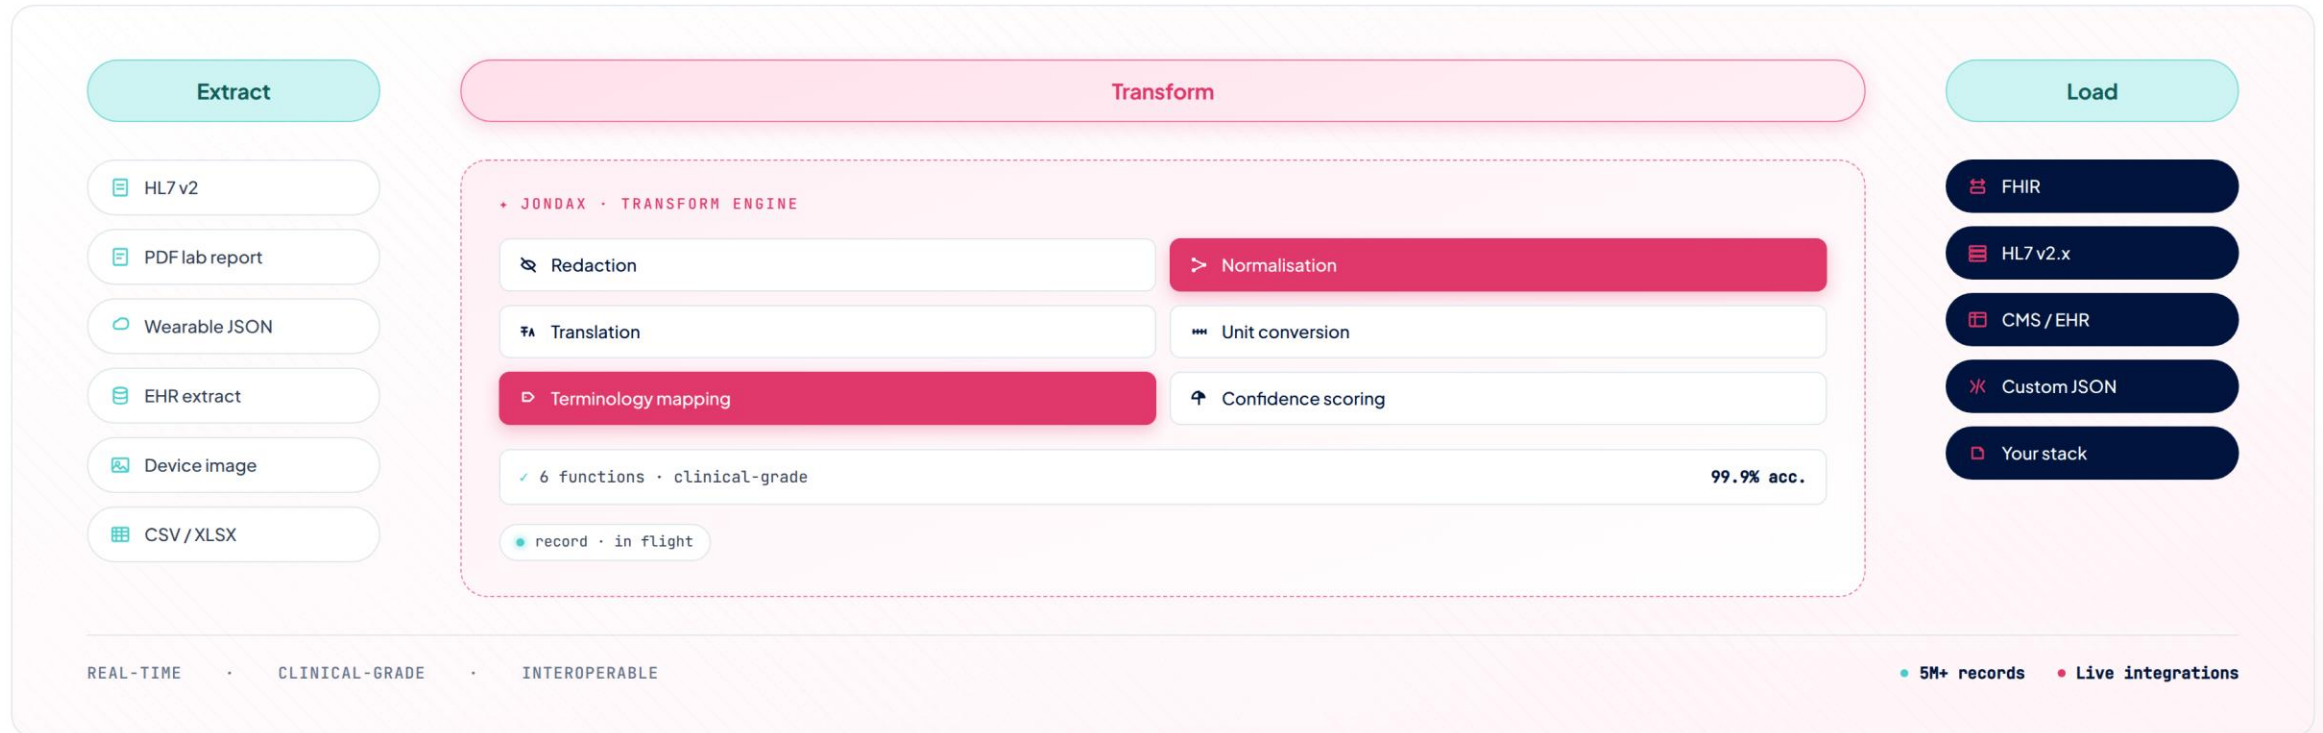
Something we don't cover? **Tell us what you're building** and we'll see if there's a path to a new module.

— HOW IT WORKS

Extract. Transform. Load.

Whatever your input. Whatever your output. JondaX is the engine in between.

06 · CAPABILITY MAP · ETL FLOW



— WHAT JONDAX DOES

Six functions. One pipeline.

Every file that enters JondaX runs through this sequence. **Each step is auditable, traceable, and built for clinical-grade output.**



redact()

Removes **identifiable patient information** while preserving every clinical signal.




normaliseFile()

Turns **any input format** into a single canonical structure JondaX can reason about.



translate()

Reads **33+ languages** and clinical shorthand. Every source becomes one common language.



convertUnits()

Detects and converts between **any unit system**. Knows when to flag, when to convert, when to reject.



mapTerminology()

Aligns synonyms, abbreviations, brand names, and regional terms to **a single canonical vocabulary**.



assignCodes()

Assigns the **right medical code** for every observation. Validates against context, not just dictionaries.

— REAL NUMBERS, IN PRODUCTION

Not a pilot. A platform.

JondaX is processing live clinical data right now. **500K+ files**. **5M+ data points**. **3M+ reference graph**. Trained on real-world clinical messiness, not lab samples.

FILES PROCESSED

500K+

Documents transformed

Across PDFs, HL7, FHIR, CSV, scans, device streams.

DATA POINTS

5M+

Clinical observations structured

Each one cleaned, coded, and validated by JondaX.

REFERENCE GRAPH

3M+

Mappings in our knowledge graph

Synonyms, codes, units, ranges across regions.

READINESS

TRL₈₋₉

Active commercial deployments

Defensible in real-world operating environments.

★ InterSystems Startup Impact Award winner

★ Eagle Ventures competition winner

↗ Active partners: Anzer · BeHealthcare

— ALREADY IN PRODUCTION

Building **together** with healthcare leaders.

JondaX is in active commercial deployment with EHR providers, healthcare platforms, and IT consultancies. **One integration with us. Many problems solved at once.**

EHR PARTNER

Anzer

Live integration. JondaX powers their data ingestion across multi-format clinical inputs.

EHR PARTNER

BeHealthcare

Live integration. Clean, structured data flowing into their EHR for downstream workflows.

PROGRAMME PARTNER

Google for Startups

Backed and supported through the Google for Startups programme. JondaX runs on GCP infrastructure.



AWARD

InterSystems Startup Impact Award



AWARD

Eagle Ventures Competition Winner



DISTRIBUTION

EHR & Health IT consultancies

— JOBS TO BE DONE

Four jobs. One engine.

Different teams hire JondaX for different jobs. **One integration. Every job gets done.**

DIAGNOSTICS LAB

THE JOB

Deliver to any client, in any format they need.

“Every clinic we onboard wants a different output format. Custom integrations don’t scale, and PDFs are no longer enough.”

— WITH JONDAX

Send us your results in any format you have. JondaX **cleans, codes, and outputs to whatever each client needs:** FHIR, HL7 v2.x, CSV, JSON, EHR-specific.

→ One ingestion, every client format covered.

RESEARCH & CLINICAL TRIALS

THE JOB

Get from raw multi-site data to analysis-ready in weeks.

“Trial data comes from 30+ source systems. Harmonising it eats months before our team can even start the analysis.”

— WITH JONDAX

Pipe in data from any site, any platform, any format. JondaX returns a **single, harmonised, research-ready dataset.** EDC platforms and study teams get clean input from day one.

→ Weeks to analysis, not months.

HOSPITAL & CLINIC GROUPS

THE JOB

See the whole patient, in-hospital and remote.

“Lab results in one system, vitals in another, wearables nowhere. We treat what we can see, not what’s actually happening with the patient.”

— WITH JONDAX

Labs, devices, wearables, prior records: ingested, coded, and unified. **Every relevant data point in one clinical view, whether the patient is at the bedside or being monitored at home.** Delivered through your EHR or CMS.

→ One patient, one record, every signal.

DIGITAL HEALTH STARTUP

THE JOB

Let engineers build product, not parsers.

“We have 12 different lab and EHR feeds to support. Our roadmap is parsing, not product. Every new partner kills another sprint.”

— WITH JONDAX

Point JondaX at every source you support. **One contract, one schema, one integration.** Engineers stop maintaining parsers and get back to features.

→ 1 API replaces 12 integrations.

— BUILT FOR PRODUCTION

Engineered for the real world. Day one.

We've done the regulatory and infrastructure work so your team doesn't have to. Plug JondaX in. Stay focused on building product.

✓ HIPAA-aligned

✓ GDPR-ready

✓ ISO-grade controls

✓ BAA & DPA available

✓ Multi-jurisdictional



Privacy by design

PII redaction baked into the pipeline. Data held 30 days for issue resolution, then permanently deleted.



GCP infrastructure

Core platform on Google Cloud. Audited, scalable, with regional deployment options across markets.



Audit-ready output

Every transformation is traceable, versioned, and inspectable. Your compliance team will love you.



Human-in-the-loop

Edge cases route to licensed reviewers for verification. Your data quality bar stays high.

FLEXIBLE ARCHITECTURE


REST API. FHIR-native output. SDKs available. Works with any modern data stack, from greenfield startups to legacy hospital systems.

CASE STUDY

For a single GP clinic

JondaX turns wasted admin hours into patient care and growth.

 **Future-ready:** Can support Satusehat via FHIR + LOINC mapping




For a single primary care clinic
estimated based on research

Annual time saved	Annual revenue unlocked
~33 days	~USD 35K
	<i>from redirecting time from manual data entry and data retrieval to additional patient care</i>

— CASE STUDY

For a longevity platform

JondaX turns wasted development time to address data harmonization into time focused on creating value for their clients with their algorithms and at the same time enabling scalability



Spanish · Indonesian · German · French · Dutch · and more

Data transformed automatically based on requirements

By removing technical and logistical barriers to lab ingestion and normalisation, we've empowered our clients to unlock deeper insights and provide better outcomes for their patients.

—A European longevity platform, JondaX client

— IF THIS RESONATED

Three ways to start.

Whether you're sketching architecture or shipping next quarter, there's a path in. Pick the one that fits where your team is today.

01

15 MINUTES

A technical conversation

Send us your messiest data sample. We'll show you, in your own data, what JondaX would output. No commitment, just clarity.

[Book a session →](#)

02

2 WEEKS

A scoped pilot

Pick one data source giving your team the most pain. We integrate, transform, and deliver clean output for it. You measure the difference.

[Start a pilot →](#)

03

WHEN YOU'RE READY

Self-serve access

Try the API on your own. Sign up, drop in your first file, see structured output back in minutes. No call required.

[Try jonda.health →](#)

— THANK YOU

You came to build **the future.** Let's give your team **the data** to do it.

From data chaos to data clarity. We'd love to build that with you.

SPEAKER

Suhina Singh

Founder & CEO · Jonda Health

GET IN TOUCH

hello@jonda.health

jonda.health



jondahealth