DiveScale
Your Technology Partner

embedded world North America 2026
Anaheim, September 22 to 24

Your hardware is ready. We build the software around it.

Cloud platforms, fleet consoles, companion apps and AI features for connected products, built by an AI-native engineering team that has shipped them to production.



15+
production platforms designed, built and launched

23%
lower electricity costs for a verified energy client

75%
faster incident reporting on an industrial safety AI platform

Lunul console and homeowner app, built by DiveScale

An AI-native engineering team that ships the software layer of connected products

DiveScale is a San Diego software company with an engineering team in Lahore, Pakistan. We design, build and run production software: the cloud, data, apps and AI that sit on top of your hardware.

15+ production platforms

Shipped and running for clients in the US and Europe, from IoT fleets to HIPAA-grade healthcare.

AI-native by default

Our engineers build with AI every day, and we ship AI features that hold up in production.

Cloud to app, one team

Cloud and IoT, backend, web, mobile, data and AI engineers who work as one delivery team.

With you after launch

Support and optimization retainers keep your platform running and improving once it ships.

Industries we have shipped for: smart home and lighting, energy, industrial safety, healthcare, veterinary, hospitality, B2B data and e-commerce.

HE

Hassaan Elahi

Founder and CEO, DiveScale LLC

8+ years building production software, including leading engineering teams at multinationals

Founded DiveScale in 2023 and has led the delivery of 15+ production platforms since

AWS Certified Cloud Practitioner and ACM programming competition winner

hassaan.elahi@divescale.com

From the device to the dashboard, we build everything after the board

Your team owns the hardware. We build what comes after it: secure connectivity, the data layer, and the apps operators, installers and customers use.

Your team

What DiveScale builds



Device

Boards, sensors, controllers and firmware designed by your hardware team.

Connectivity and identity

MQTT and AWS IoT Core, X.509 certificates, fleet provisioning, gateways and pairing flows.

Cloud and data

APIs, device registry and shadows, time-series storage, rules, alerts and firmware and OTA management.

Apps and consoles

Fleet and operations consoles, installer commissioning apps, customer web and mobile apps.

AI and automation

Edge ML, anomaly detection, assistants over your manuals and service data, vision checks.

Connected products fail at the seams: device and gateway, gateway and cloud, prototype and fleet. We design these layers together so the seams are engineered, not assumed. Need extra hands instead? Our engineers can also join your team directly.

What we build for hardware and connected product teams

Provisioning and device identity

Secure onboarding for every unit that leaves the factory, with certificates, claim flows and a device registry your support team can search.

X.509 and mTLS, JITP and fleet provisioning, QR and BLE claim flows

Telemetry and fleet data

Ingestion that holds at fleet scale, time-series storage with retention tiers, and rules that turn readings into alerts and actions.

MQTT, AWS IoT Rules, TimescaleDB, Timestream, InfluxDB, EventBridge

Firmware and OTA management

A version registry per board and core, OTA campaigns with staged rollout and rollback, and per-device version visibility.

AWS IoT Jobs, S3 firmware storage, canary to beta to general rollout

Commissioning and installer apps

Guided pairing for controllers, dimmers and sensors, job workflows, and live status for field teams during installation.

React Native and native mobile, offline-tolerant flows, support requests

Fleet consoles and customer apps

Operations dashboards for device health and project status, and the web and mobile apps your customers use every day.

React, Next.js, real-time updates, multi-tenant access and roles

Edge and product AI

Presence and anomaly models at the edge, predictive maintenance, and assistants that answer from your manuals and service history.

Edge ML, computer vision, LangChain, RAG with pgvector, Claude and OpenAI

Lunul: intelligent home lighting, from installation to move-in

Client: Lunul (Ensemble) / Smart home lighting and controls / New construction and open-wall renovation

Lunul unites premium lighting, dimming engines and wall controls in one system that architects, designers and builders specify. It learns each household's rhythms and adjusts light locally, even offline.

The challenge

- Handover between installers and homeowners was complex and manual
- Pairing controllers, dimmers and sensors with the app was hard in the field
- No real-time view of devices and projects during installation

What DiveScale built

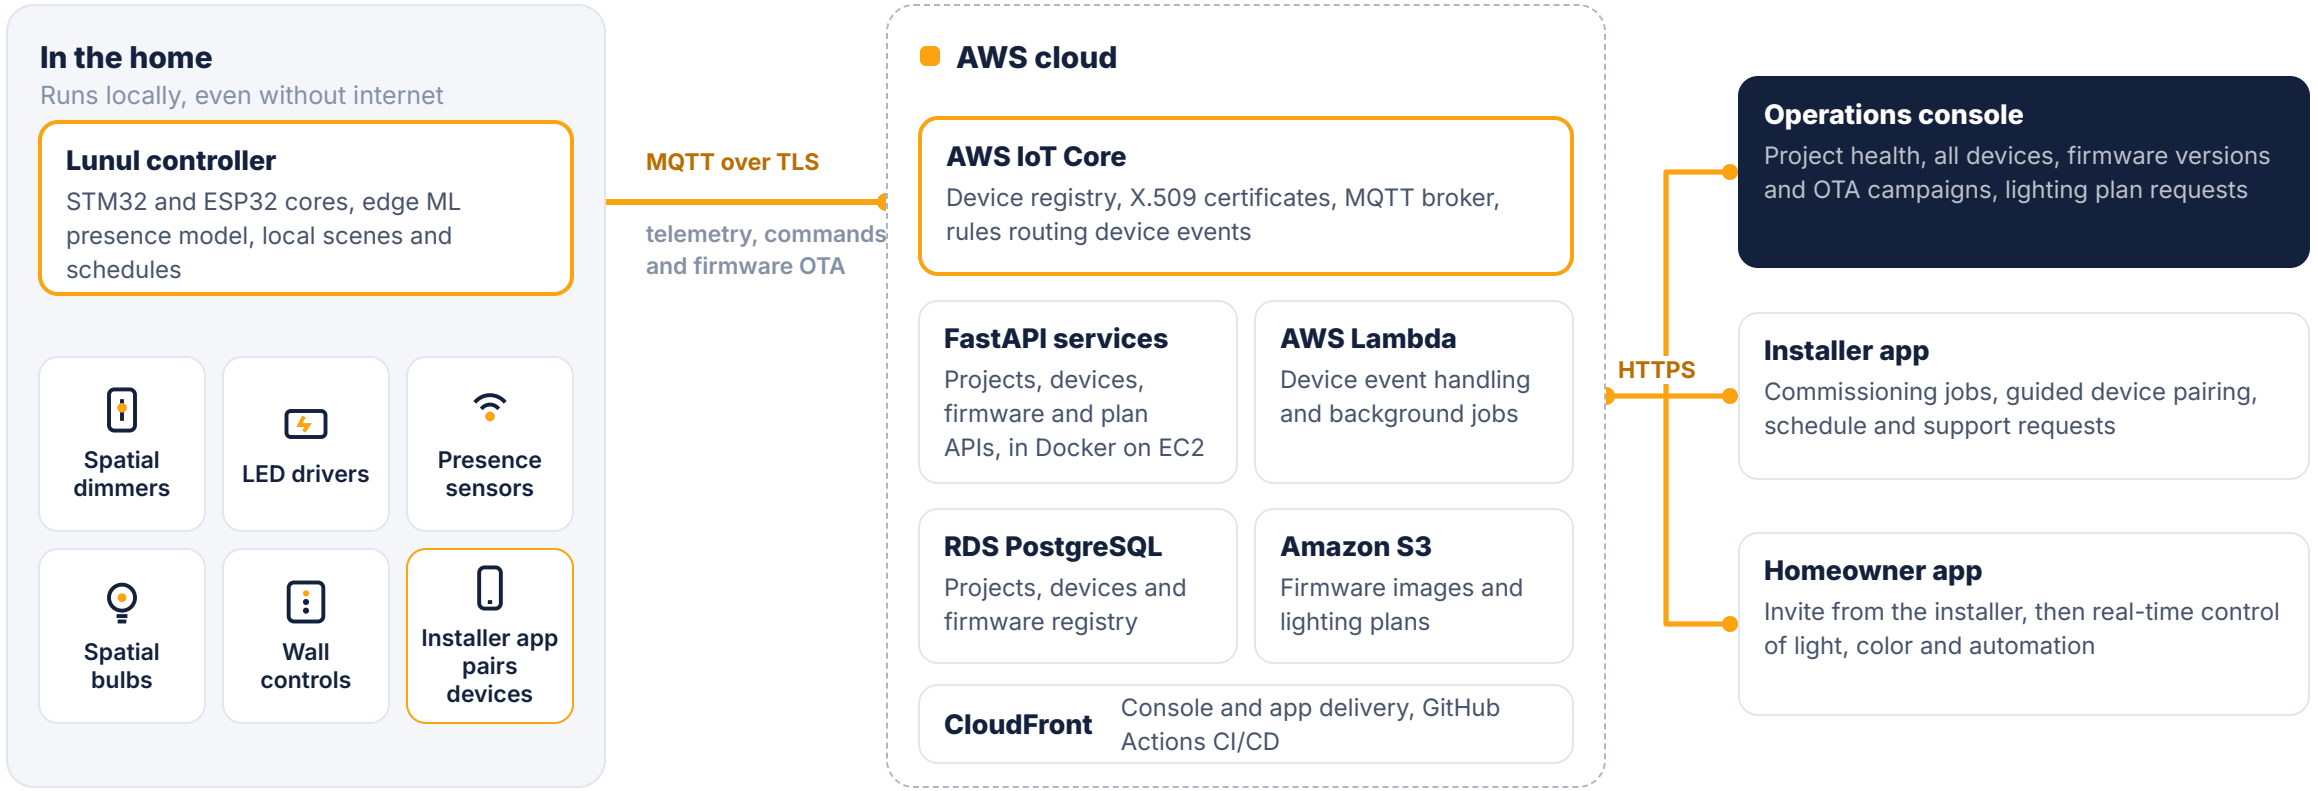
- **Operations console** for projects, devices, firmware and lighting plans
- **Installer mobile app** with a guided pairing and commissioning flow
- **Homeowner experience** with automated invites and real-time control of light, color and automation
- **AWS IoT Core backend** linking every device to its project



React | TypeScript | Python and FastAPI | PostgreSQL | AWS IoT Core | Lambda | EC2 and RDS | S3 and CloudFront | GitHub Actions

How Lunul fits together, from the wall to the cloud

Scenes and presence run locally in the home. The cloud runs projects, devices, firmware and apps.



Every project tracked through Planning — Shipped — Started — Commissioned — Moved in

Simplified view of the production system

Hotelly: in-room hardware, managed across every hotel

Client: Hotelly / Hotel room technology: Apple TV IPTV systems, iPad compendiums and digital signage

- **Multi-property operations console** with every hotel site, its users, billing and orders in one place
- **Hotel setups and equipment orders** tracked from order to install, with a hardware catalog that includes IPTV headends
- **QR-based device enrollment** and automated resets of in-room entertainment when guests check out
- **Support tickets and alerts** that resolve common device issues before they reach a technician

70 to 95%
less manual work resetting in-room entertainment

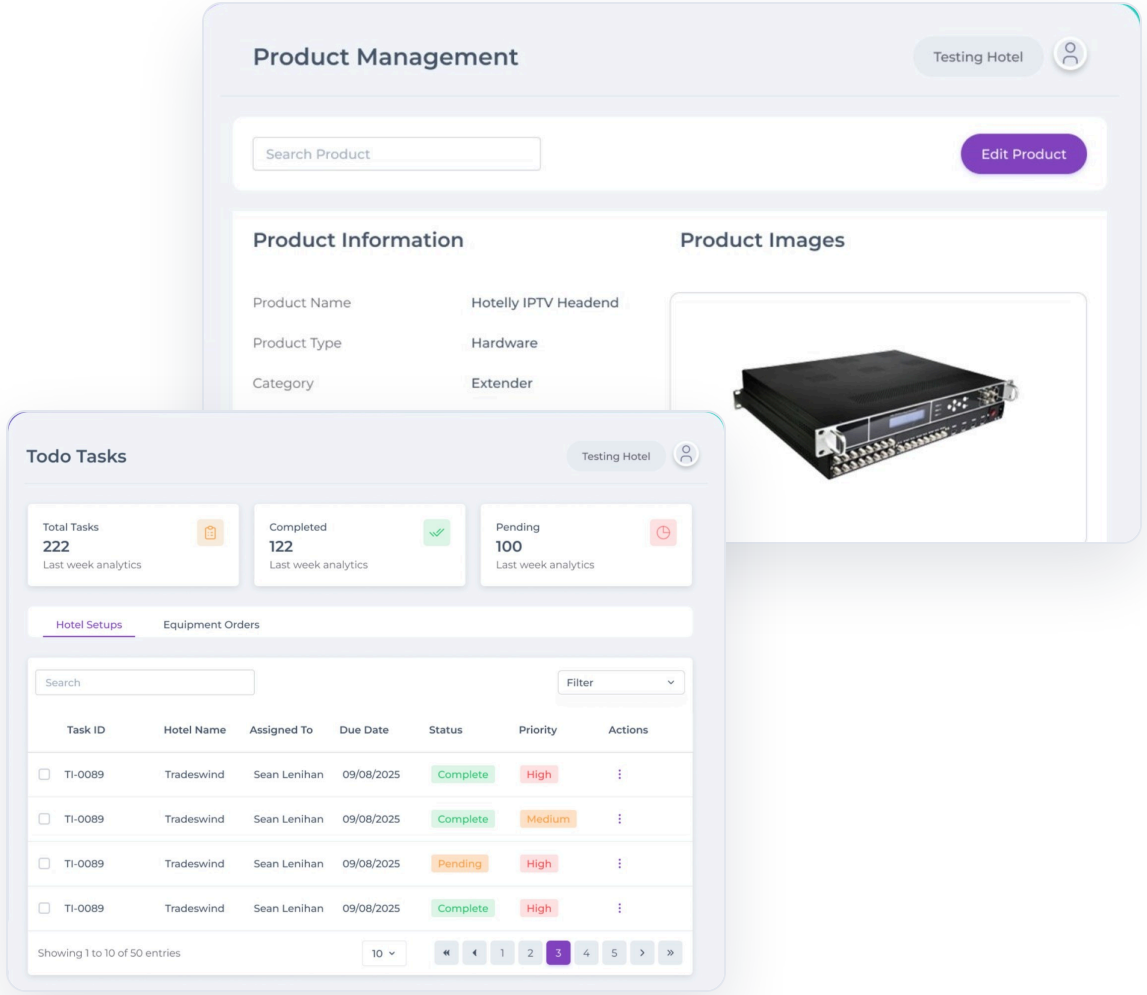
50 to 80%
faster device provisioning with QR enrollment

50 to 75%
of common issues resolved before escalation

60 to 90%
less admin time with unified billing

Estimated ranges from the Hotelly case study, based on published industry benchmarks.

Stack: Next.js, NestJS, Prisma, PostgreSQL, Redis, AWS ECS on Fargate, RDS Multi-AZ, S3, CloudFront, Stripe



Trawa: running every energy asset as one fleet

Client: Trawa / Energy platform for commercial and industrial electricity users in Germany

- **Intelligent asset control:** CHP units, batteries, EV charging stations and cooling systems coordinated on demand, grid signals and price forecasts
- **Load shifting, peak shaving and price arbitrage** that run automatically instead of by hand
- **Consumption analytics** by site, time and usage type, replacing monthly bills with real-time dashboards
- **Green energy procurement** with direct PPA offers from producers

23%
lower electricity costs for a verified customer within 6 months

10 to 30%
average electricity savings across customers

60 min
to a final offer, down from 2 to 4 weeks

about €620
a year in arbitrage value per 10 kWh of battery

Stack: AWS IoT, MQTT, TimescaleDB, PostgreSQL, Python, Node.js, React, D3 and Chart.js, Docker



Assets the platform coordinates


CHP units

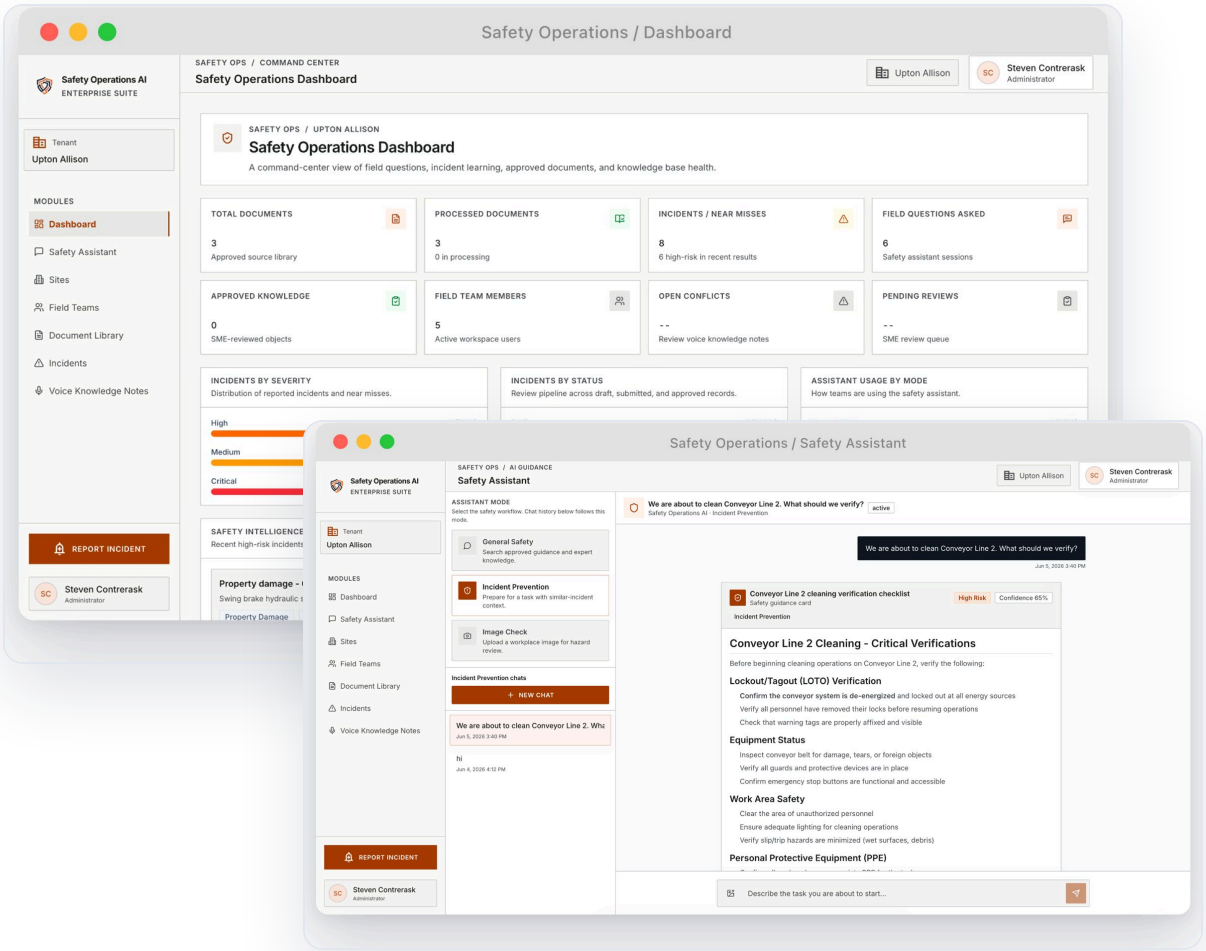

Cooling systems


EV charging


Heating and AC

Safety Operations AI: trusted guidance for field teams

Industrial safety platform / Built from SOPs, incident reports and site photos for manufacturing, construction, oil and gas, and logistics



- **Field worker copilot** that answers from approved procedures, with source citations
- **Image-based hazard checks**, pre-work risk briefs, PPE guidance and stop-work triggers
- **Expert approval and conflict detection** before guidance reaches the field, with a full audit trail

75%
faster incident reporting, 2 hours to 30 minutes

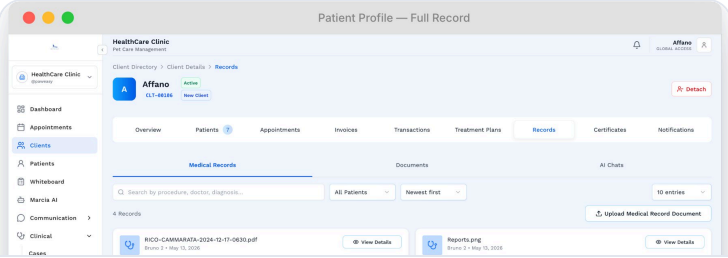
87%
faster toolbox talk prep, down to 15 minutes

70%
less time on audits and inspections

99.9%
of guidance traceable to its source document

The same pattern works for your products: turn device manuals, service logs and field photos into an assistant for your technicians and support team.

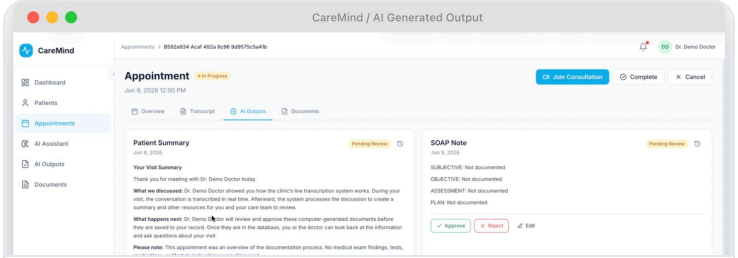
More platforms we have designed and launched



PetHealthcare AI OS

Veterinary hospital operating system

An AI-native platform for clinics that cannot go offline, on Amazon EKS with an Aurora global database.



CareMind AI

HIPAA-grade telehealth

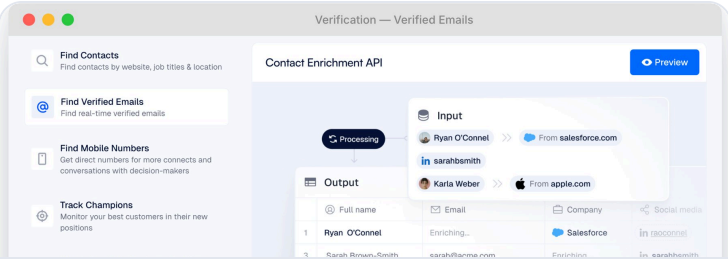
Video visits with live AI transcription and clinician-approved notes, cutting documentation time by 60%.

1.5 days to 2 hours

Sub-Ready

Property management operations

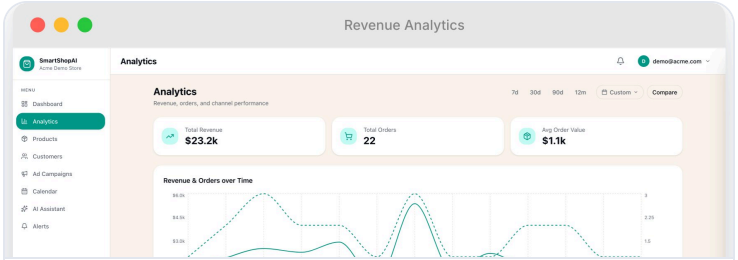
AI onboarding for subcontractors: W-9 and COI checks, MSA signing in DocuSign, and vendor setup in QuickBooks Online.



Waterfall

B2B data platform

One API over 30+ data vendors, with ETL pipelines on AWS Lambda, Step Functions and ECS.



Smart Shop AI

E-commerce operations

Shopify, Google Ads and Meta Ads in one dashboard, with a real-time AI analyst that flags what to act on.

Real-time voice agents

Voice interfaces for products and services

Voice booking and hands-free guidance with natural conversation, built on LiveKit and Deepgram.

Hardware, protocols and platforms we integrate with

If your device can talk, we can connect it to a cloud, an app and the people who need its data.

Device platforms

- ESP32 and ESP32-S3
- STM32
- Nordic nRF52 and nRF53
- ARM Cortex-M
- Raspberry Pi
- Embedded Linux
- NVIDIA Jetson
- Google Coral
- Arduino
- Industrial PCs and gateways

Wireless and networking

- Wi-Fi
- Bluetooth LE
- Matter
- Thread
- Zigbee
- Z-Wave
- LoRaWAN
- LTE-M and NB-IoT
- Ethernet and PoE

Industrial and building

- Modbus RTU and TCP
- OPC UA
- BACnet
- CAN bus
- KNX
- DALI
- DMX512
- 0 to 10V dimming
- OCPP for EV chargers

Messaging and APIs

- MQTT
- CoAP
- AMQP
- WebSockets
- Server-sent events
- REST
- GraphQL
- gRPC
- Webhooks

IoT cloud

- AWS IoT Core
- Device Shadow
- Fleet Provisioning
- IoT Jobs for OTA
- AWS IoT Greengrass
- Azure IoT Hub
- Azure IoT Edge
- EMQX and HiveMQ

Data and time series

- TimescaleDB
- Amazon Timestream
- InfluxDB
- PostgreSQL
- Kinesis
- EventBridge
- Grafana
- Snowflake and BigQuery export

Security

- X.509 certificates
- mTLS
- JITP and JITR
- Certificate rotation
- Signed OTA images
- Role-based access
- Audit logging

Edge AI and vision

- TensorFlow Lite
- ONNX Runtime
- OpenCV
- Google Vision
- Anomaly detection
- LLMs with RAG
- Voice with Deepgram

How a connected product build runs with us

We prove the whole path early, then scale it. A real device, real connectivity and a real dashboard come first, before we add features.

1. Scope the system

Device, gateway, cloud, apps and operations designed together, with the outcome defined up front.

2. Prove the seams

A real device reporting to a real dashboard within the first milestone.

3. Build for the fleet

Provisioning, telemetry storage, alerts, and OTA with staged rollout and rollback.

4. Pilot and harden

A small field pilot, heavy instrumentation, and fixes for what the field surfaces.

5. Operate and evolve

Fleet operations, release cadence, and a roadmap driven by the data the platform produces.

How we keep quality high

Output defined first

We agree what "working" means, then write the tests around it before we build.

We close the loop

Every failure found in testing or the field feeds back, so the same error does not return.

Built to hand over

Infrastructure as code, CI/CD on every change, and delivery into your repositories and cloud accounts.

What clients say about working with us

> Hassaan and team are true professionals... well organized without being rigid, very clear communication, apply software development best practices end-to-end... and deliver high quality solutions quickly. They will be our first call for full stack work in the future.

Brian Hendrickson
CEO and Co-Founder, Lunul

"I highly recommend Hassaan for any web, cloud, or AI project. He is an excellent engineer with profound expertise in modern technologies, including setting up AWS and Azure platforms for our ChatGPT app."

Daniel, CEO of Humaitrix

"Hassaan is an exceptional developer and a pleasure to work with... Hassaan excelled as a technical project manager, leading the React/Node development team with confidence and skill."

Abdullah Sabih, CEO of Masaf KSA

"They are reliable, proactive, and genuinely care about creating value. The quality of the work and the smooth communication made the whole process easy and productive."

Said Aadmsaid, CEO of AI Maths Tutor

DiveScale
Your Technology Partner

Let's build the software around your hardware

Tell us what your device does and who needs its data. We will come back with a plan, a timeline and a fixed quote or a team proposal.

Option one

Fixed-scope build

A platform, console, app or AI feature delivered against agreed milestones and a fixed price.

- First milestone proves device to dashboard
- Scope, milestones and price agreed up front
- You own the code and infrastructure

Best for a defined product or feature with a launch date

Option two

Dedicated engineers

Senior cloud, IoT, full-stack and AI engineers who join your team and work in your tools.

- Start with a 2-week trial
- Scale the team up or down as you grow
- Backed by a firm that owns delivery

Best for teams that need more engineering capacity now

Talk to the founder

Hassaan Elahi

Founder and CEO, DiveScale LLC

Book a call

Scan to share your project and book a 30-minute call.

hassaan.elahi@divescale.com
+1 (302) 308-2674 / divescale.com