SENSORLOGIX · FORWARD DEPLOYMENT

Hardware and software specification

A collector beside each sensor, and a forward node where the analysis has to keep going with no link. This is what each one is built from, what it runs, and roughly what it carries.

The forward node runs the whole DataLogiX platform with SensorLogiX, under the same rules as the rear. It installs from media and never calls home.

AT A GLANCE

Three tiers, sized to where they go

	TIER 1 Edge collector	TIER 2 Forward node	TIER 3 Forward node, large
WHERE	Vehicles, drone ground stations, masts and aircraft.	Ships, forward bases and incident command posts.	Larger ships and headquarters elements.
WHAT IT DOES	Captures and buffers sensor data through a lost link, with optional detection on the box. It doesn't run analyses.	The whole platform for one site, fully disconnected for days.	The whole platform with more video streams, longer retention and failover.
HARDWARE	A small sealed rugged box with an NVIDIA Jetson Thor module.	One short-depth rugged 2U server with an NVIDIA RTX PRO 6000 Blackwell Max-Q, 96 GB.	Two rugged 2U servers, each with one 94–96 GB GPU: NVIDIA H100 NVL or RTX PRO 6000.
LANGUAGE MODEL	gpt-oss-20b, optional	gpt-oss-120b	gpt-oss-120b
POWER	About 220 W	About 1.0–1.2 kW	About 2.5–3 kW

Collectors upload to a forward node or the rear. Forward nodes sync both ways with the rear over a network, on removable media or through your cross-domain device.

The hardware, tier by tier

A minimum and a recommended build for each tier. The chassis and its GPU slots are the parts to check first: they decide what fits.

Edge collector

TIER 1

ITEM	MINIMUM	RECOMMENDED
Compute module	NVIDIA Jetson T4000: 64 GB, 40–70 W	NVIDIA Jetson T5000 (AGX Thor): 128 GB, 40–130 W
For cold or a 10-year life	NVIDIA Jetson AGX Orin Industrial: 64 GB ECC, module rated −40 to 85 °C	The same
Storage	1–2 TB NVMe, self-encrypting (TCG Opal), FIPS 140-3	4 TB NVMe, the same
Network	1–2 × GbE; the radio or modem is yours	2 × 2.5 or 10 GbE
Power input	18–32 VDC vehicle power with MIL-STD-1275F transient protection; 28 VDC aircraft power to MIL-STD-704F	The same
Enclosure	Sealed, fanless or filtered; IP65–67 for outdoor masts	The same
Power draw	About 150 W	About 220 W at full load
Security	TPM 2.0 or the module's secure boot; encrypted storage	Adds a USB FIPS 140-3 Level 3 HSM where keys sign at the edge

Forward node

TIER 2

ITEM	MINIMUM	RECOMMENDED
Chassis	Short-depth rugged 2U, about 400–475 mm deep, with two double-wide GPU slots rated 300 W or more, front service and filtered airflow	The same
CPU	32 cores or more, x86-64, one or two sockets	32–48 cores
Memory	256 GB ECC	256–384 GB ECC
GPU	1 × NVIDIA RTX PRO 6000 Blackwell Max-Q: 96 GB GDDR7 ECC, 300 W, up to 4 MIG partitions	The same, with the second slot kept free
Data storage	8 TB usable on mirrored NVMe, self-encrypting (TCG Opal), FIPS 140-3, about one drive-write a day	15 TB usable, the same
Boot	Mirrored M.2 pair, encrypted with LUKS, key sealed to the TPM	The same
Network	2 × 10/25 GbE (SFP28) and a management port	2 × 25 GbE and 2 × 10 GbE
Security	TPM 2.0 and a FIPS 140-3 Level 3 HSM, USB or PCIe	PCIe HSM
Power	1.0–1.2 kW at full load; dual supplies; 100–240 VAC, with a 28 VDC option for vehicles	The same
UPS	Online double-conversion, 1.5–2 kVA, 15–30 minutes, signaling a clean shutdown	The same
Time	GPS receiver with PPS; no network time source needed	The same
In the kit	Transit cases, CAC/PIV readers, FIPS 140-3 encrypted removable drives for sync, a local switch	The same

Forward node, large

TIER 3 · TWO NODES

ITEM	PER NODE
Chassis	Short-depth or standard-depth rugged 2U with at least two double-wide GPU slots; 400 W or more for the H100 NVL
CPU and memory	48–64 cores; 384–512 GB ECC
GPU	Node A: NVIDIA H100 NVL (94 GB) or RTX PRO 6000 Max-Q (96 GB). Node B: the same, for video, overflow and model failover
Storage	30–60 TB usable across the pair, mirrored and replicated between the nodes
Network	2 × 25 GbE per node, plus a direct 25 or 100 GbE link between the nodes
Security	TPM 2.0 and a PCIe FIPS 140-3 Level 3 HSM in each node
Power	About 1.3–1.5 kW per node, 2.5–3 kW in all; a UPS sized for both

Why two nodes: one 94–96 GB GPU already runs gpt-oss-120b. The second node adds failover, keeps video work apart from the model and adds capacity.

GPUs we build with

NVIDIA

GPU	MEMORY	POWER	MIG	USED FOR
RTX PRO 6000 Blackwell Max-Q	96 GB GDDR7	300 W	Up to 4	Forward node
H100 NVL	94 GB HBM3	350–400 W	Up to 7	Large forward node
L40S	48 GB GDDR6	350 W	None	Smaller builds, with gpt-oss-20b
L4	24 GB	72 W	None	Video only, or a small model
Jetson T5000 / T4000	128 / 64 GB	40–130 / 40–70 W	Yes	Edge collector

2 · WHAT RUNS WHERE

What each tier does, and what waits

Forward, the checks are the cloud's: the SQL guard, number bindings, recomputation, the verifier, traceability, markings and dissemination run exactly as they do in the rear. A capability that can't meet its standard on local hardware is switched off, never run in a weaker form.

CAPABILITY	COLLECTOR	FORWARD	LARGE	FORWARD NODE, NO GPU
Capture, buffering, store-and-forward	Yes	Yes	Yes	Yes
Ingest and parsing, entity search, maps, dashboards, review, dissemination, markings and traceability	—	Yes	Yes	Yes
New analyses and reports, written by the model	—	Yes	Yes	Queued, with a notice, until a GPU or a link is available
Video detection and tracking	Optional	Yes	Yes	Queued
OCR of scanned pages	—	Yes	Yes	Tesseract only, less accurate
Semantic and image search	—	Yes	Yes	Keyword and entity search only

Connected extras, such as Microsoft 365, Purview, Google Drive and hosted models, run on a forward node only when a link is up and policy allows them. In FIPS mode, features whose algorithms sit outside FIPS 140-3 are off on every tier: post-quantum signatures, OpenPGP, C2PA content credentials and X25519/Ed25519.

The same software everywhere it runs

The forward node runs the same signed build as the rear. It needs no network to install, run or update.

Platform

LAYER	WHAT RUNS
Operating system	Red Hat Enterprise Linux 9, or Ubuntu 24.04 LTS with Ubuntu Pro, in FIPS mode from install. Disks encrypted with LUKS, the key sealed to TPM 2.0.
Kubernetes	RKE2 in FIPS mode, hardened to the DISA STIG and the CIS benchmark, with Pod Security “restricted” and audit logging always on. One node, or two for the large forward node.
Ingress	Traefik
GPU	NVIDIA driver, container toolkit and GPU operator, with MIG partitions.
Database	PostgreSQL 17 with pgvector and PostGIS
Object storage	SeaweedFS, S3-compatible
Workflows	Temporal, self-hosted
Keys and secrets	OpenBao, sealed by the HSM over PKCS#11 or by Shamir key shares. Envelope keys are rooted in the HSM.
Identity	Keycloak with CAC/PIV smart cards, OIDC and SAML, in FIPS mode. A local directory, mapped to the rear’s on the first sync.
Registry and updates	zot with cosign signatures and TUF-checked content. Bundles on media, delta updates and rollback. No call-home.
Monitoring	OpenTelemetry collector, Prometheus, Grafana and Loki
Backup	pgBackRest and Velero, to encrypted removable media
Time	chrony and gpsd, with PPS from the GPS receiver

DataLogiX on the node

AREA	WHAT RUNS
Platform	The DataLogiX API, web app and workflow engine, webhooks and the relay gateway.
Documents	Conversion of Office, PDF and older formats; OCR; Word, PDF, PowerPoint and Excel reports; citation styles; XBRL and C-CDA readers; malware scanning with offline signature updates.
Security	OpenTDF key access, a local prompt screener that fails closed, and a code sandbox with no network.
Ingest	The relay, which on a collector captures video, splits out KLV metadata and stores and forwards; SFTP drop; watched folders; connectors.

Models

PURPOSE	EDGE COLLECTOR	FORWARD NODE	LARGE FORWARD NODE
Planner, writer and checker	gpt-oss-20b, optional	gpt-oss-120b (MXFP4, about 65 GB), or gpt-oss-20b	gpt-oss-120b
Second checker, from another model family	—	An open model of about 8–30B parameters, such as Qwen3 or Mistral Small	The same, larger

PURPOSE	EDGE COLLECTOR	FORWARD NODE	LARGE FORWARD NODE
Embeddings	—	A local text embedder; SigLIP 2 for images	The same
Video detection and tracking	RF-DETR and ByteTrack on ONNX Runtime	The same, plus a captioner. Face detection is off unless policy turns it on	The same
OCR	—	PaddleOCR-VL on the GPU; Tesseract as fallback	The same
Speech to text	—	Local speech recognition	The same
Forecasting	—	Chronos-2 and statistical models, on the CPU	The same

Models ship inside the bundle and are checked by hash before they load. All are open-weight, under permissive licenses such as Apache-2.0. The checker is never from the writer's model family.

4 · CAPACITY

Roughly what each tier carries APPROXIMATE

All figures here are approximate and meant for sizing. Real numbers depend on the hardware, the questions, the data and how busy the site is, so we size each build with you.

Users and work

	FORWARD NODE	LARGE FORWARD NODE
Named users	25–50	75–150
People working at once	10–15	30–50
Analyses running at once	2–3	5–8
Question to verified report, typical	10–25 minutes	6–15 minutes
Entity and document search	Under 1 second, across about 10 million passages	The same, with more headroom

Search, maps, review and dashboards don't use the GPU; the language model sets the pace. A report takes about 30,000–60,000 generated tokens across the planner, writer and checkers.

Model speed

MODEL AND GPU	ONE REQUEST	MANY AT ONCE
gpt-oss-120b on H100 NVL	80–150 tokens/s	Several hundred tokens/s
gpt-oss-120b on RTX PRO 6000 Max-Q	50–110 tokens/s	200–400 tokens/s
gpt-oss-20b on RTX PRO 6000 or L40S	120–250 tokens/s	500+ tokens/s
gpt-oss-20b on Jetson T5000	20–50 tokens/s	—

Video and tracks

	EDGE COLLECTOR	FORWARD NODE	LARGE FORWARD NODE
Live 1080p streams with detection, at about 5 sampled frames a second each	1–2, on the box	4–6, while the model serves users	10–16

	EDGE COLLECTOR	FORWARD NODE	LARGE FORWARD NODE
Recorded video	—	About 4–8× real time per GPU partition	About twice the forward node
Track and message updates on the live map (CoT, AIS, ADS-B)	Forwarded	Within about 5 seconds of receipt	The same

Detections arrive within about 2 minutes of each 30-second video segment closing.

Documents and data

	FORWARD NODE	LARGE FORWARD NODE
Scanned pages through OCR	1,000–2,000 an hour, off-peak	2,500–5,000 an hour
Born-digital documents	3,000–6,000 pages an hour	6,000–12,000 pages an hour
Structured loads (CSV, database extracts)	10–50 million rows an hour	20–100 million rows an hour

Storage

	EDGE COLLECTOR	FORWARD NODE, 8–15 TB	LARGE FORWARD NODE, 30–60 TB
Video	Hours to days of buffer; the oldest dropped only by policy	About 2–4 weeks of raw video from four 1080p streams, at about 2.7 TB a week, unless policy sends raw video to the rear and keeps clips	3–6 months at four streams, or more streams
Models and software	20–50 GB	150–250 GB	150–250 GB
Data, documents and indexes	—	Typically 1–3 TB	Typically 5–10 TB

Running it

- **Install** From media, with no network, in about 30 minutes.
- **Cold start** About 5–10 minutes from power to serving, plus the HSM or operator unseal.
- **Updates** Signed delta bundles on media, checked before they install, with rollback.
- **Disconnected** 72 hours and more with no link at all.
- **Zeroize** Keys destroyed in seconds, leaving the data unreadable; then the disks are wiped.
- **Sync with the rear** Resumable, by priority and within a bandwidth cap. Every record is checked against markings and release rules on both sides.

5 · STANDARDS

Standards the kit is built to

We don’t certify hardware. The kit is specified to these standards, the hardware vendor’s certification of the chosen model qualifies it, and accreditation stays with you.

STANDARD	COVERS	TIERS	WHO HOLDS IT
MIL-STD-810H	Environment	All	Hardware vendor
MIL-STD-461H	Electromagnetic interference	All	Hardware vendor
MIL-STD-1275F	28 VDC vehicle power	1, 2	Hardware vendor
MIL-STD-704F	Aircraft power	1	Hardware vendor
MIL-DTL-901E, MIL-STD-167-1A	Shipboard shock and vibration	2, 3	Hardware vendor and shipyard
RTCA DO-160G	Airborne equipment	1	Hardware vendor
IEC 60945, IEC 60529, IEC 62368-1, FCC Part 15B, NEBS	Marine equipment, ingress protection, safety, emissions, telecom sites	As needed	Hardware vendor
FIPS 140-3	Cryptographic modules	All	The OS, drive and HSM vendors
DISA STIGs	Hardening of RKE2, Kubernetes and the OS	2, 3	We configure; your assessor checks
TEMPEST, CSfC, authority to operate	Emanations, classified networks, accreditation	2, 3	You
TAA, NDAA §889	Country of origin, covered telecommunications	All	Hardware vendor; EVAS by representation

6 · SECURITY ON THE BOX

Locked down before it ships

The node keeps the rear’s controls with no link at all, and it’s built to be inspected.

- **Encrypted at rest.** Self-encrypting drives under LUKS, per-tenant data keys wrapped by the HSM, and sealed columns encrypted just as in the cloud.
- **Audited.** Every action is recorded as it is in the cloud, and the audit trail syncs to the rear.
- **Signed where it’s made.** Reports and traces are signed with the node’s own key and timestamped from GPS time. The rear checks each signature on sync.
- **Smart-card sign-in.** CAC/PIV, with no shared accounts.
- **No outbound network.** Egress is denied by default. Connected extras run only by policy, and every attempt to reach out is recorded.
- **Evidence for your assessor.** Software bills of materials, scan results, STIG checklist results and a control mapping come with every build.

Tell us where it has to run

We’ll size the kit with you, from the sensors and links you have to the people who will use it.

evas.live/contact