Integration of Foundation Models with Up-/Down-stream Applications and Geospatial Digital Twins

Authors:
Aziza Ibrayeva, CEO, CELSTRAT (Celestial Strategy)
David Aronchick, CEO, Expanso

CELSTRAT

expanso

Importance of Geospatial Digital Twins for NASA & ESA

Mission-critical apps:

- Climate monitoring
- Disaster response
- Logistics
- Military reconnaissance

Current challenges:

- Limited bandwidth
- Latency
- Intermittent connectivity
- High security

expanso

The Geospatial Data Explosion

Market growth: Geospatial analytics to reach $226B by 2030 (12.6% CAGR)

Growth drivers: AI, IoT, remote sensing, smart cities, urban planning

Data volume: 100 TB/day from weather data; global data to exceed 175 ZB by 2025

Edge processing critical for relevance, efficiency, and security

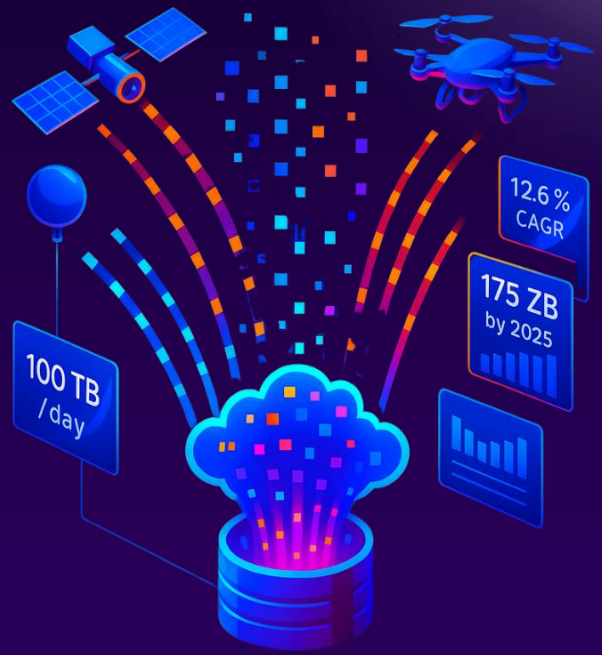

expanso

Satellite Operation Constraints

Typical LEO satellite orbit period: **~90 minutes** (14–16 orbits/day).

Ground station passes per day: **5–8** (8–12 min each).

Real-world downlink speeds: X-band realistic speeds **10–20 Mbps** (theoretical 100–200 Mbps).

Data downlink limitations: Only **10–20%** of daily generated data transmitted.

Centralized **Log Processing**

And many others.

Typical centralized architectures require transferring massive volumes of data to Earth, causing delays, high costs, and missed opportunities in time-sensitive applications.

Why Edge Processing Matters

- Local filtering of low-value/noise data
- Prioritize anomalies or model-predicted targets
- Reduces congestion in ground data pipelines

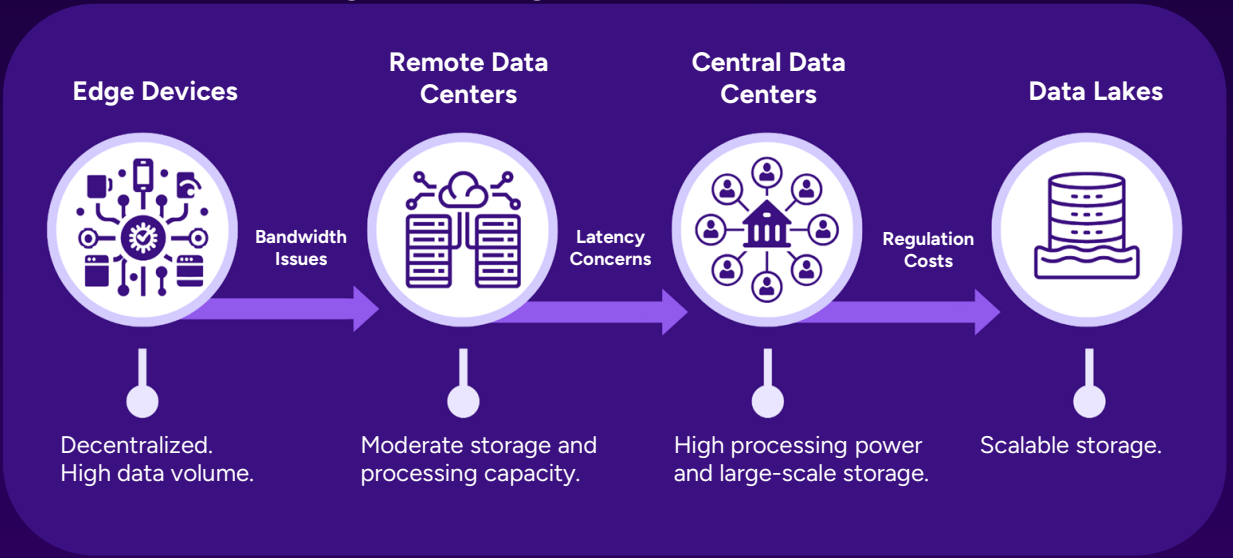

In orbit,
stages past the edge are **INACCESSIBLE**,
making onboard filtering critical

expanso

Foundation Models on the Edge

MODEL DEPLOYMENT APPROACH

- Foundation Models (e.g., YOLOv5 for object detection) containerized and deployed on edge devices.
- Bacalhau executes jobs locally using WASM/Docker, allowing inference in disconnected environments.
- Results are stored, queued, and synced post-reconnection to ground systems.

Benefits of Edge AI Deployment

BENEFITS

- **Low Latency:** Local inference avoids cloud round-trips, keeping latency under ~5s.
- **Bandwidth Optimization:** Only high-value, summarized results transmitted—>90% data reduction.
- **Privacy by Design:** Raw data never leaves the edge device, aligning with secure mission requirements.
- **Scalability:** Distributed architecture supports multi-node inference (fleet-level or constellation-scale).

Research Framework - Edge Architecture

- Distributed, containerized inference jobs/tasks
- Operates in disconnected environments (onboard, remote nods/payloads)
- Queues tasks locally and syncs only essential results post-connection
- Open-source, audit-capable, and vendor-neutral

Implication:
Achieving gains in efficiency and cost WITHOUT ripping and replacing.

Submarine as a Space Proxy

- **Operational Similarities:** Submarines and satellites both operate in bandwidth-constrained, isolated environments with limited real-time connectivity.
- **Edge AI Deployment:** Bacalhau powers local AI workloads directly on unmanned undersea vehicles (UUVs), enabling inference without persistent comms.
- **Resource Optimization:** Efficient use of onboard compute - no new infrastructure required.
- **Security and Compliance:** Meets U.S. Navy security standards, aligning with space mission requirements.

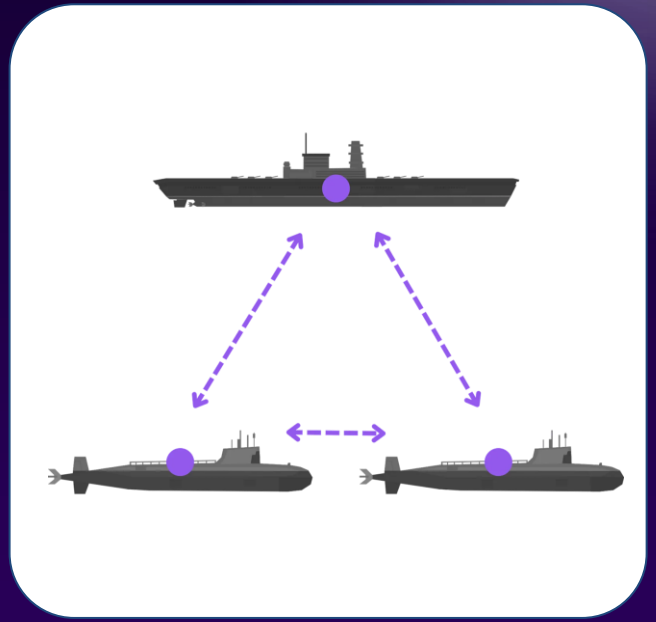

Bacalhau runs **where the data is**

Raw Data	Ingestion	Storage	Transform	Load	Integrate	Visualize

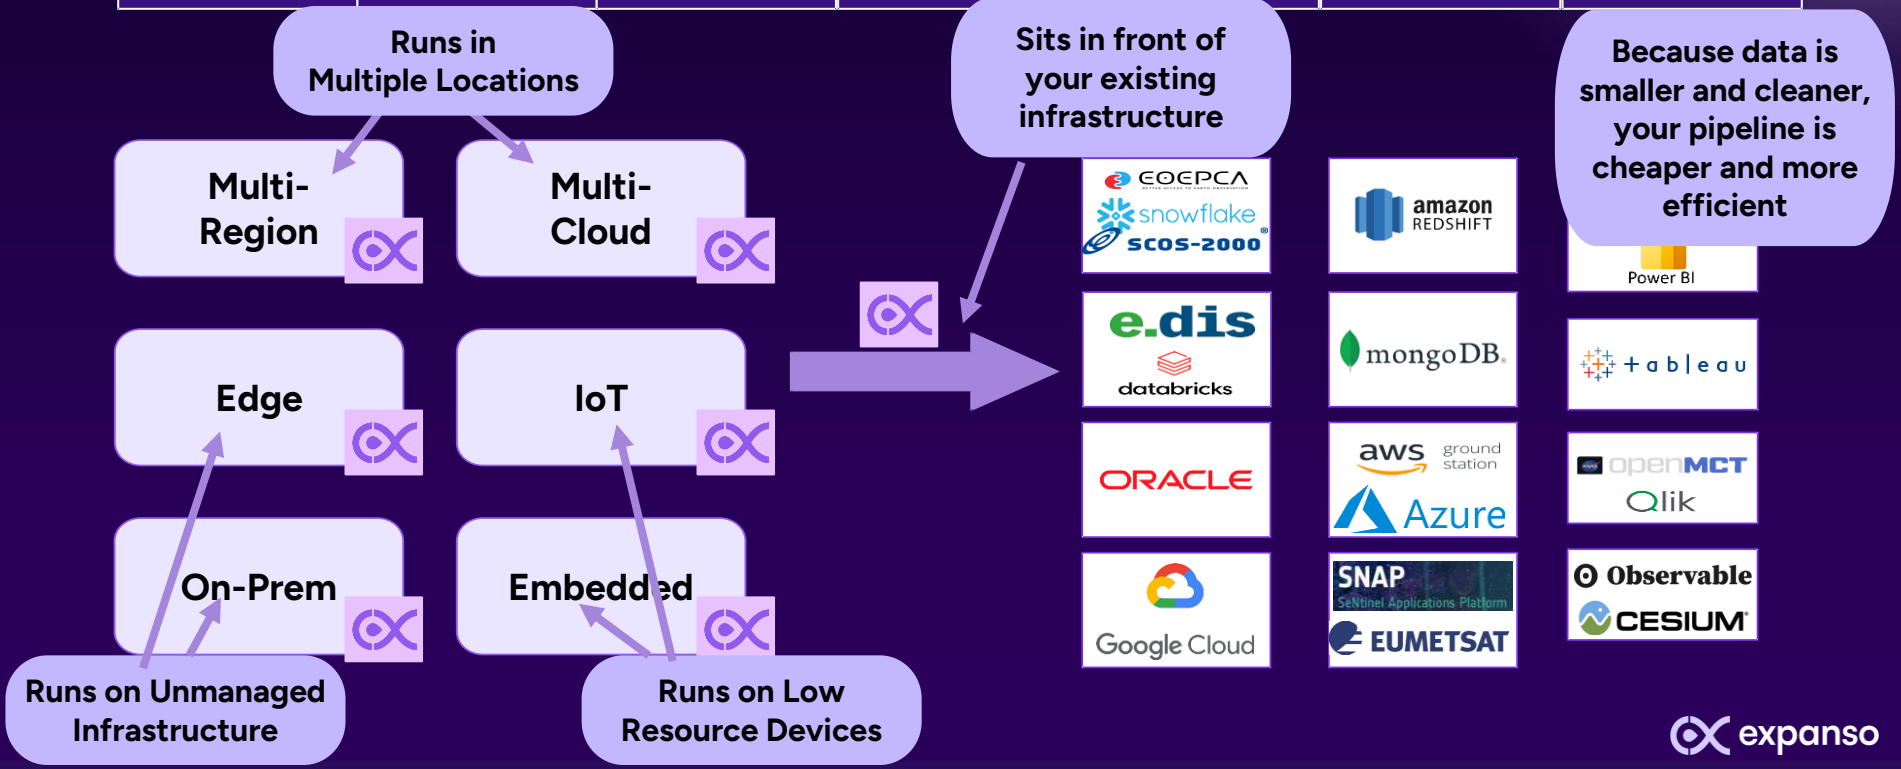

expanso

Results from Analog Testing

Metric	Value / Range	Context / Note
Inference latency	~<5 seconds (internal test)	Fast detection in resource-constrained environments via local inference
Bandwidth savings	>90% (internal test)	Substantial reduction by filtering raw data before sync
Data transmitted per session	~10–20% of original volume	Only high-value results transmitted; rest processed at source
Job queue uptime	100%	All jobs retained and executed post-connection
Success rate under disconnect	Near 100%	Reliable offline execution with automatic sync upon reconnect

Future Roadmap

Orbital test campaign under planning

Goals: Benchmark on-orbit inference, measure latency, validate reconnection

Seeking partners for payload hosting and telemetry evaluation

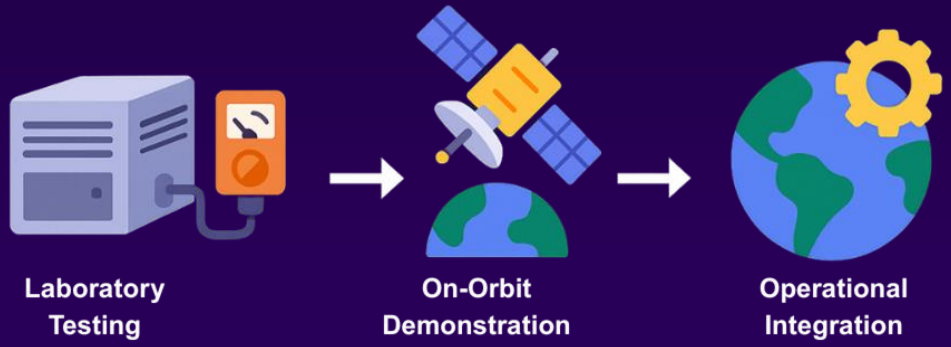

expanso

Strategic Fit & Call for Collaboration

ALIGNMENT WITH NASA & ESA PRIORITIES

- Supports NASA ESTO and ESA FutureEO goals: onboard analytics, data minimization, operational autonomy
- Edge compute = reduced downlink, faster response, and enhanced mission security

RESEARCH INVITATION

- Early testing confirms feasibility in constrained environments (submarine analog)
- Strong potential for orbital deployments and Digital Twin applications
- Open to payload partners and mission collaborators

expanso

Q&A and Thank you!

We welcome collaboration - please get in touch.

Aziza Ibrayeva: aziza.i@celstrat.com
David Aronchick: aronchick@expanso.io

CELSTRAT

expanso