

Why AI-enabled drone swarms need mesh ranging radios

Centurion Tech, White Paper v0.02, June 2026

Alejandro Hernandez

Introduction

As the AI boom expands in full swing, innovative techniques to solve the demands for communication at higher-and-higher data rates must be developed. In data centers, this is already visible: the industry is moving toward high-radix Ethernet fabrics, optical interconnects, co-packaged optics, and carefully scheduled traffic because model training and distributed inference are increasingly network-bound. In heavy edge-compute networks where wired optical links are impractical, it will fall upon wireless link architectures to innovate.

“The network is ultimately responsible for AI performance.” - NVIDIA^[1]

This exact bottleneck is arriving at the frontier of academia. Autonomous drones, robots, and sensor teams increasingly depend on local AI models, visual-inertial odometry, LiDAR mapping, collaborative SLAM, cooperative perception, and distributed planning. These systems do not merely exchange commands. They exchange features, descriptors, keyframes, map fragments, model outputs, poses, covariances, and health telemetry. A swarm is therefore not just a group of aircraft. It is a moving, wireless, compute cluster.

Beyond high datarate links, there are other constraining factors in the innovation of drone swarm technology. Scalability: how do you efficiently communicate in a network of hundreds of drones where each node begins to distort its many neighbors?

The U.S. Department of Defense’s Replicator initiative was built around the premise that thousands of attritable autonomous systems may need to be fielded quickly across domains. ^[3]

At the same time, the positioning assumptions behind many drone systems are breaking. GNSS denial and spoofing are becoming ever-present operational realities in contested and industrial environments.

Open reporting on the Black Sea and Ukraine region describes persistent, dynamic GNSS interference campaigns that degrade navigation for aircraft, ships, and drones. ^[2]

In order for drone swarms to effectively coordinate, both local and relative positioning have to be accounted for, without GNSS. The state-of-the-art ranging solution lies in collaborative SLAM; however, it relies on high-throughput wireless links and fails catastrophically in degraded visual conditions.

The conventional response is to use UWB for ranging. This comes with several key limitations:

1. Each UWB radio comes as a detriment to cost and SWaP from hardware overhead in board area, antennas, power consumption, and integration complexity
2. UWB pulses introduce distortion into adjacent communication transceivers, creating RF coexistence problems that become significantly more severe with scaling
3. UWB suffers from significant latency constraints from scheduled polling cycles, to the point where reliable ranging begins to fail at 10+ UAVs

The key takeaway: UWB scales poorly.

The Idea

Integrated Sensing and Communications (ISAC) collapses this stack. Instead of carrying data on one radio and measuring range on another, the swarm radio itself becomes a source of range, timing, channel, and uncertainty measurements.

Integrating the range sensing layer into the communication layer is a novel way to eliminate these drawbacks. Simultaneously achieving sub-meter precision ranging and multi-Mbps throughput will be challenging; despite this, Centurion Tech has devised Phalanx, a RF hardware-to-protocol-level stack to subvert these challenges and solve this problem with a cross-domain solution. The time is now. Geopolitics is pushing the pace of AI and drone innovation to an explosive sprint. Centurion Tech Inc. will pioneer the radio technology that will serve as the crux of the incoming autonomy age.

Figure 1. Wireless data demand is rising while autonomy adds local burst traffic.

Metric	Value	Relative scale
Global mobile data traffic, 2025	146 EB/month	
Total mobile network traffic incl. FWA, 2025	203 EB/month	
Global mobile data traffic, 2031 forecast	328 EB/month	
Total mobile network traffic incl. FWA, 2031 forecast	515 EB/month	

Ericsson forecasts total mobile data traffic of 328 EB/month in 2031 and total mobile network traffic including fixed wireless access of 515 EB/month. The drone-swarm problem is smaller in absolute traffic, but harsher in latency, mobility, RF contention, and link volatility. [5]

Table 1: The macro trend is clear: wireless networks are being asked to carry more data, more intelligence, and more control-critical traffic.

Why a Mesh Ranging Radio?

ISAC Radios offer several fundamental benefits compared to other sensing and communication stackups, especially when deployed at scale. Significant cost savings can be achieved by embedding sensing into the already-necessary high bandwidth communications layer without the need for additional UWB/sensor hardware. This architecture can leverage lower latency by estimating propagation delay directly from modified data packets rather than running independent scheduled polling events. ISAC Radios can integrate seamlessly into the existing communications infrastructure, offering the same plug-and-play capability of modern WiFi modules. Perhaps most importantly, they can be scaled massively beyond 10+ drones without the throughput or latency constraints faced by other sensing and communication stacks, utilizing an integration-first approach to deploy them rapidly. For advanced users/researchers, ISAC Radios can expose key signal metrics unavailable in traditional radio, providing several rich DoFs for use in model-based estimators. These advantages are explained below and considered against the consequential challenges and additional costs associated with integrating sensing and communication.

Current State-of-the-Art

System	What it does well	What it costs	Where it breaks
RTK GNSS	Outdoor global position	Receiver, antenna, correction link or base station	GNSS-denied, indoors, urban canyons, jamming, spoofing
UWB ranging	Dedicated decimeter-class pairwise range	Extra radio, antenna, MAC schedule, integration path, comms interference	Scaling beyond 10+ UAVs, polling latency, jamming, Multipath Fading
Visual / LiDAR SLAM	Rich geometry and local odometry	Cameras/LiDAR, compute, features, keyframes	Data-rate limitations, Low texture, darkness, occlusion, poor overlap, repeated scenes
Collaborative SLAM	Shared map and pose-graph consistency	Compute, coordination, descriptors, keyframes, states	Data-rate limitations, Network delays, visual overlap, feature quality
MANET radio	Mesh peer-to-peer data transport	Mesh networking occurs at IP-layer, abstraction at cost of performance	Absent positional data, spoofing, jamming, Multipath Fading
ISAC radio	Scalable High-throughput mesh + low-latency precision RF ranging factors	Calibration, Complex PHY/MAC design, estimator integration	HPM jamming, Multipath Fading

Table 2: The current autonomy stack solves communication, ranging, and mapping separately. ISAC Radio collapses part of that stack by moving ranging into the radio already needed for high-throughput communication.

Reduced Hardware Expenses

A dedicated UWB module sounds cheap and simple. But at swarm scale, the cost is not only the transceiver. It is the second antenna, placement constraints, RF coexistence, power, firmware integration, ranging scheduler, timestamp reconciliation, test process, and the fact that a separate localization radio cannot carry the high-throughput data that collaborative autonomy already needs.

A localization-native mesh radio asks a different question:

If the swarm must already carry a high-throughput SDR/FPGA mesh radio, what is the marginal cost of making that radio emit range factors?

The answer is not zero. It requires calibrated timestamps, waveform design, pilot allocation, channel estimation, clock-skew correction, and multipath rejection. But the marginal cost may be far lower than the full system cost of a second ranging radio on every drone, especially when the network must scale beyond small demonstrations.

Scalability

Fleet size	All-pairs links	Implication for separate ranging	PoseLink approach
10 drones	45 links	Polling is manageable	Sparse range graph plus data packets
20 drones	190 links	Scheduling starts to dominate	Nearest-neighbor plus long-baseline edges
50 drones	1,225 links	All-pairs ranging becomes impractical	Estimator-driven graph scheduling
100 drones	4,950 links	Dedicated polling cycles do not scale	Clustered mesh, range-on-packet, adaptive pilots

Table 3: The number of potential pairwise links grows as $(N(N-1)/2)$. At 20+ drones, the MAC scheduler is as important as the range sensor.

A swarm does not need every drone to range every other drone every frame. It needs enough well-conditioned edges to estimate a stable relative frame, maintain collision-relevant local geometry, and prevent map/formation drift. That suggests a sparse, adaptive ranging graph:

$$L(N) = \frac{N(N-1)}{2}$$

but

$$L_{\text{scheduled}} \ll L(N)$$

The scheduler should select links by expected estimator value:

- near-neighbor links for collision avoidance and formation maintenance,

- long-baseline links to reduce global graph drift,
- anchor or leader links when a global frame is available,
- cross-cluster links to merge local maps,
- and degraded-vision links when visual SLAM confidence drops.

This is where ISAC becomes more than waveform design. The radio, estimator, and planner must be coupled.

Speed of Deployment

A localization-native radio should be deployed like a modern mesh module, not like a custom research stack. The goal is a plug-in system with:

- power, RF, and compute interfaces,
- ROS 2 and MAVLink messages,
- calibrated factory delay parameters,
- timestamped range-factor outputs,
- optional raw CSI/CIR export for researchers,
- and a default scheduler that works before custom tuning.

The value is not that every customer writes their own estimator. The value is that the radio exports both packets and navigation factors in forms that autonomy stacks already understand.

Support for Collaborative SLAM and Swarm Communication

Collaborative SLAM already proves the demand for high-bandwidth, low-latency drone-to-drone communication. D^2 SLAM reports representative message sizes of 20.1 kB for complete stereo keyframes and 56.3 kB for complete omnidirectional keyframes at 5 Hz, along with compact keyframes and pose-graph updates. [4]

D2SLAM traffic type	Size	Rate	Payload rate
Complete stereo keyframe	20.1 kB	5 Hz	0.80 Mbps
Complete omni keyframe	56.3 kB	5 Hz	2.25 Mbps
Compact stereo keyframe	1.2 kB	5 Hz	0.048 Mbps
Compact omni keyframe	4.8 kB	5 Hz	0.192 Mbps
Example VINS update	3.5 kB	5 Hz	0.14 Mbps
Example PGO update	81 kB	1 Hz	0.65 Mbps
RF range factor	32-128 B	10-50 Hz/link	kbps/link

Table 4: Collaborative visual SLAM can be Mbps-class per active UAV. RF range factors are tiny by comparison, but can stabilize the same pose graph.

The radio therefore has two jobs:

1. carry the keyframes, descriptors, and state updates that collaborative autonomy already needs;
2. produce cheap RF range factors that keep the relative graph constrained when visual conditions degrade.

Market Opportunity

Multiple trends are converging.

Trend	What changed	Why it matters for Phalanx
AI networking	AI infrastructure is increasingly network-bound; hyperscale vendors now optimize full-stack fabrics around AI traffic. [1]	Edge swarms will also need deterministic, high-rate local links for features, maps, and coordination.
Drone scale	Defense and commercial autonomy are moving from single platforms toward many attritable or cooperative systems. [3]	Hundreds of nodes cannot be treated as a small lab WLAN.
GNSS denial	Interference and spoofing are operational realities in contested regions. [2]	Relative, infrastructure-light localization becomes necessary.
Visual SLAM maturity	Collaborative SLAM can deliver strong relative localization, but depends on visual overlap, features, compute, and communication. [4]	RF range factors become a complementary sensor, not a replacement.
ISAC validation	Recent OFDM ISAC prototypes demonstrate real-time video, range/angle sensing, and radio SLAM. [7]	The concept is no longer purely theoretical.
SDR/FPGA edge compute	Compact SDRs, FPGAs, and Jetson-class compute make custom PHY-to-SLAM integration practical.	A cross-layer radio is now buildable on UAV-class hardware.

Table 5: Phalanx sits at the intersection of AI networking, swarm autonomy, GNSS-denied operation, collaborative SLAM, and SDR/FPGA hardware.

Centurion Tech will become the telecommunications powerhouse of the Autonomous era

System Architecture

Centurion Tech's ISAC Radio is a full-stack radio architecture, not only a waveform.

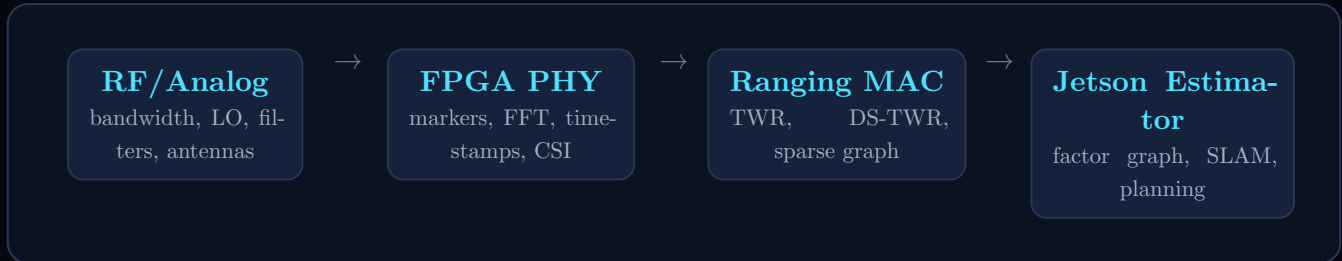


Table 6: Architecture separates nanosecond-sensitive PHY work from graph-level autonomy. The FPGA timestamps and extracts channel observables; the Jetson fuses them with VIO, LiDAR, and SLAM.

Physical architecture

Each node contains:

- SDR RF front end appropriate to the operating band;
- stable oscillator, preferably TCXO or better;
- FPGA fabric for correlation, FFT/IFFT, timestamping, pilot extraction, and CSI/CIR output;
- antenna system designed for the airframe;
- Jetson Orin Nano or similar companion compute;
- calibration memory for RF-chain delay, antenna delay, and temperature compensation;
- ROS 2 / MAVLink / DDS integration;
- optional PPS, RTK, or wired sync input for anchor mode.

Network architecture

The network operates in three layers.

1. Data layer. Normal mesh traffic: commands, telemetry, keyframes, descriptors, maps, and state updates.
2. Ranging layer. FTM-like TWR/DS-TWR exchanges embedded into normal packets or scheduled as mini-slots.
3. Estimator layer. Range factors, visual factors, inertial priors, loop closures, and optional anchors are fused into a shared swarm state.

Message	Data payload	Localization payload
Poll	mesh header, optional data	TX marker timestamp ID
Response	ACK, routing, optional data	RX/TX timestamps, CIR quality, SNR
Follow-up	optional payload	DS-TWR correction or clock-skew update
Broadcast state	pose, velocity, covariance	range factors, NLOS flags, link quality
SLAM burst	keyframe, descriptors, loop closure	coincident RF link metrics and range priors

Table 7: The same packet exchange can carry mesh payloads and ranging metadata.

Cost and System Inversion

ISAC Radio is tailored to these conditions:

- high-throughput mesh networking;
- relative localization without GNSS;
- collaborative SLAM or cooperative perception traffic;
- low SWaP;
- rapid deployment;
- and scalable scheduling.

Cost item	Mesh + UWB stack	ISAC stack	System-level implication
High-throughput data radio	Required	Required	No difference; swarm data still needs a mesh
Relative ranging radio	Separate UWB module	Embedded in mesh radio	Potential hardware and integration reduction
Antenna system	Separate or carefully co-located	Shared or integrated design	Lower SWaP and RF co-existence burden
Ranging schedule	Independent polling cycle	Merged with data/MAC schedule	Potential latency and air-time savings
Estimator interface	Range-only, often module-specific	Range, variance, CIR, phase, NLOS, CSI	Richer factors for SLAM and graph optimization
Regulatory path	UWB-specific constraints	Band-dependent non-UWB or licensed/authorized operation	Requires careful RF strategy either way

Table 8: ISAC Radio’s advantage is not cheaper silicon alone. It is the collapse of communication, timing, ranging, and estimator observability into one radio stack.

Performance Targets and Baselines

The central thesis

The radio will become the next navigation sensor. Drone swarms already require a high-throughput wireless mesh to exchange commands, features, keyframes, state updates, and map fragments. Centurion Tech's localization-native radio turns those same packets into calibrated range, channel, timing, and uncertainty measurements for swarm-level navigation.

A recent 5G-NR-aligned OFDM ISAC prototype reports real-time high-definition video transmission, range/angle sensing, and radio SLAM, with 0.3 m range RMSE, 2.3 degree angle RMSE, and 0.25 m radio-SLAM localization error in its real-world tests. [7] This is not the same as a small-drone mesh radio, but it shows that communication-centric OFDM ISAC can produce useful sensing observables while carrying real data.

Centurion Tech's target is not to beat every dedicated ranging technology on day one. The near-term target is system-level value:

- high-throughput mesh communication,
- sub-meter RF range factors in practical LoS conditions,
- decimeter-class targets under calibrated, favorable conditions,
- robust covariance and NLOS detection,
- sparse graph localization,
- and fusion with visual/LiDAR SLAM.

Conclusion

ISAC radios are ambitious because they force more meaning into the same spectrum. A localization-native radio must communicate, range, estimate channel quality, support swarm graph localization, and cooperate with visual/LiDAR SLAM without collapsing under latency or calibration error. That is a difficult engineering problem, but it is aligned with the direction of the industry: AI workloads are becoming network-bound, wireless data traffic is rising, autonomy is moving to the edge, GNSS denial is operationally relevant, and collaborative SLAM already depends on robust communication.

Centurion Tech's position is that drone swarms will not be unlocked by communication or localization alone. They require the fusion of both. The winning radio will not simply move packets. It will expose the physical world: range, timing, channel structure, uncertainty, and graph value.

References

1. NVIDIA, “Spectrum-X Ethernet Networking Platform.” <https://www.nvidia.com/en-us/networking/spectrumx/>
2. Spire, “GNSS interference report: Crimea and the Black Sea region.” <https://spire.com/blog/space-reconnaissance/gnss-interference-report-crimea-and-the-black-sea-region/>
3. U.S. Department of Defense, “DoD Innovation Fact Sheet,” August 2024. <https://media.defense.gov/2024/Aug/07/2003519333/-1/-1/0/DOD-INNOVATION-FACT-SHEET-AUGUST-2024.PDF>
4. Hao Xu, Peize Liu, Xinyi Chen, Shaojie Shen, “D2SLAM: Decentralized and Distributed Collaborative Visual-inertial SLAM System for Aerial Swarm,” arXiv:2211.01538. <https://arxiv.org/abs/2211.01538>
5. Ericsson, “Mobile data traffic outlook,” Ericsson Mobility Report, June 2026. <https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast>
6. Qorvo, “DWM3000 Ultra-Wideband Transceiver Module.” <https://www.qorvo.com/products/p/DWM3000>
7. Z. Wang et al., “ISAC Prototype System for Multi-Domain Cooperative Communication Networks,” arXiv:2410.22956. <https://arxiv.org/abs/2410.22956>
8. eCFR, 47 CFR 15.521, “Technical requirements applicable to all UWB devices.” <https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-15/subpart-F/section-15.521>
9. Android Developers, “Wi-Fi location: ranging with RTT.” <https://developer.android.com/develop/connectivity/wifi/wifi-rtt>
10. IEEE Standards Association, “IEEE 802.11az Standard Improving Wi-Fi Location Accuracy.” <https://standards.ieee.org/beyond-standards/newly-released-ieee-802-11az-standard-improving-wi-fi-location-accuracy-is-set-to-unleash-a-new-wave-of-innovation/>
11. 3GPP, “New Radio positioning support.” <https://www.3gpp.org/technologies/nr-positioning-2>
12. Silicon Labs, “Bluetooth Channel Sounding.” <https://www.silabs.com/wireless/bluetooth/channel-sounding>
13. Silvus Technologies, “StreamCaster MANET Radios.” <https://silvustechnologies.com/products/streamcaster-radios/>
14. Doodle Labs, “Mesh Rider Radios.” <https://doodlelabs.com/products/>
15. Lumen Orbit / Starcloud, “Why we should train AI in space,” White Paper v1.03, September 2024. <https://starcloudinc.github.io/wp.pdf>