BleedIO Tech — NWA Due Diligence Q&A

July 2026

NWA - BleedIO Tech Questions and Answers

Questions received: June 30, 2026 · **Responses prepared:** July 2026

BleedIO Tech, Inc. — Confidential. Prepared for the NWA due-diligence team; not for distribution beyond that team. Figures are as of the dates noted in each answer.

Data-room citations are given as *dataroom.bleedio.com* → *folder* → *document* and resolve in the BleedIO investor data room.

Corrections to question premises (index)

Several questions contain premises that do not match current facts. Each is corrected in context in the relevant answer; the material ones are indexed here for the reviewer's convenience:

Question	Premise in the question	Current fact
Q2	Non-provisional patents "due May 12 and October 16, 2026 — have they been filed?"	Not yet converted — a deliberate, layered IP strategy; conversion decisions managed with counsel (Q2)
Q8	Sub-meter accuracy data "from the Snap-on deployment"	No Snap-on deployment exists yet (agreement in progress). Current locMESH accuracy is 1–1.5 m; sub-meter/centimeter-level arrives with BLE 6.0 Channel Sounding (Q8)
Q10	"Subnet segmentation by floor"	The mesh standard segments by **groups**, not subnets; per-floor/zone grouping is supported today (Q10)
Q11	"3+ months zero failures; field deployment in July"	The locMESH lab evaluation is in a **planned pause**; Lab POC Round 2 reopens Q3 2026, field planned Q3–Q4 2026; a netMESH phase follows later, not yet dated (Q11, Q24)
Q13	"35 installs" at breakeven	≈30 install units in the base case, dominated by ≈25 fire departments (Q13)
Q14	"$900K split equally — $300K each"	**60% GTM ($540K) / 25% Engineering ($225K) / 15% Legal & Ops ($135K)** (Q14)

Question	Premise in the question	Current fact
Q17	"$9M valuation cap; $130K raised"	**$6M post-money cap** (Post-Money SAFE); the ≈$130K was **pre-seed capital (closed)** — no commitments closed yet on the current seed round, which is in active DD with multiple angel groups (Q17)
Q26	"ROFR for Series A at a minimum of $500K"	Right of First **Offer**; $500K is the Designated-Allocation **ceiling** (greater of pro-rata or $500K), one-time at the next ≥$2M preferred round (Q26)
Team	Four cofounders (older materials)	**Five cofounders** — Stan, Mike, Lee (original founders, common stock); Brad and Rodion (adopted cofounders, options) (Q32)

Technology & IP Defensibility

1. What has BleedIO built on top of the open wireless mesh standard, and what can a customer do with netMESH that they could not do by implementing the open spec themselves?

> **Question as asked:** "BLE Mesh is an open Bluetooth SIG standard. Walk us through exactly what BleedIO has built on top of the standard. Specifically, what can a customer do with netMESH that they could not do by implementing the open BLE Mesh spec themselves?"

BleedIO is not selling the open wireless mesh standard itself. We are building **DRAW**, our proprietary firmware and platform ecosystem, on top of an open, non-proprietary mesh foundation. The standard provides basic mesh communication primitives. DRAW turns that foundation into operational products customers can deploy, secure, monitor, automate, and integrate into real industrial, public sector, and defense/autonomy workflows.

Two products, one ecosystem. BleedIO offers two products, both part of the DRAW ecosystem:

- **locMESH** — RTLS (real-time location) and data collection based on BLE. This is our positioning and location-awareness product, delivering personnel/equipment tracking, zone alerts, and environmental/telemetry data.
- **netMESH** — data collection and position awareness based on BLE Mesh. This is our decentralized, self-healing mesh product, adding multi-hop mesh connectivity, store-and-forward offline operation, and scale beyond what point-to-point BLE provides.

A key advantage of the DRAW architecture is that these products share a common firmware and platform base: **locMESH readers can be upgraded over-the-air (OTA) to netMESH by adding a simple hardware dongle.** This means a customer can start with locMESH RTLS/data collection and grow into full mesh capability without replacing their installed reader base — a low-friction upgrade path that protects the customer's hardware investment and lets us expand within an account over time.

Why this matters — DRAW's advantage versus building it yourself. The reviewer's underlying question is really "what stops a customer from just implementing the open spec?" The answer is that the open standard only provides raw mesh communication primitives — the equivalent of an engine with no vehicle around it. Everything that makes a mesh *usable in production* still has to be built on top, and that is precisely what DRAW already delivers as a finished, field-proven product.

A customer choosing to build it themselves would face the full cost and risk of recreating that entire stack — firmware models, routing behavior, provisioning, identity, enhanced security, telemetry, SaaS management, dashboards, alerts, gateway/cloud integration, deployment tooling, and application-specific workflows. That is a multi-year engineering program with no operational track record at the end of it, versus deploying DRAW and getting a working system on day one. The specific capabilities DRAW provides — and that a self-build customer would have to reproduce from scratch — are:

- **Proprietary firmware models** for sensors, equipment, and other mesh node types. The open standard is weighted toward basic lighting/control models and does not provide the full set of sensor, motor, industrial-device, and operational models BleedIO needs for real deployments.
- **Enhanced security**, including BleedIO's AES-256 extension. The underlying standard uses AES-128 security; BleedIO extends the security layer for customers that require stronger protection.
- **Decentralized, self-healing network behavior** with no single point of failure.
- **Store-and-forward operation** so the system can continue collecting and relaying data when cloud or gateway connectivity is intermittent.
- **SaaS management layer** for deployment, monitoring, node state, telemetry, customer operations, and fleet visibility.
- **Edge-AI intelligence layer** that runs at the network edge, not only in the cloud: nodes and gateways make the *initial* decisions locally — rules, filtering, anomaly detection — and push only the **final numbers to the cloud/operator**. The mesh therefore carries distilled results rather than raw streams (which multiplies effective capacity — Question 12), decisions continue to be made when the uplink is down (offline continuity), and edge decision latency stays **under 100 ms**.
- **AI-assisted operating layer** that is independent of the open standard and makes the SaaS easier to operate, more robust, and better able to support configuration, monitoring, anomaly detection, and operational decision-making.
- **Application workflows** layered on top of connectivity, including location awareness, zone alerts, personnel/equipment tracking, environmental sensing, and operational dashboards through locMESH.
- **Domain-specific orchestration and IP**, including the "Network in Advance" patent family and invention disclosures around mesh routing, edge AI, identity, positioning, and orchestration.

So the customer benefit is not that BleedIO invented the underlying radio standard. The benefit is that BleedIO has built the proprietary operating layer required to make that standard useful in the field. A customer could theoretically start from the open spec, but they would still need to spend years building the firmware models, enhanced security layer, AI-assisted SaaS operations, device management, routing/orchestration logic, and vertical-specific workflows that DRAW/netMESH already provides.

Our defensibility is therefore the full DRAW ecosystem: proprietary firmware, expanded device models, AES-256 security extension, AI-assisted SaaS layer, deployment tooling, and operational know-how from real customer environments. The open standard is the substrate; DRAW/netMESH is the product.

A useful analogy: the MANET model. The same relationship exists one layer up the wireless stack, in the tactical world. MANET (mobile ad-hoc networking) is an open, well-established *category* of infrastructure-less self-healing mesh — not any single company's invention. Silvus Technologies did not invent MANET; it built a proprietary, hardened, deployable *implementation* on top of it (its MN-MIMO waveform, resilience to jamming, self-forming/self-healing mesh, and the software that made it usable across drones, ground vehicles, and dismounted units). That implementation-and-usability layer — not ownership of the MANET category — is what Motorola Solutions acquired for $4.4B in 2025.

DRAW is the same model applied to the open mesh standard: BleedIO does not own or sell the underlying standard; we build the proprietary firmware, security, orchestration ("Network in Advance"), SaaS, and application layer that turn it into a product customers can actually deploy. Just as a customer could not replicate Silvus by reading the MANET literature, a customer cannot replicate DRAW by reading the open mesh spec — in both cases the value is the implementation, the operational hardening, and the software ecosystem built on top of an open category. (This comparison is developed further in Question 6.)

Sources. Data room: Technology Thesis (dataroom.bleedio.com → 01_Company_Overview → BleedIO

Tech Technology Thesis.pdf), Product Vision (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Product Vision.pdf), Platform Capabilities (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Platform Capabilities.pdf), netMESH Data Sheet (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech netMESH Data Sheet.pdf), locMESH Data Sheet (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech locMESH Data Sheet.pdf). Comparable (Silvus/MANET → Motorola $4.4B), developed in Question 6.

2. What is the status of the non-provisional patent filings due May 12 and October 16, 2026?

> **Question as asked:** "What is the status of the non-provisional patent filings due May 12 and October 16, 2026? Have they been filed? What specific claims are being protected — particularly around the"Network in Advance" feature and the decentralized orchestration engine?"

Current filing status. BleedIO has two provisional patent applications on file, prosecuted through our patent counsel, Patent 360 LLC (Irvine, CA):

- **Provisional #1 — DN1855 "Network in Advance"** (App. No. 63/804,380), filed May 12, 2025. This application covers the core "Network in Advance" concept and the decentralized orchestration behavior at the heart of the DRAW platform.
- **Provisional #2**, filed October 16, 2025, covering additional subject matter in the same family.

We have **not yet converted these to full non-provisional filings.** The May 12, 2026 and October 16, 2026 conversion dates are decision points we are actively managing with counsel; we are deliberate about how and when we commit to full prosecution rather than treating the deadlines as automatic conversions.

Why this is a deliberate choice, not a gap. We do not believe a small startup's moat should rest on issued patents. Patents are slow, expensive, and — for a company at our stage — provide limited practical protection: a well-capitalized incumbent can litigate a young startup into the ground regardless of who is right on the merits, and a single patent rarely covers the full surface of a real product. We therefore treat patents as **one layer** of a consolidated, multi-layer IP strategy rather than the foundation of it.

Our IP strategy has several reinforcing layers:

- **Provisional patent filings** on the highest-value, most novel inventions — the "Network in Advance" orchestration concept above — to establish priority dates while we validate which claims are worth the cost of full prosecution.
- **A pipeline of invention disclosures** (8 currently in the pipeline: DN1856–DN1860 and DN1875–DN1878) covering mesh routing, edge AI, identity, and positioning. These let us document and date inventions internally and decide later which graduate to provisional or non-provisional filings.
- **Registered copyright** on core code (DN1892, Wi-Fi Mesh Gateway Code, Reg. #1-14883902841, registered March 17, 2025).
- **Trade secrets and operational know-how** — the proprietary DRAW firmware models, AES-256 security extension, AI-assisted SaaS operations layer, and the deployment and tuning knowledge accumulated from real customer environments. Much of what makes DRAW hard to replicate is not patentable subject matter at all; it is the accumulated engineering and field experience embedded in the platform.

The real moat. Our defensibility is the full DRAW ecosystem and the compounding advantage of being deployed in real industrial, public sector, and defense/autonomy environments — not any single patent. The provisional filings protect the specific inventions where a priority date matters most; the broader moat is the integrated product, the proprietary firmware and security work, the operational data, and the customer relationships a competitor cannot copy from a patent document.

We are happy to walk through the specific claim language in the provisional filings under NDA, and to share our counsel's view on the conversion decisions.

Sources. Data room: IP Patent Status Summary (dataroom.bleedio.com → 04_IP_&*Compliance → BleedIO Tech IP Patent Status Summary.pdf), IP Milestone Timeline (dataroom.bleedio.com → 04_IP*&*Compliance → BleedIO Tech IP Milestone Timeline.pdf), IP & Compliance Summary (dataroom.bleedio.com → 04_IP*&_Compliance → BleedIO Tech IP & Compliance Summary.pdf).

3. How is Stan's IP structured within the company?

> **Question as asked:** "Stan holds the IP. How is that IP structured within the company? Is it assigned to BleedIO Tech, Inc., and are there any employment or consulting arrangements that could create IP ownership ambiguity?"

All IP is owned by the company. BleedIO Tech, Inc. (a Delaware C-Corporation, incorporated May 2024) is the owner of the company's intellectual property. The provisional patent applications, invention disclosures, copyrighted code, firmware, and trade secrets described above are held by the company, not by Stan personally.

No ownership ambiguity. BleedIO Tech has standard agreements in place with its cofounders and contractors under which IP rights are assigned to BleedIO Tech, Inc. This means that inventions and work product created for the company — including work originating with Stan, the other cofounders, and outside contractors — are assigned to the company rather than retained individually. There are therefore no employment or consulting arrangements that create IP ownership ambiguity: the assignment agreements are specifically designed to consolidate all rights in the company.

Our corporate and securities counsel is Duane Morris LLP, and our patent prosecution and IP strategy counsel is Patent 360 LLC. We are happy to make the underlying assignment agreements available for review under NDA as part of confirmatory diligence.

Sources. Data room: IP & Compliance Summary (dataroom.bleedio.com → 04_IP_&*Compliance → BleedIO Tech IP & Compliance Summary.pdf), Certificate of Incorporation (dataroom.bleedio.com → 03_Legal*&*Governance → BleedIO Tech Certificate of Incorporation.pdf) (03_Legal*&_Governance also holds bylaws, board consents, and cap table). Assignment agreements available under NDA.

4. If Silicon Labs or Texas Instruments builds a managed mesh orchestration SaaS product, what is BleedIO's specific defense?

> **Question as asked:** "If Silicon Labs / Texas Instruments decides to build a managed BLE Mesh orchestration SaaS product, which their Series 3 platform and TI acquisition suggest they could. What is your specific defense? What would they have to replicate, and how long would it take?"

Silicon Labs and TI are hardware companies, and their expansion is a tailwind for us, not a threat. Silicon Labs and Texas Instruments make chips. Their platforms — including the Silicon Labs Series 3 platform, which combines BLE, mesh, and other protocols on-chip — are hardware platforms. What they do *not* do is build the software layer on top: provisioning, data collection, device management, orchestration, and application workflows. They deliberately delegate that to software companies. That division of labor is exactly where BleedIO lives.

Because DRAW is hardware-agnostic and runs on any compatible hardware, **every new chip and platform these vendors ship expands the set of hardware that can run DRAW.** BleedIO already works with STMicroelectronics and Nordic Semiconductor; bringing Silicon Labs and TI into the mix would be a great expansion of our supported hardware base. TI in particular has historically treated Bluetooth mesh as a secondary SDK feature on its SimpleLink line (a Zephyr-based add-on) rather than a first-class product push, so a deeper TI move into the BLE-mesh market makes the DRAW offer *stronger*, not weaker — it adds another maturing silicon source our platform can run on. Series 3 bringing more capable, lower-cost, multi-protocol chips to market means more devices DRAW can operate on. We *want* as many hardware vendors competing in this market as possible — it commoditizes and improves the substrate our software sits on, and it de-risks us from any single silicon supplier.

What they would have to replicate, and why it is not their business model. For Silicon Labs or TI to compete with BleedIO directly, they would have to build and operate the entire software product

we described in Question 1:

- Proprietary firmware models spanning sensors, equipment, and industrial device types beyond the basic control/lighting models in the open standard.
- An enhanced security layer (our AES-256 extension over the standard's AES-128).
- A decentralized, self-healing orchestration engine with store-and-forward offline operation ("Network in Advance").
- A full SaaS management layer — provisioning, monitoring, telemetry, fleet visibility, dashboards, alerting.
- AI at both ends of the architecture: an **edge-AI inference layer** (initial decisions made locally at the edge, with only distilled results pushed to the cloud/operator) and an AI-assisted cloud operations layer.
- Vertical-specific application workflows (RTLS, zone alerts, personnel/equipment tracking) and the deployment know-how earned in real customer environments.

This is a multi-year software and operations effort in a business that chip vendors have historically chosen *not* to enter, because software services carry different margins, support obligations, and go-to-market motions than silicon. They monetize by selling more chips, and a healthy ecosystem of software partners like BleedIO helps them do exactly that. Our relationship with the hardware layer is complementary, not competitive.

Where the real competitive pressure comes from. We do not anticipate Silicon Labs or TI moving into the software market — that is not their business model. If competition comes from anywhere, it is more likely from other *software* incumbents than from the chip vendors. The most relevant same-model software competitor is **Wirepas** (a Finnish decentralized-mesh software company whose stack runs on third-party silicon such as Silicon Labs SoCs, deployed on 10M+ chipsets); the hyperscalers — principally **Amazon (AWS IoT Core / Sidewalk)** and **Google** — compete on the device-management and cloud layer. We detail this landscape in Question 6. And even against those players, our defense is not one or two patents on BLE mesh; it is the whole DRAW ecosystem: the integrated software stack, proprietary firmware and security work, AI-assisted operations, and the operational data and customer relationships from live deployments. That is the point we keep returning to — the moat is the ecosystem, not any single filing.

Our defense, concretely. (1) We are hardware-agnostic, so we benefit from — rather than compete with — every vendor's platform expansion, and we already run on STM and Nordic silicon. (2) The moat is the integrated DRAW software stack, our security and firmware work, our AI-assisted operations, and the operational data and customer relationships from live deployments, none of which is replicable from a datasheet or a reference design. (3) The natural end-state with the chip vendors is partnership: they want a healthy software ecosystem, and DRAW makes their silicon more useful to industrial, public sector, and defense/autonomy buyers.

Sources. Data room: Defensibility & Moat Analysis (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Defensibility & Moat Analysis.pdf), Competitive Landscape (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Competitive Landscape.pdf), Partner Integration Architecture (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech Partner Integration Architecture.pdf). External: Silicon Labs + Wirepas (10M+ chipsets, Jun 2025); see Question 6 for the full competitive landscape and citations.

5. How does netMESH handle brownfield environments where most installed sensors are not BLE-enabled?

> **Question as asked:** "You mention backward compatibility with older non-mesh BLE devices. How does netMESH handle brownfield environments where the majority of installed sensors are not BLE-enabled at all? Is there a gateway or bridge product on the roadmap?"

We handle brownfield environments through a tiered approach that matches the integration method to what the customer already has installed, plus a gateway for data egress. There are three cases:

- **Sensor is already BLE mesh–capable:** we simply re-provision it into the DRAW network. No new hardware is required — the existing device joins the DRAW mesh directly.

- **Sensor is BLE but not mesh:** we deploy specialized DRAW mesh nodes that read these non-mesh BLE sensors and relay their data into the mesh. The legacy sensor keeps operating as-is; our node acts as the mesh-side reader, so the customer does not have to rip and replace working hardware.
- **Sensor is not BLE at all:** for the fully non-BLE installed base, DRAW runs as an overlay alongside the existing infrastructure rather than replacing it — the practical brownfield model where a customer stands up a DRAW mesh for the data and coverage they need while legacy systems continue running. Where a customer wants that non-BLE data brought together with the mesh, the integration path is through gateway-level data collection rather than direct mesh membership.

Gateway / bridge for egress. Yes — there is a gateway product. DRAW includes a Wi-Fi Mesh Gateway (the subject of our registered copyright, DN1892) that serves as the bridge for data egress from the DRAW mesh to Wi-Fi/cloud infrastructure. This is what moves mesh-collected data out to the customer's systems and our SaaS layer, and it is the integration point for pulling brownfield data together with DRAW telemetry.

Practical positioning. In real deployments we rarely need to convert an entire installed base at once. The mesh-node-reads-legacy-sensor path and the overlay model let us deliver value on day one — standing up DRAW coverage and the sensors the customer actually cares about — without a forced, site-wide hardware replacement. This is closely related to our broader brownfield GTM strategy and the Lufthansa-style zone-overlay model discussed in Questions 18 and 19.

Sources. Data room: netMESH Data Sheet (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech netMESH Data Sheet.pdf), Partner Integration Architecture (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Partner Integration Architecture.pdf), Architecture Diagram (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Architecture Diagram.pdf). Wi-Fi Mesh Gateway: registered copyright DN1892 (see IP & Compliance Summary (dataroom.bleedio.com → 04_IP*&_Compliance → BleedIO Tech IP & Compliance Summary.pdf)).

6. How does BleedIO differentiate or position versus Texas Instruments, Motorola, Semtech, and their emerging edge IoT strategies?

> **Question as asked:** "There are specific companies building edge IoT integration. Specifically comment on your differentiation and/or strategic positioning versus Texas Instruments, Motorola, Semtech, and their emerging strategies?"

Our model is already validated by a recent $4.4B outcome. The clearest way to understand BleedIO's positioning is the Silvus Technologies precedent. Silvus is a software-defined MANET (mobile ad-hoc network) provider — it did not invent the underlying MANET category; it built a resilient, deployable *implementation* on top of it and delivered a solution usable across autonomous systems, military, law enforcement, and enterprise customers. Motorola Solutions acquired Silvus for **$4.4 billion** (announced May 2025, closed August 2025), with Silvus on track for roughly $475M revenue at ≈45% margins plus up to $600M in earnouts.

BleedIO Tech is the same model, one layer down the stack: **we do not invent the wireless standard — we embrace it and deliver a proprietary, deployable solution (DRAW) that different parties can actually use.** Silvus proved that the value in this market accrues to the *solution and implementation provider* that makes a mesh standard operationally useful, resilient, and integratable — not to whoever owns the underlying radio technology. That is exactly BleedIO's thesis for industrial, public sector, and defense/autonomy mesh.

How Silvus played MANET. MANET — mobile ad-hoc networking — is a broad, well-established category of infrastructure-less wireless mesh, not a proprietary Silvus invention. Its defining properties are that nodes form and re-form the network dynamically with no fixed base stations, so the network self-heals as nodes move, drop out, or are jammed. Silvus built a software-defined radio implementation on top of this category and hardened it for the most demanding environments: high-throughput data, video, and voice; resilience to jamming, detection, and interference; and operation across drones, ground vehicles, and dismounted units without any fixed infrastructure. Their differentiation was never "we own MANET" — it was the quality, resilience, and deployability of their *implementation*,

plus the software that made it usable by autonomous-systems makers, militaries, and law enforcement. That implementation-and-usability layer is what Motorola paid $4.4B for.

How BleedIO plays DRAW. BleedIO applies the identical playbook to a different, complementary layer of the wireless stack. The open wireless mesh standard underneath DRAW provides the basic mesh primitives — the equivalent of the raw MANET category. DRAW is our proprietary implementation and extension on top of it: proprietary firmware models beyond the standard's basic control models; an enhanced security layer (AES-256 extending the standard's AES-128); a decentralized, self-healing orchestration engine with store-and-forward offline operation ("Network in Advance"); a full SaaS management and AI-assisted operations layer; and vertical application workflows including precision RTLS (1–1.5 m today, with sub-meter/centimeter-level accuracy on the BLE 6.0 Channel Sounding roadmap — Question 8). Just as Silvus turned the MANET category into a resilient, deployable product for the defense and autonomy world, BleedIO turns the open mesh standard into DRAW — a resilient, deployable product for industrial, public sector, and defense/autonomy customers. The technical envelopes differ (Silvus targets long-range, high-bandwidth tactical RF; DRAW targets high-density, low-power indoor/industrial mesh with precise location), but the *business model is the same*: embrace an open mesh category, build the proprietary implementation and software that makes it operationally useful, and become the layer customers actually deploy.

The Silvus outcome is therefore not just a comparable transaction — it is direct evidence that the market rewards exactly the model BleedIO is executing, and that large strategics acquire the implementation/solution provider rather than trying to build that layer themselves.

Positioning versus each named player:

- **Texas Instruments and Silicon Labs (chip vendors):** As covered in Question 4, these are hardware companies that deliberately delegate the software layer to partners. They are a tailwind — more silicon for DRAW to run on — not competitors. We already run on STM and Nordic silicon and would welcome SL/TI chips into the supported base.
- **Motorola:** Motorola's own $4.4B acquisition of Silvus demonstrates that a large strategic sees implementation-layer mesh companies as high-value assets to *buy*, not commodities to compete away. That validates both our model and a credible strategic-exit path. Motorola operates at the mission-critical radio and public-safety infrastructure tier; BleedIO delivers the deployable DRAW software solution layer, which is complementary to — and potentially acquirable by — players at that tier.
- **Semtech (LoRa):** Semtech's strength is long-range, low-bandwidth connectivity (LoRa/LoRaWAN), a different technical envelope than DRAW's high-density, low-latency, self-healing mesh with precision RTLS and rich telemetry. These serve different use cases; where they overlap, our differentiation is the full DRAW operating layer — provisioning, orchestration, security, AI-assisted operations, and application workflows — not just the radio link.

The broader software competitive landscape. The question names hardware players, but the more relevant competitive set for DRAW is the *software* layer — mesh provisioning, device management, orchestration, and data collection. We track it in two segments:

Industrial IoT mesh software:

- **Wirepas** (Finland) is the closest same-model competitor: a decentralized, self-healing mesh sold as software/firmware (provisioning, device management, positioning, gateway software and APIs) that runs on third-party silicon, including Silicon Labs SoCs. It has real scale — deployed on 10M+ chipsets — and is our most direct industrial comparable. Our differentiation against Wirepas is the full DRAW operating layer for our target verticals: precision RTLS/locMESH, AES-256 security extension, AI-assisted SaaS operations, and vertical-specific workflows for industrial, public sector, and defense/autonomy, rather than a horizontal connectivity stack.
- **Amazon (AWS IoT Core / Sidewalk)** and **Google** compete at the hyperscaler device-management and cloud layer. Notably, Amazon Sidewalk is a *shared, community-sourced* low-power network (BLE plus LoRa/FSK) that piggybacks on existing Echo/Ring devices as gateways — it is explicitly *not* a customer-owned, dedicated, deployable mesh. That is a fundamentally different (and, for industrial and defense buyers, unsuitable) model from a private DRAW deployment the customer owns and controls.

- **Memfault** and similar firmware+SaaS observability/OTA platforms are commercial-model analogues to DRAW's SaaS layer but are not mesh solutions; they reinforce that the firmware-plus-SaaS model is a proven category, not a competitive threat to the mesh itself.

Drone swarm / MANET:

- This segment is dominated by **combined hardware-and-software radio vendors** — Silvus (now Motorola), Persistent Systems (Wave Relay / MPU5), Doodle Labs (Mesh Rider), and Rajant (Kinetic Mesh / InstaMesh) — plus higher-layer autonomy-orchestration software such as Anduril's Lattice. These are defense-grade tactical RF systems at a different technical envelope and price point than DRAW.
- **The economics are the wedge.** MANET is a **high-throughput, high-cost network** — the right tool for high-value assets, such as $1M-class heavy drones, where a $5K–$30K radio per node is a rounding error on the platform cost. But defense procurement is now explicitly asking for **under-$10K drones** — and a swarm of under-$10K drones needs **sub-$1,000 mesh nodes**. **MANET economics do not sustain that**: the radio would cost as much as, or more than, the aircraft it flies on. DRAW's sub-$100-per-node economics (Question 25) are built for exactly this attritable-swarm class, which the MANET vendors structurally cannot reach without abandoning their own cost architecture.
- The strategically important point: this space is served almost entirely by *hardware-coupled* radio stacks priced for high-value platforms. There is a relative gap for a **pure software/firmware-plus-SaaS mesh** that is hardware-agnostic, deployable across commodity silicon, and priced for attritable swarms — which is exactly DRAW's positioning, and the same implementation-provider gap Silvus filled on the tactical side before its $4.4B exit.

The through-line: across all of these, our differentiation is not a radio standard or a patent. It is that BleedIO delivers the complete, proprietary DRAW solution that makes the underlying mesh usable by real customers — the same implementation-provider model the market just rewarded with a $4.4B outcome.

Sources. Data room: Competitive Landscape (dataroom.bleedio.com → 06_Market_&*Traction → BleedIO Tech Competitive Landscape.pdf), Comparable Exits & Acquisition Landscape (dataroom.bleedio.com → 06_Market*&*Traction → BleedIO Tech Comparable Exits & Acquisition Landscape.pdf), Defensibility & Moat Analysis (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech Defensibility & Moat Analysis.pdf). External — Silvus/Motorola $4.4B: Motorola Solutions press release, SEC 8-K; Wirepas scale: Silicon Labs + Wirepas 10M+ chipsets, Wirepas Mesh; Amazon Sidewalk model: AWS IoT Core for Sidewalk; Silvus swarm/defense focus: Silvus unmanned systems.

Product Performance

7. What is the full infrastructure requirement per deployment, and can BleedIO share a TCO model for a 10-site enterprise rollout?

> **Question as asked:** "The solution is dependent on BLE gateways for data egress, which can become costly to deploy and maintain at scale. What is the full infrastructure requirement per deployment? Can you share a total-cost-of-ownership model for a 10-site enterprise rollout?"

Yes. Below is the infrastructure requirement per site and an illustrative 10-site TCO model. The unit economics are grounded in our actual figures; the site-count assumptions are typical mid-to-large enterprise sites and would be tuned to the specific rollout.

Infrastructure requirement per site. A typical enterprise site is served by:

- **100–200 DRAW nodes/readers** — the mesh endpoints that provide coverage, positioning, and data collection. Node count scales with facility size, coverage density, and how many assets/zones the customer wants tracked.
- **2–4 gateways** — the Wi-Fi Mesh Gateway units that egress mesh data to the customer's network and our SaaS layer. Gateway count scales with site footprint and redundancy requirements.

- **DRAW SaaS subscription** — the management, monitoring, telemetry, and application layer (annual, recurring).

A defining cost advantage is that our **per-node cost is under $100** — materially lower than competing RTLS and industrial-mesh hardware. Because the mesh is self-healing and multi-hop, we also avoid the dense fixed-anchor/gateway infrastructure that drives cost in many competing RTLS systems: nodes relay for each other rather than each requiring a direct line to fixed infrastructure.

Per-site cost structure (illustrative):

Component	Basis	Per-site (illustrative)
DRAW nodes/readers	100–200 nodes × <$100/node	≈$10K–$20K (one-time hardware)
Gateways	2–4 units	Modest one-time hardware add
Deployment & commissioning	Install, provisioning, calibration	Included in a per-site deployment line
DRAW SaaS	Annual subscription per site	**≈$48K–$60K/year (recurring)**

The node/gateway hardware plus deployment lands a fully commissioned site in the **≈$100K–$150K one-time** range (consistent with our per-deployment economics on comparable engagements), with the DRAW SaaS subscription (≈$48K–$60K/site/year) as the recurring layer on top.

Illustrative 10-site enterprise TCO (3-year horizon):

Line item	Basis	10-site total
Site deployment (hardware + install)	≈$100K–$150K × 10	≈$1.0M–$1.5M (one-time, Year 1)
DRAW SaaS	≈$48K–$60K/site/yr × 10 sites × 3 yrs	≈$1.44M–$1.8M (recurring)
3-year TCO	Deployment + 3 yrs SaaS	**≈$2.4M–$3.3M**

This puts the **per-site 3-year cost of ownership in the ≈$240K–$330K range**, which is consistent with the 3-year ACV band we model per deployment. A phased rollout (a subset of sites in Year 1, expanding as the customer validates) further smooths the capital profile — and the OTA-upgrade path described in Question 1 lets a customer start with locMESH RTLS at lower initial node cost and grow into full netMESH mesh without replacing the installed reader base.

Important caveats. These figures are an illustrative model to show structure and order of magnitude, built on real unit economics (per-node <$100; per-deployment ≈$100K–$150K; SaaS ≈$48K–$60K/site/yr). Actual numbers depend on facility size, node density, coverage and redundancy requirements, integration scope, and contract structure (e.g. a negotiated multi-site enterprise agreement rather than simple 10× site pricing). We are happy to build a bottoms-up TCO tailored to a specific 10-site rollout — with the full bill of materials, deployment plan, and SaaS terms — under NDA.

A note on why these figures differ from the per-zone pricing in Question 19. Pricing varies by deployment class, and the numbers here reflect a **higher-cost class of site** (the per-site economics above are drawn from our maritime deployment model). Maritime is materially more expensive per site than a standard industrial zone: it requires different, ruggedized marine-grade readers and gateways, different (marine) sensors, and more onboard edge/offline capability because reliable cloud connectivity is often unavailable at sea. A standard industrial brownfield zone follows the lighter per-zone/per-device economics in Question 19 (roughly $25K–$60K+ per-zone deployment plus $3–$5/device/month SaaS). Both are the same product and pricing philosophy — one-time deployment fee plus recurring per-device/per-site SaaS — scaled to the hardware, environment, and connectivity realities of the vertical. We would use the industrial per-zone economics for a typical 10-site industrial rollout, and the maritime economics for a fleet deployment.

Sources. Data room: Financial Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf), Financial Assumptions & Definitions (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Assumptions & Definitions.pdf), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market_&_Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf). Unit economics (per-node <$100; per-deployment $100K–$150K; SaaS $48K–$60K/site/yr; 3-yr ACV $150K–$330K) per the financial model. Per-site node/gateway counts per BleedIO engineering.

8. Can BleedIO share raw positioning accuracy data for locMESH, including deployment conditions and limits of the sub-meter claim?

> **Question as asked:** "BleedIO claims sub-meter RTLS positioning accuracy for locMESH. The VOS COO has not observed that level of performance. Can you share raw positioning accuracy data from the Snap-on deployment — including device density, node spacing, and environment conditions? What are the limits of that accuracy claim?"

First, a clarification on the premise. The question references a "Snap-on deployment" as the source of accuracy data. To be precise: Snap-on is a channel/OEM agreement in progress, not a live paying deployment, so there is no Snap-on field dataset to share. We want to be careful not to imply performance data from an engagement that is not yet deployed. We are glad to share our actual positioning test data and to be transparent about the real limits of the accuracy claim — including the VOS COO's observation, which is consistent with how we characterize the technology below.

Current locMESH accuracy. locMESH delivers positioning accuracy in the **industry-standard 1–1.5 meter range** for BLE-based RTLS. Accuracy is not a single fixed number; it depends primarily on two environmental factors:

- **RF noise / interference** — the electromagnetic environment of the site (metal structures, machinery, competing 2.4 GHz traffic, multipath reflections).
- **Geometry** — node/anchor spacing, density, and the placement geometry relative to the tracked asset.

We have made significant improvements in our positioning math over the last several years, which has tightened accuracy and robustness within that band, but we do not claim performance beyond what the underlying BLE physics currently supports. (This math-model work is ongoing, including a signed research collaboration with Villanova University focused specifically on improving mesh-positioning models — see BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf).) If a party has heard a "sub-meter" figure presented as a guaranteed spec, that is more aggressive than we would state today: **1–1.5 m under typical conditions is the honest characterization**, with the low end achievable in favorable RF/geometry and the high end in noisy or poorly-geometried environments. This is consistent with what the VOS COO observed, and we would rather set that expectation accurately than oversell it.

Where sub-meter / centimeter accuracy is genuinely headed. The BLE ecosystem itself is closing this gap. **Bluetooth 6.0 Channel Sounding** — a Bluetooth SIG standard feature using phase-based ranging (PBR) and round-trip time (RTT) — is designed to deliver centimeter-level distance measurement on standard 2.4 GHz BLE radios, with early implementations already demonstrating roughly ±10–20 cm at short range (Bluetooth SIG). Because DRAW is built on the open BLE foundation and is hardware-agnostic, **as BLE 6.0 Channel Sounding silicon reaches the market, locMESH is positioned to move from meter-level to centimeter-level accuracy** without a change to our platform architecture. That is our roadmap for high-precision positioning, and it is grounded in a ratified industry standard rather than a proprietary claim we are asking investors to take on faith.

Important context: RTLS is not the core of netMESH. It is worth separating the two products here (see Question 1). Precise positioning is central to **locMESH**. For **netMESH**, RTLS plays a *minor* role — netMESH is primarily a **data-collection** product. It can provide RTLS, and we are developing positioning awareness for drone-swarm coordination (position awareness among nodes in a swarm), but the attributes that matter most to netMESH customers are:

- **Low latency** — timely delivery of collected data.

- **Resiliency** — reliable operation in contested, noisy, or degraded RF environments.
- **Network portability** — the network moves and re-forms with the deployment rather than depending on fixed infrastructure.

This matters competitively: the current incumbents on the market tend to force a trade-off — they deliver **either low latency or high resiliency in contested environments, but not both.** netMESH's design target is to hold low latency *and* resiliency together, which is the differentiated value for data collection, and a distinct question from RTLS precision. So the sub-meter-accuracy concern, while a fair question for locMESH, does not bear on the core value proposition of netMESH.

What we can share. We are happy to provide our positioning test data — accuracy distributions against node spacing, density, and RF conditions — under NDA, along with the environmental assumptions behind each figure, so the accuracy claim can be evaluated on real numbers rather than a headline spec.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf), locMESH Data Sheet (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech locMESH Data Sheet.pdf), locMESH Precision Data Platform (dataroom.bleedio.com → 05_Product&_Technology → BleedIO Tech locMESH Precision Data Platform.pdf). Industry standard (BLE 6.0 Channel Sounding, centimeter-level ranging): Bluetooth SIG — Channel Sounding, Minew — Bluetooth 6.0 Channel Sounding.

9. How does netMESH handle rapid topology changes in dynamic environments, and what is the re-convergence time?

> **Question as asked:** "VOS raised concerns about performance in highly dynamic environments — construction sites, for example — where physical layouts change constantly. How does netMESH handle rapid topology changes? What is the re-convergence time after a significant physical change to the environment?"

This is the problem our platform was built to solve — not an edge case. Dynamic, portable environments are the *core design target* of the DRAW platform, not a limitation we work around. Legacy wireless (Wi-Fi, LoRa, Zigbee, cellular) anchors connectivity to fixed infrastructure — a hub, gateway, or cloud uplink — so it breaks exactly where the layout keeps changing: fire scenes, construction sites, disaster response, live events, contested airspace. DRAW's entire premise is the opposite: every device relays for every other device, so **the network travels with the operation, not the building.** A changing physical environment is the condition we designed for.

locMESH — proven in the most dynamic environment there is. locMESH is specifically designed for dynamic and portable deployments. It stands up in minutes with no RF surveying, no fixed anchors, and no pre-planning, and delivers real RTLS on the spot. The clearest proof is that BleedIO is already selling locMESH to fire departments (signed, contract onboarding paid — including hardware) for exactly this: real-time location on a live fire scene — a portable environment where the layout is unknown, changing, and hostile, and where GPS fails indoors (BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf); locMESH Data Sheet (dataroom.bleedio.com → 05_Product&_Technology → BleedIO Tech locMESH Data Sheet.pdf)). Firefighter location and accountability inside a burning structure is a problem that had no reliable wireless solution before; locMESH solves it today. If the technology performs in a fire scene, ordinary industrial dynamism like a construction site is well within its envelope.

netMESH — portable by default, settles as the environment changes. netMESH takes this further. It is portable by default and requires no site-specific preparation: nodes form the mesh on power-up and the network continuously re-settles as the environment changes — as nodes move, are added or removed, or as RF paths shift. The self-healing mechanisms that do this are (netMESH Data Sheet (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech netMESH Data Sheet.pdf)):

- Automatic node exclusion and re-inclusion as nodes drop out or rejoin
- Network partition detection and reseeding
- Dynamic path repair and hop-count optimization
- Congestion detection with adaptive relay back-off

- Continuous path-quality scoring for route selection

Re-convergence time. Our design and lab validation target is **re-convergence in under 3 seconds under 10% simultaneous node loss, with under 1% packet loss under stress** (netMESH Data Sheet (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech netMESH Data Sheet.pdf)). *For context on why that threshold matters: refinery SCADA control loops time out and trigger safe-shutdown if a failure isn't rerouted within about 3 seconds, so <3 s is the bar for mission-critical industrial use (1000-Node Stress Test Analysis (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech 1000-Node Stress Test Analysis.pdf)). Silicon Labs' published Bluetooth Mesh 1.1 benchmark — a 256-node network in a live commercial office (active Wi-Fi and BLE present) — measured >99% reliability and <200 ms latency across up to 6 hops for ≤16-byte payloads (netMESH Data Sheet (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech netMESH Data Sheet.pdf)).

Honest scoping. We want to be precise about maturity. locMESH is live and proven in the field, including the portable fire-scene use case. netMESH is in development (patent pending), and the <3 s reconvergence figure is **lab-validated and field-consistent, but not yet proven at large scale in production.** We are candid that at very high node counts, in heterogeneous hardware and RF-hostile conditions, naive mesh implementations can see reconvergence stretch to tens of seconds — which is precisely why we are running a **1,000+ node stress test (targeted Q3–Q4 2026)** whose explicit pass criteria include holding <3 s reconvergence under 10% node loss, 100% of the time, across multiple failure scenarios (1000-Node Stress Test Analysis (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech 1000-Node Stress Test Analysis.pdf)). Rather than quote a scale figure we have not yet earned, we present the scoping as it stands. The stress-test white paper and the netMESH data sheet with the full mechanism set are available in the data room.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf), locMESH Data Sheet (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech locMESH Data Sheet.pdf), netMESH Data Sheet (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech netMESH Data Sheet.pdf), 1000-Node Stress Test Analysis (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech 1000-Node Stress Test Analysis.pdf). Industry context on multi-hop BLE mesh self-healing: Bluetooth SIG — The Case for Bluetooth Mesh (dataroom.bleedio.com → 05_Product*&_Technology → Bluetooth SIG - The Case for Bluetooth Mesh.pdf). External — Silicon Labs' published Bluetooth Mesh 1.1 benchmark (>99% reliability, <200 ms to 6 hops, 256-node office network): Silicon Labs AN1424.

10. How does netMESH address multi-level structures where signals cross floor boundaries?

> **Question as asked:** "How does netMESH address multi-level structures, mezzanines, multi-floor warehouses, hospitals, where BLE signals do not respect floor boundaries? Is subnet segmentation by floor currently supported, or is that a 2026 roadmap item?"

Multi-level structures are supported today — through two distinct capabilities that are often conflated, so we separate them explicitly:

1. locMESH provides both 2D and 3D mapping.

- **2D mapping** resolves position on a plane (X/Y) — the standard configuration for a floor.
- **3D mapping** resolves elevation as well, by including the vertical dimension in the trilateration. The practical trade-off is reader count: **3D mapping requires almost twice as many readers as 2D**, because the reader geometry must span different elevations for the trilateration to resolve height. 3D is the right choice for tall single-volume spaces — warehouses with high racking, mezzanines, atriums — where "how high" matters as much as "where."

2. Floor segmentation is a different capability — and it works with either mapping mode. Determining *which floor* a person or asset is on does not require 3D mapping: floor segmentation is achieved with **both 2D and 3D mapping**, with one hard deployment requirement — **in a multi-floor structure, readers are required on each floor.** Each floor's readers resolve position on that floor (2D per floor, or 3D where the use case needs elevation within a floor), and the per-floor reader

placement is what disambiguates which floor a reading belongs to — rather than treating the building as a single flat plane where signals bleeding across floor boundaries would confuse the estimate.

In other words, the fact that BLE signals do not respect floor boundaries is handled by **reader placement and geometry**, which is exactly the lever RTLS systems use for vertical resolution. Signals crossing floors are not a failure mode for the mesh — for the *connectivity* layer, cross-floor propagation is actually helpful: it provides redundant vertical relay paths, which strengthens the self-healing mesh. For the *positioning* layer, the per-floor (and, in 3D, multi-elevation) reader design is what disambiguates which floor/height a reading belongs to.

On segmentation by floor. The question asks about "subnet segmentation by floor." One clarification on terminology: the underlying open mesh standard does not partition addressing into subnets the way IP networks do — it uses **groups** (group addresses with a publish/subscribe model) as the native mechanism for logically organizing a network, including by floor or zone. DRAW uses this today: we segment multi-floor and multi-zone deployments through a combination of **group** configuration and per-floor/zone anchor design, so a large multi-floor building is organized into logical partitions (per-floor or per-zone groups) rather than one undifferentiated mesh. That gives the customer floor/zone-level addressing, monitoring, and control without requiring separate physical networks. We are glad to walk through the specific group/zone segmentation approach for a given building topology as part of technical diligence.

Accuracy envelope. As with our other positioning answers (see Question 8), locMESH accuracy remains in the industry-standard 1–1.5 m range and is dependent on RF noise and anchor geometry; multi-floor and single-floor-3D positioning follow the same accuracy envelope, driven by how densely and how well-distributed the anchors are placed across floors and elevations. The capability to resolve floor and height is present today; the precision within that resolution is governed by the same physics and anchor-geometry factors described in Question 8, and improves as BLE 6.0 Channel Sounding hardware reaches the market.

Sources. Data room: locMESH Data Sheet (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech locMESH Data Sheet.pdf), locMESH Precision Data Platform (dataroom.bleedio.com → 05_Product*&*Technology → BleedIO Tech locMESH Precision Data Platform.pdf), Architecture Diagram (dataroom.bleedio.com → 05_Product*&*Technology → BleedIO Tech Architecture Diagram.pdf). Group/zone segmentation mechanism (industry standard): Bluetooth SIG — The Case for Bluetooth Mesh (dataroom.bleedio.com → 05_Product*&_Technology → Bluetooth SIG - The Case for Bluetooth Mesh.pdf). Positioning accuracy envelope per Question 8.

11. What specific Chevron performance metrics are being tracked, and what are the next field-testing conditions?

> **Question as asked:** "Chevron has been in lab operation for 3+ months with zero failures. What are the specific performance metrics being tracked: uptime %, packet loss rate, latency? And what are the conditions of the next phase (field deployment in July)?"

A clarification on the current state, stated plainly. Chevron remains a **paying customer** and the relationship is active. The engagement is a **locMESH** evaluation — real-time location (RTLS) plus collection from BLE legacy sensors — in a Chevron **lab/office environment**, not a live refinery. That lab stage is in a **planned pause**: by mutual agreement, Chevron asked us to come back with improvements in **positioning accuracy and UI/UX** before proceeding, and we are working on exactly those. This is a planned, scoped pause with a defined path back — not a stalled or lost account. We would rather report that candidly than overstate a "months of flawless uptime" narrative.

To be precise about two points the question presumes:

- The current evaluation is **locMESH** (RTLS + BLE legacy-sensor data collection). A **netMESH** evaluation with Chevron is planned as a later phase — no date is committed yet.
- The **refinery field-testing timeline has moved.** It is not a July / Q2-2026 live-refinery deployment. The working plan is to **reopen lab testing in Q3 2026 (Lab POC Round 2)** and, gated on those results, run **field testing in Q3–Q4 2026.** We treat these as working targets rather

than hard commitments, because under-promising beats a failed deployment in a safety-critical environment.

The next step is a second lab round — "Lab POC Round 2" (Q3 2026) — before any field testing. We are explicit about the sequence: we reactivate in the lab first, not in the field. Lab POC Round 2 is the customer-run **locMESH** lab phase that resumes the evaluation once we deliver the accuracy and UI/UX improvements Chevron requested — those improvements, plus operational stability, are its pass criteria. Field testing follows in **Q3–Q4 2026**, gated on the lab results.

On the specific metrics the question asks about — uptime %, packet loss rate, latency. Those are the headline pass metrics for the **netMESH** evaluation phase, which is planned with Chevron as the step after the locMESH sequence and is **not yet dated**:

- **Uptime %** — sustained availability of the netMESH network over the test window.
- **Packet loss rate** — data delivery reliability across the mesh, including under the RF conditions of an industrial environment.
- **Latency** — end-to-end delivery time for collected data and location updates.

These are the same attributes we hold as netMESH's core value (low latency, resiliency, reliable data collection — see Questions 8 and 9), so that phase will be the direct, customer-run validation of those claims when it is scheduled.

Why this sequence. A refinery is a safety-critical, RF-hostile environment. Moving to the field before the lab results and usability improvements are proven would be the wrong risk profile for both Chevron and us. The sequence is therefore: (1) deliver the accuracy and UI/UX improvements Chevron asked for → (2) Lab POC Round 2 (locMESH, Q3 2026) → (3) field testing (Q3–Q4 2026, gated on lab results) → (4) the netMESH evaluation phase (uptime, packet loss, latency) and commercialization thereafter, on a timeline Chevron and BleedIO set together. We are happy to share the test plan and success criteria under NDA.

This is precisely what the investment is for. The Chevron planned pause is not a stalled customer — it is a snapshot of the real work: finishing the commercial product — the locMESH accuracy and UI/UX now, and bringing netMESH to a hardened commercial implementation next. That research and software-development effort is genuinely hard, and it is not easy for *any* company — which is exactly why **there is no meaningful commercial implementation of a decentralized, self-healing, low-latency data-collection mesh on the market today.** The problem is complex and labor-intensive to solve well; that difficulty is the barrier to entry and the reason the opportunity exists. We are raising this round specifically to fund the engineering that takes netMESH the last distance — the accuracy, UI/UX, and scale-hardening that turn a promising lab result (and an engaged Fortune 500 customer like Chevron) into a deployable product. In other words, the gap the question is probing is the gap the capital closes, and it is defensible precisely because it is difficult.

Sources. Data room: Customer & Pipeline Status Matrix (dataroom.bleedio.com → 06_Market_&*Traction → BleedIO Tech Customer & Pipeline Status Matrix.pdf), Customer Traction & Pilot Portfolio (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Customer Traction & Pilot Portfolio.pdf), netMESH Data Sheet (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech netMESH Data Sheet.pdf). Positioning-accuracy context per Question 8; netMESH performance targets per Question 9. Test plan and success criteria for the next lab phase available under NDA.

12. What is the largest deployed or stress-tested network, and what are the realistic bandwidth and scale limits?

> **Question as asked:** "You claim support for up to 32,000 devices per network. What is the largest network you have actually deployed or stress-tested? At what node count have you observed performance degradation? Further, BLE networks have a natural bandwidth limitation. How should we view maximum scale deployments from a bandwidth perspective? What are the realistic limits in number of devices and data transmission per sensor. What are the sensor types / use cases / markets beyond reasonable scope of this solution."

What the 32,000 figure is — and is not. The ≈32,000 number (precisely 32,767 nodes per network) is the **theoretical maximum a single DRAW network can address**, not a size we have deployed

or stress-tested. We want to be explicit about that distinction. It is an architectural ceiling, and we would never present it as a proven deployment figure.

Why the theoretical ceiling is not the operationally interesting number. In practice, nobody runs a network of that size as a single flat mesh, and it would be poor engineering to try. Real-world scale is not about how many nodes you can flood into one undifferentiated network — it is about **routing efficiency and subdividing the network into logical groups/zones** (using the group mechanism described in Question 10). A well-designed large deployment is a set of coordinated group/zone partitions with efficient routing between them, not one giant flat broadcast domain. So the meaningful engineering questions are "how well does routing scale?" and "how cleanly can we partition into groups?" — not "can we hit 32,767 in one flat network?"

Where we stand on hard data. We **do not have large-scale stress-test data today.** We are not going to quote a node count at which we have observed degradation, because we have not yet run that test at scale and we will not manufacture a number. What we can say is grounded in the mesh protocol testing we do have (see Question 9): Silicon Labs' published Bluetooth Mesh 1.1 benchmark (a 256-node office network) showed >99% reliability and <200 ms latency across up to 6 hops for small payloads. Beyond that, large-scale behavior is exactly what we intend to measure, not assert.

The planned test. We plan to run a **1,000+ device stress test in Q3–Q4 2026.** Its purpose is precisely to answer this question with data: to characterize routing performance, group/zone subdivision, reconvergence, packet loss, and latency at scale, and to identify the node counts and conditions where performance degrades — so we can state real limits rather than theoretical ones. The pass criteria and methodology are described in our stress-test white paper.

On BLE bandwidth and realistic limits. The question rightly notes that BLE has inherent bandwidth constraints, and it does — BLE mesh is optimized for many small, frequent messages (sensor telemetry, location updates, control), not high-throughput streaming. That shapes where DRAW fits and where it does not:

- **In scope / well-suited:** sensor telemetry, RTLS/location, status and control messaging, alerts, and other low-to-moderate data-rate operational data — where the value is reliability, low latency, resiliency, and device density rather than raw throughput.
- **Out of scope / not the right tool:** high-bandwidth applications such as live video streaming, bulk file transfer, or high-rate continuous data per node. For those, a mesh built on the BLE physical layer is the wrong choice, and we would not position DRAW there.

Edge AI multiplies effective capacity. A structural part of how DRAW lives within BLE's bandwidth envelope is its **edge-AI layer**: nodes and gateways make the initial decisions at the edge — filtering, thresholding, anomaly detection — so the mesh carries **final numbers, not raw streams**, and the cloud/operator receives distilled results. The right bandwidth question is therefore not "how much raw sensor data can BLE carry?" but "how many *decisions and results* can it carry?" — a far smaller payload per sensor, which is exactly how a low-throughput mesh serves dense sensor populations (see the edge-AI layer in Question 1).

Realistic scale is therefore a function of **per-node data rate × node count × hop depth**, partitioned across groups — with edge processing shrinking the per-node data rate — not a single headline device count. A deployment of many low-rate sensors partitioned into well-routed groups scales very differently from one trying to push high data volumes through every node. We are happy to work through the scale envelope for a specific use case, and the Q3–Q4 2026 stress test is designed to replace these principled bounds with measured ones.

Sources. Data room: netMESH Data Sheet (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech netMESH Data Sheet.pdf) (32,767-node architectural ceiling; mesh protocol test results), 1000-Node Stress Test Analysis (dataroom.bleedio.com → 05_Product*&*Technology → BleedIO Tech 1000-Node Stress Test Analysis.pdf) (planned Q3–Q4 2026 test, methodology and pass criteria). Group/zone subdivision per Question 10; mesh performance context per Question 9. Industry context: Bluetooth SIG — The Case for Bluetooth Mesh (dataroom.bleedio.com → 05_Product*&_Technology → Bluetooth SIG - The Case for Bluetooth Mesh.pdf).

Financials & Use of Proceeds

13. What assumptions support the 2027 breakeven projection at $1.2M ARR and 35 installs?

> **Question as asked:** "Your breakeven projection is 2027 at $1.2M ARR and 35 installs. Walk us through the specific assumptions: how many contracts need to close, at what average contract value, and what is your current probability-weighted pipeline against those targets?"

First, reconciling two numbers in the data room. Our model contains two distinct 2027 scenarios, and it is worth separating them so the figures are not confused:

- **Base / breakeven case — ≈$1.2M ARR (2027).** This is the conservative case the question references: roughly $1.2M ARR on ≈$2.1M revenue, reaching approximately cash-flow breakeven (≈$57K EBITDA-positive) at a 1.1× burn multiple. This is the number we stand behind as the near-term target.
- **Optimistic / maritime-scaled case — $7.1M–$8.9M ARR (2027).** A separate, upside scenario in which the Lufthansa maritime channel scales to 15–20 ships in 2027. This is the aggressive case; it is not the breakeven case, and we do not treat it as the base plan.

Everything below concerns the **$1.2M base/breakeven case.**

On "35 installs" — mapping the number to our actual mix. The $1.2M is not 35 identical installs; it is an ARR blend across revenue streams, and the install count is a *mix* dominated by fire departments. Mapped out:

Revenue stream	2027 base-case contribution	Approx. installs / customers	ACV basis
Fire departments (Public Sector)	≈$750K	≈25 paying (of ≈50 targeted; ≈50% conversion)	≈$10K–$30K each
Chevron (renewal)	≈$150K–$250K	1 renewed site (base case)	$150K–$250K/site
Snap-on OEM channel	≈$50K–$150K	1 OEM channel (royalty/embed, not a site install)	embed/royalty
Lufthansa maritime + Carbon Reform + other industrial	≈$100K–$200K	1–2 industrial + first maritime deploy	$25K–$100K
Total (base case)	**≈$1.2M–$1.3M**	**≈30 install units**	—

So "≈35 installs" is a reasonable order-of-magnitude proxy, but the honest picture is **≈25 fire-department installs plus a handful of industrial/Chevron sites plus an OEM channel** — the fire departments are the *volume* driver, and the enterprise/industrial sites are the *value* driver. It is not 35 uniform enterprise deployments.

Contracts that need to close, and the probability-weighted view. The base case depends on a specific, trackable set of closes:

- **≈25 paying fire departments** (from ≈48–50 in pipeline at ≈50% conversion) — the primary breakeven driver.
- **Chevron renewal** of the current site (≈$150K–$250K ACV).
- **Snap-on OEM contract** signed and first units shipping.
- **First Lufthansa maritime deployment** and 1–2 additional industrial pilots converting.

On a probability-weighted basis, our **12-month-forward weighted pipeline is ≈$839K** (probability-adjusted; versus ≈$48K current ARR — a high pipeline/ARR ratio that we flag as carrying early-stage uncertainty), against a total unweighted pipeline of ≈$1.56M. (For clarity across our documents: the ≈$1.2M "weighted pipeline ACV" in the Financial Projections is the *total* weighted pipeline including

longer-dated maritime/OEM scale-up; the **≈$839K is the 12-month-forward slice**, and it is the number we use here.) In other words, roughly 70% of the breakeven target is covered by *already-identified, probability-weighted, 12-month-forward* pipeline — the remainder comes from converting the longer-dated pipeline we can already name, not from net-new demand we have not yet sourced.

Crucially, breakeven does not depend on Chevron expansion. This matters given Chevron's current status (the locMESH lab evaluation is in a planned pause pending accuracy/UI-UX improvements, with Lab POC Round 2 in Q3 2026 as the next step — see Question 11). The **$1.2M base case is reachable on fire departments + the Chevron *renewal* + Snap-on**, i.e. without counting any Chevron multi-site expansion revenue. Chevron adding 2–3 sites (≈$400K–$600K at run-rate) is real upside on top of the base case, not a load-bearing assumption underneath it. We deliberately do not lean the breakeven case on the paused expansion.

Both scenarios, stated openly. We would rather show the range than a single point:

Scenario	2027 ARR	Path	Series A timing
Base case	≈$1.2M	≈25 FDs + Chevron renewal + Snap-on + first maritime	≈2H 2027
Downside (Chevron stalls entirely and/or FD conversion <25%)	≈$500K–$800K	fire departments carry a reduced base; enterprise slips	slips to Q4 2027–Q1 2028, valuation from lower base
Upside (Chevron multi-site expansion + maritime scales)	up to $7.1M–$8.9M	maritime channel 15–20 ships + Chevron 2–3 sites	earlier / larger round

The plain read: fire departments are the engine of the base case (short 1–2 month sales cycle, repeatable, IAFF-accelerated — see Question 20), Chevron renewal and Snap-on are supporting contributors, and Chevron expansion + maritime are the upside. The single biggest risk to $1.2M is fire-department conversion rate, which is why our 2026 execution milestones are built around proving that conversion (e.g. first 5 paid FDs by Q3 2026, 25 pilots in field by Q4 2026).

Sources. Data room: Financial Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf) (base vs. optimistic cases, burn multiple, EBITDA), Financial Assumptions & Definitions (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Assumptions & Definitions.pdf), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market_&*Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf), Customer & Pipeline Status Matrix (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech Customer & Pipeline Status Matrix.pdf) (weighted pipeline ACV). Chevron status per Question 11; vertical sequencing per Question 20.

14. What milestones does the $900K seed round unlock by allocation, and which allocation could be cut by 30% if necessary?

> **Question as asked:** "The seed round is $900K split equally across engineering, sales, and G&A ($300K each). What specific milestones does each bucket unlock? If you had to cut one allocation by 30%, which would it be and why?"

One correction on the premise first. The question describes the round as "split equally across engineering, sales, and G&A ($300K each)." That is not our allocation. The **$900K seed is split 60% / 25% / 15%**, weighted heavily toward go-to-market, not divided equally:

Allocation	Amount	What it unlocks
Go-to-market — 60%	**$540K**	Convert pilots to paying contracts; scale the fire-department pipeline (IAFF / Ignite Conference channel); hire the first solutions engineer. This is the bucket that drives the $1.2M breakeven case (Question 13).
Product & Engineering — 25%	**$225K**	Hire 2–3 engineers; advance netMESH toward commercial readiness (the accuracy/UI-UX and scale-hardening work behind Questions 11–12); edge-AI development.
Legal & Operations — 15%	**$135K**	Non-provisional patent filing decisions (the May and October 2026 deadlines, Question 2); compliance; operating infrastructure.

Why the weighting is toward GTM. At the seed stage our binding constraint is not building more technology — locMESH is live and netMESH has a working evaluation — it is *converting identified demand into paying contracts.* The $1.2M breakeven case is a conversion story (≈25 fire departments plus enterprise/OEM), so the capital is weighted to the function that closes and deploys those contracts. The round is milestone-tied, not just runway: at post-hire burn of ≈$25K–$35K/month it provides ≈18–24 months to Series A.

Milestones each bucket unlocks:

- **GTM ($540K):** first ≈25 paying fire departments (from ≈50 in pipeline); Chevron renewal converted and supported; Snap-on OEM moved to shipping; first Lufthansa maritime deployment; the repeatable deployment playbook and integrator-onboarding motion that make the model scale without proportional headcount.
- **Engineering ($225K):** the accuracy and UI/UX improvements that unblock the Chevron evaluation and enable Lab POC Round 2 (Question 11); netMESH scale-hardening validated by the 1,000+ node stress test (Questions 9, 12); edge-AI features.
- **Legal & Ops ($135K):** patent conversion decisions on schedule; compliance posture matured; clean operating infrastructure for enterprise procurement.

If forced to cut one allocation by 30%, it would be Engineering — and here is the reasoning. A 30% cut is roughly $68K off the $225K engineering bucket. We would cut there, not GTM or Legal & Ops, because:

- **GTM cannot be cut without breaking the base case.** GTM funds the fire-department conversion that *is* the $1.2M breakeven engine (Question 13). Cutting the revenue engine to save cash is self-defeating at this stage — it would push out breakeven and Series A, the opposite of what the round is for.
- **Legal & Ops is small and largely non-discretionary.** At $135K it is already the smallest bucket, and much of it (patent deadlines, compliance, corporate operations) is time-bound or obligatory. A 30% cut saves little and risks missing hard deadlines.
- **Engineering is the most deferrable because it funds the *upside*, not the base case.** netMESH commercialization is the platform/upside play; the near-term breakeven does not depend on it (it runs on locMESH + fire departments + Chevron renewal). Trimming engineering 30% slows netMESH maturation and the stress-test timeline — a real cost to the *upside* and to Defense/2028+ optionality — but it does **not** break the near-term revenue plan. We would manage the cut by prioritizing the specific engineering work that unblocks Chevron (accuracy/UX) over

longer-horizon features, and by leaning on fractional/partner engineering rather than full-time hires.

In short: we protect the revenue engine (GTM) and the obligatory spend (Legal & Ops), and we absorb a cut in the bucket that funds future upside (Engineering) — accepting slower netMESH progress as the price of extending runway, because the base-case path to breakeven survives it.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf) (use-of-proceeds 60/25/15, milestones, burn/runway), Financial Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf), Seed Round Term Sheet (dataroom.bleedio.com → 02_Financials → BleedIO Tech Seed Round Term Sheet.pdf). Breakeven dependency per Question 13; netMESH readiness per Questions 11–12; patent deadlines per Question 2.

15. How has the financial model changed since the Keiretsu DD report in June 2025?

> **Question as asked:** "How has the financial model evolved since the Keiretsu DD report in June 2025? That report projected $350K in 2025 revenue. What actually happened, and what drove the revision to the current model?"

What the Keiretsu report projected, and what actually happened. The independent Keiretsu Forum due-diligence report (June 2025) projected approximately **$350K in 2025 revenue.** Actual 2025 revenue was **$35,975** — all of it recognized in **Q4 2025** (Q1–Q3 2025 were pre-revenue at $0), per our Puzzle.io cash-activity records. In other words, **2025 came in at roughly 10% of the Keiretsu projection**, and first revenue only began in the final quarter of the year. We are not going to paper over that: the 2025 revenue target was substantially missed.

The revision in one sentence: rather than a broken business, the shift is best understood as an **approximately 12-month rightward move in the revenue timeline** — first revenue arrived in Q4 2025 rather than across 2025, and the ≈$350K revenue level the Keiretsu model expected in 2025 is roughly what our current model projects for **2026 (≈$360K)** — driven primarily by the fact that **both the software development and the business development proved harder and slower than we expected**: more technical barriers in the SW build than the 2025 model anticipated, and industrial contracts that took far longer to sign than modeled (Snap-on, for example, was in motion for roughly 10 months before reaching its current approved state).

What drove the revision. The primary driver, stated directly: **both the software development and the business development proved harder and slower than we expected.** We hit more technical barriers building the platform than the 2025 model assumed, and industrial contracts took far longer to sign than modeled. Everything else follows from those two:

- **Software development was harder than anticipated.** Building DRAW — the proprietary firmware, the decentralized self-healing mesh, the accuracy and orchestration work — ran into more engineering barriers than the June 2025 plan expected. netMESH in particular is at an early technology-readiness stage (TRL 3-class), and bringing it toward commercial deployment is genuinely hard R&D and software work (see Questions 11–12). The 2025 model implicitly assumed a faster development path than the engineering reality allowed.
- **Business development / industrial contracts were slower than anticipated.** Large industrial enterprises are slow to sign and fulfill. As a concrete example, **Snap-on was in motion for roughly 10 months** — "crawling" through evaluation and internal process — before reaching its current approved state (Question 24). That is normal for enterprise/industrial procurement, but it is materially slower than the 2025 model assumed, and it pushed revenue recognition to the right.
- **locMESH became the near-term commercial foundation.** With the full platform and the enterprise contracts taking longer, we made locMESH — live, generating revenue, deployed with real customers — the near-term revenue engine while netMESH matures. This is a risk-management architecture, not a pivot away from netMESH: locMESH funds the company and generates hostile-environment operational data that informs netMESH development.
- **Deliberate pacing compounded the timing.** We also set a deliberate pace to avoid over-committing ahead of product readiness (see Questions 11 and 20), which we discuss further

under the pipeline explanation below.

Why revenue is small now but steps up sharply — and why that step is credible. The jump from ≈$36K (2025) toward the ≈$360K (2026) and ≈$1.2M breakeven (2027, Question 13) is not a hockey-stick assumption; it reflects the nature of the revenue we have booked so far and the shape of the contracts ahead:

- **2025 and Q1–Q2 2026 revenue is from paid pilots — real, but pilots.** These were **paid** engagements, which is meaningful validation, but they ran on our **MVP, not the commercial product.** Pilot-scale, MVP-stage revenue is intentionally small; it proves customers will pay before we have the full commercial product, not that we have hit commercial scale.
- **The large industrial contracts are slow to sign and fulfill — but each one is large.** Enterprise/industrial procurement (Lufthansa, Snap-on, Chevron) takes time to close and deploy. The upside of that same dynamic: **a single contract with Lufthansa or Snap-on can, on its own, cover the entire 2025 ARR gap and roughly half of 2026.** The revenue is lumpy and back-loaded because it is concentrated in a few large enterprise contracts, not spread across many small ones — so it looks like a "jump" precisely because one or two closings move the whole number.
- **The pipeline was deliberately stalled until the commercial product is ready.** We intentionally held back on pushing pilots to full commercial contracts (the same discipline as the Chevron planned pause, Question 11) rather than sign customers onto a product not yet ready to deliver at commercial quality.

This is exactly why the raise is needed. The gap between pilot revenue today and the pipeline ahead is a **capacity and product-readiness gap, not a demand gap.** We need more people at BleedIO to (a) pay down the technical debt and finish the commercial product, and (b) execute the pipeline we deliberately stalled while the product matured (see Questions 28–30 on hiring). The revenue step-up is the conversion of that held-back pipeline once the workforce and the commercial product are in place — which is what this round funds.

How the current model differs from the Keiretsu-era model:

Dimension	Keiretsu-era model (June 2025, per the report)	Current model (2026)
2025 revenue	$350K projected	**$35,975 actual** (all Q4 2025) — ≈10% of projection
2026 / 2027 revenue	$2.8M / $4.5M projected	≈$360K (2026) / ≈$1.2M base case (2027)
Round being raised	**$2.5M SAFE at an $11M post-money cap**	**$900K SAFE at a $6M post-money cap** — the raise itself was right-sized alongside the model (Question 17)
≈$350K revenue level	Expected 2025	Now projected ≈2026 (≈$360K) — ≈12-month shift
Primary near-term product	Less differentiated between products	**locMESH-first**; netMESH as platform/upside
Breakeven	(per Keiretsu-era assumptions)	≈$1.2M ARR / ≈$2.1M revenue, ≈2027 (Question 13)
netMESH commercialization	Assumed sooner	Explicitly staged; funded by this round (Questions 11–12, 14)

The comparison shows a consistent pattern, not a selective one: we revised **revenue, timeline, and the size of the raise itself** downward to match engineering and sales reality. A model that cuts its own ask from $2.5M at an $11M cap to $900K at a $6M cap is being honest with itself, not just with investors.

Why this makes the current model more credible, not less. The revised model is grounded in what has actually happened: locMESH is deployed and generating revenue, the netMESH timeline is

scoped honestly rather than optimistically, and the breakeven case (Question 13) now rests on the repeatable fire-department motion plus Chevron renewal rather than on assumptions that did not hold. We would rather present a model that survived contact with reality — and be candid about the 2025 miss — than defend the original projection.

Sources. Data room: BleedIO Tech Due Diligence Report by Keiretsu Group June 2025 (dataroom.bleedio.com → 01_Company_Overview → BleedIO Tech Due Diligence Report by Keiretsu Group June 2025.pdf) (verified against the report: $350K 2025 / $2.8M 2026 / $4.5M 2027 projections; $2.5M SAFE at $11M post-money cap), Cash Activity Report — Jan 2025 to Mar 2026 (dataroom.bleedio.com → 02_Financials → BleedIO Tech Cash Activity Report-Jan 2025 to Mar 2026.pdf) (2025 actual revenue: $30,975 subscription + $5,000 services = **$35,975, all Q4 2025**; Q1–Q3 2025 = $0), Financial Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf), Financial Assumptions & Definitions (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Assumptions & Definitions.pdf) (current model; actuals from Puzzle.io). Breakeven detail per Question 13; netMESH maturity per Questions 11–12. Primary revision driver (software development and business development both harder/slower than expected — e.g. Snap-on ≈10-month cycle) per BleedIO management; Snap-on status per Question 24.

16. At what ARR and customer milestone will BleedIO initiate Series A conversations?

> **Question as asked:** "At what ARR and customer milestone do you plan to initiate Series A conversations? What is the expected Series A size, timeline, and target investor profile?"

ARR and customer milestone. We plan to initiate Series A conversations at **$1M–$2M ARR**, targeted for **2027 — most likely the second half** (our planning window is Q2–Q4 2027). The ARR threshold is paired with a set of concrete customer/product milestones that make the round fundable at a strong step-up:

- **50+ fire-department deployments** (the public-sector volume driver — see Questions 13, 20)
- **Snap-on OEM in production** (embedded firmware shipping)
- **netMESH commercially deployed** (the platform maturing past the current in-development stage — Questions 11–12)
- Chevron progressing through Lab POC Round 2 → field → commercialization (Question 24)

In other words, we do not initiate Series A on ARR alone; we initiate when ARR *and* the proof points (public-sector scale, OEM in production, netMESH commercial) line up, so the round is raised from evidence rather than promise.

Size. The Series A target is **$3M–$8M depending on ARR at raise** — the term sheet's scenario table (Question 17) frames $3M–$5M at the conservative case (≈$1.0M ARR) and $5M–$8M at the base case (≈$1.5M ARR); a separate optimistic/maritime-scaled model contemplates a larger raise if the Lufthansa channel scales faster (see Question 13). Use of Series A proceeds: scale the GTM team, expand the integrator certification/onboarding program (Question 27), fund netMESH commercial launch, and international expansion via the Lufthansa maritime channel. It also funds the next org layer — dedicated customer success and the formal CTO hire that completes the separation of Stan's roles (Questions 28, 30).

Timeline. Seed close → ≈18–24 months of runway → **Series A in 2H 2027.** The seed round is explicitly sized to reach the Series A milestones above.

Target investor profile. We are targeting a blend: **institutional venture investors focused on industrial IoT, deep-tech, and connectivity infrastructure**, alongside **strategic / corporate-venture investors.** The strategic angle is already live — we are engaged with **Chevron's CVC** and **Lufthansa is making CVC introductions** (Question 24) — which positions the Series A to include strategic capital from partners who are also commercial customers, not only financial VCs. The comparable exit landscape (Cradlepoint/Nozomi; Motorola/Silvus at $4.4B) also makes the category legible to strategic acquirers, which supports both Series A strategic interest and the longer-term exit thesis.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Q&A.pdf) ($1M–$2M ARR at Series A, 2H 2027; $3M–$5M target; milestones), Financial*

Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf), Comparable Exits & Acquisition Landscape (dataroom.bleedio.com → 06_Market&_Traction → BleedIO Tech Comparable Exits & Acquisition Landscape.pdf). Milestones per Questions 13/20; CVC engagement per Question 24; org build-out per Questions 28/30. Exit comparables verified: Ericsson–Cradlepoint $1.1B (2020); Mitsubishi Electric–Nozomi Networks ≈$1B (closed January 2026).

17. Is the round currently open, and what is the current committed total and current terms?

> **Question as asked:** "The one-pager (February 2026) shows $130K raised against a $900K target with a $9M valuation cap. Is the round currently open? What is the current total committed and on what terms?"

Yes — the round is open and active. BleedIO is raising a **$900K seed round**, and we are currently **in due diligence with multiple angel groups.** The round is progressing through those groups' diligence processes rather than being closed or oversubscribed; this NWA diligence is part of that active process.

Current terms — from our proposed (indicative, non-binding) Seed Round Term Sheet (April 2026). The specific terms are set out in the term sheet in the data room. Correcting one figure in the question — it references a "$9M valuation cap," but our cap is **$6M**:

Term	Detail
Round / target raise	Seed — **$900,000**
Instrument	**Post-Money SAFE** (YC standard template)
Valuation cap	**$6,000,000 post-money** (not $9M)
Minimum investment	$25,000
Pro rata rights	Yes — proportional right to participate in the next priced round
Early Investor Program	**10% discount → $5.4M effective cap** for early commitments; applies to the **first $200K committed within 30 days** of the term sheet date (Founding Investor $100K+; Early Investor $25K–$99K)
MFN	If the Company later issues SAFEs on more favorable terms before conversion, this round's investors may adopt them
Information rights	Investors of $50K+ receive quarterly financial updates, an annual summary, and material-event notifications
Board / governance	Founder-managed; **no board seats offered in this round**; formal board anticipated at the first priced round

Existing instrument on the cap table: one prior SAFE — **$100K from Prox SG at a $2M cap** (July 11, 2024). The current $6M cap is roughly a **3× step-up** from that SAFE, which the term sheet attributes to the Chevron deployment, the Snap-on OEM channel agreement, and the two patent provisionals.

Excerpt from the term sheet (Offering Summary):

> *"Round: Seed. Target Raise: $900,000. Minimum Investment: $25,000. Instrument: Post-Money SAFE (YC standard template). Valuation Cap: $6,000,000 post-money. Pro Rata Rights: Yes — proportional right to participate in the next priced round."*
>
> *"How the discount works: Early investors receive a SAFE with a $5.4M post-money cap (10% below the standard $6M cap)… The 10% discount applies to the first $200K committed within*

30 days of the term sheet date. Once the $200K early allocation is filled, all subsequent investments are at the standard $6M cap."

Committed to date — and a correction on the "$130K raised." The **≈$130K the question cites (from the February 2026 one-pager) was pre-seed capital, and that pre-seed is closed** — it is not a commitment against the current $900K seed round. On the current seed round, **no commitments have closed yet**: the round is in **active due diligence with multiple angel groups** (this NWA process among them), and definitive SAFE agreements are executed per-investor as those diligence processes conclude.

Status note. The term sheet is expressly marked **"Indicative — Non-Binding"** and "for discussion purposes only" — it frames the round for the angel groups currently in diligence. We are happy to provide live status on the diligence processes directly.

Ownership context. Post-seed, founders retain approximately **82% ownership** — above the ≈56% seed-stage median (Carta 2026) — reflecting disciplined capital management and minimal pre-seed dilution (see Question 32 for the cofounder equity structure; note that the term sheet's "4 cofounders" and Chevron-traction highlight lines predate the current statements in Questions 11 and 32 — where the term sheet and this document differ, this document is current).

Sources. Data room: Seed Round Term Sheet (April 2026, indicative/non-binding) (dataroom.bleedio.com → 02_Financials → BleedIO Tech Seed Round Term Sheet.pdf) (all terms and excerpts above — $900K / $6M post-money SAFE cap / $25K min / 10%–$5.4M early program on first $200K / pro rata / MFN / info rights / no board seats), Cap Table (dataroom.bleedio.com → 02_Financials → BleedIO Tech Cap Table.pdf), BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf). Round-open status and current commitments (in DD with multiple angel groups) per BleedIO management as of the response date.

Go-To-Market and Device Compatibility

18. What is BleedIO's explicit strategy for industrial brownfield deployments?

> **Question as asked:** "Many industrial environments have large installed bases of sensors that are not BLE-enabled. What is your explicit strategy for brownfield deployments? Do you recommend hardware replacement, gateway bridging, or positioning primarily in greenfield?"

Our brownfield strategy starts from a clear principle: DRAW is complementary, not rip-and-replace. locMESH and netMESH are not sold as a replacement for a customer's existing wireless or wired infrastructure. They are a **complementary connectivity and location layer** that customers add where their current infrastructure cannot deliver — without tearing anything out. This is deliberate: it removes the biggest source of friction in industrial sales (the cost, risk, and downtime of ripping out installed systems) and lets us deliver value alongside what the customer already runs.

The specific problem we complement: where Wi-Fi makes things worse, not better. The clearest brownfield case is Wi-Fi coverage in industrial environments. In metal-dense, RF-hostile facilities — refineries, manufacturing floors, warehouses — Wi-Fi coverage is often inadequate, and the conventional fix makes it worse. Adding more access points to chase coverage **adds more interference**, drives **co-channel contention and broadcast storms**, and the net result is **rising latency and declining reliability** — the opposite of what the customer wanted. Past a point, more Wi-Fi hardware degrades the very coverage it was meant to improve.

There is also a hard density limit. Access points support only a bounded number of clients before performance falls off — Cisco's own design guidance places the practical clients-per-AP figure well below its theoretical radio maximums (typical realistic planning is on the order of a couple hundred devices per AP at most, and lower for demanding traffic). Dense industrial sensor deployments — hundreds or thousands of sensors in a facility — blow straight past that, so covering them with Wi-Fi APs is both expensive and structurally fragile. netMESH is purpose-built for the opposite profile: a **low-latency, low-throughput network with practically no dead zones**, designed to carry large numbers of small, frequent sensor messages reliably rather than to stream high-bandwidth data.

This is not just a theoretical argument. In our discussions with C-level operations leaders, the recurring, unprompted concern is that **Wi-Fi is not stable enough and cannot cover everything** — and that piling more sensors onto each AP makes it worse, not better. DRAW is designed for exactly this failure mode: a decentralized, self-healing mesh where every node relays for every other node, so coverage improves as you add nodes (more nodes = more relay paths) rather than degrading. Even when a customer already has nominal coverage, the stability and density problems are why they keep looking for an alternative. We do not fight the customer's Wi-Fi; we cover the gaps it structurally cannot fill.

How this maps to the device-integration tiers (see Question 5). Brownfield is not only about coverage; it is also about the installed sensor base. Our explicit approach, matched to what the customer already has:

- **Do not replace what works.** BLE mesh–capable devices re-provision into DRAW directly; BLE-but-not-mesh sensors are read by specialized DRAW mesh nodes; fully non-BLE systems keep running while DRAW operates as an overlay, with data brought together at the gateway. (Full detail in Question 5.)
- **Bridge to and from the wired network for large, complicated sites.** For big or structurally complex environments, we can **bridge netMESH into a customer's wired network and bridge back out to netMESH on the far side** — using the wired backbone as a fast transit link between mesh segments. This lets us span long distances, cross physical barriers, and stitch multiple mesh zones together without forcing every hop to be wireless, which is important for sprawling industrial facilities where no single mesh domain would be sensible.
- **Deliver value on day one** by standing up DRAW coverage and the specific data/location the customer needs, without a forced, site-wide conversion.
- **Expand within the account over time** — including the OTA locMESH→netMESH upgrade path (Question 1) — so the footprint grows as the customer validates results.

Positioning summary — greenfield vs. brownfield. We are not "greenfield only." Greenfield deployments are simpler for us, but our differentiated wedge is precisely the brownfield case: environments where the customer has already tried to solve coverage/connectivity with conventional infrastructure and hit its structural limits. We position DRAW as the layer that sits **underneath and alongside** existing systems — complementary to the customer's IoT platforms and wired/wireless infrastructure, competitive only against other pure connectivity-infrastructure approaches. The Lufthansa-style dedicated zone-overlay model is the concrete commercial expression of this brownfield positioning, and its pricing is addressed in Question 19.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf) (connectivity-infrastructure-layer / complementary positioning), Competitive Landscape (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Competitive Landscape.pdf), Partner Integration Architecture (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Partner Integration Architecture.pdf), netMESH Data Sheet (dataroom.bleedio.com → 05_Product&*_Technology → BleedIO Tech netMESH Data Sheet.pdf). External — Wi-Fi client-density limits: Cisco Meraki — Approximating Maximum Clients per Access Point. Device-integration tiers per Question 5; OTA upgrade path per Question 1; overlay pricing per Question 19.

19. Is the Lufthansa-style dedicated overlay model the primary brownfield positioning, and how is it priced?

> **Question as asked:** "The Lufthansa model — dedicated BLE Mesh overlays per operational zone, running alongside rather than replacing existing infrastructure — seems like a practical approach for environments with mixed installed hardware. Is this your primary brownfield positioning, and how do you price it?"

Yes — the dedicated zone-overlay model is our primary model for expansion in industrial installations. The Lufthansa-style approach — a dedicated DRAW mesh overlay per operational zone, running alongside rather than replacing existing infrastructure — is exactly how we expand inside industrial accounts. A customer stands up a DRAW overlay in the zone where their current connectivity falls short (see the Wi-Fi coverage/density problem in Question 18), proves value there, and then adds zones. Each zone is a self-contained, self-healing overlay, and the deployment grows

zone-by-zone rather than through a single site-wide project.

What the overlay is scoped for — and what it is not. netMESH in this model is a **dedicated sensor network for low-volume, real-time data**: sensor telemetry, status, location, alerts, and control messages. It is deliberately **not** designed for high-volume payloads such as video streams or heavy compute/data-intensive traffic (see the bandwidth discussion in Question 12). Scoping the overlay to low-volume real-time sensor data is what lets it deliver the reliability, low latency, and near-zero-dead-zone coverage that make it valuable — and it keeps the pricing clean, because value scales with devices and zones, not with bandwidth.

How it is priced. The overlay maps naturally onto a **two-part, software-first pricing model**, consistent with our existing deployment economics:

1. **Per-zone deployment fee (one-time):** covers design, node/gateway provisioning, install, and commissioning of the dedicated overlay in that operational zone. Our documented deployment economics are **$25K–$60K+ ACV per deployment** for a typical 50–200-node zone; larger or more complex zones sit at the upper end or above. Hardware is partner-supplied (sub-$100/node), so this line is services-and-integration, not hardware margin.
2. **Per-device monthly SaaS subscription (recurring):** the DRAW management, monitoring, telemetry, and application layer, priced at **$3–$5 per device per month** (roughly $10,800–$24,000/year for a typical zone, scaling with device count). This is the primary recurring-revenue engine and carries **80%+ gross margin**, since hardware is partner-supplied.

This structure has three properties that fit the overlay model well:

- **It aligns price with value and with the product's design.** Because netMESH is a low-volume sensor network, per-device pricing tracks exactly what the customer is buying — coverage and reliable data from a defined set of sensors — rather than metered bandwidth. A customer who adds sensors or zones pays proportionally more; a customer running heavy data would be mis-served and mis-priced, which is why we scope the overlay to sensor data.
- **It makes land-and-expand natural.** A first zone is a modest, low-risk purchase ($25K–$60K deploy + per-device SaaS). Each additional zone is an incremental deployment fee plus incremental per-device SaaS, so account revenue compounds as the overlay spreads — and the recurring SaaS base grows with every device added.
- **It is predictable for the customer.** Per-device, per-zone pricing is easy for an operations buyer to model and budget, with no surprise bandwidth or overage charges.

For a multi-zone or multi-site enterprise, we would wrap these into a negotiated framework agreement (volume-tiered per-device SaaS, standardized per-zone deployment pricing) rather than pricing each zone in isolation — the same way the 10-site TCO in Question 7 is structured. We are happy to model specific zone counts and device densities for a given facility under NDA.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf) (SaaS $3–$5/device/month; deployment $25K–$60K+ ACV; 50–200-node typical deployment; 80%+ SaaS gross margin), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf), Financial Assumptions & Definitions (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Assumptions & Definitions.pdf). Brownfield/complementary positioning per Question 18; bandwidth scope per Question 12; multi-site framework per Question 7.

20. Which verticals are in scope for active sales in 2026-2027, and which are 2028+ targets?

> **Question as asked:** "Which verticals are in-scope for active sales in 2026–2027, and which are 2028+ targets? How are you prioritizing given a 5-person team and limited GTM capacity?"

We operate across three verticals — **Industrial**, **Public Sector**, and **Defense & Autonomy** — and we sequence them deliberately given a lean team.

Active sales in 2026–2027:

- **Public Sector (near-term wedge).** Fire departments and first responders are the fastest path to revenue: GPS fails indoors and there is no reliable incumbent for firefighter location/accountability, the evaluation cycle is short (**1–2 months**), and the IAFF relationship gives us a channel into 27,000+ U.S. departments. We have 2 fire departments signed (contract onboarding paid, including hardware; deployments active in the field), with **50 targeted by end of 2026.** This vertical is in active, prioritized sales now because the cycle is short and the sales motion is repeatable.
- **Industrial (highest-value wedge).** Our highest-value revenue per account, sold now: O&G refinery telemetry (Chevron, paying customer), industrial tooling (Snap-on, OEM agreement in progress), maritime/cruise IoT (Lufthansa Industry Solutions, channel partner — contract onboarding paid), and live-event operations (Oracle Red Bull Racing, agreement in progress). Industrial is in active sales in 2026–2027, expanded primarily through the zone-overlay model (Questions 18–19) and through channel/OEM partners rather than direct headcount.

Market-entry / building in 2026–2027, scaling 2028+:

- **Defense & Autonomy.** This is at **market-entry stage** today — drone-swarm coordination and contested-environment operations, with netMESH as the platform target. The pull is real (the $13.4B FY2026 DoD autonomy budget, and DRAW's ITAR-free open-standard foundation enabling direct federal/allied procurement), but it depends on netMESH maturing to commercial readiness (see Questions 11–12). We are seeding relationships and positioning now; meaningful revenue scale is a **2028+** horizon, gated on netMESH commercialization.

Explicit 2028+ targets (next-wave verticals):

- **Healthcare and logistics** are identified next-wave verticals but are **2028+**, because both carry significant regulatory requirements (HIPAA, FDA, maritime certifications) that gate entry and are not the right use of a 5-person team's capacity today. Our regulatory-readiness path for these is addressed in Question 22.

How we prioritize given a 5-person team. We are deliberately lean, and prioritization is a survival constraint, not just a preference. Three principles govern it:

1. **Sequence by sales-cycle speed and repeatability.** Public Sector first (short 1–2 month cycle, repeatable motion), Industrial in parallel for account value, Defense & Autonomy and the next-wave verticals later when the product and the team can support longer, more complex cycles.
2. **Multiply reach through channels and OEM, not headcount.** Lufthansa IS (maritime reseller), Snap-on (OEM embedding), and the system-integrator network let us cover more market than a 5-person team could sell directly. This is why we can be "in" Industrial broadly without a large direct sales force.
3. **Don't over-commit ahead of readiness.** We intentionally pace pipeline to avoid overpromising on a vertical (e.g. Defense, healthcare, logistics) before netMESH and our regulatory/GTM capacity are ready to deliver — the same discipline reflected in the Chevron sequencing (Question 11). The seed round's GTM allocation (Question 14) funds the first sales/solutions hires that expand this capacity.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Q&A.pdf) (vertical scope, statuses, sales cycles), Customer & Pipeline Status Matrix (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Customer & Pipeline Status Matrix.pdf), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf), TAM, SAM, SOM Analysis (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech TAM, SAM, SOM Analysis 2026-05-23.pdf). Regulatory path per Question 22; use-of-proceeds/hiring per Question 14; netMESH readiness per Questions 11–12.

21. Is licensing netMESH to chip manufacturers a near-term strategy or a long-term aspiration?

> **Question as asked:** "You mentioned the possibility of licensing netMESH to chip manufacturers like Nordic or STM — embedding it upstream in the hardware ecosystem. Is this a

near-term commercial strategy or a long-term aspiration? Have you had conversations with Nordic, STM, or other chipmakers?"

Licensing to chip manufacturers is a long-term aspiration, not a near-term strategy — and we want to be candid about that. It is a real part of the long-range vision (embedding DRAW upstream in the hardware ecosystem), but it is not something we are pursuing or pricing into the near-term plan, and we have not represented it as imminent.

To be direct about the current state: we have **no current conversations with STMicroelectronics, Nordic, or other chipmakers about licensing.** We should be clear about why: **BleedIO is too small today to represent the volume a chip manufacturer needs to justify a licensing arrangement.** Licensing upstream is a scale game — a chipmaker embeds a stack when there is enough downstream demand to move silicon — and we have not yet reached that scale. Claiming active licensing talks would overstate where we are.

(For clarity, this is distinct from our hardware relationships: as noted in Question 4, DRAW already *runs on* STM and Nordic silicon today. That is supply and compatibility — DRAW being hardware-agnostic across their chips — not a licensing deal in which they embed netMESH. The two should not be confused.)

What has to happen first — and why the timing works. There is a concrete technical reason licensing becomes more compelling over time. The open mesh standard was originally designed for **lighting applications — switches and dimmers** — so its built-in device models are weighted toward that world. It does **not** include built-in models for **motors and many industrial sensors**, which is exactly the gap DRAW's proprietary firmware models fill (see Question 1). As industrial installations grow and more third-party manufacturers need mesh-enabled motors, industrial sensors, and equipment, the demand for a richer, industrial-grade model library grows with them. That is the point at which DRAW's models become valuable to embed upstream, and at which BleedIO would carry the volume and the reference deployments to make a licensing conversation credible.

So the sequence is: (1) win industrial and public-sector deployments now, building the proprietary industrial model library and the deployment volume; (2) as third-party industrial manufacturers adopt the ecosystem, reach the scale where chipmakers have a reason to embed DRAW; (3) only then does upstream licensing to chip manufacturers become a near-term commercial motion rather than an aspiration. Presenting it in that order reflects where we actually are — chipmaker licensing is not on the table today, and we do not imply otherwise.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Q&A.pdf), Product Vision (dataroom.bleedio.com → 05_Product*&*Technology → BleedIO Tech Product Vision.pdf), netMESH Data Sheet (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech netMESH Data Sheet.pdf) (proprietary firmware/device models beyond the standard's lighting-oriented models). Hardware-agnostic relationship with STM/Nordic per Question 4; DRAW device-model differentiation per Question 1.

22. What is BleedIO's regulatory readiness path for healthcare and logistics?

> **Question as asked:** "Healthcare and logistics are identified as next-wave verticals. Both have significant regulatory requirements (HIPAA, maritime certifications, FDA). What is your path to regulatory readiness in those markets, and at what product maturity level?"

The key point first: netMESH is a data carrier, not a data collector. The question assumes BleedIO would carry the full regulatory burden of healthcare and logistics. But netMESH is the **transport layer** — it moves data across the mesh; it is not, by itself, the system that collects, stores, or is the custodian of protected health information or regulated logistics records. Most of the regulatory weight (HIPAA, FDA) attaches to the collection, storage, and processing of the data and to the medical device or the record system — not to the wireless carrier that transports messages. So as a data carrier, **netMESH does not carry a direct HIPAA or logistics regulatory requirement** in the way the question implies. This is an important scoping distinction, and we would make it clearly to any healthcare buyer.

If and when BleedIO moves into data collection for healthcare, the path exists. We can extend from carrier into data collection for healthcare use cases, and we have the building blocks to do so responsibly:

- **HIPAA-eligible cloud.** Our platform runs on **Microsoft Azure and AWS**, both of which offer HIPAA-eligible services under a Business Associate Agreement (BAA). Building on that foundation means the cloud/storage layer can be operated within a HIPAA-compliant architecture when a healthcare data-collection use case requires it — rather than us having to build a compliant cloud from scratch. (To be precise: Azure and AWS being HIPAA-eligible is the enabling foundation; full HIPAA compliance is a property of the whole system and process, which we would complete for a specific healthcare deployment.)
- **Encryption in transit.** netMESH data collection is encrypted — **AES-128 or AES-256 depending on the netMESH version** (our AES-256 extension over the standard's baseline; see Question 1). Encryption in transit is a core control HIPAA expects, and it is built into the product rather than bolted on.

Product-maturity gating. Consistent with our answer on vertical sequencing (Question 20), healthcare and logistics are **2028+ next-wave verticals, not 2026–2027 active markets.** The regulatory path is therefore staged and deliberately not front-loaded:

1. **Today:** netMESH is positioned as a data carrier, where the direct regulatory burden is limited, and we do not market it into HIPAA/FDA-regulated data-collection roles it is not yet certified for.
2. **As netMESH reaches commercial maturity** (Questions 11–12) and a concrete healthcare data-collection opportunity is real, we complete the compliance work for that specific use case on the HIPAA-eligible Azure/AWS foundation, with encryption already in place, and add the BAA, access controls, audit logging, and documentation a specific deployment requires.
3. **Logistics** carries lighter regulatory requirements than healthcare in most cases; where maritime or other certifications apply, they are handled at the deployment/vertical level (e.g. through channel partners such as Lufthansa IS for maritime).

The accurate scoping: carrier now; compliant data collection when the product and a real opportunity are ready — not a claim of healthcare/FDA readiness we have not yet earned.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf) (cloud on Azure/AWS/GCP), netMESH Data Sheet (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech netMESH Data Sheet.pdf) (AES-256 end-to-end encryption), IP & Compliance Summary (dataroom.bleedio.com → 04_IP*&_Compliance → BleedIO Tech IP & Compliance Summary.pdf) (early-stage compliance posture). External — HIPAA eligibility of the cloud foundation: Microsoft Azure HIPAA compliance, AWS HIPAA Compliance. Vertical sequencing per Question 20; encryption/device models per Question 1; netMESH maturity per Questions 11–12.

Commercial Traction and Pipeline

23. What is current ARR as of June 2026, which contracts generate revenue, and what is the pricing structure?

> **Question as asked:** "What is your current ARR as of June 2026? Which specific contracts are generating revenue, and what is the pricing structure for each?"

Current ARR — stated plainly. Current recurring revenue is approximately **$3,000–$4,000 MRR (≈$48K ARR)** as of June 2026. We are not going to inflate this: the company is early-revenue, and the ≈$48K is the honest current figure. The story is the trajectory from here — the ≈$1.2M breakeven case (Question 13) is a conversion of the identified pipeline, and the ≈$48K is the proven, paying base underneath it.

Which contracts generate revenue today — with the recurring vs. one-time distinction made explicit.

- **Chevron — paying customer** (locMESH lab evaluation — RTLS + BLE legacy sensors; the lab stage is in a **planned pause** pending accuracy/UI-UX improvements, with Lab POC Round 2 in Q3 2026 as the next step — see Question 11). Revenue-generating and active.

- **Lufthansa Industry Solutions — channel partner; contract onboarding paid.** Lufthansa IS has **paid for contract onboarding** (payment received May 7, 2026); **recurring revenue-share has not yet started** — it begins as cruise-line deployments go live (Question 24). BleedIO sells to Lufthansa IS, which resells to cruise lines.
- **Two fire departments — signed; contract onboarding paid (including hardware).** Both departments **paid for contract onboarding, which includes the hardware**; deployments are active in the field (Question 9). **Recurring SaaS payments have not yet started** on these accounts.

Snap-on and Oracle Red Bull Racing are agreements in progress / OEM in review (not yet revenue-generating — see Question 24); Carbon Reform is signed but at hardware-development stage.

Pricing structure. Our model is software-first, two-part (consistent with Questions 7 and 19):

- **SaaS subscription: $3–$5 per device per month** (the recurring engine; ≈80%+ gross margin, since hardware is partner-supplied).
- **Network deployment: $25K–$60K+ ACV per deployment** (one-time install/integration for a typical 50–200-node deployment).
- **Channel/OEM variants:** Lufthansa maritime is a revenue-share channel model (per-ship economics differ — Question 7/19); Snap-on is an OEM embed (per-device royalty at shipment scale). Public-sector/fire-department SaaS sits near a floor (≈$900/month, ≈$10.8K/yr minimum) and scales with devices.

What this base demonstrates. Current ARR is small and we present it as such; what it demonstrates is a **live, paid, multi-segment base** — a Fortune 500 industrial customer paying today, a maritime channel partner that has paid onboarding, and public-sector customers that have paid onboarding (including hardware) with active field deployments — across all three verticals. Customers commit real money at every tier before the recurring engine turns on, which is the proof that the commercial model works in the field before we scale it against the pipeline.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Q&A.pdf) ($3K–$4K MRR; Chevron paying; Lufthansa IS onboarding payment May 7 2026; SaaS $3–$5/device/mo; deployment $25K–$60K ACV; recurring vs. one-time characterization per BleedIO management), Customer & Pipeline Status Matrix (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech Customer & Pipeline Status Matrix.pdf), Financial Assumptions & Definitions (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Assumptions & Definitions.pdf). Chevron planned-pause status per Question 11; pricing detail per Questions 7/19; account pipeline per Question 24.

24. What are the next steps, timelines, contract value ranges, and blockers for Chevron, Snap-on, and Lufthansa?

> **Question as asked:** "Walk us through the conversion pipeline: for Chevron, Snap-on, and Lufthansa — what are the specific next steps, decision timelines, and contract value ranges? What are the blockers to signing?"
>
> Contract values below are **modeled ACV ranges**, consistent with the Target Contract Pipeline table in the Seed Round Term Sheet (Question 17). None of the three contracts is signed at these values yet — the per-account blockers are stated plainly.

Below is the conversion pipeline for each of the three accounts — next step, timeline, contract value range, and the specific blocker to signing.

Snap-on (industrial OEM / VAR).

- **Where it stands:** The **decision is made and approved.** An approved showroom installation is scheduled **ASAP (July/August 2026)**, and the **VAR (value-added reseller) agreement is in review.**
- **Next step:** Complete the showroom installation and execute the VAR agreement.
- **Timeline:** Showroom install July/August 2026; VAR agreement finalizing in parallel.

- **The only blocker is BleedIO team readiness** — i.e. our own execution capacity to deliver the install and support, not any decision or approval gap on Snap-on's side. This is exactly the kind of near-term capacity the seed round's GTM allocation funds (Question 14).
- **Contract value range (modeled ACV): Snap-on '26 (showroom + first VAR shipments): $50K–$150K → Snap-on '27+ (OEM embed at production volume): $300K–$500K ACV**, with $1M+ lifetime value per OEM partner as shipments scale.

Lufthansa Industry Solutions (maritime channel).

- **Where it stands:** Lufthansa IS is a signed channel partner that has **paid for contract onboarding** (payment received May 7, 2026); **recurring revenue-share has not yet started** — it begins with live cruise deployments. The **Miami office installation is complete.**
- **Next step:** The **first joint sales meeting — Lufthansa and BleedIO together — plus a demonstration with Carnival Cruise Lines, is scheduled for July/August 2026.** The step after that is a **pilot on a ship.** Lufthansa is also **making CVC (corporate venture capital) introductions** on our behalf.
- **Timeline:** Joint Carnival meeting/demo July/August 2026 → ship pilot following a successful demo → first cruise deployment targeted Q4 2026.
- **Blocker: Loss of momentum** — the risk here is not a "no," it is letting the Carnival demo-to-pilot sequence stall. Sustaining momentum through the joint sales motion is the key.
- **Contract value range (modeled ACV): Lufthansa '26 (pilot ship): $100K–$150K deployment + $48K–$60K/yr SaaS per ship (≈$150K–$330K 3-year ACV per ship) → Lufthansa '27 (2–3 ships): $200K–$400K channel ACV** (per-ship economics per Questions 7 and 19). Upside for scale: Carnival Cruise Line alone operates ≈29 ships — a full-fleet rollout of that single brand models to ≈$4.4M–$9.6M over three years.

Chevron (industrial O&G).

- **Where it stands:** Paying customer; the locMESH lab evaluation is in a **planned pause** pending the accuracy and UI/UX improvements we are completing (Question 11).
- **Next step and sequence: Lab POC Round 2 (Q3 2026)** — reactivate in the lab first — then **field testing (planned Q3–Q4 2026)**, with **commercialization thereafter**; a **netMESH evaluation phase follows later (not yet dated).** We are also **already engaged with Chevron's CVC.**
- **Timeline:** Lab POC Round 2 Q3 2026 → field Q3–Q4 2026 (gated on lab results) → commercialization after successful field validation.
- **Blocker: Technical failures** — the gating risk is technical: whether locMESH meets the accuracy/UI-UX bar in Lab POC Round 2 and the field, and later whether netMESH meets the uptime/packet-loss/latency bar in its subsequent evaluation. This is a product-execution risk, not a commercial or relationship risk, and it is the risk the engineering spend (Question 14) is aimed at retiring.
- **Contract value range (modeled ACV): Chevron '26 (renewal, current site): $150K–$250K → Chevron '27–'28 (expansion, 2–3 sites): +$400K–$600K at run-rate** (2–3 sites × $150K–$250K per-site ACV), with a 5-site path modeling to $750K–$1.25M ARR. For reconciliation with our revenue-bridge model: the bridge's 2027 Chevron figure of $400K–$600K is the *recognized-revenue total including the renewal*, reflecting partial-year ramp of the expansion sites — the run-rate increment stated here is the same model viewed at steady state.

Cross-cutting note — CVC engagement. Two of the three accounts are also strategic-investor conversations: we are already engaged with **Chevron's CVC**, and **Lufthansa is making CVC introductions.** These are potential sources of both strategic capital and deeper commercial commitment, and they reinforce that these are not arm's-length pilots but relationships the counterparties are investing in.

Summary table:

Account	Next step	Timeline	Blocker	Contract value
Snap-on	Showroom install + VAR agreement	July/Aug 2026	BleedIO team readiness (our execution)	$50K–$150K (’26) → $300K–$500K ACV (’27+)
Lufthansa	Joint Carnival demo → ship pilot	Jul/Aug 2026 → Q4 2026	Loss of momentum	≈$150K–$330K 3-yr ACV/ship → $200K–$400K (’27)
Chevron	Lab POC Round 2 → field → commercialize	Q3 → Q3–Q4 2026	Technical failures	$150K–$250K renewal → +$400K–$600K expansion (run-rate)

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Q&A.pdf) (Lufthansa channel partner, onboarding payment May 7 2026; first cruise deployment Q4 2026), Customer & Pipeline Status Matrix (dataroom.bleedio.com → 06_Market&Traction → BleedIO Tech Customer & Pipeline Status Matrix.pdf), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf), VAR Agreement (Customer Contracts) (dataroom.bleedio.com → 08_Customer_Contracts → BleedIO Tech, Inc._VAR Agreement_09-24-2025 - Executed.pdf). Chevron status/sequence per Question 11; maritime pricing per Question 19; use-of-proceeds per Question 14. Contract values are **modeled ACV ranges** per the Seed Round Term Sheet’s Target Contract Pipeline table (Question 17) and the financial model — not signed contract values. Carnival fleet size (≈29 ships): public fleet listings (carnival.com/cruise-ships). CVC details per BleedIO management.

25. What is the status of the defense contractor drone swarm opportunity?

> **Question as asked:** “You mentioned a defense contractor drone swarm contract opportunity that came up the morning of our meeting. What is the status of that contract — size, timeline, probability of close, and who is the contracting entity?”

Here are the specifics on the drone-swarm opportunity referenced in our meeting:

- **Contracting entity: Draganfly Inc. (NASDAQ/CSE: DPRO).** Draganfly is an established, publicly-traded U.S.-linked UAV/autonomous-systems company with genuine defense traction — its Commander 3XL platform was selected by a branch of the U.S. Department of Defense (2025), and it holds defense-partner relationships with DLA prime contractors. This is a credible, named counterparty, not a speculative lead.
- **Size:** a **$750K proof-of-concept (POC)**, with an **eight-figure opportunity for full rollout** if the POC succeeds.
- **Probability of close: 80%+** (management estimate) on the POC.
- **Timeline / structure:** POC first, full-rollout scale contingent on POC success — the standard defense adoption path (prove the capability, then scale to program volume).

Why BleedIO fits this opportunity. Drone-swarm coordination in GPS-denied or contested environments is precisely the problem netMESH is designed for, and DRAW brings structural advantages that the incumbent tactical-radio options (Silvus, Persistent Systems) cannot match at the relevant price point:

- **Cost per node: sub-$100 vs. $5K–$30K for MANET radios.** Attritable swarms require a radio that costs a fraction of the drone; a $30K radio cannot ride on a $5K drone. This is the core of the defense thesis (see Question 6 and the Silvus comparable).
- **Low RF emission = low RF signature.** BLE’s low transmit power means each node emits a fraction of a MANET radio’s RF energy, reducing the swarm’s detectability to passive RF/direction-finding sensors — a tactical advantage that is not purchasable from higher-power tactical radios.

- **Decentralized, self-healing coordination without a ground-control link** — the netMESH design target (Question 9).
- **ITAR-free open-standard foundation** lowers export friction for allied-nation and federal procurement (program-specific export review may still apply).

Maturity and scope. This is a **real, named, high-probability POC opportunity**, and it validates the demand thesis for defense/autonomy. But it is a **POC**, and Defense & Autonomy is our **market-entry / 2028+ scaling vertical** (Question 20) — the eight-figure rollout depends on netMESH continuing to mature to the required performance and on the POC succeeding. Consistent with how we present the rest of this diligence, we treat the $750K POC as the near-term, high-probability item and the eight-figure rollout as significant upside optionality, not as base-case underwriting. We do not build our breakeven case (Question 13) on defense revenue.

Sources. Contract specifics (entity, size, 80%+ probability) per BleedIO management. Data room — drone-swarm technical/defense positioning: Drone Swarm Operations (dataroom.bleedio.com → 05_Product_&*Technology → BleedIO Tech Drone Swarm Operations.pdf), Defense Contractor Requirements for Autonomous Mesh Networks (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Defense Contractor Requirements for Autonomous Mesh Networks.pdf), BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech Q&A.pdf). External — Draganfly is a public company with DoD adoption: Draganfly — Drones for Military & Government, Draganfly SEC filings. Cost-per-node/defense thesis per Question 6; netMESH maturity per Questions 11–12; vertical sequencing per Question 20. DoD FY2026 AI/autonomy budget ($13.4B, first dedicated allocation): DefenseScoop. MANET radio price magnitude: public procurement records (e.g. StreamCaster 4480 at $13,500/unit, City of San Diego PO, 2018).

26. What are the specific Capital Factory equity and ROFR terms?

> **Question as asked:** "Capital Factory recently took a small equity stake and a right of first refusal for Series A at a minimum of $500K. What are the specific terms — equity percentage, ROFR conditions, and what milestones trigger or expire the ROFR?"

The specific terms are set out in the executed **Equity Issuance Agreement** between BleedIO Tech, Inc. and **2026 Fund, a series of Capital Factory All Access, LP** (dated May 22, 2026, signed by Stan Podolski for BleedIO and Jamie Serio, Managing Director, for Capital Factory), available in the data room.

Equity stake. Capital Factory received **8,113 shares of Common Stock at $0.00001 per share, equal to exactly 1% of BleedIO's fully-diluted capitalization** (811,273 fully-diluted shares immediately following issuance). Two important characterizations:

- This is an **accelerator-membership equity grant, not a cash investment.** The purchase price for the shares is applied as BleedIO's membership fee for the **Capital Factory All Access Membership Program** (paid to Capital Factory Ventures, LLC). Capital Factory is not putting new cash into the balance sheet; the 1% is the cost of the accelerator membership.
- The fully-diluted 1% calculation **excludes outstanding SAFEs/convertible notes** (e.g. the Prox SG SAFE), per the agreement's definition.

The "$500K" and the investment right — it is a Right of First *Offer* (ROFO), not a Right of First *Refusal*. A precise reading of the agreement matters here, because the "$500K minimum" and "ROFR" in the question are close to, but not exactly, the contract language:

- **What Capital Factory has** is a **Right of First Offer** to invest in BleedIO's **next Qualified Equity Financing** (defined as a preferred-stock round raising ≥ $2,000,000). It is a right to *participate by purchasing* new preferred shares — not a right to match or block a third party's terms, and not a right tied specifically to "Series A" by name.
- **The $500K is the "Designated Allocation."** Capital Factory may purchase up to the **greater of (a) its pro-rata share or (b) $500,000** of the preferred stock being sold for new money in that financing. So $500K is the *floor on the size of the allocation Capital Factory may elect to buy*, not a minimum it is obligated to invest and not a minimum round size.

- **Trigger and mechanics:** the right attaches to the next Qualified Equity Financing (≥ $2M preferred). BleedIO must give Capital Factory **20 business days' prior written notice** of a Qualified (or Non-Qualified) Financing, furnish the financing documents/term sheet and a pro-forma cap table, and give Capital Factory at least **20 business days** to elect whether to exercise.
- **Expiration / termination:** the investment right (including the ROFO) **terminates upon the earlier of** (1) consummation of that next Qualified Equity Financing (provided BleedIO complied with the notice provisions) — i.e. it is a **one-time, next-round right that is used up at the next qualifying raise**; or (2) an IPO, a qualifying merger/consolidation, or Capital Factory no longer holding any BleedIO securities.

Other terms worth noting for diligence.

- **No board seat.** The agreement grants investment/information rights, not board representation.
- **Put right.** Capital Factory holds a put right to require BleedIO to repurchase all of its securities for an aggregate **$1.00** — a standard accelerator "clean exit" mechanic.
- **Competitor-investment freedom.** Capital Factory (as a VC/accelerator) is expressly free to invest in competing companies.
- **Lock-up** (180-day IPO lock-up) and standard reps; governed by **Delaware** law, venue **Travis County, Texas**.

Net for an incoming investor. The Capital Factory position is a small (1%), non-cash, accelerator-membership equity stake with a one-time right of first *offer* to buy up to the greater of pro-rata or $500K of the next ≥$2M preferred round, that right expiring once that round closes. It is a light, standard accelerator arrangement — not a blocking right, not a board seat, and not a standing ROFR over all future financings. We note that the question's "ROFR for Series A at a minimum of $500K" is a reasonable paraphrase, but the precise mechanics are the ROFO/Designated-Allocation structure above, and we would characterize it that way to avoid overstating Capital Factory's rights.

(Terminology note: this Capital Factory Right of First Offer is entirely separate from the Company's Right of First Refusal on employee stock-option share transfers under §9A of the 2024 Stock Option Plan — the two ROF-rights are unrelated and should not be conflated.)

Sources. Data room: Capital Factory All Access — 1% Equity Issuance Agreement (Signed) (dataroom.bleedio.com → 08_Customer_Contracts → All Access 2026 - 1% Equity Issuance.docx - Signed.pdf) (all terms above — 1% / 8,113 shares, Designated Allocation = greater of pro-rata or $500K, ROFO on next ≥$2M Qualified Equity Financing, 20-business-day notice, termination at that financing/IPO/merger, put right for $1.00, no board seat). Unrelated employee-transfer ROFR: 2024 Stock Option Plan (dataroom.bleedio.com → 03_Legal_&_Governance → BleedIO Tech 2024 Stock Option Plan.pdf) §9A.

27. How many of the 100+ channel partners are actively in diligence or piloting versus passive interest?

> **Question as asked:** "You have 100+ channel partners on a waitlist. How many are actively in diligence or piloting vs. passive interest? What does conversion to an active channel partner look like in practice?"

The honest breakdown. We have 100+ system integrators on the waitlist, and we want to be straight about what that number means: **the large majority are in a holding pattern — passive interest, not active diligence or piloting.** They have signed up to the waitlist and represent real demand signal, but they are waiting while we improve and streamline our onboarding and partner-activation process. The accurate characterization is a qualified waitlist — not 100+ partners actively selling today.

The activated / in-process partners are deliberately few — and deliberately chosen. The select partners we have activated or have in-process are **all experienced commercial integrators.** That is by design: at this stage they are not just distribution, they are **product partners** — they are actively providing feedback on the platform, and that feedback is exactly what we are using to build a repeatable, scalable onboarding and activation process. In other words, we are intentionally running

a small number of high-quality integrators first to harden the onboarding motion, rather than trying to activate 100 partners against a process that is not yet streamlined.

What "conversion to an active channel partner" looks like in practice. The progression is:

1. **Waitlist (passive interest)** — signed up, demand signal, not yet engaged in activation. This is where most of the 100+ sit today.
2. **Activation / in-process (active)** — an experienced integrator is onboarded, trained on the platform, and begins providing feedback and working toward deployments. This is the small, curated group today.
3. **Active channel partner** — the integrator can independently scope, deploy, and support DRAW for their own customers, with BleedIO as the infrastructure layer underneath (the same complementary positioning as Questions 18–19). At that point the partner's customers become BleedIO's addressable market without proportional BleedIO headcount.

Why the waitlist is throttled on purpose — and how the round changes it. The gating factor today is not partner demand; it is our onboarding and activation capacity as a lean team. We are deliberately not flooding activation ahead of a streamlined process, because a bad onboarding experience at scale would waste the demand. The seed round's GTM allocation (Question 14) funds the solutions-engineering capacity and the structured integrator-onboarding program that let us convert the waitlist from passive to active in volume. So the sequence is: harden onboarding with a few expert integrators now → launch the structured onboarding program → convert the waitlist at scale. The 100+ is the pipeline that program is built to unlock, not a claim of 100+ active partners today.

Sources. Data room: Forum Interest Report (dataroom.bleedio.com → 06_Market_&*Traction → BleedIO Tech Forum Interest Report.pdf) (100+ system integrators on waitlist), BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product&Technology → BleedIO Tech Q&A.pdf) (integrator network / distribution model), Partner Integration Architecture (dataroom.bleedio.com → 05_Product*&_Technology → BleedIO Tech Partner Integration Architecture.pdf). Complementary channel positioning per Questions 18–19; onboarding capacity funded per Question 14. Exact active/in-process partner count per BleedIO management.

Team

28. What is the plan to distribute Stan's CEO, CTO, and IP-holder responsibilities as the company scales?

> **Question as asked:** "Stan, you are simultaneously CEO, CTO, and the primary IP holder. What is the plan to distribute these responsibilities as the company scales? Are there specific engineering or technical leadership hires planned in the seed round?"

This is exactly why we are raising. We have already laid down the pipeline (see Questions 20, 24, 25, 27). The constraint is not demand and it is not strategy — it is **workforce to execute against the pipeline we have built.** The seed round is fundamentally about adding the people who let Stan hand off the hats he is currently wearing and focus on where he is irreplaceable.

Where things stand today — candidly. Stan currently carries multiple roles at once, which is a normal early-startup reality but also the organizational gap we are addressing head-on:

- **Engineering:** we have **2 developers on frontend and backend (Stan and Rodion)** plus **1 firmware developer on a shared basis.** That is a very lean engineering team for the commercial product we need to ship.
- **Finance & operations:** Stan also carries **CFO and COO** responsibilities. Stan and Lee both hold MBAs and can cover finance to a degree, but specialized financial expertise is preferable to founders wearing the CFO hat.

We are not going to pretend this is anything other than what it is: too few people, too many hats. The point of the round is to fix it.

The immediate hiring plan. The seed round funds a specific, prioritized set of hires that redistribute Stan's load and let us execute the pipeline:

Hire	Purpose	Notes
Operations / implementation team (start with 1)	Pre-sales and post-sales delivery; owns customer deployments	Ideally scales to ≈1 person per large account (e.g. Lufthansa, Snap-on)
+1 Sales	Additional sales capacity beyond the CRO	Our Chief of Revenue is flying ≈90% of the time; we need more closing capacity
+2 Developers	Commercial-product delivery, bug-fixing, ongoing engineering	**One at team-lead / CTO level** — begins taking technical leadership off Stan
Fractional CFO	Specialized financial leadership	Replaces founders-as-CFO; fractional first, given stage

This directly offloads Stan: the developers (especially the lead/CTO-level hire) take day-to-day engineering; the operations hire takes deployment/delivery (the COO-flavored load); the fractional CFO takes finance. What Stan keeps is where a founder's leverage is highest.

How AI leverage stretches the lean team — and where it stops. We are heavy, sophisticated users of LLMs and agents, and it materially extends what a small team can do:

- We built an **internal agent that tracks our documents, drafts emails, and suggests follow-ups** — in fact, this due-diligence response was written with the help of that agent.
- We run a **complex coding-agent setup** that accelerates development.

But we are clear-eyed about the limits: **the best LLM cannot follow a live conversation and close a sale, and it can help *with* development but cannot *run* development.** AI lets us punch above our headcount; it does not remove the need for the operations, sales, and senior-engineering hires above. We say that plainly rather than claim AI eliminates the hiring need.

Stan's end-state role. As the company scales, Stan settles into **CEO and IP holder, owning software design, architecture, and product direction.** He will **not be involved in day-to-day software development** — that moves to the engineering team and the team-lead/CTO-level hire, maturing toward a formal CTO separation as we approach Series A. The IP-holder function is already structurally clean (all IP is assigned to BleedIO Tech, Inc. — see Question 3), so "distributing the IP-holder role" is really about deepening the engineering bench beneath Stan rather than reassigning ownership.

The through-line: the founder-concentration the question probes is real, acknowledged, and is precisely the gap this round closes. The pipeline exists; the raise buys the workforce to convert it, and to move Stan from doing four jobs to doing the one only he can do.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf) (organizational gap: Stan carries CEO + product/engineering; seed-funded team + CTO separation), Team Overview (dataroom.bleedio.com → 07_Team → BleedIO Tech Team Overview.pdf), Management Bios (dataroom.bleedio.com → 07_Team → BleedIO Tech Management Bios 2026-05-15.pdf). Use-of-proceeds funding these hires per Question 14; IP ownership structure per Question 3; sales-hire specifics per Question 29. Current team composition and immediate hiring plan per BleedIO management.

29. What sales roles, timing, and hiring criteria are planned from the seed round?

> **Question as asked:** "You acknowledged you are slow-walking pipeline opportunities to avoid over-promising. With $300K allocated to sales hires in the seed round, what are the specific roles, timing, and hiring criteria? How quickly can you realistically onboard a sales team that can close enterprise deals?"

The seed round's 60% go-to-market allocation ($540K — see Question 14) funds the customer-facing hires that let us execute the pipeline we have already built. This is the GTM-specific view of the broader hiring plan in Question 28.

The sales / go-to-market roles:

Role	Timing	Why now	Hiring criteria
Operations / implementation (pre- and post-sales), start with 1	First GTM hire (near-term)	Our Chief of Revenue and founders are carrying deployment and customer delivery on top of everything else; this role owns it	Experienced commercial deployment / solutions-engineering background; can own a large account's delivery end-to-end. Ideally scale to ≈1 per large account (Lufthansa, Snap-on)
Additional salesperson (+1)	Near-term, alongside operations	Our **Chief of Revenue is flying ≈90% of the time** — closing capacity is the bottleneck on converting pipeline	Enterprise/industrial sales experience; able to run a full cycle in our verticals; comfortable in a technical, complementary-infrastructure sale (Questions 18–19)

Why these roles, and why in this order. The near-term revenue engine is conversion, not lead generation — we have pipeline (Questions 20, 24, 27); we need capacity to *deploy and close* it. So the first GTM hires are (1) an operations/implementation person to take delivery off the CRO and founders, and (2) an additional closer to relieve a CRO who is already at ≈90% travel. The operations role also unlocks the channel-partner onboarding program (Question 27) and reduces founder time on delivery, which is the same load-shedding logic as Question 28.

Timing and pace. These are near-term seed hires (funded on close), sequenced ahead of the deeper 2028+ scaling because the base-case breakeven (Question 13) depends on converting the current pipeline, and that conversion is capacity-limited today. We onboard deliberately — a small number of high-quality hires who can operate in a technical, multi-stakeholder industrial/public-sector sale — rather than scaling a sales team faster than we can support it.

On how fast a sales team can close enterprise deals. We are realistic: enterprise/industrial cycles run longer (Chevron-class procurement), while the public-sector/fire-department cycle is short (1–2 months — Question 20). So the sales hires are pointed first at the faster-converting public-sector pipeline and at supporting the already-advanced Snap-on and Lufthansa motions (Question 24), where the relationships and next steps already exist, rather than starting cold. That is how a small, new sales team produces near-term closes rather than waiting out long cold cycles.

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&*Technology* → *BleedIO Tech Q&A.pdf) (GTM allocation, solutions-engineer hire, distribution model), Sales Pipeline & Revenue Forecast (dataroom.bleedio.com → 06_Market*&_Traction → BleedIO Tech Sales Pipeline & Revenue Forecast.pdf), Management Bios (dataroom.bleedio.com → 07_Team → BleedIO Tech Management Bios 2026-05-15.pdf). Use-of-proceeds per Question 14; sales-cycle length per Question 20; account next-steps per Question 24; full hiring picture per Question 28. Specific roles/timing per BleedIO management.

30. What will the organizational structure look like 12 months after seed close?

> **Question as asked:** "What does the organizational structure look like 12 months from now, assuming the seed round closes? Who are the first 3–5 hires, in priority order?"

Where we start. Today the team is **5 cofounders** with a very lean execution bench: **2 developers on frontend/backend (Stan and Rodion), plus 1 firmware developer on a shared basis**, with

Stan also carrying CFO and COO (see Questions 28 and 32). Twelve months after the seed close, the goal is a company that has shed Stan's excess hats and has dedicated capacity in engineering, delivery/operations, sales, and finance.

The first 3–5 hires, in priority order. These follow directly from the load-shedding logic in Questions 28–29:

1. **Operations / implementation lead (pre- and post-sales).** First hire. Owns customer deployment and delivery — taking the COO-flavored delivery load off Stan and the CRO. Unblocks large-account execution (Lufthansa, Snap-on) and the channel-onboarding program (Question 27).
2. **Senior developer — team-lead / CTO-level.** Begins taking day-to-day engineering leadership off Stan; the seed-stage seed of the eventual CTO role. Highest-leverage engineering hire.
3. **Second developer.** Adds commercial-product delivery and bug-fixing capacity so the roadmap (and the Chevron accuracy/UI-UX work — Question 11) moves without Stan in the critical path.
4. **Additional salesperson.** Relieves a Chief of Revenue who is traveling ≈90% of the time; adds closing capacity against the pipeline (Questions 20, 24).
5. **Fractional CFO.** Specialized finance leadership, replacing founders-as-CFO. Fractional first, given stage.

Priority order reflects the binding constraints: delivery capacity and senior engineering come first (they unblock revenue *and* offload Stan most), then additional closing capacity, then finance specialization.

The org 12 months out. The resulting structure is a still-lean but role-complete organization — roughly a doubling of headcount into defined functions rather than founders-wearing-many-hats:

Function	Today	≈12 months post-seed
Engineering	2 devs (Stan + Rodion) + 1 shared firmware	Stan (architecture/product) + team-lead/CTO-level + Rodion + new dev + firmware; Stan out of day-to-day dev
Operations / delivery	Founders (Stan/COO hat)	Dedicated operations/implementation lead (scaling toward 1 per large account)
Sales / GTM	CRO (Brad) + founders	CRO + additional salesperson; supported by the operations hire
Finance	Stan/Lee (MBAs, part-time)	Fractional CFO + founder oversight
Leadership	Stan (CEO/CTO/CFO/COO), Lee (CMO), Brad (CRO), Mikhail (HW/Eng), Rodion (Eng/co-inventor)	Stan (CEO, IP holder, product/architecture); roles distributed; CTO separation maturing toward Series A

This maps to our planned FTE growth (from a lean ≈2026 base toward a low-double-digit team by 2027) and stays deliberately capital-efficient — we scale through channel/OEM partners and AI leverage (Question 28) rather than proportional headcount. The defining change is qualitative, not just numeric: **Stan moves from doing four jobs to doing the one only he can do (CEO, IP, product/architecture)**, with dedicated owners for engineering delivery, operations, sales, and finance.

Series A trajectory. Beyond the 12-month window, Series A funds the next layer: expanding the sales team, dedicated customer success, and the formal CTO hire that completes the separation of Stan's product/engineering role from the CEO role (Question 16).

Sources. Data room: BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf) (hire sequence: solutions engineer Q2 2026, 2–3 engineers Q2–Q3, sales/BD Q3; post-Series A CTO), Financial Projections & Benchmark Analysis (dataroom.bleedio.com → 02_Financials → BleedIO Tech Financial Projections & Benchmark Analysis 2026-05-23.pdf) (FTE headcount plan),

Team Overview (dataroom.bleedio.com → 07_Team → BleedIO Tech Team Overview.pdf). Full hiring rationale per Questions 28–29; use-of-proceeds per Question 14; Series A per Question 16. Current team and priority order per BleedIO management.

31. Which advisors, board members, or domain experts is BleedIO targeting?

> **Question as asked:** "Are there key advisors, board members, or domain experts you are targeting to onboard? The pitch deck mentions advisor onboarding as a near-term milestone. Who specifically, and in what capacity?"

We are actively recruiting advisors, and advisor onboarding is a near-term milestone as the pitch deck notes. Current status: **two advisors signed and onboarded, both actively serving.**

Advisors (signed and onboarded):

- **Bill Goodwin — onboarded (design / ecosystem).** Bill is local and active in the Philadelphia startup ecosystem, well-networked, and a design artist by trade. In practice he has already contributed the way we want advisors to contribute: **introductions to investors and to commercial opportunities.** He is the model for the "foundational" advisor relationship — network access plus domain input (design).
- **Guylain Roy MacHabée — onboarded and actively serving (RTLS / wireless).** Guylain is a Strategic Advisor under an executed Consulting Agreement and NSO grant (**24,030 NSO shares; 4-year vest, 1-year cliff; granted March 25, 2026**), and he is actively serving in the advisor role. (We quote option grants as share counts rather than percentages: option-grant percentage calculations use a different capitalization basis than the common-stock issuance percentages elsewhere in the data room, so share counts are the unambiguous figure.) His distinctive value is **deep experience in RTLS and other wireless endeavors**, directly relevant to locMESH/netMESH, alongside the foundational contributions — **access to investors and commercial opportunities.** He has already contributed substantively: much of the strategic/technical framing sharpening in our materials (e.g. the DRAW-vs-open-standard moat articulation and the brownfield/complementary positioning reflected in Questions 4, 6, and 18) came from his review.

What we are targeting next. With **design and technical/RTLS advisors now in place**, the next set of advisors we are recruiting should bring **commercial and defense networks for partnership** — people who can open doors in our industrial/enterprise commercial channels and in the defense/autonomy vertical (the Draganfly-type opportunities in Question 25 and the 2028+ defense scaling in Question 20). In other words, we are sequencing advisor recruitment by capability gap: design and technical first (in hand), commercial and defense-network advisors next.

In what capacity. These are advisory relationships (not board seats), providing network access, domain expertise, and commercial/investor introductions. Advisors are compensated through the equity/advisor framework under the 2024 Stock Option Plan (which expressly permits advisor grants). Board composition remains founder-controlled at this stage; formal independent board additions would come with the Series A.

Sources. Data room: Team Overview (dataroom.bleedio.com → 07_Team → BleedIO Tech Team Overview.pdf), BleedIO Tech Pitch Deck (dataroom.bleedio.com → BleedIO Tech Pitch Deck.pdf) (advisor onboarding milestone), 2024 Stock Option Plan (dataroom.bleedio.com → 03_Legal_&_Governance → BleedIO Tech 2024 Stock Option Plan.pdf) (advisor equity mechanism). Guylain Roy MacHabée strategic-advisor terms (24,030 NSO / 4yr-1yr cliff, Consulting + NSO Agreements executed March 25, 2026) per the executed advisor agreements on file. Bill Goodwin status and next-advisor (commercial/defense-network) targeting per BleedIO management. Defense-network targeting rationale per Questions 20 and 25.

32. What is each cofounder's current employment status, time commitment, and outside obligation profile?

> **Question as asked:** "Confirm cofounder's current employment status, time commitment to BleedIO, and any outside roles or obligations that survive the seed close."

All cofounders are full-time on BleedIO, with zero competing outside roles or obligations that would survive the seed close. We can state that cleanly for every cofounder:

Cofounder	Role	Status	Time commitment
Stanislav (Stan) Podolski	CEO / Cofounder	Full-time	100% BleedIO (currently also carrying CFO/COO/engineering — see Question 28)
Mikhail (Mike) Kolobov	Cofounder / Hardware & Engineering	Full-time	100% BleedIO (Prox SG relationship noted below)
Lee (Leonardo) Beup	Cofounder / CMO	Full-time	100% BleedIO
Brad Niems	Cofounder / CRO	Full-time	100% BleedIO (traveling ≈90% on sales — see Question 29)
Rodion Rodionov	Cofounder / Engineering (co-inventor)	Full-time	100% BleedIO

There are **no competing employment arrangements, side ventures, or conflicting outside obligations that persist past the seed close.** (The one related-party relationship — Mike and Prox SG, our own hardware partner — is aligned and disclosed below.)

On the equity structure (for clarity). BleedIO has **five cofounders**, in two groups. **Stan, Mike, and Lee are the original founders** and hold founder common stock. **Brad and Rodion are adopted cofounders** — they joined after formation and were brought in as cofounders — which is why their equity is structured as **options** (4-year vest, 1-year cliff) rather than founder common stock. This is the expected, clean explanation for why two of the five cofounders hold options: it reflects when and how they joined, not a difference in commitment. All five are full-time.

The one relationship to disclose clearly — Mike and Prox SG. Mike shares responsibilities with **Prox SG**, and we want to be fully transparent about it because it is a related-party relationship — but an *aligned* one, not a conflict:

- **Prox SG is BleedIO's hardware partner** — it builds the prototypes/hardware for BleedIO Tech. Mike's Prox SG involvement is therefore in service of BleedIO's own hardware development, not a competing external job.
- **The relationship is disclosed and structured**, not informal: Prox SG is a **pre-seed investor** in BleedIO (a $100K SAFE at a $2M cap, July 2024), and the Prox SG ↔ BleedIO arrangement is governed by **IP-licensing and invention-filing/licensing agreements** on file in the data room. IP created for BleedIO is assigned to BleedIO Tech, Inc. (see Question 3), so the hardware-partner relationship does not create IP-ownership ambiguity.
- **Net effect:** rather than an outside obligation that competes for Mike's time, the Prox SG relationship is a strategic advantage — an aligned, invested hardware partner led by a technical cofounder, with the commercial and IP terms already papered.

Summary. All cofounders full-time, 100% committed, no surviving outside obligations. The only related-party relationship (Mike/Prox SG) is fully disclosed, contractually structured, aligned with BleedIO's interests (Prox SG is our HW prototyping partner and a pre-seed investor), and clean on IP. We are happy to make the Prox SG SAFE and the IP-licensing/invention agreements available for confirmatory diligence.

Sources. Data room: Management Bios (dataroom.bleedio.com → 07_Team → BleedIO Tech Management Bios 2026-05-15.pdf), Team Overview (dataroom.bleedio.com → 07_Team → BleedIO Tech Team Overview.pdf) (full-time cofounders), BleedIO Tech Q&A (dataroom.bleedio.com → 05_Product_&_Technology → BleedIO Tech Q&A.pdf) (Prox SG SAFE $100K/$2M cap; Mike = Founder/CEO Prox SG; Rodion Rodionov co-inventor, holds options). Prox SG IP-licensing and invention-filing/licensing

agreements available in the data room (Legal & Governance / Customer Contracts). IP assignment to BleedIO Tech, Inc. per Question 3.