Scott Sanders, chief growth officer at Forterra, on the defense tech startup playbook

TEAM

Jan-Erik Asplund
Co-Founder
jan@sacra.com

DISCLAIMERS

Scott Sanders, chief growth officer at Forterra, on the defense tech startup playbook

By **Jan-Erik Asplund**

EXPERT INTERVIEW

Scott Sanders

Chief Growth Officer
Forterra

Background

Scott Sanders is chief growth officer at Forterra (previously RRAI) and he was an early employee at Anduril. We talked to Scott to better understand how defense tech startups innovate on business model in contrast to incumbent defense primes, how companies like RRAI gain leverage through their product-centric approach, and how startups today can sell into and navigate the DoD's procurement processes.

Interview

What is RRAI and what inspired you to join the company?

After college, I was commissioned as a ground intelligence officer in the Marine Corps. I did five tours to combat zones, and became a Marine Special Operator. That's where the story starts.

When I was in Special Operations Forces (SOF), one of the things that was happening in the geopolitical scene was the war in Iraq and Afghanistan. One of the big things that the SOF Enterprise would do is partner with the indigenous forces and go with them and do stuff. Around 2014, SOF was losing the authorities to directly partner with local forces on the ground as there was no tolerance for US casualties.

If you think about the equipment that the SOF enterprise had, it was to remotely work with local forces. Just imagine teaching your parents to reset a WiFi router over the phone. Now, think about doing that with bombs in a foreign country with limited communication and a language barrier. I could barely get it done with my English-speaking parents who are somewhat tech-savvy.

So, you have this huge problem of "Okay, I can't go with you to combat anymore, but you need me to have the same effect."

When we were gearing up for our first deployment to a new mission in the Philippines, we weren't going to get the authorities to actually go with our force. We said, "How do we solve this?"

We didn't know what we were doing when we started, but what we did is we built what is now a program of record inside of SOCOM called the Remote Advise and and Assist Virtual Accompaniment Kit (RAA-VAK) which is based on the Governments Team Assault Kit (TAK) software system. It's a bunch of COTS and GOTS tooling and some home-brewed software that enables forces to digitally partner with their partner forces.

I didn't think a whole lot of it at the time. When I got out of the military, I was done. I'd done nine and a half years, and I really wanted to go work in tech. Through the process of trying to figure out where I wanted to work, I came across Anduril Industries.

Around the first time I talked to Trae Stephens, Anduril's co-founder, the company was six people—the five founders and one engineer.

I was like, "Oh! I didn't want to work in defense ever again. I had no desire to work in the DoD at all." Brian and Palmer and

the crew were like, "We're going to change how the technology gets delivered to operators, soldiers, and Marines. Do you want to come along for the ride?" I was like, "Okay - this sounds awesome."

I joined there as a very early employee and then, did our first defense contract at Anduril—a $12.5 million contract with the Marine Corps—and then, continued to go on and build some of the biggest programs of record that they have. It was a great experience and it's an awesome team over there.

We actually work with them pretty closely here, because in order for CJADC2 or any of the battlefield things to actually work, you need autonomous ground vehicles. It's a great relationship that we have with them and we look forward to working with them in the future.

You started your career in defense tech at Anduril. What were the critical elements of Anduril's success in selling to the Pentagon as a startup founded by someone who was from the VR/gaming space?

I'd say when I came in they had this mission from day one. Palmer actually said it very well in one of podcasts he did, that "Anduril is built on a mission to change the way technology is fielded because it matters."

Now, it's great if you create a new VR headset or you develop the next big SaaS application, but you're not going to make the world a fundamentally different place because of that. The defense problem was something that Palmer's cared about since he got started in VR. He started doing that in the lab space before he even started Oculus. Really, it was like, "This is the mission; this is what we're going to do." It attracted a lot of people who cared about that mission.

It's very similar to what drives people here at RRAI. We see these problems on the battlefield and the tech adoption by the defense community and traditional acquisition method is just so slow. People's lives are impacted by that. If you're slow to field autonomous transportation vehicles in the DoD, the consequence is that humans die, or they lose their legs, or they have multiple compressed vertebrae in their spines.

I don't think I could get up every day and go to work at a Salesforce or a SaaS company no matter how much money I

was paid because I just don't care. Here, we care and we attract people who care. I think Anduril is the same way. People are driven there by the mission. Yeah, it's great to be an early employee in a venture-backed company, but no one was sitting there checking their Carta account when they're working fifteen-hour days like, "Ah! One day this might be worth something." Hit refresh. "Did we get a new 409A?" No one cared.

We were thinking about the impact we could have on people and how we could maintain Western democracy as the de facto best version of government. Western democracy has flaws, but it's pretty good. I like living in America—I'm a big fan.

It's the same thing that we believe here, and we can speed this up by taking a much different vector into the DoD space than a traditional company can and we can just move so much faster, because we're smaller. Even now, they have 2,500 people today. That's still a fraction of the size of a major defense prime.

Can you give us a high-level breakdown of how a new defense technology gets from the R&D stage to being deployed in the real world?

Anytime you want to build a new aircraft carrier, new submarine, or some big capital project, the DoD has a budget system—the Planning, Programming, Budgeting & Execution Progress (PPBE) and Joint Capabilities Integration and Development System (JCIDS) and all this structure—which is a centrally managed system which is objectively communist in nature to create innovation.

The (Federal Acquisition Regulation) FAR exists because there's not a market. If you have a market, the market sets prices, if you have a single buyer, a single seller, a monopsony and monopoly relationship, it's really hard to determine a fair price for both parties. So, the FAR is actually constructed to be fair for both, and you have to do that because that way the government can't extract too much value from the commercial side and vice versa.

So, when you build those big systems and those big projects, you have to do it in a semi-deliberate way. JCIDS and PPBE

makes sense for building an aircraft carrier. If you built the wrong thing, you just plowed a couple billion dollars into a thing you're going to scrap. We still do that as a country, but it wasn't because the problem wasn't analyzed enough. When you look at innovative technology, autonomy systems, battlefield networking, radios, these are things that don't require billions of dollars to build the first prototype. They require tens of millions of dollars or hundreds of millions of dollars.

You still have to work through the process of getting an approved requirement, getting a program executive officer and a PM to drive these things through. It's really not that different from early adopter sales into commercial enterprise businesses. You're going to have the same problems of getting people brought on. It's just different when you have Congress as this other entity that has to exist to fund those projects.

We run the same playbook here as we did at Anduril which hinges on getting a different vector into the space. It's very similar to SpaceX—you have to have a high conviction that not only is the thing you're building a step change better or a step change lower cost, but you have to be able to build it and have the conviction to build it, get it fielded, and get it tested. You still have to work through all those elements in some fashion without waiting for the DoD to say, "Hey! I need this thing. I need a new widget. Here's the performance specification of what I want. Go build exactly what I tell you." That's a slow five to six year on-ramp period to get that technology in the field.

It's very different when you show up and say, "You need a ground autonomy solution? Great! We have this application kit you can integrate on the factory line today. We've invested hundreds of millions of dollars in building it. This is the price. Also, we sell this commercially so you can follow commercial pricing rules from FAR 12 or FAR 15."

There's about 15 DoD statutes and parts of US Code Title 10 that instruct the DoD on how to do procurement, and you can show up and you can eliminate really the first three or four years of that tail where the government's paying you to develop it and they're telling you exactly what to build.

The best example is SpaceX vs. United Launch Alliance. Both have a heavy lift mission. Who's actually launching rockets that are doing something today? You can knock SpaceX all you

want for some of those things blowing up, but how many Artemis missions have flown and how much did it cost? Those were 10X the cost, and that was Congressionally-directed technology.

You probably don't want Congress telling you how to build your tech.. The second you're going to tell NASA what technology to use, that's crazy.

So you can't go build things on spec that you don't have high conviction about. Whether it was here or Anduril, we have had very high conviction that the products we're building are meaningfully different—a dramatic technology change or dramatically cheaper than what's out there. You can't just build a better mousetrap.

We see this with Google and Microsoft trying to compete to own the DOD email market. Google's great, but they haven't made a dent in Microsoft's ownership of the DoD market for enterprise email. It's not even close.

If there was a company which has the dry powder to go attack that problem, it's Google, and DoD is the single biggest buyer of IT solutions in the US. So, you'd think that they'd be able to go make a bigger dent in that market.

DoD sales are a 4D chess game, where you have to have users who want your product. You have to have technology that's very different from what the current incumbent has. You have to map requirements. You have to get program executive officers on board with it. You have to work through the solicitation processes. You have to deliver that technology ahead of your competitors. You have to get Congress to buy off on funding that over the long-term, and you have to get all those things working together against a cohesive capture strategy to deliver that technology to the field.

If you and I got together and we were the best engineers in the world, and we built a fifth-generation fighter that cost $5 million bucks a print, and we parked it in front of the Pentagon and did a test flight, it would still take them two years to buy it. You just can't have a better product. You also have to know how the entire system works, and it's a very complicated system.

There seem to be roughly two different business models that modern defense startups are using to innovate

around the cost-plus contract—the SpaceX/Relativity Space approach of heavy vertical integration, and the Anduril platform approach. Does that map to reality in your experience? How might you position it? Where does RRAI fit?

The good examples are SpaceX and Palantir. They took about five to seven years to get their first substantial contract of more than a million dollars. They had patient capital. They had billionaire founders. They could play that game.

At Anduril, we were very lucky that Palantir had already broken down a lot of the walls, so we were able to get our first sales within the six months of founding. I remember Matt Grimm screenshotting the cash in our checking account and going, "We got traction." It was like $30,000 or something. It was nothing. But it mattered.

Then we got our first defense contract for $12.5 million about a year into the company existing, and that was huge. That was a step change. That allowed us to demonstrate that this approach worked.

We're similar here where we have this Generation 3 autonomy kit, which is an integrated hardware-software stack product that enabled us to get on four different programs of record within two years of launching this strategy.

Our company is very different. We've existed for a long period of time, but mostly as an applied research company. We pivoted the entire business about two years ago to go single product, multi-market, multi-application, and to go generate that traction very quickly through running a very similar strategy playbook of working closely with those programs that have a need for it and have funding against either legacy programs or incoming programs like Robotic Combat Vehicle (RCV).

But to your point, if we were to go start a ground autonomy company today, it would probably take you five to seven years to get your first prototype in the field. There's not that many investors who are willing to wait that long. Companies like Hermeus are in a similar position. They're not flying hypersonic jets right now, but they've got to generate revenue and there's ways to do that through close cooperation with the DoD and giving them access to certain data sets or simulation training

and things like that. We're lucky that we don't have to do that here, but it is a very hard road. It's why I'm actually not bullish on defense tech.

I think there's very few companies that can pull this off and generate actual meaningful traction. SBIRs don't count. They get you almost nothing. Operations and Maintenance (O&M) dollars don't get you very much either. You've got to be targeting these big programs and you have to have the capital or the recurring cash flow from other activities to take that two-year process to get a program of record. It's very hard. It's significantly harder than enterprise sales.

At Anduril, we built the first prototype in a couple months and were able to capture the first order very quickly. We iterated against that second product multiple times. When we got the order from the Marine Corps, we shipped like 30 days later.

We made a big bet in the same way we did here at RRAI on ground autonomy—all these programs are up and running. If you have a product you can ship today into testing, you have a huge advantage with the customer versus, "Here are my PowerPoint slides and my Excel breakdown of how long it'll take and how much it'll cost to build the first version."

It's very different when you do a live demo. It's like, "Here it is. It's on a truck. It works now. Is it perfect? No, it's still going to take some work."

In our case, we have such a head start and that's what you have to prove to be able to get those first contracts. So you've got to bet the farm on your first product.

Was there a particular strategy to the southern border surveillance as the introductory product or the wedge or was that incidental?

Legend has it that when they went out to CBP and said, "What's your biggest problem? Give us your biggest problem." Then,the founders and engineering ideated and said, "I think we can solve this for a dramatically lower cost."

If you do some research on Boeing's Secure Border Initiative Network (SBI-Net), it was a $9 billion project to effectively try to do the same thing that Anduril did. SBI-Net covered three miles for about a billion dollars with still radars, cameras, a pan

tilt unit, and some computers. Anduril's Lattice and Sentry Tower systems were not dramatically different from a sensor methodology several years later.

The difference with Anduril was that six years later, a thing called a GPU existed that you could put in a box.

If it wasn't for the gaming community, that wouldn't have been possible. So, Anduril was able to deliver a product to a customer who had a burning need, had the funding to go solve it, and was willing to put a product into prototyping very quickly to address that need. I think that's the big difference between an Anduril versus a Shield or Varta or these other companies is that Anduril had a product from day one that had product market fit on a recurring license basis with a very motivated customer. That's just a hard playbook to run. You better be right on your first product.

With RRAI moving toward a multi-product approach, does it imply a shift towards doing products vs. research? Does it have anything to do with larger shifts in the space that maybe Anduril and SpaceX helped pave the way for startups to more easily sell into the government, DoD, etc.?

When I started in defense tech, there were two startups doing it.

The fact that SpaceX and Palantir both sued their customers to enforce US law was pretty helpful. If you didn't have a Palantir and a SpaceX, you wouldn't have an Anduril. If you didn't have an Anduril, you wouldn't have defense tech today.

They didn't change any of the laws. They just compelled the DoD to actually follow the ones they already had and also showed it would work.

Silicon Valley has its roots in defense. It's not crazy that there's a pivot back to it. It was also very controversial in 2017 and 2018 to be working on defense given the political climate in the US. Now, it's cool. It's wild to me. That one shifted pretty quick.

But the reason we're taking a product approach is you want to run a scale business with high gross margins. The only way you do that is by delivering a scale product on a recurring

revenue basis where your cost basis doesn't change or ideally, gets much lower with scale.

It's no different than SaaS. You just have an element of hardware associated with it. If you're trying to scale a company just on research, you're going to get paid effectively by the hour, on a cost plus contract or a time-materials contract.

So, you look at a Lockheed, a BAE, an Oshkosh Defense—any large defense prime—their contribution margins are extremely low, because they're saying, "Hey! Government. We'll do this research or build this system for you. Here's the amount of people working on it. The only way you scale that is through massive amounts of headcount."

It's also the reason why there is no incentive for a defense prime to deliver faster, or at lower costs, because it's just a fixed amount of margin—8 to 10% on the total billing of the contract. It's obviously oversimplified, but that's effectively the math. You can see that in SEC filings for every one of these companies. They only can scale by increasing the size and duration of contracts. There is no incentive to move faster or deliver a lower cost product.

That means they just have to get bigger and bigger and unless you have the ability to go scale a company for 40 years to try to compete with Lockheed, it's never going to work on an applied research basis. You're going to have to go build a product that generates cash flow and then scale into it.

That's a way to think about why to take this approach, because there's tens of thousands of small defense contractors doing applied research. The odds of any of them getting as big as a Lockheed is zero. You're just never going to be able to grow fast enough to do that.

Given the long timelines in contracting and the capital requirements of doing R&D upfront, how do companies in this space manage the capital risk involved in building a product before selling it?

It comes down to risk. We've poured $300 million in investment into building our stack at RRAI. Anduril has probably put $2.5 billion into developing theirs and their systems.

We're a little different in that we actually sell our systems commercially. So, we sell a commercial item, and by law, we have to sell it at the same price to the DoD that we sell it commercially and vice versa. The rate is the rate.

If it costs us $1 or $100,000 to deliver something that costs $100,000, that's on us. That's the risk that we're taking. If they don't think it's a fair price, they don't have to buy it.

When you're doing cost plus contracting, you have to show every nut and every bolt and what everything costs. Then you get a fixed margin on top of it. There's two ways to not build products like this in defense.

At RRAI we sell our products as commercial items because we have a commercial business. There's another way to do this called non-developmental items. That's more like what Anduril and some of the new tech startups do where they're not actually selling it commercially, because you have to have proof of commercial sales to do a fair and reasonable assessment. You can still be a commercial item, but the pricing justification is actually separate from commerciality.

Non-developmental items generally apply the same rules, where it's entirely funded under private research and development and then you negotiate on price like you would anywhere else. So, it's much more aligned with the commercial space where your iPhone's a good example. It's like $110 in parts. You pay $1,300 for it, because that's the market price. That's what the market supports. It doesn't mean you get to price gouge—it's actually the opposite.

The reality though is unless the DoD is going to pay for all that research and therefore have all the rights to it, which they're likely not going to do, then you get to price it appropriately and your gross margins follow.

Can you talk about how the business models of companies like Anduril and SpaceX have changed how the government thinks about procurement and incentives? Where are traditional contractors threatened? Where are they most durable given this shift?

Anduril's trying to build a software platform that then has multiple products that feed off of it. When Anduril started

working on counter-UAS, it was just a ground tower with the mast now pointing up but using very similar software.
We didn't have to go recreate a whole new thing to start that vertical. It's why having your first product be a very key fundamental thing like a sensor fusion platform is pretty critical because you can go do that in other places.

I argue there's a third model on defense tech, which is a tightly integrated product that enables a platform versus being the platform itself.

Our long-term goal is not to be an Anduril. What we want to do is be able to capture as much of the broader autonomy market as possible, but doing it through the serviceable markets today such as defense.

There is no serviceable market today for middle mile autonomy. No one's taking the driver out in fleet operations anytime soon. Robo-taxis has recently proven my thesis on this that the public isn't ready for that technology. I believe robo-taxis to be much safer than a human driving down the road and texting, but the public isn't ready for it.

Our approach is about building a platform in the autonomy space.

Our platform is integrated hardware and software that enables it to connect into any battlefield C2 system, our own C2 system, commercial C2 systems for yard operations, things like that.

Other people can build on top of that platform, because we're fielding it before anyone else's, but it's a tightly coupled product that enables a larger network effect. We're taking a much more horizontal approach to that market.

With the way we build our autonomy systems, if a DoD customer shows up tomorrow and says they need a thousand of these things on a Joint Light Tactical Vehicle (JLTV) in the next year, we can go do that, because we have a unified core product base that is repeatable and scalable.

If the DoD never talked to us again and a yard customer showed up and said, "I want a thousand," we would take the exact same product and do it on a yard truck.

So we can address multiple markets and multiple go-to-market strategies with a single platform that enables them to build on top of it. It's more of a hybrid approach. Our goal is actually to sell it for a dramatically lower price than what the DoD would do for their own research and development.

For comparison, our system is about a tenth of the cost of the DoD's own funded internal ground autonomy platform on a per vehicle basis. The joke I always tell Congress is, "I'm the US Government. I show up to your house, and I tell you, I'm going to lower your taxes or build you a self-driving car. I'd probably believe you on the taxes." It is just crazy that the government thinks they can compete in enterprise software, because they've proven they can't. The cost basis that you can find through public documentation justifies that.

You're spending 10 times as much for a thing that's never been fielded. The big advantage to a commercial product in the defense space is that while you still have to listen to your customers to build the thing that they want, you can make driven engineering choices to say exactly what you want to do to get the outcome you need.

That's taken the chains off of the government S&T process where they want to direct industry to do a very specific thing. That's really what SpaceX and what we've done is prove that you can do that and you can get a product to fielding and you can work closely with your government partners to get to that outcome.

The horizontal platform is the way we think about it, but you want to create SaaS-like margin profiles with the thing that involves hardware and the only way to do that is to have very low hardware costs.

Can you talk about margins in the defense contracting world? Is there a world where companies like Anduril can get to 40-50% margins selling into the government—particularly selling recurring software subscriptions—or does that require, for instance, selling into commercial use cases as well?

Part of it is anchoring around that kit which we sell for a firm fixed price per year and it's effectively a license or a lease. It's a SaaS-style approach. Then governments are like, "Hey! We

want you to improve the ability for that to work in GPS environments" or "We want you to have it do mine clearance" or something you just wouldn't do in the commercial space. Then they can pay you to test and develop that. They can say, "Hey, can you integrate your AutoDrive autonomy stack with another C2 system?" There's a bunch of things that you can create and our logic is that the power of the platform is that others will build value on top of it.

We want to enable—on the DoD's side—the weapon systems to have more precision when they're launched and we want to make sure that we can integrate with the next vehicle type. There's all these other things that go along with it instead of just selling straight licenses.

What we don't want to do though is create an upsell mentality with the product of "Oh! You've got the basic vehicle autonomy, but you really want a super advanced version" We want to have the best thing in the field at all times that is the most capable version of the product. We want our products to be sticky because of performance, not black box vendor lock.has to be better than anyone else's thing, at the most comparable price, and you can only do that through scale.

That's why we've taken this horizontal platform play because that's the only way to do it in this space to get the scale you need. If you just do yard trucks or you just do on-road, you are completely wed to the market dynamics that you are entering. Whereas "Hey! I'm glad we're on an on-road truck company right now", we could go do that. We have run trucks on the road. We don't think it's particularly hard, but the market's not there. So, we don't have to go invest in a bunch of research somewhere else to go say, "What can we do with this, next?"

There are endless applications for autonomy, from things that are the size of small drones to large ground vehicle platforms and we can attack all of them with a single product in the same way that Salesforce attacks B2B enterprise.

At RRAI, what problems have you had to solve in terms of sourcing chips, cameras, radar, LiDAR, sensors, and other components? Is there an edge to be gained through differentiated sourcing tactics or customization? We have seen autonomy companies seeking cost/scale advantages by optimizing systems to run through using consumer-

grade chips/components. Or, alternately, companies using a lot of custom components.

We're buying across multiple DoD and commercial programs that are running the exact same stack. Granted we're not even remotely close to the scale of a Tesla, but this was Tesla's approach. It is why the chip shortage didn't impact them as badly as it did Ford. They use the same chips on every single system. Similarly, we're able to buy quantities in the hundreds or thousands for stuff, whereas if we had a different product on every single vehicle, we wouldn't be able to do that.

We also had a huge advantage in that the company didn't take venture funding until 2021, but we have existed since 2002 and had been doing ground autonomy for the DoD this entire time.

All the things that are really hard edge cases, the DoD tried solving first—How do you work in the snow? How can you work in the rain? How can you find the military crest of a hill with an autonomy system? If autonomy doesn't work in all climates, in all domains, the DoD is not interested.

Look at where commercial autonomy applications are in the public sector now. No one's running Detroit as their beta testing area for robo-taxis. You have snow, you have rain, you have horrible roads. You have this 3D problem. So, we built our stack from the ground-up to be tightly coupled with that hardware because we never had the advantage of just throwing cash at a problem from a development standpoint until two years ago.

So, a lot of the time, the discipline around "You build the lowest cost system from a hardware perspective" means having very tightly coupled ASICs and GPUs and using off-the-shelf stuff where you can. We're not trying to create our own LiDARs; we don't think we need to. We know a lot of autonomy companies are trying that. It's very expensive and that enables us to have a tightly coupled supply chain. We don't have to go procure a bunch of new stuff for every vehicle we spin up. We can spin up a new platform as a beta in about 60 days from getting the platform to us. The traditional timeline for that is usually like five to six years.

So, we have this huge advantage on the tech side where we built a platform designed to do this and tightly coupled it with

low-cost hardware to have the maximum efficacy at the lowest cost. That allows us to attack a lot of markets that just wouldn't be addressable by a traditional commercial on-road company. **Can you talk about the opposed approaches of 1) vertical-specific autonomy companies like Outrider (logistics) and Kodiak (trucking), and 2) horizontal autonomy companies like RRAI (logistics + trucking + defense + other commercial use cases) and Shield AI (software for drones + jets)? What are the benefits and trade-offs of each approach?**

We raised capital in 2021. We raised a $200 million round from SoftBank really to consolidate around this single-product approach. When I started here, we had 17 legacy contracts that we had to effectively spin down and close out. Simultaneously we had to scale up product-based revenues into programs of record on the DoD side and emerging pilots on the commercial side.

I think we're a little bit different in that we can show our actual product revenue versus development revenue from an OEM. Lots of the commercial autonomy companies are doing small pilots to show traction, but they're highly negative contribution margin because you're having to babysit them without the scale to back them up.

Not everyone has access to large amounts of capital and can light a billion dollars per quarter on fire to get their systems into the field. So, it creates a different discipline, but our goal as a company is to be cash flow positive in the next 18 months off of our current product offering, because we want to be able to selectively choose, "Okay, if we're going to raise more money, what verticals are we going to accelerate to do that?" versus just using that to sustain operations.

That's a big difference and I don't think you can do that to your point with a single market vertical approach, because what happens if the bottom falls out of the freight industry like it's currently doing? Those companies pull back on investment and emerging technology and say, "Yeah, we're just going to wait it out." Those are low gross margin businesses. They don't have a lot of free cash flow just to throw at tech that's not quite ready to be completely driver-out.. There's hard safety case requirements to go after those things.

On the DoD side, you still have to build a safe platform, but they have a pressing need for the national defense to have autonomous vehicles in the same way they do in the air domain. That's why Shield AI is getting the traction that they have.

We are a country of roughly 330 million people. Our peer competitors combined, have billions. If we want to go into a shooting war, we literally couldn't draft enough people to go fight that war. You have to do it with robots. Ukraine and the scenario in Taiwan have only reinforced that. We just don't have an interest in putting US troops into combat. The only way you win that is somehow get much bigger, much quicker and draft a huge population to go create that deterrence or you get 10X the firepower out of the humans that you currently have and the only way to do that is robotics.

That has led to the drive from the department to field more robotic systems as fast as humanly possible. We're not building a new tank right now. The next programs that are coming out in this space all have an autonomy requirement such as ATV-S/GEARS, RCV, CTT and ROGUE.

ROGUE Fires just went into fielding. It's a $2.6 billion program that we're working on with our partner Oshkosh Defense to put autonomous missile launchers in the hands of Marines. There's already a substantial number of those fielded and they'll continue to field more. Other parts of the department have projects in that bucket because you just can't get enough people. We could barely recruit to the force design we have today. You certainly couldn't increase the size of the force by 10% overnight, let alone the 100% you need to increase it to go fight a real shooting war with China.

That's the difference—DoD is here today, we can go attack that. If all of a sudden there's the biggest driver shortage in the world in the yard terminal space and those companies are motivated to move quickly, we've laid the foundation to go accelerate that on a single product. We wouldn't have to pivot. That's the big advantage and you foundationally have to build a different set of technologies and a different company to address a problem like that.

To our knowledge, there's no one else who's trying to do this. So, either we're very right or we're very wrong. We'll find out in

the next five years. So, we're betting on the right, but it just makes sense to us to de-risk your market entry point.

How have more recent advances in AI changed the autonomy landscape in helping solve problems related to the feasibility and economics of deploying these systems commercially? For example, companies using multimodal LLMs to improve autonomous navigation, especially with regard to labeling and reasoning around edge cases (Ghost Autonomy).

Speaking very narrowly to the ground autonomy piece first—ground autonomy for defense is very, very different from on-road in the sense that it's a 3D problem. You are traveling off-road, and on a road or surface you've never traveled before, in an environment you've never seen. You likely can't do that well with generative AI. Because the datasets do not exist.

LLMs are awesome. There's a great tool called FARGPT, which is just exclusively trained on all the different Federal Act regulations of which there's actually 16. Most people don't know this. There's not just the one. I use that all the time and it's wrong about 15% of the time.

It is only trained on government-issued documentation and it's a great way to get started. You're like, "Okay, where was this regulation again? I don't want to go comb through 15,000 pages of the FAR", but you have to check it every time. "Okay, it said this citation. Let me actually read that citation on the FAR. Ah, yeah, that's incorrect, FARGPT. This is the actual thing."

So think about that in ground autonomy. You have to have a massive model of the off-road environment for basically the entire world and every modality in every different type of climate to build a model that could address that. What Tesla's doing is not crazy, but I don't think it'll happen in the next five years, because the data sets that you need and the precision that you need are not even close.

The second part is safety. To understand what the intended function is for a system, you have to be able to get a set amount of inputs and they have to give you the set amount of outputs every single time. Think about asking a question to ChatGPT. You ask the same question on two different days. You're going to get two different answers. You might get two

different answers within 30 seconds. So, if I can't certify, especially the government or regulators, that in a given scenario a vehicle is going to behave exactly the same way every single time, you'll never get a safety certification.
It's useful for other things. "Can I have it debug my code? Can I create new simulated environments that a generative AI algorithm can create?" It's a helpful tool, but if you try to do that in the ground autonomy space, it's just not going to work. You might get a 90% solution, but in ground autonomy, the last 10% is 99.9% of the work. You have to be consistent.

It has to exactly do what you want the system to do every time. The irony is that our CEO Alberto Lacaze, who's an incredible engineer, was teaching this stuff 20 years ago before it was cool, before he jumped into the build-a-business phase of his life. If there's ever more a bigger advocate for generative technologies, it's Alberto, but we don't have that approach in our code base, because we don't think it closes. We don't think you can get a safety case from it, and we have pretty high conviction on that.

If RRAI succeeds, what does the world look like in 15 years? How are things different?

I think number one is fully autonomous robotic combat formations—the drone wars are coming, and we want to be at the forefront of that.

We think that's better for humanity. I certainly don't want to go into a shooting war personally with another country. Combat sucks. Anyone who tells you otherwise is lying to you. We think that's coming, and in the next war—because there will be the next war—these technologies will enable dramatically fewer US casualties.

What we're going to see is the same kind of shift in warfare that we saw with the emergence of things like jets and missile systems.

I also think that by that point in time, you'll have fully automated yards and off-highway drayage runs on private roads. That will become mainstay much like it took about 15 years for the inside of warehouses to be fully automated. That took a while and that's a very controlled space.

Fifteen years from now, you'll have fully automated yard facilities, transportation depots, things like that. Around that time is when you'll have the first truly unmanned scale deployments of on-road vehicles—like truly driver-out and autonomous. It'll take that long to get the public adoption, not the technology. There's a public perception and regulatory problem that, I think, will just be very slow to mature.

Disclaimers