

DARPA Information Innovation Office (I2O) Office Wide Proposers Day

Q&A

Question: We've heard rumors that ACO will come out with a Transition Council. Can you explain a little bit more about that?

Answer: So, ACO is the Adaptive Capabilities Office, so it's a part of DARPA that doesn't have a broad public exposure. So maybe many of you aren't familiar with it. It's generally been our kind of special projects office. Sometimes, DARPA does things because we're uniquely positioned to do it for the Department of Defense, not because it's really a DARPA program. These tend to be large-scale joint cross-cutting activities where we're driving disruption. But because of often it's models and processes and organizational disruption more than the underlying tech disruption. And so, what we found over history is that it's better to, to carve these things out away from our core tech offices because they need a different form of management.

Within ACO, we created something called ACO-T, which is a transition support function for program managers across the agency, to better help PMs navigate what, what is a much more complex situation because we're not just accepting the traditional program of record pathway. Transition occurs on a spectrum that can range from the program of record, a specific military transition, all the way to purely commercial activities, and everything in between.

The Transition Council is a decision-making process to help drive that. We have Tech Council where we approve funding for new programs. Transition Council approves funding for existing programs that have satisfied the DARPA "technical miracle" and are now really in that "how do we scale the impact" conversation and they may need some additional resources.

Question: I had a question about the resilient software slide. I'm familiar with all those programs and all the great work that went into them. Where do you think they fell short of disrupting that cycle that you had outlined in the upper left graphic?

Answer: This links to the last question strongly. If we think about transition purely as a program manager activity, then the scale of what's kind of the reasonable level of transition that can be achieved is kind of limited, right? Because, you know, the program manager is saying, hey, you know, I've developed this magical technology. How do I get to "it also disrupts something?" They're going to talk to a program office. They're going to start going to the acquisition side. They're going to do all the things you have to do to line up the requirements, the money, for the operational side to take it. But it takes a tremendous amount of energy and effort to go drive all those conversations and get somewhere. It's to be expected that only one or two programs can push that all the way through. I think that's what we see, good PMs that have done great work and then ended up with these transitions. We're trying to take a broader, more integrated view of agency-scale transitions. A DARPA-wide activity to push this transition through. They're

a conversation between the DARPA Director and, you know, the Deputy Secretary of Defense about how we move this through the process from the top down.

Question: What is I2O's hiring strategy for the next year?

Answer: We have a lot of people who express interest in being program managers. Every time we go and talk to anybody, we say that we are always hiring program managers. At DARPA, every program manager and everyone who has decision authority has an expiration date on their badge. We turn over roughly a quarter of the office every year. We are always looking for people who have ideas. I'll meet with someone and talk about possible ideas. If the candidate has ideas that seem interesting, sometimes aligned with the strategy, sometimes not, if the ideas are surprising and something in the DARPA sphere and roughly in the domains that I2O has an interest in. If we like the ideas, we'll ask the candidate to put together a sample pitch that answers the first three of the Heilmeier Catechism questions in 1-2 short paragraphs. What are you trying to do? How is it done today? This demonstrates a broad understanding of the state of the art. And the third question is what is new about your approach and why do you think it might succeed? You must have a glimmer that the problem might be solvable. As an example, transporter technology from Star Trek, it's easy to answer the first two questions, but nobody has that grounding that it might possibly be solvable. We then iterate on that draft, diving in to create well-grounded answers to those questions.

Part of what we're assessing is whether the candidate is willing to take feedback, because nobody comes to DARPA with all the skills they need to be a successful program manager. Being able to take feedback is critically important to being successful, while also assessing communication skills. Because you as a program manager don't do any of the technical work yourself, you have to be able to motivate other people to do the technical work. So, if you can't communicate, you're not going to be successful. Then, when the candidate is please with the quality of their written material, they come back and talk to me. I often then ask about what other ideas they have that they think might be of relevance to assess their ability to create ideas that are not so "one trick pony."

If they're good with that, then we bring them in for an interview. They meet with a bunch of PMs, they meet with me, they meet with the deputy of our office. We have a set of criteria for evaluating them, which is the things that you might expect based on what I just told you. If they do a good job on those interviews, they go meet with the director of the agency.

This answer is more of the process than the strategy. Now, we're shortest on AI PMs, as you might imagine, given the demand for people who understand AI across like every single organization on the planet. We are most focused on trying to get high-quality people who understand AI and national security ramifications. Not as much focus on the "apply AI to a problem" area. We're pretty good at applying AI to problems, but like AGI, how do we defend against AGI? How do we think through the ramifications of

AGI? Is the existential risk of AGI a thing or not a thing? How do we think about agents run amok? How do we defend against that?

Question: Besides quantum, AI, cyber, what is the next disruptive technology you find challenging to explore?

Answer: Well, I mean, to some extent, I want you to tell me. There's like, 350 of you in this meeting and only one of me. I do think one problem that is hard to solve is the PAI/CAI surveillance ecosystem of all the publicly available, commercially available data that is getting generated, like every time you touch your phone, every time you drive a car, every time you go on a commercial flight. Adversaries can buy data that can track every single person here, every single important person. Why bother with high tech surveillance devices when you can just go buy this data? How do we defend against that when maybe the right answer is that the government put in laws that say you can't aggregate the data, but the \$2.2 trillion industry, that's a massive lobbying effort to not regulate that, not make it illegal.

So, if we're not going to be able to just make that illegal, how do we make it so that adversaries can't track all of us perfectly? That's maybe not a new technology, but it's a consequence of existing technology that is, I'd love a solution for that.

Question: For the last 20 years, we've been hearing about formal methods. If there's ever a system that had the money and the time and everything else, but it was still attacks, how do we get past that problem?

Answer: Yeah, I think it's going to be a multi-pronged approach. I think, one possible way is authority to operate. The normal approach for authority to operate is a super labor intensive, super paper-oriented process that takes forever. It doesn't accomplish what it's supposed to accomplish. And everyone hates it. There is an alternative, which is a formal methods-based approach, which is much faster and it actually much more closely accomplishes what you want it to accomplish. Right? It's tied to the actual artifacts instead of like the training of the people that were, you know, carrying out the code and is tied into the formal methods-based process, which then could be tied to kind of like the leading edge of the formal methods train, as it were.

So you adopt a formal methods-based ATO process. We will get you paperwork generated in, you know, a day instead of in like half a year. And when you change the artifact, we will regenerate your paperwork in a day instead of making you take another six months and we can tie it to the actual code. So that ATO now maybe it's tied to an architecture description that was generated from the thing and that architecture. Now you can analyze it to figure out where is the most important place to go, like use generated crypto that is proven to be functionally correct and replace the parser with a generated parser. And like is the foot in the door as one piece. And then I think like taking red teaming of cyber way more seriously. And maybe we do a demonstration where we like actually, instead of just like white carding things or, just writing written reports about like, cyber could have done this, actually go demonstrate the effect that cyber could have done. Like maybe we need to make a plane actually fall out of the sky

and blow up, right? Like, because people don't seem to believe when we say cyber could have made this plane crash like they seem to take it seriously when you actually make a plane crash. Maybe we might have to do something like that. I mean, China might help us out, right? But the goal is to get us, you know, to get us to like, take it seriously before China does something.

Question: One of the issues that we've been facing is acquiring and maintaining the talent to do these things. As you mentioned, industry right now can outstrip us in terms of compensation packages, and that has to do with funding. We can razor thin our margins, and we're still not competitive in some of these cases. And further, you have to go to your potential talent and say, as you said, the protection we often go for is an air gapped network. So, you must do your work in a cave and people want to work from home. And what, if anything, are we doing to approach this problem where we have these great ideas and we may even have potential solutions, but we don't have the people to execute?

Answer: Those are really good points. And those are things that we are thinking about quite hard. Some of the things that we're doing, like whenever we are doing an effort, we, we look at the classification level of the effort. So we try to do as much as possible at the unclass level. So you don't have to do it in the cave if you if you don't want to. The DARPA Connect efforts are to try to, recruit as many, like, as many people as possible to make it so that it's not always the same people. So reaching more people to engage. Not everyone is motivated by money. Obviously, we're all here, and we probably could all be in a different room where we could be getting paid a lot more money. So, reinforcing that the mission-driven pieces, we are trying to work on, understanding and streamlining the security requirements and the, in the IT requirements. Obviously, that's a work in progress, and we still have a long way to go. So, if you have specific, pain points that you would like to make us more aware of, please don't hesitate to reach out. And I can talk through those specific pieces as I started with the opening, like, you guys are a strategic resource and, I appreciate that. It is not always the easiest to work with DARPA, and we would like to make it as easy as possible, given that, we don't have perfect control, like, certifying SCIF space and things like that is not always something that that we have exact control over. But if you have specific things, I'd love to hear about them, offline.