RAAF ACDR John Haly:

All right. Ladies and gentlemen, welcome. Good to see the room filling up and thanks for joining us. If you are not here to talk or listen to a discussion about enabling manned and unmanned teaming, then you are in the wrong place, but you should totally stay because it's going to be great.

It's my great pleasure to introduce our three expert panelists that we have today. We have Mr. John Clark, Vice President and General Manager at Lockheed Martin Skunk Works, Dr. Ben Strasser, Principal Lead, Mosaic Autonomy Research with General Dynamics Missions Systems, and Mr. Mike Atwood, Senior Director, Advanced Programs Group, General Atomics Aeronautical Systems. Gentlemen, welcome.

As I said, we're here to discuss how we can go and enable crewed and uncrewed platforms and systems to work together for a common purpose and common missions. And so, the format of today probably won't surprise you, but we're going to invite each of our experts to talk and to give some opening comments and then we'll launch into some questions from there. John, why don't you take us away?

John Clark:

Thank you. John Clark. I've had the fortunate, I guess, career up growth, if you will, of working in unmanned systems for the better part of 20 plus years, almost 25 now, of doing unmanned aircraft. What was interesting and how things have navigated is that there's always this dimension when you're talking about these unmanned or uncrewed systems, how the human interacts with them.

And so, as we talk through this crewed/uncrewed teaming or manned/unmanned teaming and how we integrate these adjunct systems to the crewed systems, a lot of things really drive back to that human machine interface, how you interact with them, and it takes me back to very early S&T research that was I working with the Office of Naval Research now 20 years ago. We were putting a lot of autonomy software in place and having users go through and evaluate it.

We had this capability where, right out of the gate, you'd see the users ask, "Why did the system do that? I really wish it would've done that and instead." And I think that as we go forward with a lot of this manned/unmanned teaming and crewed/uncrewed teaming, we need to make sure that we don't replicate those types of experiences. We have to have the human actually understand why and what those other systems are doing and how it benefits them because they've already got a heavily task saturated job operating in a cockpit. And if we're inundating them with additional tasks and requirements to understand what the system that's supposed to be helping them is doing then, then we're failing as an industry.

RAAF ACDR John Haly:

Yeah, thanks very much there. And just for a note, you came to a 40-minute presentation but the clock hasn't started. So look out, this is going to go forever. Ben, I'll turn over to you to build on that and to provide some words as well.

Dr. Ben Strasser:

Excellent. Thank you, John. I want to thank the AFA for giving me the honor to speak today. At General Dynamics, our north star for autonomy is really about taking us from a highly orchestrated, tightly controlled paradigm into an adaptive one, featuring multiple manned and unmanned agents collaborating together on complex, open-ended missions. In particular, we want them to respond well

to unanticipated tactical situations and be able to function and make those responses without particular necessary input from humans on the loop.

If we want to have manned/unmanned teaming, we want to envision this, we think there's four key capabilities that need to be born out. First, intent, commander's intent. This goes beyond what are my orders for the particular mission and into what is the mission ecosystem? What are the overall objectives? If I see something unexpected, how should I respond? And before we can have unmanned assets responding well, we need to make sure they understand the situation and are empowered to make the same kinds of decisions a human might make in response to unanticipated tactical situations.

Second is your role. In response, once you have that mission ecosystem model, you need to determine at any one moment what is my role and how is that going to be changing over time? How am I supposed to be functioning in this system? The third is really the tactics, then you need to optimally execute your role to complete specific tactics. There's a lot of really great work in using autonomous systems to execute very specific roles. And so, we need to identify what role to we need to execute.

And finally, and most importantly for this panel, is trust of course, trust and transparency. How do we give enough insight to humans on the loop as to what the autonomy system is doing and why it's doing it? And how do we build that trust that that system is going to make good decisions? That's what I hope to get to today. Thank you.

RAAF ACDR John Haly:

Perfect. Thank you. And Mike?

Mike Atwood:

Yeah, it's interesting to hear John and Ben speak, being in the middle of that. I, too, like John, started at General Atomics almost 20 years ago and my first task was to support the Hellfire installation on the MQ-1 Predator, and that set my life on a trajectory that I could have never expected. Fast forward through the Afghanistan conflict and the weapons engagement, build forward to the F-35 operational deployment, I was part of the operation where we lased a target to launch the first JDM an F-35.

So for me, the manned/unmanned teaming something, a concept is something that I've been living with for quite some time and I feel like we're at a really interesting point in our technology portfolio right now where we can truly innovate by integrating. We have so much technology from the driving car industry, from the high power compute gaming industry, we're finally at the point of realizing some level of closed loop artificial intelligence. And I think we're now taking that and abstracting that to a level that a human can actually interact with that.

Kendall has talked about OI3, air dominance, family of systems. I think we're on the precipice of something very, very special with the collaborative combat aircraft, and very excited to be a part of that and be in the testing program with things like Skyborg to go through and try to understand what the opportunities are for our war fighters to use that technology.

RAAF ACDR John Haly:

Mike, could you help me out with ... you used the expression closed loop artificial intelligence. Could you explain that?

Mike Atwood:

Yeah. A lot of people in this room have probably operated an MQ-1 or MQ-9 Predator. It's very open loop. You slew an EOR to it, you look at a target, you laser designate it, you pull a trigger, there's two

pilots and it's open loop to the human. The human's constantly closing the loop for the system to engage the target or provide the ISR. What we've realized on the Skyborg program is we put the human on the loop. So we've closed the autonomy loop.

And what we're doing now is we're setting objectives and constraints. We're saying, "Fly over there, find this thing, don't fly over here, stay away from that thing." The robot system, the autonomy engine, is basically solving that problem for us. And so, that closes the decision making loop. That closes some of the [inaudible 00:07:38] with the objectives and constraints changing based on the human outer loop. And so, the human operates on the outer loop, while the machine operates on that inner closed loop.

RAAF ACDR John Haly:

Okay. And that comes back to the discussion, Ben, that you were talking before about it being empowered to make those decisions. When we were talking about this earlier, you spoke about this in the context of explainability of what's happening. Could you take us through that?

Dr. Ben Strasser:

Absolutely. I think there's two components, and one is really about making sure we bring in both the semantic elements of the mission, as well as the physical elements. The semantic description, this is what comes from the written op board. "Here is the role I'm performing, here's my job, here are the things I'm interacting with." And then of course this interacts directly with the physics in a complicated way. Both of these talk to each other and both influence each other.

But when you start to bring in the semantic thing and you make that an essential part of the control loop, then you almost have built in explainability because the machine can always say in words what it thinks it's doing. And so, that's I think one part of it, bringing the semantics. The second part, there's best of breed academic research says what we need is simpler heads up displays that can communicate information like purpose, performance, and process. What is the thing doing, how well is it doing, and how was the decision made for it to do that?

And if you start to add in more information about its confidence and its ability to accomplish its mission, it's uncertainty about what it is it's supposed to do, that actually further improves metrics of human understanding. And the final thing I'll say before I ramble too long is that this is somewhere where AI can actually come into play, not necessarily for flying the aircraft, but for helping to curate the data that gets presented to a manned teammate, because if you have one human on the loop and one unmanned aircraft, that's one thing. You can present a lot of information.

When you start to have five unmanned systems, 10, 50, 100, at some point, there's cognitive overload. So I think the idea of an AI curator can actually help improve the explainability and reduce the cognitive load.

RAAF ACDR John Haly:

Yeah, that's interesting. I get cognitive overload with one platform under my control. So I can understand how multiple would do that. From Lockheed Martin perspective, particularly from a Skunk Works perspective I guess, I'd be interested to hear about the ethical considerations that I know you've worked through and your concept of purpose driven work.

John Clark:

As we've navigated through that early science and technology related research and then systematically went to go build up systems that do stuff that we're talking about, that element of explainability and

determinism had to be intrinsic in it. One of the things that we built in based on that human interaction such that the user would understand what the system was doing, is that we built what we called a flexible autonomy framework where the user could go in and dial in and say, "All right, for this decision, the system is authorized to make that decision. And for this other activity, for instance, dropping a weapon, dropping a weapon is not authorized and it will require a human in the loop intervention to facilitate the weapon actually being released."

And so, we systematically have gone through and put this framework in place such that the user can customize the decisions that it's authorized to make. And in that spirit of, as Mike had highlighted the on the loop construct, that on the loop construct is allowing ... At the end of the day, we very much look at Air Force doctrine and understanding there's an accountability element for every weapon system that we put out there and how can that accountability be traced back to individual users?

And so, in that construct, we've put that decision framework in place such that the ethics of how the system is operating is actually driven by the user that configured it to do what it needed to do in that day. And as we've continued to advance those capabilities, as we explore adding machine learning and artificial intelligence on top of it, we've actually started to evaluate ways in which you can take model based simulation and explore, "All right, we're going to take these systems into a dense threat environment. You're going to configure them a certain way, this is the response that comes out."

And then you run those on a recursive basis so that you can start to see AI is not deterministic in all cases. And so, you can start to understand how you're going to get a response profile and then that response profile can then be driven by the user of the operator to understand, "All right, I'm okay with this response profile and authorizing the system to make these types of decisions, but I'm not okay with this response profile." So ultimately, you're going to have a more overwhelmed user taking on more tasks to then authorize these other adjunct systems to support.

Mike Atwood:

Yeah, I want to jump in there because I think John's starting to hit on something that maybe not a lot of people talk about when we talk about machine learning, but I think is essential to understanding it. Many years ago, we all had Google Photos, start patterning our face and tell us who we were on our phones. That's called supervised learning. Then all of a sudden, you had this thing on your computer where you had to show how many cars were in a scene, Captcha, to get through your thing. That's called unsupervised learning, where they actually look at entity types and start patterning things. That's analogous to the Maven program with Google and then the Smart Sensor program with the Joint Artificial Intelligence Center.

What John's talking about now is a modeling framework that allows an emerging technology called reinforced learning. And this is very exciting because what you do in this case is you define a state space, a world that your machine can operate in and then a set of actions and you let the machine run through all those permeations and basically self-learn all the different behaviors it can find within that state space.

And so, I found actually comfort in using the reinforced learning model as we go into flight testing with group five unmanned systems because the state space, the maximum extent of what that machine can do. It's much like your cruise control in your car. You start by a closed loop control of it throttling back and forth, then we trust a radar to come into the equation and we grow that trust in the system and add that. Then we see a steering wheel that can turn itself and we slowly become a constant of that, but we know it can't turn the radio dial, it can't turn the air conditioning on and off. And so, we're comfortable with it moving laterally and longitudinally in that state space.

And what we're finding now in the manned/unmanned teaming is the squadrons are ready to start accepting more degrees of freedom to the system. Not just flying in a circle, but maybe queuing mission systems, maybe doing electronic warfare, doing comms functionality. And we're building upon the framework, that flexible framework, that John talked about to really let the war fighters develop these really exciting TTPs within those reinforced learned state spaces.

Dr. Ben Strasser:

Actually, reinforcement learning is a really interesting idea because ... sorry to jump in if I can jump in on top of you here. It's a really interesting idea because, as Mike was saying, it allows us to really define pretty open-ended problems where you can define pretty broad objectives and say go and hope the machine learns. And I think the real challenge is in incorporating trust into that system and making sure there are enough constraints so the training actually converges, so you actually get there.

And that's where I think introducing some hierarchies in the reinforcement learning. Are you really letting the AI pilot think about the whole controls to accomplish an end-to-end mission? Are you confident in that or do you want to introduce maybe separate decision layers too. What is my objective that I should be pursuing? How should I go about pursuing it and then let me execute that. And one thing I like about the hierarchical paradigm is it lets you mix AI technologies when you trust them with traditional optimization, when that works better. So I don't know if you guys want to ... Oh, sorry.

John Clark:

One of the things that I'm going to pose out to the audience as you're navigating through your day to day challenges and think about what our adversary's doing, the question posed to me was centered around ethics and how we apply ethics in our employment of artificial intelligence. And as it was just highlighted, there's a lot of advancements here. Those advancements that are not unique to the US. That technology and that capability is available to both the nation state actors as well as those folks that now we just say that they're the peer adversaries.

And candidly, I'm not so certain that they're going to be having these types of conversations about ethics and their AI, and what ramifications does that have for us as we're going in and navigating those types of fights? And as folks that build capabilities for all of you, understanding how do we navigate that space. Having specifically worked with the user community in fielding capability where we explicitly for decisions of Air Force doctrine and safety, we pulled capabilities back and explicitly did not allow them to be available to the end user because of the higher ground that we take.

And I'm not advocating that we don't take that higher ground. However, we have to think through the implications of what the adversary is going to do. And that's going to cause us to maybe have to think asymmetrically about how we fight these reinforcement learning capabilities. If a peer adversary is putting a large amount of aircraft out coming after us and they're having emergent behavior that we've not seen and how we combat it, that's going to be a challenge if we don't have a system in place that can be as adaptive or more adaptive than what the adversary puts forward.

RAAF ACDR John Haly:

Ethical considerations aside, where it's a consideration of trust in terms of our either senior leaders or our political masters trust in the system in order to allow us to go and do things, what policy constraints do you think are surmountable, aside from the ethical things, and how do you think we, as essentially a government industry team, take that part of our population, ultimately our leadership in some cases, on that journey?

Mike Atwood:

Yeah, I'll start there. I had a chance to sit with General Kelly for a little bit and talk about his belief in how we approach what we want these machines to do. And he's a huge advocate of these Adversary UX program that in the vein of what John was just talking about, we go out, we give it to the weapons schools and we use it as red. And I think we'll quickly realize how capable these systems can become in the hands of our adversaries. And I think that will be maybe the Sputnik moment of cultural change where we realize when we saw flyer F-22s and F-35s in the range, how challenging it is to go against that.

And I think that will spark a debate of needing to create an inverse to that, where when we put weaponized product in the battle space that's unmanned, the adversary treats it just like it has the lethality of the manned aircraft. So I think the answer somewhere lies in an experimentation program like Adversary UX to start exposing these underlying policy ethics ROE issues that you raise up.

RAAF ACDR John Haly:

Yeah, I agree. Adversary UX is likely to be a really good forum for that sort of development. The other one that springs to mind would be in what we would describe as shaping or phase zero type of activities that are short of conflict. Any thoughts down that line?

John Clark:

Yeah, I'll take a swing at it. I think that this is one of the things that, especially with our times right now and one of the things that I emphasize with my team, is that any of these types of capabilities that as a collective enterprise, working both contractor as well as the Air Force and the broader DOD, the more that we can show and highlight and publicize, I think the better off we are going to be in the near term, in particular with China. The fact is that they're very aggressive, but culturally, they will look for that guarantee that they're going to win and we need to be showing off things that maybe give them pause, give them that little bit of reticence that maybe today's not the day as it's been quoted in the past. And so, I think that there's a lot of opportunity for us collectively.

We've even talked internally within our Lockheed Martin team of partnering with other contractors as a mechanism, much like we did in the '40s when it was World War II and we're looking at Manhattan Project. Having that type of collaboration where everything operates in a [inaudible 00:20:34] environment and we figure out a way to go get some capabilities out there, no kidding in the next 12 or 18 months, to just change the game with respect to the adversary.

The environment is not quite to that point, but it's maybe one day or one event away from having that sort of environment. So inside the Skunk Works, we've already been thinking about that, about how we would prepare ourselves to navigate in that sort of environment to meet some urgent national need, even though in my opinion, it's already here.

Mike Atwood:

Yeah, I think there's a sub narrative that I was picking up a little bit in John and disruptive tactics in how we employ these weapon systems. I'll just share an experience I had on the Skyborg program out in the Edwards test range. We were doing some manned and unmanned teaming with some F-16s from Lockheed and the MQ-20 platform by General Atomics, and we were using them as loyal wingmen, so blue aircraft. The F-16 would send a links 16 command over to the MQ-20s and we try to prosecute some other F-16s that were acting as adversaries.

And what we found really quickly is the F-16 would run out of gas and our blue wingman manned aircraft and it would go off to Mammoth ski area to go get gas from a KC-135, sit there for a while, while the UADs stayed on station and kept surveilling the target, holding custody, watching it fly through the range so when the F 16 could come back. And the F-16 pilot calls on the radio and goes, "Why do I need to come back? I'll just stay on the tanker, you guys stay over there and we'll all look from these multi-static perspectives at these targets out there."

And we realize that the way in which we think about air tactics, air combat maneuvers, BVR, beyond visual flight range engagements is fundamentally going to change for us. And I think some of the things that we're going to start experimenting with are these fundamentally different expansive tactics that are possible now through this wide mathematical trade space that Ben and his team work on.

Dr. Ben Strasser:

Well, I'll take that cue. Thank you very much. Yes, of course, we at GD are very proud of the work we've done to build, we call our commanders algorithm toolbox. And I'm sure Lockheed also has a toolbox of really advanced tactical algorithms for unmanned assets. I think your point about collaboration earlier is a really good one because there's not a monopoly of good ideas out there. And one thing we have to figure out is, as you walk around the trade floor, one floor below, there are a lot of autonomy companies doing really interesting work coming up with very tailor-made advanced tactics to solve very particular problems.

And if you want to go to our north star, being able to accomplish end-to-end missions and having autonomy assets that can do more than one thing, we're actually going to need to combine tactics from different vendors. And the buzz words about open architecture and infrastructure as code paradigm start to really matter when we need to combine novel capabilities from multiple providers in order to get an autonomous capability that's more powerful than the sum of its parts. So that's something that we're really championing at GD, is how can we lower the barrier of entry, build architectures and development pipelines that are easy to incorporate third parties, easy to incorporate good ideas from smaller vendors to ultimately deliver the best autonomy capabilities to the war fighters as quickly as possible?

RAAF ACDR John Haly:

What do you think is likely to hold that back from being achieved? Will it be the competitive development between different vendors? Will it be the requirements that are set and defined by the services that aren't conducive to that? What do you think the impediments are likely to be?

Dr. Ben Strasser:

I think the two main obstacles are you're at the competitive element. Everyone wants to have their own proprietary system that is the best in the world. And I think I had another component. I'm drawing a blink on my second component, I apologize.

John Clark:

Yeah, I'll jump in. I think that what I look at right now having navigated through this in the past is that there's definitely a cultural dimension. The technology has come a long way. There's a lot of capabilities that are out there. We could do a show of hands. Who in here would go hop in a Tesla right now and immediately press the auto drive button and let it take you back to Reagan National with your eyes closed? There's not going to be very many folks. And so, there's that element of trust in the system.

And so, that's going to be no different than our user community that you've been trained thoroughly and thoughtfully through how you employ a fighter aircraft. And I'll share a little anecdotal story that emphasizes this. And so, we have the Alpha Dog Fight program that we participated in with DARPA. And for those of you that recall, our Lockheed Martin team, we placed second. Well, why didn't we place first? Obvious question that I posed to the team. And we went back through things and we evaluated. We evaluated against the team that won, and what we observed is that we had worked with some of our user community in our reinforcement learning process and we had put some specific behaviors in our algorithms that were driven by the pilots in terms of how they would prosecute a mission.

We went and looked at how the competitor that placed first, how they had executed it, and they were not constrained by those same tactics that were put in there. In fact, they were actually outside of doctrine that was authorized based on how aircraft are used. And so, at that point, that's the question is that, which one was better? Was it better to follow the doctrine or was it better to win? And so, I think that that's going to be one of those things that culturally we're going to have to navigate our way through. There's that acceptance and moving through, understanding the technology, embracing the technology, and then what risks are we willing to take as a nation?

Many of you I've talked with in the past, I think that's one of the places that collectively we can get a whole lot better. And this is a prime space for us to go explore that, is take more risks. We need to go explore this trade space a whole lot more thoughtfully and we're not going to do it by analyzing it on paper. We actually have to go experiment. We're going to have to get some aircraft in the air. We're going to have to fail a few times and learn from those failures to then move forward with this is the right way to go do it. And so, I think that that's the number one thing that's keeping us from being able to really make the leap forward with this technology. I don't think it's the technology.

RAAF ACDR John Haly:

Switching gears slightly but not very much, how do you think envisage that we are likely to mission plan for these types of capabilities? Do you see this as being something that's done in exquisite vaults well ahead of time, or are these likely to be not mission planned ahead of time, but actually just employed within their capabilities by whomever has custody of them?

Mike Atwood:

I'll take a lead on that. In my generation of working with the MQ-9 Reaper, we essentially do no mission planning. The plane has such long endurance, it launches into the range and it's dynamic. You get a task from a JTAC, from an airborne battle manager, and you're doing some function. And so, GA has had to think really hard about how do we look at this battle space where there's not that dynamic human, those objectives and constraints that I talked about earlier.

And the best way I can describe it right now is something like Waze on your phone that you're going to say, Hey, I want to go from here to here or do this objective at this point in time in this place and it's going to give you all these purple lines. It's going to say, "This is the fastest drive, this is shortest drive, this is the most beautiful drive to make you smile." And the people on that outer loop that we talked about earlier are going to have to implement that doctrine in a way that looks at the risk acceptance posture, looks at the ROE and makes the best human subjective decision about the multiple courses of action that can be executed in that trade space.

Once we set that level of risk acceptance and time of arrival and all these things, we're going to have to trust some level of closed loop automation to solve that problem when it gets there. Data links are constantly being challenged every day. We need more assurance when we're not connected and we need to live in worlds of sporadic connectivity, we need more edge processing. And I think that mission

planning will happen probably not in a vault because it's in an objective constraint level, but ultimately, be put on the machine to close the actual execution of that within the objectives and constraints.

John Clark:

Just to provide maybe a complimentary, not contrarian view, but I think that what you're going to find is that given the way in which we can model things, there's going to be a lot of activity that happens in a vault and that's going to be basically simulating all these different permutations, going back to this modeling environment that we talked about earlier, where we're going to explore the different ways in which the mission could unfold with different things that pop out.

Based on that, that's the type of thing that's going to be used to facilitate what is traditionally the ATO. I mean we still have to get the airplanes to the right airfield, those airplanes have to have the right amount of gas and understand how that's all going to be orchestrated. And so, you're going to have that level of mission planning, but you're not going to have that same level of mission planning of this is exactly ... you're going to have 79 way points in your mission and you're going to take an image on this one and you're going to drop a weapon on this one, more akin to some of the traditional LO mission planning that has happened in the past, where you just follow the line.

I think that once you've gone through that model based evaluation in the vault and understand all the capabilities that can be brought to the environment, then you're going to execute with that initial ATO construct with the airplanes coming in. And then it's going to be very dynamic, very adaptive. You're going to do all that processing at the edge. And likely one of those permutations you explored will manifest, but it's not going to look exactly like it was in the vault.

RAAF ACDR John Haly:

What do you think, Ben?

Dr. Ben Strasser:

Just to add onto the conversation, I mentioned earlier commander's intent is really important and being able to capture the mission ecosystem, which isn't just the written words on the op board or the exact objectives we want to accomplish, but really the tactical context. Mission planning today, there's a lot of analysis and a lot of artifacts that we develop and a lot of processes that exist for a reason to help plan missions.

One of the things that we think has to be an important part of this conversation is really a human centered design approach so we can leverage all the work that's already being done to plan missions for manned assets and capture that in a computer understandable format for the unmanned assets, because at the end of the day, this is not just a panel on autonomy, it's a panel on manned/unmanned teaming. And so, we have to make sure we have that understanding of the manned mission, even if we're allowing some level of improvisation or adaptation or even just on the fly planning based on the objectives and constraints.

RAAF ACDR John Haly:

Yeah, interesting. Do you think it's a truism or it just can't be achieved, the concept of this balance between cost versus the lethality/survivability thing? Is there a knee in the curve that we should be all aiming at and exploiting, or is that a falsity?

Mike Atwood:

Yeah, that's an interesting question and one that I think about a lot as I make material decisions for where General Atomic sets with technology. Being part of this for so long, I've seen the X-45, X-47, J-UCAV programs, which were giant flying wings, 40,000 pounds, 8,000 foot runways, and I think we tried to make those loyal wingmen in the 2012 era and we realized they're a little too monolithic, they're a little too big. And we realized the adversaries were going to swarm tactics and attrition mass mattered, and we needed some level of attritability.

And so, the pendulum swung over to the target drone community and growing those up into more capable UAVs. That's the Valkyrie program, XQ-58A. And as that executed, we realized it was hard to man, train and equip that. The operational readiness of radio bottles and parachutes, it was hard to provide a phase zero deterrent capability with that. And I think what you've seen emerge with Secretary Kendall is this middle class. And it's not that different than an RQ-170 or an MQ-9 type platform that have stood the test of time for the last 15 to 20 years because they've been adaptable to the changing mission in the Middle East.

And so, I think we've consolidated as a war fighting community around this 10 to 20,000 pound class of utilitarian, adaptable CCAs. And I think they embody a fundamentally different philosophy than just a small F-35. And it's been exciting to see how we compliment the manned aircraft and we bring sensors that are offset and disaggregated, and not just make a small fighter or make a target drone that's so attritable that it doesn't have much capability.

John Clark:

I'll build on what Mike highlighted there. As we went through and evaluated these types of ideas and concepts, they're not new. We've been talking about some of these things for quite some time. As that pendulum swung and it was down in this a triable class and you're looking at things, we did a lot of operations analysis trying to explore is there something where you can get sufficient capability into a contested environment and have a meaningful impact while understanding what that sustainment and logistics tale was going to look like to support them?

And the candid answer is we couldn't find anything, just as Mike highlighted. As you continue to look at anything that was really in that lower class vehicle, in the end, they became really exquisite targets and they would just be shot down on day one and you didn't get them back. So they weren't attritable, they were truly expendable. And so, at that point, when it's expendable, everything's about just getting as much cost out of it as possible. Attritable, you want it back.

And so, you start to have this dilemma that, "All right, I've just put this really sophisticated IRST on there. All right, the price just went up. Now, I want that airplane back every time. Well, now that I want it back, every time I'm going to have to put this additional survivability content on there. Maybe I've got to put a jammer on it. Now I've got to put some [inaudible 00:34:30] materials on it. Now that price point keeps going up and now it just snowballs on itself."

And so, I genuinely agree that where that middle class has emerged, I think that there is a sweet spot that we can find where you're going to have a class of vehicle that has the right amount of survivability, the right amount of sensors to actually compliment the fighters. And I'll close with this little dimension. I've shared with a few folks that have come in and looked at our operations analysis. And so, when we play chess, the pawns are the front end. Our unmanned systems, I think we can all argue would be the pawns in it. They're the ones that are going to stimulate the activity.

You don't play chess by putting all your pawns behind the systems or the pieces on the board that matter. And so, as you go explore that, you've got to have the systems, those pawns, actually be able to get close enough to the adversary to make an impact and do something to stimulate the behavior that you want in this teaming construct. And so, these ideas of having unmanned aircraft that are way

behind the fighters, all of our OA says that that's not a really good value proposition and it's not actually helping, because at that point, you just bring more F-35s or F-18s to the fight because it's not actually making an impact for the humans that are putting their lives at risk. We need to put the unmanned aircraft out in front and they have to actually persist long enough to make an impact.

RAAF ACDR John Haly:

Speaking of the humans, that's a really great segue. Ben, I'll throw this one to you initially. What do you think we need different in our people to effectively be teamed with these uncrewed aircraft? Are we producing the right sorts of people already, or is there something that we should either be looking at or cease looking at in order to be able to do this role effectively?

Dr. Ben Strasser:

That's a really great question, and I think as I mentioned before, we're looking a lot to the research community and the ideas of psychology to try to really reduce the cognitive load. I think we have some incredible pilots out there. And of course, several panelists have talked in other discussions have talked about the importance of STEM education and making sure we have that level of understanding and literacy so when we want to communicate what our unmanned systems are doing, there's a level of training and understanding for what these semantic descriptions of those algorithms mean.

Mike Atwood:

Yeah, it's interesting for me. I've watched the training curriculum of the MQ-1 and the MQ-9 and the pilots that have come through and all the training that goes into being a weapons qualified officer in the platform. I've had the pleasure to work with the 26 Weapon Squadron up in Nellis Test and Training Range, and I've realized the strength of our Air Force is in our people. The ingenuity that I've seen these war fighters do with the platform that is now almost 20 years old ...

There was just some exercises done out in Valliant Shield where the operators of the MQ-9 did things that I never thought were possible with the systems that I designed in my younger engineering career. And so, yes, I do think the Airmen of today and the Guardians really have the ingenuity, the innovation, the capability to take our war fighting systems and do things as the designer that I never imagined with them. And that really excites me for the future, and especially if we do more aggressive manned/unmanned teaming and we bring more automation and capability to the war fighter. I'm just so excited to see what they can do with it.

John Clark:

Yeah, I'll share a funny story out of that. We were going through a human factors experiment with a set of users that came into our system where we were putting them through this new technology and this new capability that we were putting through its paces to go to the field, and we actually had some MQ-9 former operators that came in to go look at our system. And what we were doing, we were eliminating the sensor operator as a part of the autonomy where there was no longer a sensor operator. But folks that had been MQ-9 sensor operators came in. And so, we put them through a mission and evaluated how well they performed the mission. And apologies to the pilots in the audience, but the sensor operator kicked all their butts. And so, she specifically did a fantastic job of using the tools and the technology in ways in which the pilots were not accustomed to.

And so, I think that that's going to be an interesting dimension of how we go through that training equip process, looking at the individuals that we're going to be training them, putting them in new circumstances and figuring out how to indoctrinate them to using the tools in a new way because that

experience, it's something that stuck with me where, by a large margin, that sensor operator, she completely executed the mission successfully where other pilots were still messing around pushing on buttons and trying to make things do exactly what they wanted, where she trusted the system and it did a lot of stuff. So I think that getting that training in there and helping people understand new dimensions, it's going to be incredibly important.

RAAF ACDR John Haly:

Well, I think you're probably not the first and you won't be the last person to realize that pilots have limitations. What actually did happen in the time since when we started talking is they did start the clock and we've come to the end of our time, unfortunately. But would you join me in thanking Ben, Mike and John representing General Dynamics, General Atomics and Lockheed Martin.

Mike Atwood:

Well, thank you, John. It's always a pleasure to talk about this. I have a deep passion for not only the planes and the autonomous stuff we build, but the war fighters. Maybe I thank you to the audience and all these ... it's nice for me being a civilian to walk around and everyone in their flight suits and the military. I feel really honored to be an American and be able to give you guys technology that help us fight this war as it ever increases against the peer adversaries. So I really enjoy this and thank you for everything that you do for our nation.

RAAF ACDR John Haly:

Thanks very much.