LrrK Letters

Vol. 1 · Issue 6 · March 2026

# The Coordination Deficit

*Border Orders, Classified Incidents, Airspace Risk, and the Race to Unify Drone Defense*

**Dateline: March 13, 2026**

## Abstract

Counter-drone capability can fail between a correct observation and an authorized response. The March record makes that institutional problem difficult to ignore: reported Texas incidents involving government aircraft prompted senatorial criticism, while an enterprise contract promises a common technical backbone and an issuer announces a large autonomous border project. These events concern different programs, but together they expose the distance between acquiring components and coordinating a public mission. This essay examines shared information, legal authority, prime-contractor responsibility and the transfer of Ukrainian experience. It distinguishes a contract ceiling from spending, an order from accepted operation and a common screen from a common decision. The argument is that coordination must be specified, practiced and maintained as a capability of its own, with institutions retaining responsibility for the choices that software makes easier to execute.

## When protection needs protection

A counter-drone system can work in a narrow technical sense and still produce the wrong public outcome. It can detect an aircraft, act on it and leave the institutions involved confronting the consequences of a mistaken identification or an unresolved safety boundary. The important failure may be neither an invisible object nor an incapable interceptor. It may be a decision made by one organization without the knowledge another possessed. That is the coordination deficit: capability distributed across institutions that have not made the relationships between their capabilities dependable.

Reuters' March 4 account of Texas incidents gives the problem a concrete form. It reports that the military mistakenly brought down a government drone on February 25 using a laser-based system, and describes airspace restrictions connected with counter-drone activity near Fort Hancock and El Paso. Senators criticized interagency coordination after a classified briefing. Those are the reported incident facts and public reactions. The briefing's classified contents and underlying sensor records are not available here for an independent reconstruction.[1]

This is an important limit on what can responsibly be inferred. A report of a mistake does not tell an outsider exactly which observation, order or assumption failed. It does establish a reason to ask how a protective mission can create another risk that aviation authorities must manage. Speculation about hidden motives would add little. The public question concerns responsibilities that should have been clear before equipment was employed: who knew which aircraft were operating, who understood the relevant hazards and who could resolve a disagreement in time.

The failure is easy to caricature as bureaucracy obstructing useful technology. Some procedures may indeed be slow, fragmented or badly designed. Yet coordination with civil aviation is not an ornamental requirement. A

protective action can affect people and operations outside the institution carrying it out. Removing consultation without resolving the underlying safety question may make an action faster while making the overall system less competent. The aim should be to shorten the route to a sound decision, not simply to remove every actor who might ask whether it is sound.

The opposite caricature is to treat uncertainty as a reason never to act. Defensive work often requires decisions before every fact is known. An institution can reasonably accept some risk when the alternative is worse. What matters is whether that acceptance belongs to an authorized decision-maker with a usable understanding of the situation. A hurried operator should not be left to invent the agreement among agencies at the moment an aircraft appears. That agreement is part of the capability the public expected to purchase.

There is an emotional pressure behind both shortcuts. Officials want to show that a new threat can be met. Operators want tools that do not leave them helpless. The public wants protection without discovering that the protection has disrupted ordinary life or endangered another government operation. Those expectations are compatible in principle and difficult in practice. A sensible account should make the difficulty visible without converting every delay into cowardice or every intervention into a success. Competence includes knowing when to withhold an action and how to explain why.

The lesson is therefore broader than the particular equipment described in Texas. Any arrangement connecting sensors, analysts, decision-makers and physical response can fail at an institutional boundary. A shared system can help if it makes the relevant facts and responsibilities easier to understand. It can worsen the problem if it gives a mistaken conclusion wider reach. The quality of the connection matters as much as the speed. Integration should be assessed through decisions improved and mistakes prevented, not merely the number of devices connected to a screen.

## The shared screen and the shared decision

Today's War Department account presents a common command-and-control platform as a response to fragmented defenses. JIATF 401 describes an enterprise agreement centered on Lattice, intended to support data sharing and coordinated responses across organizations. Its expectation is that a common technological backbone can connect sensors and responses more effectively. That is a substantial institutional proposition. The announcement does not by itself establish the quality of every connection, settle each organization's authority or demonstrate the outcome of a joint operation.[2]

A shared screen can provide real value. People may spend less time reconciling different tracks, translating identifiers or finding another agency's point of contact. A common system can preserve information that would otherwise be lost when an observation crosses an organizational boundary. It can make a decision history easier to review. These are practical benefits, rather than merely attractive graphics. They should be specified as functions that users can test, because a common visual appearance can otherwise conceal differences in what each participant believes the information means.

Consider an illustrative track seen by two sensors. One observation may be old, another uncertain; both may refer to the same aircraft or to different ones. If the display merges them without retaining that history, the resulting symbol can look more authoritative than the evidence supporting it. If it leaves every observation separate, the operator may struggle to understand the scene. Fusion is a judgment about correspondence and confidence. A useful system should help the responsible person understand that judgment without demanding a forensic investigation during every routine decision.

Identity poses a related problem. The relevant question can include whether an aircraft is known, permitted, government-operated or behaving in a way that warrants concern. Those are different attributes. A system should

not silently turn the absence of one attribute into proof of another. An unidentified object is not automatically hostile, and an authorized identity does not guarantee harmless behavior. Institutions need agreed meanings and procedures for updating them. A technical interface can carry those decisions; it should not manufacture the legal conclusion through an ambiguous label.

Authority must remain distinguishable from visibility. Two organizations may see the same track while possessing different powers to respond. The permissions can depend on location, mission and circumstances. Sharing a platform does not merge the agencies into one legal actor. It should instead make the relevant decision route easier to follow: who is responsible, who can approve a response and what information must be considered. If the system gives someone the ability to click a control, that ability should correspond to an actual delegation, rather than stand in for one.

A common platform also changes the consequences of error. A local misunderstanding can be distributed rapidly. A software update can affect several organizations at once. An access-control problem can give the wrong person influence over information used elsewhere. These are reasons to govern shared infrastructure seriously, not reasons to reject it. The benefits of integration are strongest when the institutions understand what becomes shared and what remains local. They need an account of degraded service, mistaken data and recovery that is as practical as the account of normal operation.

The commercial relationship belongs inside that account. The platform provider may contribute expertise that government institutions do not possess internally. It may also become difficult to replace if interfaces, training and records depend heavily on its services. A public buyer should examine how information can be retrieved, how another system can connect and what happens when terms or suppliers change. The question is not whether one vendor is inherently untrustworthy. It is whether a public mission retains choices after adopting useful infrastructure.

The strongest defense of a common backbone is that fragmentation has its own cost and dependency. Many separate arrangements can create duplicated purchasing, inconsistent security practices and gaps no participant owns. Common standards and a well-supported platform may reduce that burden considerably. The appropriate test is comparative: which risks are reduced, which become concentrated and which safeguards keep the arrangement governable? A serious integration policy should be able to discuss both sides without treating a single platform as either a universal cure or an inevitable capture of the state.

## Someone has to own the integration

Ondas' March 4 release describes another form of consolidation. The company says its Airobotics subsidiary received an initial $20 million purchase order under a multiyear national border-protection program and would lead development, customization, integration and staged deployment. Its proposed architecture includes large drone fleets, ground infrastructure and command software. The customer is described as a major government entity, not publicly identified in the announcement as the United States. It would be a mistake to merge this order with the Texas incidents merely because both involve borders.[3]

The appeal of a prime contractor is accountability. Instead of asking a buyer to make several suppliers' components work together, one organization undertakes to deliver a functioning arrangement. That can improve coordination if the prime has the authority, information and resources to resolve conflicts. It can also simplify the public story without simplifying the underlying dependencies. The buyer should know which responsibilities the prime has actually accepted and which remain with agencies, local operators or other suppliers outside the contract.

Integration has a cost that a component price does not reveal. Equipment must be configured for its location, communications must be dependable enough for the intended use and personnel must understand the procedures. An installed item may perform differently when combined with another subsystem. A system can satisfy each supplier's local acceptance test while failing the overall mission. The contract should therefore identify what counts as integrated performance, including the conditions in which that performance will be assessed and the responsibilities triggered when it falls short.

An initial order can finance a useful first stage without establishing the final national architecture. It may pay for design work, integration or staged deployment whose success determines later purchases. The distinction helps the buyer as well as the reader. A program can learn from an early phase and adjust the next one. It becomes harder to do that if the initial announcement has already been treated as proof of full deployment. The pressure to defend a grand narrative can then compete with the practical need to admit what the first installation revealed.

Scale introduces its own operational questions. A fleet requires more than aircraft counts. Availability depends on maintenance, charging or replacement arrangements, staffing, communication and the handling of faults. More simultaneous information can increase the burden on the people reviewing it. Automation may reduce that burden in some cases and create new kinds of oversight in others. A system designed around a persistent mission needs evidence about persistence, including what happens when one part is unavailable. Its intended scale is a requirement to investigate, rather than an achievement established by the order.

The public institution should retain the ability to understand its system even when it delegates integration. It needs enough information to distinguish a supplier fault from an operational procedure and to make a sensible change later. That does not require owning every design or doing every engineering task internally. It requires avoiding a relationship in which only the vendor can explain what happened. A contractor's expertise is most useful when it enlarges the customer's capacity to govern the mission instead of making that capacity permanently dependent on the contractor's interpretation.

Acceptance is the moment at which these distinctions become concrete. The buyer should establish which outcomes, conditions and failure responses matter before a delivered system is declared ready. A local test cannot settle every future condition, but it can reveal whether the agreed arrangement exists and where its limitations lie. The record should preserve those limitations for the operators who inherit the system. Otherwise a qualified acceptance can gradually become an unqualified promise as it moves through procurement reports, executive summaries and public announcements.

## Experience cannot be shipped without a receiving institution

Reuters reported on March 5 that the United States and Qatar were discussing Ukrainian interceptor capabilities with Kyiv. The account relies partly on sources familiar with the discussions and distinguishes interest from completed acquisition. It also raises a practical constraint: experienced crews and training capacity cannot be assumed to move with equipment. No particular quantity, delivered configuration or local validation should be inferred from those talks. The report establishes a demand for experience acquired elsewhere and uncertainty about how that experience can be transferred.[4]

The distinction is easy to lose because battlefield experience has understandable authority. A system repeatedly adapted under pressure may contain insights a peacetime development process would take longer to acquire. Buyers should take that learning seriously. They should also ask which part they are acquiring. An aircraft, operating method, training relationship and software service are different things. Their usefulness may depend on one another. Buying the airframe without the surrounding knowledge can give the purchaser a product with an impressive history and little capacity to reproduce its results.

The receiving environment may change the problem. Climate, terrain, communication conditions and aviation arrangements can alter how a system performs and how it may be used. Operator preparation and maintenance practices can differ. None of this erases the value of the original experience. It explains why transfer requires local work. A responsible buyer should seek the shortest sound route to understanding those differences, rather than demand that every lesson be rediscovered or assume that a reputation makes adaptation unnecessary.

There is also a responsibility toward the source of the knowledge. A force still fighting for its own security may have scarce experienced personnel. Sending those people elsewhere can impose a cost not captured by the aircraft price. Training partners can create useful longer-term capacity, but the immediate tradeoff remains real. A transfer agreement should recognize what is scarce: the product, the instructor's time, the capacity to support changes and the knowledge generated through continuing operations. Treating experience as a free accessory understates the contribution being requested.

This is another coordination problem because the parties need shared expectations about the outcome. Is the first purpose evaluation, urgent limited use or a sustained production and training relationship? Who determines whether the transferred system is ready? Which organization records failures and communicates a change back to the original engineers? If those questions remain implicit, participants can work energetically toward different definitions of completion. Explicit goals can make cooperation faster by preventing a later dispute about what everyone believed the initial delivery meant.

Data should travel with a corresponding account of meaning. An operating result from one setting can support a hypothesis about another. It cannot settle that hypothesis without attention to conditions. A partner might need a description of what was observed, how outcomes were confirmed and which limitations remain. Sensitive detail can be protected while still allowing competent evaluation. The receiving institution's task is to turn inherited knowledge into a capability it can explain and sustain, not merely repeat the original supplier's strongest claim.

The reward can be substantial. Successful transfer can diversify supply, improve defense and give the original producer resources for further development. It can connect people learning from different circumstances. The relationship becomes more durable when useful findings flow back, including failures the original environment did not expose. Cooperation then produces additional knowledge rather than simply redistributing a finished artifact. That is a better account of technological partnership than one in which a product crosses a border and is presumed to carry its entire operating history intact.

## The contract is a route, not the destination

The enterprise announcement's largest number needs a precise reading. In a separate March 13 account, the Army says the Anduril contract has a five-year base period and a five-year optional ordering period, with a maximum potential value of up to $20 billion. It expressly states that this is not an obligated amount. The framework covers commercially available software, hardware, data and computing infrastructure, and support; it is broader than one counter-drone purchase. A route for placing orders should not be reported as if the entire ceiling has already been spent.[5]

The Army gives a concrete administrative rationale: more than 120 separate procurement actions had been managed for Anduril offerings. Common terms and pricing can reduce repeated negotiations and make ordering easier to track. Those are plausible benefits. Their realized size needs subsequent evidence about time, cost and actual use. The agency also says the agreement will not substitute for competition on future programs. That statement provides a public standard against which later choices can be judged, rather than proof that every future opportunity will be equally accessible to competitors.

Consolidation can make oversight easier by providing a common point of reference. It can also concentrate commercial influence. A buyer accustomed to one framework may find the familiar route easier than evaluating an alternative. A supplier outside it may need to demonstrate more than technical merit to overcome the administrative advantage. These are possibilities inherent in procurement design, not findings that the announced contract has been abused. They suggest that the public should examine switching costs and entry routes alongside savings from common terms.

The distinction also appears in AeroVironment's space business. AeroVironment's March 3 statement described active negotiations over an amendment to its existing Space Force SCAR ground-station contract, which it said had been temporarily paused. That is a space-communications matter, not evidence of a counter-drone award. Its relevance here is the institutional distance between a commercial expectation, a negotiated contract structure and the work a customer ultimately requires.[6]

The subsequent March 10 results make the financial consequence more specific. AeroVironment reports a $151.3 million goodwill impairment in its Space reporting unit, connected to revised expectations after a January stop-work order on BADGER systems supporting SCAR. A goodwill impairment is an accounting change to recorded value, not a statement that an identical amount of cash was spent or lost during the week. Read together, the disclosures show why the March 3 negotiating statement should not be mistaken for a fresh final award or assurance that expected work will proceed unchanged.[7]

These distinctions are useful because they preserve responsibility. The company must describe its position accurately. The contracting institution must manage scope and acceptance. Investors must assess a business whose expectations can change. None can borrow certainty from the others without examining the underlying commitment. A procurement framework may improve administrative coordination, but it cannot remove the need to resolve what is being purchased, how it will be paid for and what will count as successful delivery.

The same is true of a common command platform. Its procurement can be centralized while operations remain distributed. A contract can supply interfaces without settling who has authority to use a response. A software service can reduce friction without deciding which institution bears responsibility for an error. The financial and technical layers should support that allocation of responsibility, rather than make it disappear behind the promise of integration. The most useful common system is one that makes the chain of decisions easier to follow.

## Practicing the handoff

Coordination is often spoken of as a meeting that should have happened. In a persistent technical mission it is better understood as a maintained operating capability. People need shared meanings, current contacts, tested procedures and a way to recover when the usual channel fails. Those arrangements change as staff rotate, software is updated and missions evolve. A single successful exercise cannot guarantee them forever. They require ownership and attention comparable to the maintenance given to the equipment that depends on them.

A useful exercise should therefore test handoffs, not only components. Can one organization explain the basis of an observation to another? Can conflicting information be resolved? Does the person receiving a decision understand its scope and limits? Can an authorized operation be recognized in time? These questions can be explored without assuming that every participant should have identical access or authority. The aim is to make the differences workable. A process that succeeds only when everyone has the same privileges is poorly suited to an interagency mission.

Failure should be recorded in a form that supports repair. An ambiguous alert, a delayed approval and a misidentified aircraft are different events. Their consequences should not be flattened into one performance count. An institution needs to understand whether the problem arose in sensing, interpretation, communication or

authority, and whether a proposed correction addresses that cause. Otherwise another system may be purchased to solve a failure that actually lives in the procedure. More equipment can make an unresolved coordination problem larger.

The public also needs a proportionate account. Classified information may properly protect coverage, methods or sensitive activity. It should not prevent responsible officials from explaining the categories of problem, the corrective responsibilities and the standards used to judge improvement. A closed briefing can be necessary and still leave an obligation to communicate what can be said. Without that distinction, secrecy can become the place where difficult institutional questions are stored rather than resolved. Public confidence requires something more than assurance that an explanation exists elsewhere.

The strongest case for the enterprise approach is that the state cannot meet a fast-moving problem through indefinitely fragmented purchases. Common infrastructure and clear integration responsibility can make a material difference. The strongest caution is that a large agreement can create the appearance of unity before institutions have learned to act together. Both points can be true. The question is whether procurement finances the work of coordination or becomes a substitute for demonstrating it.

As of March 13, there are meaningful commitments and serious warnings in the record. Government accounts describe a new enterprise route; an issuer describes an initial border order; partners discuss transferring Ukrainian experience; Texas incidents have prompted sharp criticism of coordination. These developments do not share one contract or one authority. Their common lesson is that capability emerges through relationships that someone must make dependable. A protective system succeeds when the right institution understands enough, acts within its responsibility and can explain the result. The contract can help provide the tools. The handoff still has to work.

## Notes and sources

1. David Shepardson, Reuters, March 4, 2026, account of Texas government-drone incidents and senators' comments after a classified briefing. Incident descriptions are attributed; underlying classified records were not independently reviewed. Specific chronology of the separate El Paso closure is not relied upon here. Contemporaneous report.
2. War Department, March 13, 2026, JIATF 401 account of the enterprise command-and-control arrangement and Lattice. Describes intended interoperability benefits, not independent operational validation. Task-force account.
3. Ondas, March 4, 2026, original announcement of Airobotics' initial $20 million purchase order and staged border-program integration. Customer nationality and accepted deployment outcomes are not established by the release; planned scale remains an issuer claim. First-party announcement.
4. Max Hunder and colleagues, Reuters, March 5, 2026, report of U.S. and Qatari discussions with Kyiv. Used for attributed talks and transfer/training questions, without asserting completed procurement or local performance. Contemporaneous report.
5. U.S. Army, March 13, 2026, enterprise-contract announcement. Explicitly identifies up to $20 billion as maximum potential value rather than an obligation, five-year base plus five-year optional period, scope and prior separate procurement actions. Army account.
6. AeroVironment, March 3, 2026, statement about a paused existing SCAR contract and negotiated amendment. An issuer account of space ground stations, not a newly awarded drone contract. Original statement.
7. AeroVironment, March 10, 2026, fiscal third-quarter results. Reports a January stop-work order and $151.3 million Space-unit goodwill impairment; an accounting disclosure, not a cash-loss or mission-performance measure. Dated results.