LrrK Letters

Vol. 1 · Issue 3 · January 2026

# Combat Data as an Industrial Input

*Mission Software, Counter-Drone Authority, and the Institutional Stack Behind Drone Dominance*

**Dateline: January 30, 2026**

## Abstract

Ukraine's plans to make wartime information available for military artificial intelligence place a difficult resource at the center of industrial competition. Combat data can reveal failures that peacetime demonstrations miss, but its value depends on collection, interpretation, access and the conditions under which a result travels. This essay examines the announced Brave1 Dataroom alongside mission-control plans, changes reported in American counter-drone governance, an FAA reorganization and the forthcoming Super Bowl airspace restrictions. These developments expose a common problem: information becomes useful through institutions able to act on it. The argument follows the movement from experience to dataset, dataset to model and model to authorized operation. It asks who receives the resulting advantage, whose experience becomes invisible, and what public obligations should accompany the commercial use of costly operational knowledge.

## The resource no factory can order

A factory can buy a camera. It cannot place an ordinary purchase order for the circumstances under which that camera fails in war. Those circumstances arrive through danger, improvisation and loss. They are recorded unevenly by people whose first obligation is to survive and accomplish a mission. When the recordings become training material for artificial intelligence, something consequential changes. Experience is being converted into an industrial input. The conversion can improve equipment and save lives. It can also conceal the choices that determine whose experience counts as knowledge.

On January 20, Reuters reported Ukrainian Defence Minister Mykhailo Fedorov's intention to let allies train military AI on wartime data. The report connects that ambition to a proposed drone mission-control system and a more explicitly measured approach to military management. The following day's Ministry of Defence announcement describes Brave1 Dataroom, built with Palantir software, as an environment for training and testing military models. Its initial emphasis is aerial-threat detection and interception, with structured visual and thermal material collected by service members. These are related developments at different stages, not evidence that every contemplated allied data-sharing arrangement is already functioning.[1][2]

The appeal is easy to understand. An aircraft photographed against a clean sky does not provide the same challenge as a distant object seen through clutter, glare or obscuration. An adversary changes tactics and equipment in response to what defenders do. An operating force encounters combinations that a laboratory might not anticipate. Data from those encounters can tell developers which neat assumptions should be abandoned. It can make a national industrial effort less dependent on what an engineer imagined a battlefield would look like.

Yet the value is not sitting inside the image like a mineral waiting to be extracted. Someone must establish what the object was, whether it was identified correctly, what happened outside the frame and which relevant conditions were recorded. The same footage can support different inferences depending on its history. A bright point can become a target label, a successful interception or a misleading example if the surrounding evidence is missing. Calling the material real-world data identifies its origin. It does not finish the work needed to make it reliable.

This distinction matters particularly when the resource becomes a bargaining asset. Ukraine's government has reason to seek practical returns from knowledge accumulated at extraordinary cost. Partners have reason to want access that improves their own systems. Companies have reason to develop products that could enter a much larger market. Their interests can align around improved defense without automatically aligning around ownership, disclosure or future dependence. The terms of access will help determine whether the exchange strengthens Ukrainian capability or merely supplies another organization's proprietary advantage.

There is no need to choose between gratitude for useful cooperation and scrutiny of its terms. A partner can contribute valuable infrastructure while acquiring a powerful position in the workflow through which data is searched, prepared and used. A government can protect sensitive information while making it difficult for outsiders to evaluate the quality of a claim. A developer can improve a model while learning little about cases that the dataset omits. The industrial opportunity is real precisely because the resource is difficult to reproduce. That scarcity makes its governance a matter of strategy.

## What the room contains, and what a label leaves out

The ministry's January 21 account gives the proposed arrangement a more specific shape than a general offer of battlefield footage. It says the platform contains structured visual and thermal datasets of aerial targets, including Shahed-type drones, and will expand them over time. Ukrainian defense developers are invited to apply, with access subject to a security-compliance procedure. That describes a controlled environment and an initial participant route. It does not describe an unrestricted public download, nor settle how the broader allied-training ambition reported by Reuters will be implemented.[2]

Controlled access is a reasonable response to the sensitivity of the material. Images can expose sensor positions, operating methods, locations or the limits of a defensive system. Even a record stripped of obvious identifiers can reveal patterns when combined with other records. Protecting those facts can be part of protecting the people who generated the data. The important question is whether access controls are designed around specific risks and legitimate uses, rather than becoming a general answer to every inquiry about how the resulting system was trained.

Inside the room, annotation is an exercise of authority. A label tells a model what a human institution has decided the example means. Some labels may rest on strong confirmation; others may be provisional judgments made under pressure. If the difference disappears during preparation, the training process can give uncertainty the appearance of certainty. An accurate record of disagreement may be more useful than a confidently incorrect category. The aim should be to preserve enough of the reasoning that a later user can recognize when the example deserves less weight.

Selection has equal importance. Footage available for training is not automatically a representative sample of everything a military encountered. Some systems record more consistently than others. Some crews preserve failures; others may lose the recording when the aircraft is lost. Events that attract attention may be overrepresented. Quiet stretches without a threat can appear less valuable even though they are essential for understanding false alarms. These are general risks in constructing an operational dataset, not allegations about the contents of Brave1 Dataroom. They should inform the questions asked of any such collection.

The missing denominator can be decisive. A collection of correctly identified targets may teach recognition while saying little about how often an alert would be wrong during ordinary surveillance. A set of successful intercepts can show what is possible while obscuring failed attempts, aborted missions or the amount of operator work involved. The data needed to improve a particular model is not always the data needed to judge whether a whole defensive arrangement is worth deploying. Those purposes overlap, but they should not be confused.

The division between training and testing deserves practical attention. Two nearly identical frames from the same flight can look like separate examples in a table. If one helps train the model and the other is used to test it, the result may reward familiarity with that encounter rather than recognition in a new one. Separating material by event, collection conditions or time can make the question more demanding. The right choice depends on the claim being evaluated. A developer trying to improve performance on an existing sensor needs a different comparison from a buyer asking whether the result transfers to another operating area.

This is why a headline accuracy figure is insufficient without a description of how the evaluation was constructed. The evaluator needs to know which examples were withheld, which changes were made after seeing the results and whether the final test remained genuinely unfamiliar. The point is not to impose one ritual on every model. It is to avoid a situation in which everyone believes a demanding test has occurred while each participant means something different by the word. A defensible result gives people enough information to disagree about its relevance intelligently.

Context also determines how knowledge travels. A model trained on a particular sensor, altitude range and operational environment can encounter unfamiliar conditions elsewhere. Another buyer may use different optics, compression, flight procedures or communication links. A change that seems minor in procurement can change what the model sees. The responsible claim is therefore narrower than tested in combat. It concerns the situations represented, the task evaluated and the limits encountered. Operational experience becomes transferable capability through this translation, rather than by borrowing the moral weight of its origin.

An accessible explanation of a model's limits need not expose protected tactics. A program can specify relevant sensor characteristics, categories of environmental coverage, treatment of ambiguous labels and methods used to separate training from evaluation. It can report failure types without publishing sensitive footage. The public may not receive the underlying dataset, but the institutions buying and authorizing a system need enough access to distinguish a robust evaluation from a promotional summary. Secrecy changes the audience for accountability. It should not eliminate the work.

## Measuring crews changes the record being measured

Fedorov's reported mission-control plan introduces a second transformation. Data would serve military management as well as machine learning. Reuters describes plans to make drone crews' performance and effectiveness more visible to decision-makers. Such a system could help commanders identify useful practices and understand resource use. It also changes the incentives of the people producing the record. Once a measurement affects recognition, resources or continued responsibility, it becomes part of the environment in which the measured activity occurs.[1]

The strongest case for measurement starts with the cost of not knowing. Scarce aircraft, ammunition, time and specialist personnel should not be allocated by habit when reliable operational information could reveal a better choice. A force may need to identify a weak process quickly, understand why one unit performs differently or stop repeating a costly mistake. A clear record can protect successful teams from being overlooked and make equipment failures harder to blame on operators. It can improve judgment by giving a commander a more accurate account of what is happening.

But a measure of effectiveness always contains a choice about the outcome that matters. Completed sorties are easier to count than useful reconnaissance. A destroyed object may be easier to show than a threat deterred or a mission prudently abandoned. A team performing difficult work can appear less productive than one operating under favorable conditions. If the management system rewards a simple count without recording those conditions, it may direct resources toward the most legible activity. The resulting dataset then reflects the incentive as well as the war.

This does not make systematic measurement impossible. It makes comparison a substantive task. An evaluation should account for mission difficulty, exposure, equipment state and the confidence attached to an outcome. It should allow people to explain an apparent failure without treating every explanation as an excuse. A report that distinguishes an equipment fault from an unsuccessful tactic can help both procurement and training. A report that merges them into one score can make everyone optimize the wrong thing with impressive numerical precision.

There is also a dignity issue inside the administrative one. The people generating combat records are not merely sensors in a collection network. They have knowledge about the circumstances that produced the data, including facts that a standardized field may not capture. They may also bear the consequences when a distant model or dashboard misreads their work. A useful system should give them a route to contest a record, correct a label and contribute an explanation. That is a way of improving the evidence, rather than a sentimental exemption from accountability.

The feedback loop can become productive when the directions are kept visible. A crew reports a limitation. Engineers change the equipment or software. A new version is issued. Operators describe whether the change helped under comparable conditions. Commanders use that account when deciding what to buy or how to employ it. The loop becomes much harder to interpret if every stage stores a different identity for the same aircraft, or if the delivered revision inherits a claim established for an earlier one. Remembering versions is part of remembering experience.

For allies, the lesson is not simply to obtain more footage. It is to learn how the material was generated and what institutions made it credible. A partner importing a dataset without the collection discipline, operator feedback and interpretation behind it may acquire an impressive archive with a weak account of meaning. Building the capacity to learn together can be more valuable than transferring files once. That requires agreements about corrections, later discoveries and the return of useful results to the people who supplied the original experience.

## Detection meets the permission to act

An AI model's improved recognition does not complete a counter-drone system. The next steps involve a person or institution deciding what the observation permits. DefenseScoop reported on January 26 that the Pentagon had consolidated counter-drone guidance signed in December, including changes to defensive areas, treatment of unauthorized surveillance, interagency sensor-data sharing and designation of covered facilities. Its account also describes an inspector-general critique of unclear and inconsistent policy. These are reported governance changes and findings, rather than evidence that every installation now has effective protection.[3]

The relationship to combat data is direct. A model can produce a more useful assessment while the organization receiving it remains unsure who may respond. Conversely, expanded authority can make an imperfect assessment more consequential. Neither technical performance nor legal permission should be judged in isolation. The practical system is the chain joining observation, interpretation, decision and physical effect. Weakness in any link can dominate the result, including when every organization can point to a successful component test.

The American domestic setting makes that chain especially demanding. Nearby aviation, communications and ordinary public activity matter to a response. An unauthorized aircraft may represent negligence, curiosity,

deliberate surveillance or an immediate threat. The available time for distinguishing those cases may be short, but the distinction does not cease to matter. A commander's confidence should reflect both the reliability of the observation and the authority to use a particular response in that location. An alert is not a legal conclusion.

Information sharing can improve that judgment. Another agency may recognize the activity, hold relevant permission or provide a second observation. Sharing also raises questions about access, retention and use beyond the original defensive purpose. A system built to protect a facility can collect information about people who are not attacking it. The appropriate safeguards should follow the actual flow of information, including the conditions under which a record can become evidence in another process. Useful coordination needs rules precise enough to guide ordinary work before a crisis.

Today's FAA and FBI announcement for Super Bowl LX gives these abstractions a concrete future event. It establishes temporary no-drone zones around Levi's Stadium and several San Francisco venues. The game-day restriction includes a thirty-nautical-mile radius extending to 18,000 feet during the specified afternoon and evening period. The notice describes federal detection, assessment, enforcement and mitigation authority. As of January 30, it is a rule and an announced security arrangement for coming dates, not an after-action assessment of how well the event was protected.[4]

This distinction matters because the event will generate another kind of data. Investigators may identify an operator, recover an aircraft or determine that a reported threat was ordinary activity. Those dispositions can teach a system more than its initial alert count. Without them, a detection technology may accumulate apparently impressive observations while its users learn little about accuracy. A restrained response that resolves a mistaken alert can be a success. A system that reports only interceptions or seizures may make that success disappear from the institutional record.

The measure of coordination should therefore include how a case is closed. Which organization resolved the track? What information changed the decision? Was a response delayed by confusion, or prudently withheld because uncertainty remained? These are questions for authorized review, not a demand to publish sensitive defensive coverage. They connect operational learning with accountability. The same record that improves a future model can help an institution explain why it acted, and can reveal when its procedures need to change before the next event.

## The organizations that turn information into a market

On January 27 the FAA announced a plan to reorganize, including an Office of Advanced Aviation Technologies and an agency-wide safety-management organization. The announcement describes a designated point for emerging aviation integration and an effort to share safety information across previously separate lines of business. It is a recognition that new technology can cross the boundaries of an older organizational chart. An aircraft may need decisions about operations, airspace and certification whose assumptions must fit together.[5]

Organizational reform can reduce the cost of finding the person responsible. It can also create another layer between an applicant and a decision. The difference depends on authority, staffing and the information that moves with a case. If an office receives responsibility without the ability to resolve conflicts, it can become a mailbox for unresolved questions. If it can reconcile assumptions and preserve evidence across reviews, it may make the whole system faster. The announcement creates a reason to watch those functions, rather than evidence that coordination has already improved.

That administrative environment is part of the industrial input alongside combat data. A company must translate a promising model into equipment that can be tested, supplied and operated. It needs access to customers and a credible account of the rules affecting them. Capital providers make judgments about how much time and money

that translation will require. A government can encourage innovation through research access while leaving the deployment path too uncertain for a smaller company to finance. Conversely, an efficient path can turn a modest technical improvement into a useful service.

Dynamic Aerospace Systems' January 20 shareholder letter illustrates how many stages can coexist inside one commercial narrative. Filed as an SEC exhibit, it discusses demonstrations, a Drone Dominance proposal, component relationships, intended procurement access and financing ambitions. Those statements reveal the company's strategy. They do not make a proposal an award, an intended listing an authorization or an MOU a completed deployment. The significance is the combination: technical development is being financed and described through expectations about institutions that control access to demand.[6]

An investor is entitled to consider that opportunity. A public buyer needs a different kind of answer. It must know what exists now, what remains under development and what resources would be needed to deliver the promised configuration. A researcher evaluating an AI result needs still another account, including the task, data and evaluation conditions. A healthy industrial system allows these descriptions to refer to one another without pretending that one document serves every audience. A shareholder letter can introduce a hypothesis about a company's future. A delivery and acceptance record answers whether a particular promise was met.

Data access can also affect the shape of competition. Firms that learn through a common environment may improve rapidly, but the platform operator can acquire influence over the methods and interfaces that become normal. This can be useful coordination or an expensive dependency, depending on the terms. Governments should consider whether participants can retain useful knowledge, move their work where permitted, understand changes in access and receive fair treatment. The risk is not limited to a malicious provider. A commercially rational change in terms can still reshape an industrial ecosystem.

The return to the data's public origin should be part of the bargain. Companies need opportunities to earn money and fund further development. Ukraine and its partners need a capacity that survives a particular commercial relationship. Agreements can protect sensitive data while requiring meaningful feedback, support and access to results necessary for continued operation. The issue is less whether a firm owns some intellectual property than whether the public institutions and operators whose experience made it possible retain practical choices afterward. A valuable partnership should enlarge their agency.

## Learning without turning war into a showroom

There is an understandable desire to make suffering yield something useful. Better detection, safer operations and more effective defense can be important returns. That desire should make the terms of industrial learning more demanding, not less. Combat experience deserves neither mystical authority nor casual extraction. It is evidence produced under conditions that require care to interpret and responsibility toward the people involved. A market that celebrates access without attending to those conditions risks making the origin of the material a promotional asset while discarding its hardest lessons.

The strongest argument for the new arrangements is that isolation wastes experience. A nation confronting urgent threats should be able to learn with partners, use capable technical infrastructure and improve equipment at a useful pace. An imperfect initial dataset can still be more valuable than a beautifully designed peacetime exercise that misses the relevant failure. A controlled environment can protect material while making it available to people able to improve defenses. Demanding certainty before any learning begins would be a costly misunderstanding of how development works.

The strongest objection concerns the terms of acceleration. A result can travel farther than its evidence, particularly when combat-tested becomes an unqualified product attribute. A performance measure can reward

legible activity over useful work. A platform can make data available while narrowing the choices of its users. A new authority can turn a better observation into a response before institutions have sufficiently considered error. These are different hazards. They call for specific controls rather than one broad declaration that AI or data is either the solution or the threat.

The work ahead should connect three accounts. First, the history of the material: where it came from, how uncertain judgments were handled and which conditions it represents. Second, the history of the system: what was trained, tested, changed and delivered. Third, the history of use: who was authorized to act, what occurred and what later investigation revealed. Keeping those accounts connected makes correction possible. It also lets a buyer recognize when an impressive claim no longer refers to the machine or environment in front of it.

As January ends, the important achievement is an opening toward more organized learning, not a demonstrated final outcome. Ukraine is asserting that its experience has strategic value. Developers are being offered a controlled route to some of it. American institutions are changing how emerging aviation is organized and how counter-drone authority is described. The test will be whether these arrangements return knowledge as usable capability, preserve the ability to correct mistakes and strengthen the people supplying the original experience. Data becomes an industrial resource through institutions. Those institutions should remain answerable for what the resource is made to do.

## Notes and sources

1. Max Hunder, Reuters, January 20, 2026, account of Fedorov's announced allied AI-training and mission-control plans. Used for proposed policy and stated management aims, rather than proof of completed rollout or model performance. Contemporaneous report.
2. Ministry of Defence of Ukraine, January 21, 2026, Brave1 Dataroom launch announcement, with metadata indicating an update January 22. Specifies initial visual and thermal data, Ukrainian developer access and security compliance. Does not publicly establish comparative model outcomes or complete dataset composition. Ministry announcement.
3. Drew F. Lawrence, DefenseScoop, January 26, 2026, reporting December policy changes and an inspector-general critique of base-defense governance. Used as attributed reporting; the complete underlying guidance was not reviewed. Contemporaneous account.
4. FAA and FBI, January 30, 2026, Super Bowl LX restriction announcement. Specifies forthcoming dates, geographic and altitude limits and federal enforcement arrangements; no event outcome is inferred. Original notice.
5. FAA, January 27, 2026, organizational-reform announcement, including planned advanced-aviation and safety-management functions. Used for declared responsibilities rather than demonstrated improvements in coordination. Agency announcement.
6. Dynamic Aerospace Systems, January 20, 2026 shareholder letter, SEC Exhibit 99.1. Demonstrations, proposals, arrangements and financing objectives are issuer representations, not independent product validation or proof of procurement acceptance. Filed letter.