Cyberhill White Paper October 2026

From Purchase Request to Audit-Ready: Closing the Government Supply Chain Visibility Gap

How connecting procurement, shipping, property, and budget records makes every dollar and asset traceable before an auditor asks.

Eric Gibson

Senior Director, Public Sector Sales

Eric Gibson is Senior Director, Public Sector Sales at Cyberhill Partners, bringing experience as a former U.S. Foreign Service Officer and hands-on logistics experience at Amazon in both corporate and operations manager roles. He helps government agencies put their data and AI to work to strengthen supply chain visibility, accountability, and audit readiness.

EXECUTIVE SUMMARY

Government supply chain programs have made real investments in accountability: IUID (Item Unique Identification) and serialized custody tracking, General Services Administration (GSA) Schedules and Blanket Purchase Agreements (BPAs) for sourcing, dedicated property and financial systems, and an audit mandate that has been federal law for more than three decades. Yet one question remains hard to answer quickly in most programs: **where does a given purchase order (PO) stand right now, and can every dollar and every asset behind it be traced end to end?**

Today a purchase request, an order, a receipt, a shipping record, and a custody record can live in as many as five separate systems, sometimes fewer, never guaranteed to talk to each other. The Department of War has failed its financial statement audit eight years running, and the Government Accountability Office (GAO) and the DoW Inspector General (IG) trace a significant share of that failure to exactly this problem: fragmented systems, incomplete reconciliation, and property and inventory records that do not tie out (see "Why it matters now").

This is not a policy gap. The Defense Federal Acquisition Regulation Supplement (DFARS), PPBE (Planning, Programming, Budgeting, and Execution), the Anti-Deficiency Act, and the GAO Green Book already require the kind of traceability this paper describes; accountability offices are already trying to deliver it, by hand, across systems built years apart. What is missing is a capability that connects what already exists.

01

Federate in place.

Zero-copy collectors connect procurement, shipping, accountable property, and budget systems without moving or duplicating data.

02

Trace every transaction.

Every purchase request, order, receipt, and dollar resolves back to the system of record it came from, so there are no black-box answers.

03

Keep the systems you have.

Cerebro adds a governed layer over existing investments instead of replacing a finance system, a property system, or a PO system. Your program keeps its system of record.

The result is a supply chain program that can answer, in seconds, the questions an auditor, an inspector general, or a program manager already has the right to ask.

The visibility gap

No single system was ever built to hold all of the data from request through disposition. Procurement, shipping, accountable property, and budget execution are each served by systems chosen for their own purpose, often years apart, by different offices. The gap this creates is not a failure of any one of them; it is a boundary of what each was designed to do. It shows up at specific points in a program's own workflow:

- **Purchase requests and receipts live apart.** A generated request and an approved order often sit in a different system than the warehouse's receiving log, so confirming an item actually arrived means opening two systems and matching dates manually.

- **Shipping and custody transfer generate no unified signal.** Internal shipping marks a hand-off; delivery confirmation for an item tagged with an IUID or a Unique Item Identifier (UII) lands in the property system, and nothing automatically ties the two records together.

- **Property records do not natively expose the budget behind them.** A custody record shows who holds an asset, not the appropriation that paid for it, so tracing a dollar from budget line to physical location and back is manual cross-referencing, every time.

- **Budget execution can lag ground truth.** Anti-Deficiency Act guardrails depend on remaining budget being accurate, but obligations recorded in a financial system can trail what has already been ordered or received elsewhere.

- **Annual physical inventories rebuild the picture from scratch.** Because no single system holds the full lifecycle, verifying a serialized asset against its record each year means manually reassembling data that should already tie together.

- **Segregation of duties relies on process, not system controls.** Requesting, approving, and receiving are supposed to stay with different people, but each step potentially lives in a separate system with no common record. Nothing stops one person from performing two of them, and no record ties the steps together to show it happened. Enforcement depends on people following the process.

- **Contractor-held property is harder to see than government-held property.** GAO testing of contractor property records found errors in 116 of 270 sampled records, and identified 917 assets, worth roughly $109 million, that were not timely directed for delivery or recorded. That gap opens whenever the accountable system depends on someone else's report.

- **Institutional knowledge walks out the door.** When the people who know which spreadsheet reconciles which system transfer or retire, so does the ability to explain the next discrepancy.

Every one of these gaps is a reconciliation problem today. Under audit, or under an IG review, each becomes a finding.

Why it matters now

The Department of War's audit record is the clearest public evidence of what fragmented supply chain systems cost. GAO and the DoW Inspector General are consistent on the cause. This is a records problem more than a fraud problem, built from thousands of legacy systems that were never built to talk to each other, incomplete documentation, and inventory that cannot be reconciled fast enough for an auditor to rely on it.

0 of 8

DoW financial statement audits passed, ever (FY2018–FY2025)

1 of 24

Agencies under the CFO Act that have never earned a clean opinion (DoW alone)

28

Material weaknesses cited in the FY2024 DoW audit

$4.1 trillion

Total DoW assets on the FY2024 balance sheet

Indicator	Figure	Source
DoW financial statement audits passed, ever	**0 of 8** (FY2018–FY2025)	Military.com
Agencies under the CFO Act (Chief Financial Officers Act) of 1990 that have never earned a clean opinion	**1 of 24** (DoW alone)	GAO-25-108191
Material weaknesses cited in the FY2024 DoW audit	**28**, up from 20 in FY2018	GAO-25-108191
Of those, weaknesses severe enough to stop auditors mid-procedure	**17 of 28**	GAO-25-108191
Confirmed defense-related fraud, FY2017–FY2024	**$10.8 billion**	GAO-25-108191
Total DoW assets on the FY2024 balance sheet	**$4.1 trillion**	GAO-25-106868
Contractor-held assets GAO found undelivered or unrecorded on time	**917 items, about $109 million**	GAO-25-106868
Contractor property records sampled that contained data errors	**116 of 270 (43%)**	GAO-25-106868
FY2024 audit findings actually closed out	**955 of 3,322 (29%)**	GAO-25-108191
Target year for a first clean DoW audit opinion	**FY2028**	Military.com

These figures fall into four groups.

- **Audit performance:** the department has never cleared a bar every other CFO Act agency has cleared at least once. Eight straight failed audits, and DoW is the only one of twenty-four agencies that has never earned a clean opinion.
- **Internal control weaknesses:** 28 distinct weaknesses were cited in the FY2024 audit, and 17 of them were severe enough that auditors could not complete their procedures at all, meaning the records were too incomplete to check, not merely in dispute.
- **Financial and property exposure:** $4.1 trillion in assets sit on the balance sheet, $10.8 billion in confirmed fraud was reported over seven years, and a sample of contractor-held property came back with data errors in nearly half the records tested.
- **Remediation:** findings are accumulating faster than they close. Less than a third of FY2024 findings were resolved within the year.

> Together, these numbers describe one plain fact: an auditor asking where a specific dollar or asset stands cannot get a fast, confident answer today.

The pattern holds below the department level too. A program office that cannot produce a clean custody trail on demand is carrying the same risk at smaller scale: a finding at inspection, a scramble before an audit, and no confident answer when a contracting officer or an IG investigator asks where a specific dollar or asset stands right now.

The mandate: policy already calls for this

End-to-end supply chain traceability is not a capability searching for a requirement. The policies that already govern government procurement, property, and budget assume exactly the kind of connected record this paper describes; they simply predate the systems that would make it easy.

- **DFARS 252.211-7003, Item Unique Identification and Valuation.** Serialized government property must carry a UII and a valuation record traceable across its life. The clause assumes that record stays coherent from marking through disposition; it does not assume the systems that touch it along the way will agree with each other.
- **PPBE: Planning, Programming, Budgeting, and Execution.** The budget process that funds a program assumes execution data reflects reality: what has been obligated, what has been received, and what remains. Fragmented systems make that assumption harder to satisfy every cycle.
- **The Anti-Deficiency Act.** Obligations are not allowed to exceed available budget authority. Enforcing that guardrail requires obligations and remaining budget to be visible together, in something close to real time, not reconciled after the fact.
- **The GAO Green Book (Standards for Internal Control in the Federal Government).** Federal internal control standards call for information that is complete, accurate, and available when needed to support accountability, a standard that a fragmented, manually reconciled system structurally struggles to meet.
- **The CFO Act of 1990 and its audit mandate.** Every agency it covers is required to produce auditable financial statements. DoW is the one agency that has never met that bar, and GAO's own reporting names disconnected systems and incomplete property records as central reasons why.

The practical conclusion mirrors what other federal technology mandates have shown: the requirement already exists. Most programs are not missing policy. What they lack is a capability that connects the systems already required to comply with it.

The approach: one governed knowledge graph

Cerebro, Cyberhill's ontology-first knowledge graph, builds a living representation of a program's business and technical processes. Zero-copy collectors virtualize data where it already lives, using RDF (Resource Description Framework, a standard way of modeling linked data) and R2RML (a related mapping language) to federate metadata from every source into one graph, and an ontology and semantic layer on top give that graph real context and constraint. Every answer traces across the graph to its source.

Cerebro's four-step approach: Connect, Map, Model, Deliver

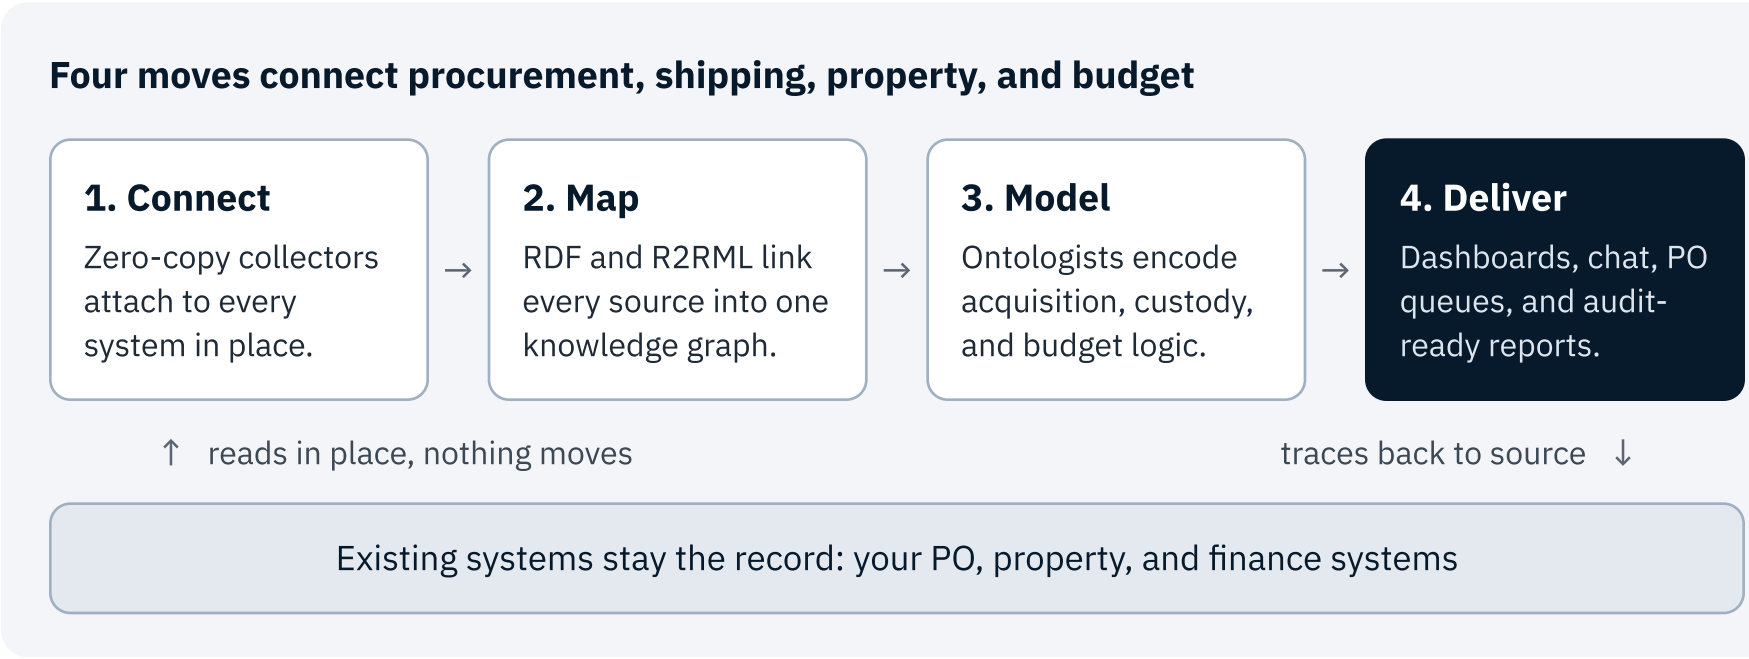

Connect and Map bring procurement, shipping, property, and budget systems into one graph without moving their data. Model gives that graph the meaning of a program's own acquisition lifecycle. Deliver puts traceable answers in front of the people who need them, while the PO system, the property system, and the finance system each stay exactly what they are today: the system of record for their part of the story.

Each step in the figure above breaks down as follows:

- **Connect.** Zero-copy collectors attach directly to the purchase order system, the shipping and receiving log, the property management system, and the finance system, wherever each one already runs today. No data is copied out or moved into a new database; the collectors read each system in place, so nothing about how a program uses those systems has to change.

- **Map.** RDF and R2RML mapping translates each source's own fields and formats into one common structure. A purchase order number in the PO system, a custody record in the property system, and an obligation in the finance system are linked together as parts of the same underlying transaction, even though each system stores that transaction differently.

- **Model.** Cyberhill's ontologists encode the program's own acquisition lifecycle, custody rules, and budget execution logic into the semantic layer sitting on top of the graph. This is what lets the graph understand, for example, that a purchase order and a budget obligation describe the same dollar at two different stages of one transaction.

- **Deliver.** The finished graph answers questions through the interfaces a program already uses: dashboards, chat, purchase order review queues, and audit-ready reports. Every answer names the system it came from, so an answer is never a black box the program has to take on faith.

The roadmap: request to disposition

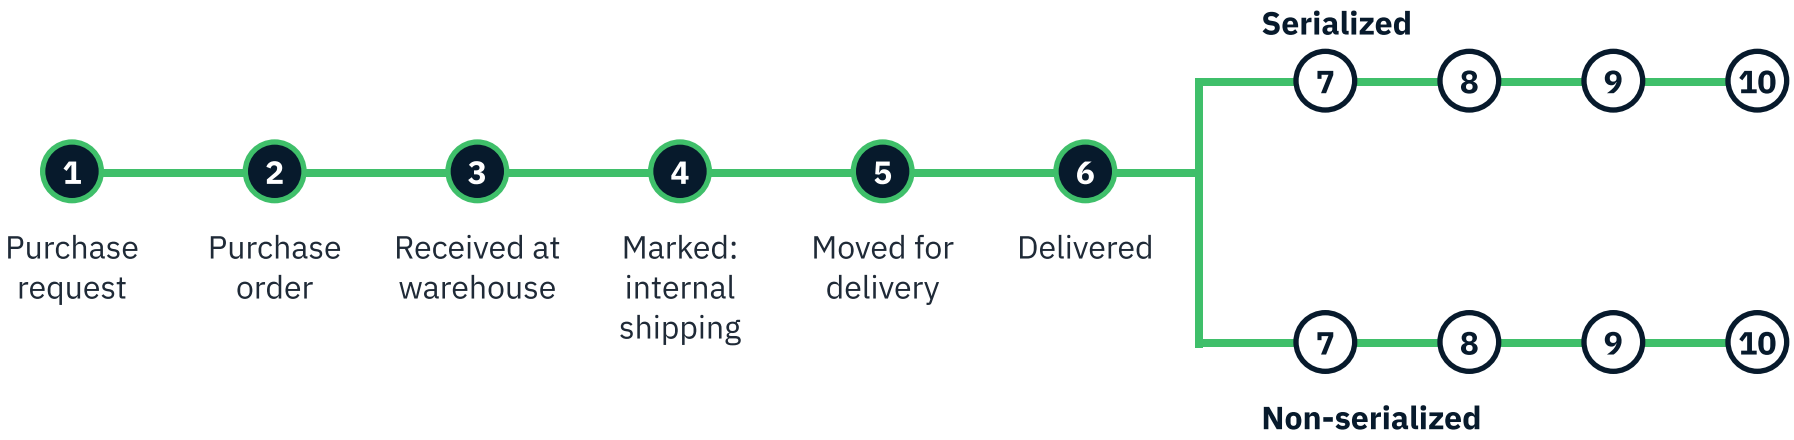

Cerebro follows a purchase through as many as ten stages, from the first request to final disposition. The first six apply to everything moving through a supply chain; the last four apply when they are relevant to the specific item: a serialized asset typically runs through all ten, while a non-serialized commodity may stop earlier. Cerebro does not replace any of the systems below; it connects them, so every stage that applies is always the same traceable record.

1. **Purchase request.** Generated from an email request or an internal system input.
2. **Purchase order.** Priced, approved, and sourced, typically on a GSA Schedule or BPA.
3. **Received at warehouse.** Physically received and inspected against the purchase order.
4. **Marked: internal shipping.** Logged and marked by the internal shipping department; a serialized item receives its IUID/UII tag under DFARS 252.211-7003.
5. **Moved for delivery.** Assigned an internal system number or a commercial shipping number, depending on destination.
6. **Delivered.** Recognized in the internal shipping system; final destination and custody are confirmed.
7. **System of record.*** A serialized item's UII and custody record enter the property management system; a non-serialized item is logged by quantity and location.
8. **Accounted and audited.*** Verified annually, or per regulation, against the property record.
9. **Repaired or repurposed.*** No longer needed where it sits: returned for repair, or reissued elsewhere.
10. **Disposition.*** Removed from accountable inventory; the internal system keeps the record for life.

*Stages 7–10 apply when they are relevant to the item; see the two examples below.

Two examples, one graph

The same delivery lifecycle applies to a serialized asset and a non-serialized commodity, then runs as far as each one's accountability actually goes. In both cases, the internal system remains the life-of-record; Cerebro is the graph that keeps it traceable.

Read the table by row, then compare the two columns at each stage. Both items move through the same first six stages: request, order, receipt, marking, movement, and delivery, with identical or near-identical handling. Stage 7 is where the two columns part ways. A serialized item receives an individual custody record tied to its own unique identifier and carries that record forward through annual verification, turn-in, and disposal. A non-serialized item is tracked only by quantity and location, and its accountability effectively ends once that quantity is logged and reconciled at the next scheduled inventory. Neither path is incomplete; the point is that Cerebro applies the correct depth of tracking to each item automatically, without forcing every purchase through the same heavyweight process or the same light one.

Stage	Serialized item: unclassified desktop computer	Non-serialized item: office furniture
1. Purchase request	IT helpdesk ticket auto-generates a request for a desktop computer replacement.	Internal system requisition for furniture for a new workstation.
2. Purchase order	Sourced on a GSA IT Schedule; priced and approved by the property custodian.	Sourced on a GSA Schedule/BPA; priced and approved by quantity, not unit.
3. Received at warehouse	Warehouse receives and inspects the unit against the PO.	Warehouse receives and counts the lot against the PO quantity.
4. Marked: internal shipping	IUID/UII tag applied (DFARS 252.211-7003); hand receipt drafted.	Logged by quantity and location, not a serial number.
5. Moved for delivery	Internal system # or commercial shipping #, by destination.	Internal system # or commercial shipping #, by destination.
6. Delivered	User signs the hand receipt; custody transfers of record.	Placed at the workstation; quantity reconciled to the internal record.
7. System of record	UII and custody record implemented in the property management system.	Tracked in the internal system by quantity and location, not by serial number.
8. Accounted and audited	Serial number physically re-verified each year against the property record.	Reconciled at scheduled physical inventory, by quantity and location.
9. Repaired or repurposed	No longer needed: item is shipped back to the warehouse (turn-in).	Reissued or consolidated at the next inventory cycle.
10. Disposition	Removed from accountable inventory; the internal system retains the full life-of-record.	Removed from the location record; no individual disposal action required.

Both examples run through the same six delivery stages and land in the internal system as the record of where the item is. From there they diverge: a serialized asset carries an individual UII and custody record forward into annual verification, then turn-in and disposal; a non-serialized commodity's record ends at its logged quantity and location, reconciled at scheduled inventory, not tracked and disposed item by item. Either way, the internal system is the record, and Cerebro is the graph that keeps it traceable end to end.

Impact by role

A single governed graph changes a different question for each role on a supply chain program, from the logistician tracking an item day to day to the auditor reconstructing a record after the fact, not a new report to read, but a faster, more confident answer to the question each of them already owns.

The value here is not limited to audit season. A logistician tracking a shipment, a warehouse technician confirming a receipt, or a program analyst chasing a stalled order runs into the same fragmented-systems problem every day, not once a year. Today, answering a simple question like where an order stands or whether an item ever arrived means checking two or three systems and often a phone call or an email to a colleague, and a different person asking the same question an hour later can get a different answer. With a connected graph, the same question returns the same answer every time, drawn from the same underlying records, in seconds instead of over the course of a day's back and forth. That is visibility into where things actually stand, speed in getting an answer instead of chasing one down, and accuracy because the answer traces to the systems of record, not to whoever happened to check last. The audit benefit follows from the same foundation, but the daily value shows up immediately, on the first order a logistician looks up.

For example:

Today

PROGRAM MANAGER

Status on a purchase order means opening the PO system, then the property system, then asking someone to check shipping, and hoping the three agree.

With Cerebro

PROGRAM MANAGER

One question, asked in plain language, answers where the order stands, what it is tied to, and what is still open: traced to each system it touched.

On the next page, the left column describes how each role works today, across systems that do not share a common record. The right column describes the same role once those systems are connected through the graph. No new system needs to be learned, only the existing systems answering together. Read each row as one role's single most common question, answered two different ways.

Role	Today	With Cerebro
Logistician / Supply Technician	Answering where an item is or whether a shipment arrived means checking the shipping log, then the property system, then calling the warehouse or a colleague, often several times in a single day for routine status questions.	One search by purchase order, item description, or shipment number shows exactly where the item is and what happened to it last, cutting a string of calls and system checks down to a single lookup.
Program Manager	Status on a purchase order means opening the PO system, then the property system, then asking someone to check shipping, and hoping the three agree.	One question, asked in plain language, answers where the order stands, what it is tied to, and what is still open: traced to each system it touched.
Accountable Property Officer	Reconciling custody records against the property system means matching spreadsheets by hand, and a mismatch is often found only at the annual inventory.	Mismatches surface as they occur, flagged for review instead of buried until the yearly count, and a serialized record is ready for verification before inventory starts.
Contracting / Budget Officer	Obligations, remaining budget, and Anti-Deficiency Act guardrails are checked against a financial system that may lag what has actually been ordered or received elsewhere.	Obligations and approvals stay visibly tied to the budget line behind them, with segregation of requesting, approving, and receiving enforced by the record, not only by process.
Auditor / Inspector General	Reconstructing a custody trail for a sampled item means requesting records from several offices and reconciling them manually. That is the work behind findings like GAO's 116-of-270 contractor property error rate.	Every request, purchase order, receipt, shipment, and hand receipt already traces to one record, searchable directly instead of requested and assembled after the fact.

None of this requires a new system of record. It requires the systems already in place to answer together instead of separately.

The same graph also gives contracting and program leadership a view they do not have today: which vendors are actually supplying which programs, how much has been bought from each one, and how much of what was bought has been received, issued, or used. Vendor spend and usage sit today in separate procurement and property records that are rarely compared side by side. Connected, they let a Contracting Officer or Program Manager see vendor concentration, duplicate purchases across offices, and underused stock before it becomes an audit finding, rather than after.

Integration with existing systems

Programs have made real, well-founded investments in their PO systems, their property systems, and their finance systems. Cerebro is designed to add to that, not replace it.

Zero-copy collectors read data where it already lives, with no migration, no duplicate storage, and no second system for anyone to maintain. RDF and R2RML mapping federate metadata from each source into the graph, and the ontology layer gives that graph the meaning of a program's own processes. Every answer the graph produces resolves back to the system that actually holds the record, so the PO system stays the system of record for purchase orders, the property system stays the system of record for custody, and the finance system stays the system of record for obligations. Program staff keep their existing logins and existing workflows; what changes is that a question spanning all three no longer requires opening all three.

Deployment follows the same principle: Cerebro is deployable into secure, classified, and air-gapped environments, and plays well with the platforms programs already run, including Databricks, ServiceNow, and IBM. Human-in-the-loop approval gates apply to every workflow. The graph surfaces answers and exceptions, but people still make the decisions.

A phased path to adoption

Cyberhill engages use-case first, architecture led, and delivers with the program's own team, the same model it runs across its public sector work, on a firm-fixed-price, milestone basis.

Each row below describes one phase, in sequence. The middle column, "What happens," is the work Cyberhill and the program team do together during that phase. The right column, "Outcome," is the specific, reviewable deliverable the program has in hand at the end of it, not a promise, but something the program can evaluate before committing to the next phase.

Phase	What happens	Outcome
1. Discover	Use cases, the current data landscape, and success measures defined with program stakeholders, covering which systems hold what and where reconciliation actually breaks down.	A scoped, prioritized set of use cases and a fixed-price proposal, before a line of pipeline code is written.
2. Design	Ontology, architecture, and a roadmap agreed with the program office.	An agreed data model for procurement, property, and budget, and a named first pilot population.
3. Build	Cyberhill engineers the graph, the pipelines, and the semantic layer against the program's real data, in weeks, not quarters.	A fixed-scope pilot in production, answering real questions for a real population: a program office, a portfolio, a single property book.
4. Enable	Use cases in production, documentation in hand, the program's own team trained to run and extend it.	The program owns the ontology, the graph, and the intellectual property, with no lock-in and a clear basis to scale.

Each phase is fixed-scope and milestone-based, deployable into secure, classified, and air-gapped environments as required, and each one delivers value on its own; an agency does not have to commit to enterprise scale to see whether the approach works for one property book or one program office first.

A single connected system is often the natural place to start. Many agencies already run a document management system, or DMS, to hold contracts, receipts, correspondence, and related procurement and property records. Because those records already reference the purchase orders, shipments, and property items the rest of this paper describes, connecting the DMS first is a natural entry point: it gives a working proof point without an enterprise commitment, built from data the agency already has in hand. Cerebro's integration layer is abstracted into flat, uniform layers, not a series of bespoke point-to-point connections, so that first connection is not a dead end. The same layer extends to the next system, and the one after that, without redesigning the underlying model. What begins as a single-system pilot around the DMS becomes the entry point for the broader enterprise deployment described above.

Conclusion

The Department of War's eight consecutive failed audits are a story about systems, not about missing money. Those systems cannot answer, quickly and confidently, where a purchase, an asset, or a dollar stands right now. Every government supply chain program carries a version of that same risk, at its own scale, every time a purchase request, a shipping record, a property record, and a budget line live in systems that were never built to talk to each other.

The policy already requires the answer this paper describes: DFARS, IUID, PPBE, the Anti-Deficiency Act, and the GAO Green Book all assume traceability that fragmented systems make hard to deliver. Cerebro provides a practical, non-disruptive way to close that gap: a governed knowledge graph that connects the systems a program already runs, traces every transaction to its source, and gives a program manager, a property officer, a budget officer, and an auditor the same confident answer, without asking anyone to give up the system of record they already trust.

> The future of government supply chain accountability lies not only in what a program can show at the next audit, but in what it can already answer, on any ordinary day, about any single dollar or asset in its custody.

See your supply chain as a graph

Start with a 30-minute Cerebro session for your program team: from first briefing to a fixed-scope pilot in weeks, not quarters.

Book your 30-minute session →

Chris Walker
VP, Public Sector
chris@cyberhillpartners.com

Caitlyn McClellan
SVP, Public Sector Sales
caitlyn.mcclellan@cyberhillpartners.com

Eric Gibson
Senior Director, Public Sector Sales
eric.gibson@cyberhillpartners.com

cyberhill.ai — The logic your AI is missing

About Cyberhill

Cyberhill is the applied AI and data modernization partner for the public sector mission, grown out of the Department of War and the Intelligence Community. Cyberhill pairs a decade of mission delivery and a deep bench of cleared engineers, PhD ontologists, data scientists, and mission subject matter experts (SMEs) with the ontology and data engineering that make AI actually work: fixing the data foundation first, rather than layering AI on top of systems that were never built to talk to each other.

Cerebro, Cyberhill's ontology-first knowledge graph engine, and Wolverine, the digital twin of the enterprise stack, are built on that same governed foundation. Cyberhill's delivery spans the Intelligence Community, the Army, the Air Force, Special Operations, and the Department of Energy (DOE) national laboratory complex, with engagements deployable into secure, classified, and air-gapped environments.

Contracting. Cyberhill Intel LLC · CAGE (Commercial and Government Entity) code 5VTS6 · UEI (Unique Entity Identifier) CQPSEH72RDB7 · Small Business