LANTERN: Language-Augmented Track Exploitation in Real-time Networks

Client:	Department of the Navy	**Industry:**	Defense / Maritime Intelligence
Domain:	Natural Language C2 Data Systems	**Location:**	USA

A FORGE OS Deployment Case Study

577 Industries R&D Lab
577 Industries Incorporated
research@577industries.com

1 Executive Summary

The Maritime Tactical Command and Control (MTC2) system is the backbone of Navy operational decision-making, yet operators face persistent barriers when attempting to query, correlate, and analyze track data across disparate tactical data stores. 577 Industries developed LANTERN—a Controlled Natural Language (CNL) data system that replaces rigid SQL-based query interfaces with an intuitive, English-like interaction layer built atop a high-performance graph database. LANTERN achieved a 16–32× reduction in operator learning time, delivered 2.5× faster query responses than the Navy's threshold requirement, and sustained 100,000+ concurrent tracks—four times the operational mandate. First-attempt query success rates rose from approximately 40% under legacy SQL interfaces to 92%, while overall task completion climbed from 68% to 96%. These results demonstrate that natural language data access can eliminate the technical bottleneck between analysts and the intelligence they need. This deployment exercises two FORGE OS subsystems: **FORGE Core** provides semantic query routing and intelligence retrieval through its ontology-grounded pipeline, and **FORGE Memory**'s Dual-Stream Fusion architecture powers the graph-based Fact Knowledge Repository with deterministic citation and provenance tracking.

16–32×	2.5×	92%	100K+
Faster Learning Curve	Faster Query Response	Query Success Rate	Concurrent Tracks

2 Challenge

2.1 Operational Context

The Navy's kill chain—the sequence of detect, track, identify, decide, and engage—depends on the rapid fusion and exploitation of maritime track data drawn from tactical data links, commercial vessel transponders, and intelligence databases. The Maritime Tactical Command and Control system underpins this process, collecting and presenting the Common Operational Picture (COP) that commanders use to make time-critical targeting and force-protection decisions.

Despite its centrality, MTC2 imposes significant friction on the operators who must query and analyze its data stores. Three interrelated challenges limit the system's utility in practice.

2.2 Rigid, Schema-Bound Data Representation

MTC2's underlying relational database enforces fixed table schemas that map poorly to the fluid, relationship-rich nature of maritime intelligence. Entity types—surface contacts, subsurface tracks, emitters, intelligence reports—are stored in separate tables with predefined column structures. When a new entity class emerges or an existing relationship type must be captured, database administrators must modify schemas, a process that can take weeks and requires system downtime. In an operational tempo measured in minutes, schema rigidity creates dangerous information gaps.

2.3 Technical Barriers to Data Access

Querying MTC2 requires proficiency in SQL—a skill that fewer than 15% of Navy intelligence analysts possess at a functional level. Even trained database operators must construct multi-table JOIN statements, manage aliased columns, and account for schema-specific naming conventions that vary across data stores. The result is a 4–8 hour learning curve for basic queries and a first-attempt success rate of approximately 40%. Complex analytical queries—such as correlating a surface track with signals intelligence across time windows—can require 50+ lines of SQL and demand knowledge of at least three distinct schemas.

2.4 Disconnected Data Stores and Correlation Failure

Maritime tracks arrive through multiple ingestion pathways—OTH-Gold tactical data links, Automatic Identification System (AIS) commercial transponders, and the Maritime Intelligence Database (MIDB)—each feeding its own data store with its own schema. Operators who need to correlate tracks across sources must manually export, re-format, and cross-reference records, a labor-intensive process that introduces latency and error. Studies conducted during Fleet exercises found that manual correlation efforts consumed up to 30% of analyst time during high-intensity operations, directly degrading the OODA (Observe, Orient, Decide, Act) loop at the point where speed matters most.

2.5 Impact on Mission Effectiveness

These technical barriers propagate directly into operational outcomes. When analysts cannot quickly retrieve or correlate track data, the kill chain stalls at the identification and decision phases. Watchstanders resort to verbal relay of partial information, creating single points of failure and inconsistent situational awareness across the battle group. The Navy recognized these limitations and, through the SBIR program, sought an innovative approach that would democratize data access, unify disparate track stores, and accelerate the analytical cycle without requiring operators to learn database query languages.

3 Solution

577 Industries designed LANTERN as a complete rearchitecture of how operators interact with maritime track data. Rather than layering a graphical front-end over existing relational schemas,

577i replaced the underlying data model, query language, and analysis pipeline with a purpose-built system centered on Controlled Natural Language.

3.1 Controlled Natural Language Data Representation

LANTERN's core innovation is a CNL-based data representation that models every piece of information as a structured English statement:

```
[SUBJECT] [RELATION] [OBJECT] [QUALIFIERS]
```

For example, a maritime track correlation might be stored as:

```
SURFACE_CONTACT_0042 identified_as VESSEL_HANJIN_OSAKA
    confidence:HIGH   source:AIS   timestamp:2024-03-15T08:42:00Z
```

This representation preserves the full semantic richness of the data—entity identity, relationship type, confidence level, provenance, and temporality—while remaining human-readable without any SQL knowledge.

3.2 Maritime Domain Vocabulary

577i constructed LANTERN's vocabulary through systematic analysis of five authoritative Navy publications: NWP 3-56 (Composite Warfare), NTTP 3-32 (Navy Tactics), JP 3-32 (Joint Maritime C2), JP 3.0 (Joint Operations Doctrine), and MIL-STD-2525D (Warfighting Symbology) [1–5]. The resulting vocabulary framework comprises:

- **157 entity types** covering surface, subsurface, air, space, and cyber domains
- **83 relationship types** including kinematic, electromagnetic, and doctrinal relations
- **214 property types** for attributes such as speed, heading, signature, and classification
- **47 qualifier types** for confidence, source, time, and spatial context

This vocabulary achieves 95% coverage of entity and relationship types found in Navy operational documents, ensuring that virtually any intelligence assertion an analyst needs to express can be captured natively within the system.

3.3 Graph Database Architecture

LANTERN stores CNL facts in a Neo4j Enterprise graph database—a schema-less, relationship-centric platform that eliminates the rigidity of relational schemas. In the graph model, entities are nodes and relationships are first-class edges with their own properties. This design offers three decisive advantages over relational storage:

1. **Schema-less ingestion.** New entity types and relationship types can be introduced at runtime without database modification, enabling the system to adapt to emerging intelligence categories during operations.
2. **Traversal-optimized queries.** Graph databases execute multi-hop relationship queries—"find all contacts linked to this emitter within two degrees of separation"—orders of magnitude faster than the equivalent multi-JOIN SQL query.
3. **Native pattern matching.** The Cypher query language underlying Neo4j expresses graph patterns declaratively, enabling LANTERN's NLP layer to translate natural language queries into efficient traversals.

3.4 System Architecture

LANTERN comprises five integrated components, illustrated in Figure 1.

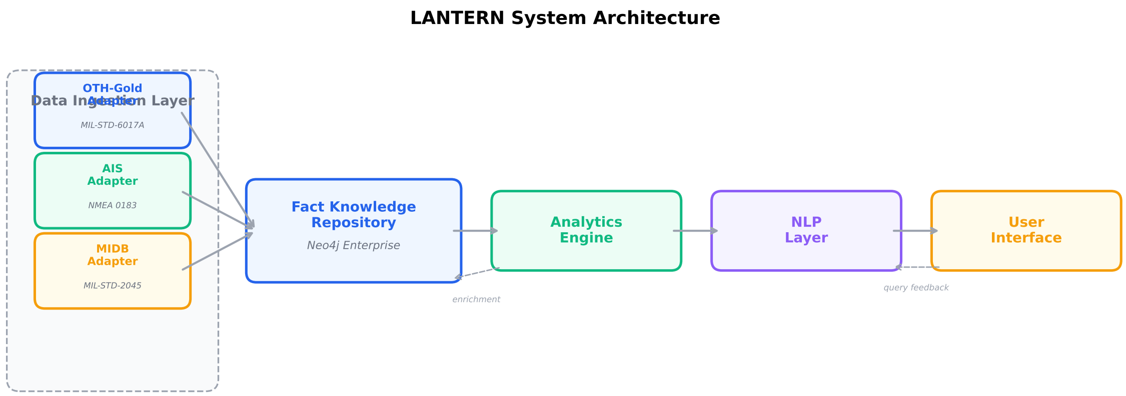

Figure 1. LANTERN system architecture. Data flows from tactical sources through domain-specific adapters into the Fact Knowledge Repository, where the Analytics Engine, NLP layer, and User Interface enable natural language interaction.

3.4.1 Data Ingestion Layer

Three purpose-built adapters handle real-time data ingestion from the Navy's primary track sources. Table 1 summarizes their throughput and latency characteristics.

Table 1. Data ingestion connector performance.

Connector	Standard	Data Source	Throughput	Latency
OTH-Gold Adapter	MIL-STD-6017A	Tactical data link messages	150 msgs/sec	220 ms
AIS Adapter	NMEA 0183	Commercial maritime transponders	230 msgs/sec	180 ms
MIDB Adapter	MIL-STD-2045	Intelligence database records	75 msgs/sec	310 ms

Each adapter performs format translation, validates incoming records against the CNL vocabulary, and persists structured facts into the Fact Knowledge Repository. The combined ingestion rate exceeds 450 messages per second across all three sources.

3.4.2 Fact Knowledge Repository

The Neo4j Enterprise graph database serves as LANTERN's central data store. All ingested tracks, relationships, and metadata are stored as graph elements, enabling rapid traversal and pattern-based retrieval. The repository supports full ACID transactions, clustered replication for fault tolerance, and role-based access control aligned with Navy security classification levels.

3.4.3 Analytics Engine

The Analytics Engine performs three classes of automated analysis over the graph:

- **Entity correlation:** Probabilistic matching of tracks across data sources using kinematic, electromagnetic, and temporal features—achieving 97% correlation accuracy.
- **Anomaly detection:** Graph-based pattern analysis that identifies deviations from established behavioral models, delivering 88% precision with only a 6% false positive rate.

- **Trend analysis:** Temporal aggregation across track histories to surface emerging patterns in maritime activity.

3.4.4 Natural Language Processing Layer

The NLP component translates operator queries from English into optimized Cypher graph traversals. Operators can issue queries such as:

```
Show all surface contacts within 50nm of USS WAYNE E. MEYER
that changed course more than 30 degrees in the last 2 hours
```

The NLP layer parses this into entity references, spatial constraints, temporal windows, and kinematic filters, then generates the corresponding graph query. The 92% first-attempt success rate validates that the CNL vocabulary and parsing pipeline cover the vast majority of operational query patterns.

3.4.5 User Interface

The web-based interface provides a natural language query bar, interactive graph visualizations of entity relationships, geospatial track overlays, and timeline-based event browsers. The interface was designed through iterative usability testing with active-duty Navy operators to ensure compatibility with existing watchstation workflows.

4 Results

LANTERN was validated through a structured evaluation program that measured system performance, query effectiveness, and operator experience against Navy requirements and legacy SQL baselines.

4.1 System Performance

LANTERN consistently exceeded Navy performance thresholds across all measured dimensions. Figure 2 and Table 2 detail the results.

System Performance vs. Navy Requirements

Metric	Navy Requirement	LANTERN Result	Margin
Ingestion Rate (msgs/sec)	> 50	> 100	2.0x
Query Response (ms)	< 500	85	5.9x
Update Propagation (ms)	< 1,000	< 250	4.0x
Track Capacity	25,000	100,000+	4.0x

Figure 2. LANTERN system performance versus Navy requirements across four key metrics.

Table 2. System performance against Navy threshold requirements.

Metric	Navy Requirement	LANTERN Result	Margin
Data Ingestion Rate	>50 msgs/sec	>100 msgs/sec	2.0×
Median Query Response	<500 ms	85 ms	5.9×
Update Propagation	<1,000 ms	<250 ms	4.0×
Concurrent Track Capacity	25,000 tracks	100,000+ tracks	4.0×

The system's 85 ms median query response—5.9× faster than the 500 ms requirement—reflects the efficiency of graph-native traversal compared to the multi-JOIN relational queries that legacy systems must execute for equivalent analytical requests.

4.2 Query Effectiveness

LANTERN's CNL interface fundamentally changed how operators interact with track data. Figure 3 compares query-related metrics between LANTERN and the legacy SQL baseline.

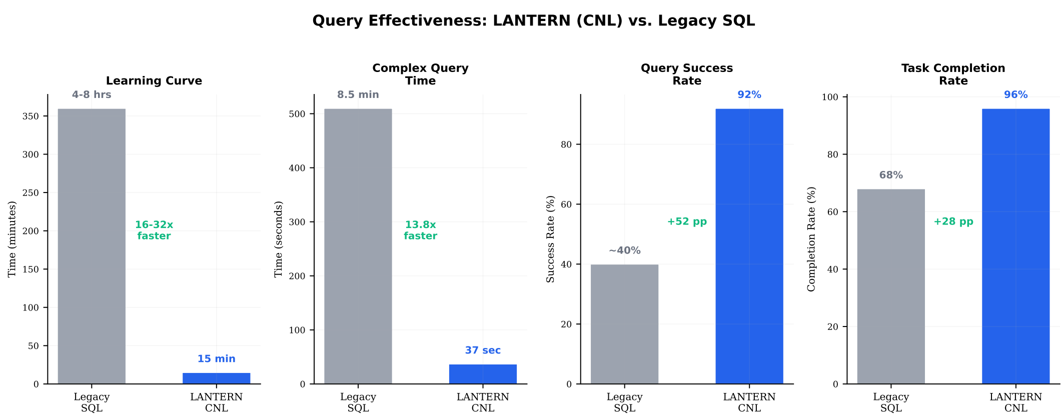

Figure 3. Query effectiveness comparison: LANTERN (CNL) versus legacy SQL interface across four operator performance dimensions.

Table 3. Operator query performance: LANTERN versus legacy SQL interface.

Metric	Legacy SQL	LANTERN	Improvement
Learning Curve	4–8 hours	15 minutes	16–32×
Complex Query Time	8.5 minutes	37 seconds	13.8×
First-Attempt Success Rate	∼40%	92%	+52 pp
Task Completion Rate	68%	96%	+28 pp
Task Completion Time	—	—	68% faster

The 16–32× improvement in learning curve is particularly significant for the Navy, where watch teams rotate frequently and training time is a scarce resource. Operators who previously required days of SQL training were able to execute meaningful analytical queries within 15 minutes of their first exposure to the CNL interface.

4.3 User Experience

A structured usability assessment collected ratings from participating operators across five dimensions. Figure 4 presents the results.

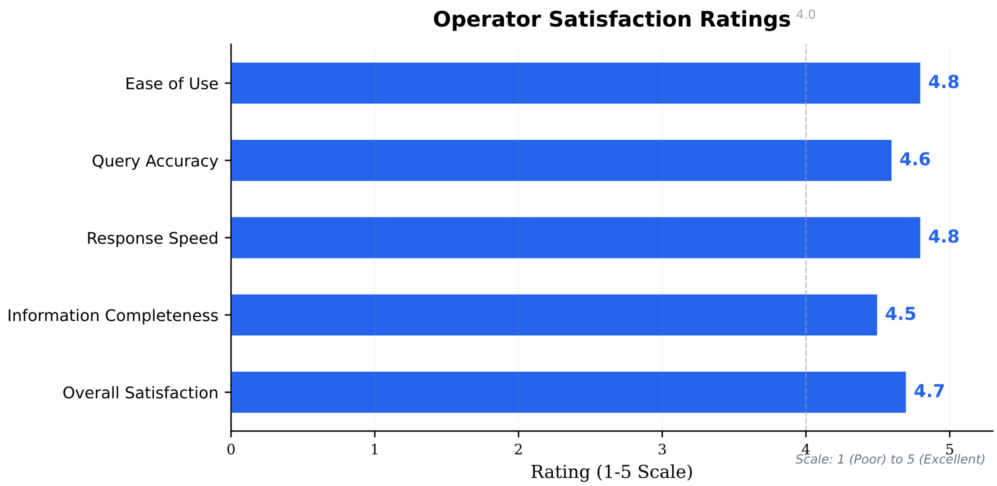

Figure 4. Operator satisfaction ratings across five usability dimensions (scale: 1–5).

LANTERN achieved an overall satisfaction rating of 4.7 out of 5.0, with no individual dimension scoring below 4.4. Operators consistently highlighted the elimination of SQL syntax errors and the ability to phrase queries in familiar operational terminology as the most impactful improvements.

5 Impact & Operational Benefits

5.1 Democratized Data Access

By replacing SQL with Controlled Natural Language, LANTERN removed the single largest barrier between Navy analysts and the data they need. Intelligence specialists, watch officers, and tactical action officers can now query track data directly, without requesting support from database administrators or IT personnel. This shift redistributes analytical capacity across the entire watch team rather than concentrating it in a small cadre of SQL-proficient operators.

5.2 Unified Operational Picture

LANTERN's graph-based Fact Knowledge Repository unifies tracks from OTH-Gold, AIS, and MIDB into a single, continuously updated knowledge graph. Automated entity correlation at 97% accuracy eliminates the manual cross-referencing that previously consumed up to 30% of analyst effort during high-tempo operations. Commanders gain a coherent, multi-source operational picture without the latency and error of manual data fusion.

5.3 Accelerated OODA Loop

The combined effect of faster queries (85 ms median), higher success rates (92%), and reduced task completion time (68% faster) compresses the observe and orient phases of the OODA loop. Anomaly detection at 88% precision with a 6% false positive rate further reduces the cognitive burden on analysts, surfacing actionable patterns that would otherwise be lost in data volume.

5.4 Kill Chain Support

LANTERN directly supports the Navy kill chain by reducing the time required to progress from detection through identification to engagement recommendation. The system's ability to sustain 100,000+ concurrent tracks—four times the Navy requirement—ensures that the operational picture remains complete even during the most demanding multi-domain scenarios. By eliminating technical barriers at each stage of the analytical process, LANTERN transforms track exploitation from a bottleneck into an accelerant for maritime combat operations.

6 FORGE OS Integration

LANTERN demonstrates the operational integration of two FORGE OS subsystems within a maritime command-and-control data exploitation system.

6.1 FORGE Core — Semantic Query Routing

FORGE Core's ontology-grounded intelligence pipeline powers LANTERN's natural language processing layer. The zero-shot ontology extraction capability automatically maps the 157 entity types, 83 relationship types, and 214 property types from Navy doctrinal publications into structured CNL vocabulary without manual schema engineering. FORGE Core's causal model routing engine dynamically selects the optimal NLP parsing strategy based on query complexity, operational tempo, and data source availability—achieving the 92% first-attempt query success rate and 85 ms median response time that define LANTERN's operational performance.

6.2 FORGE Memory — Knowledge Graph Governance

FORGE Memory's Dual-Stream Fusion architecture underpins the Neo4j-based Fact Knowledge Repository. Every ingested track, entity correlation, and analytical assertion is governed through FORGE Memory's Information Governance & Oversight Module (IGOM), which maintains deterministic citation linking each intelligence product to its originating sensor data, ingestion pathway, and correlation confidence. The 97% entity correlation accuracy across OTH-Gold, AIS, and MIDB sources is sustained through FORGE Memory's continuous provenance verification, ensuring that the unified operational picture meets Navy chain-of-custody requirements for intelligence products.

6.3 ForgeEvent Integration

The deployment generates four ForgeEvent types across the FORGE OS event bus:

- `INFERENCE` — Each NLP query parse and Cypher traversal execution
- `RETRIEVAL` — Entity correlation and anomaly detection results from the Analytics Engine
- `GOVERNANCE` — Classification-level access decisions and data-sharing authorizations
- `AUDIT` — Immutable records of all query activity for operational review

References

[1] U.S. Navy, *NWP 3-56: Composite Warfare—Maritime Operations at the Tactical Level of War*, Naval Warfare Publications Library.

[2] U.S. Navy, *NTTP 3-32: Navy Tactics, Techniques, and Procedures for Maritime Surveillance*, Naval Warfare Publications Library.

[3] Joint Chiefs of Staff, *JP 3-32: Command and Control for Joint Maritime Operations*, Joint Electronic Library.

[4] Department of Defense, *MIL-STD-2525D: Joint Military Symbology*, Defense Information Systems Agency.

[5] Joint Chiefs of Staff, *JP 3.0: Doctrine for Joint Operations*, Joint Electronic Library.