White Paper

April 2026

Semantic Knowledge Plane in Telecom

The Missing Layer for Agentic AI and Operational Intelligence

Authors: Patrick Kelly and Robert Curran

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: info@appledorerg.com• www.appledoreresearch.com

Publish date: 8-Apr-26

Cover image: Patrick Kelly

Contents

EXECUTIVE SUMMARY

Telecom operators are entering a new phase of AI adoption—moving from experimentation to operational deployment. The constraint is no longer access to data or tooling, but the ability to apply knowledge consistently across complex, multi-domain environments.

A **Semantic Knowledge Plane** (SKP) provides a practical path forward. An SKP acts as a unifying knowledge layer, complementing existing observability, inventory, and data platforms. It enables AI systems to reason across domains, explain decisions, and coordinate actions—transforming data into operational understanding.

Rather than requiring wholesale transformation, an SKP can be introduced quickly, starting with well-bounded operational use cases, and developed incrementally, expanding in scope over time. This approach allows operators to deliver immediate value, even in brownfield environments characterized by fragmented data, legacy systems, and organizational silos.

Appledore's view is that the Knowledge Plane is not a future-state architecture, but a deployable capability that can be applied today. This paper outlines how operators can begin, where value is realized first, and how vendors such as **Vitria** are operationalizing this approach in live environments.

THE MARKET CONTEXT: WHY A NEW LAYER IS EMERGING

Telecom operators are contending with an unprecedented confluence of technology change. 5G networks have expanded operational complexity across RAN, transport, and core domains. Cloud-native architectures have distributed workloads and created new abstraction layers. Data volumes are expanding rapidly, and organizational structures have become more siloed. Operations consume approximately 20% of revenue in large communications service providers, with growing pressure to automate and reduce costs.

The industry response has been predictable: massive investment in observability platforms, AIOps systems, and enterprise data fabrics. These initiatives have improved data collection, real-time visibility, and accessibility. They have not, however, addressed the underlying problem of interpretation. Operators can now detect what is happening in their networks with considerable precision. However, they still struggle, at scale, to determine why something happened or what action to take.

This gap between visibility and understanding is the primary constraint on scaling automation and realizing value from agentic AI. Data alone does not solve operational problems. Understanding does.

WHAT IS A SEMANTIC KNOWLEDGE PLANE?

A Semantic Knowledge Plane is a unified, semantically grounded, real-time model of network state, business intent, and operational policies—expressed through a formal ontology—that enables autonomous agents to observe, reason, and act with consistency across heterogeneous, multi-domain environments.

This definition requires unpacking. Unlike data planes (which transport packets) or control planes (which manage forwarding rules), a Knowledge Plane is designed explicitly for meaning. It models not just what is happening, but why, when, where, and how. It captures business context, causal relationships, dependencies, and domain semantics in machine-readable form.

An **ontology** is the Knowledge Plane's core component: a formal specification of shared domain concepts, their properties, and their relationships. In telecom, ontologies harmonize how a service, resource, fault, policy, or customer is defined across RAN, core, transport, OSS, and BSS domains. This semantic alignment is the glue that allows heterogeneous systems, tools, and agents to agree on meaning without constant, error-prone translation.

Knowledge Planes are typically implemented as **knowledge graphs** where nodes represent entities (network functions, services, customers, policies) and connections represent relationships and properties. Real-time telemetry, events, and logs are continuously ingested and mapped into the graph, keeping it synchronized with live network state. This enables agents to reason deterministically over explicit relationships. An agent can answer 'Why is this customer's throughput degrading?' or 'What is the blast radius of this failing device?' with confidence and explainability.

A useful framing: a data lake stores raw data, an OSS manages operational workflows, and a Knowledge Plane captures understanding. It unifies data across disparate sources into a coherent semantic model. It encodes dependencies between infrastructure, services, and customers. It enables programmatic reasoning rather than simple querying.

WHY THE SEMANTIC KNOWLEDGE PLANE IS THE FOUNDATION FOR AGENTIC AI

Agentic AI systems are designed to interpret context, reason across multiple variables, plan actions, and execute workflows in a coordinated, autonomous manner. In current deployments, however, most systems rely heavily on statistical inference from machine learning models and probabilistic reasoning from large language models, all operating on data that is frequently incomplete, siloed, or contextually ambiguous.

This creates three critical constraints. First, without structured representation of system relationships and dependencies, agents lack the contextual awareness needed to accurately determine service impact or identify root causes. They produce recommendations based on pattern matching, not causality. Second, agents are prone to hallucination and inconsistency because their outputs are not grounded in explicit, real system relationships. In a telecom environment, this directly translates to operational risk: incorrect actions cause outages and SLA violations. Third, the lack of explainability undermines operator trust and creates governance challenges. Regulators increasingly demand that autonomous systems can justify their actions in auditable, deterministic terms.

A Semantic Knowledge Plane addresses these constraints by introducing context-aware reasoning built on explicit, verified relationships and dependencies.

Consider the distinction: *probabilistic AI systems ask 'What pattern does this resemble?' A contextual AI system grounded in a Knowledge Plane asks 'What is actually happening, given the state of the network?'*

This shift has profound implications. It moves the operational paradigm from reactive statistical inference toward predictive and prescriptive models grounded in domain semantics. It enables multi-agent coordination across domains because all agents operate on the same canonical semantic model. It supports closed-loop automation—where observations immediately trigger reasoning, which drives actions, which generate outcomes that feed back into the Knowledge Plane, continuously refining understanding.

Consider a practical scenario: a network experiences a sudden surge in traffic to a streaming service. A detection agent observes the anomaly through telemetry and queries the Knowledge Plane to identify impacted services, affected customers, and relevant operational policies. A planning agent reasons over candidate remediation actions by modeling their impact within the Knowledge Plane. An orchestration agent executes the chosen action. All agents read from and write to the same Knowledge Plane, ensuring coordination and preventing conflicting actions. Without this foundation, each agent operates in isolation, with no shared context—risking duplication, conflicts, and suboptimal outcomes.

The Knowledge Plane enables deterministic context, semantic consistency across agents, causal reasoning for impact analysis, explainability grounded in domain logic rather than opaque neural networks, and real-time feedback loops that close the observe-reason-act-measure-learn cycle. These are not optional features for mature autonomous networks. They are mandatory.

WHERE THE SKP FITS: DISTINGUISHING THE KNOWLEDGE LAYER

The market already includes observability platforms, network inventory systems, and data fabrics. Understanding where an SKP differs from these is essential for correct investment decisions.

Observability vs. Knowledge

Observability platforms—such as time-series databases, distributed tracing systems, and log aggregators—excel at collecting and visualizing what is happening. They capture metrics, events, logs, and traces with high fidelity and in real time. They answer the question: 'What is occurring right now?' They do not, however, address causality or enable reasoning about relationships. An observability system detects that a service is slow; a Knowledge Plane identifies why—because an upstream resource has exhausted its memory buffer, which impacts three downstream services, which will breach SLAs if not remedied in the next five minutes.

Network Inventory vs. Knowledge

Network inventory and topology systems maintain structured representations of assets—what exists and how it is physically or logically connected. These systems are the authoritative source for network configuration and asset relationships. They are static or slowly evolving. A Knowledge Plane operates at a different level: it is dynamic, continuously updated with real-time state, and focused on semantic relationships and business context, not just physical topology. It answers questions about operational impact and business risk, not just network connectivity.

Architectural Integration

In practice, these systems are complementary, not competing. Observability platforms feed real-time telemetry into the Knowledge Plane. Inventory systems provide the baseline for network topology and asset relationships. The Knowledge Plane sits above both, providing the semantic glue and causal reasoning layer that makes their outputs actionable. A well-architected Knowledge Plane ingests data from observability and inventory systems, maps it into a unified ontology, applies inference and reasoning engines to derive insights, and exposes those insights via APIs to OSS, AIOps, and AI systems.

This architecture typically consists of five key components. First, a graph database that stores entities (network functions, services, customers, faults) and their relationships using either RDF or property graph models. Second, an ontology layer that defines shared semantics—typically aligned with TM Forum SID standards or telecom-specific extensions. Third, data ingestion and mapping pipelines that continuously integrate telemetry, events, logs, inventory data, and ticket information. Fourth, a reasoning engine that applies rules and inference to derive insights about impact, causality, and recommended actions. Finally, an API and query layer that exposes the Knowledge Plane to downstream systems. The result is not a static snapshot but a continuously evolving, real-time representation of the network and its behavior.

WHERE OPERATORS CAN START: PRACTICAL REALITIES OF AI DEPLOYMENT

In the last two years, operators have become convinced that AI, in various forms, is now essential to their efforts to streamline, modernize, automate and grow. The ready accessibility of AI tooling and platforms has made it possible to experiment fast – revealing the challenges that telcos face in turning AI explorations into at-scale deployments.

But the realities that operators face are not barriers to entry. They define where and how to begin. What seems initially like an insurmountable set of problems turns out to have a strong common thread: the encapsulation, application and lifecycle management of knowledge.

SKP Delivers Value First Where Data Is Fragmented

Data fragmentation is a defining characteristic of telecom environments—and a primary driver for introducing a knowledge layer. Rather than requiring full data unification upfront, a Knowledge Plane can be applied selectively, mapping relationships across a subset of systems to deliver immediate operational insight.

RAN data is often locked in radio equipment EMS/NMS silos with proprietary data models and limited export capabilities. Core network data is distributed across EPC, 5GC, and cloud-native platforms. OSS and BSS systems remain disconnected from network data. Customer data, billing, and finance live on separate islands. The result: dozens of heterogeneous data sources with incompatible schemas, APIs, authentication mechanisms, and governance models. Many operators lack the data mesh infrastructure or governance maturity to unify these sources, let alone map them into a coherent semantic model.

Integration Complexity

Telecom environments are inherently heterogeneous, spanning legacy OSS/BSS platforms, cloud-native systems, vendor-specific network management tools, and real-time telemetry streams. Rather than attempting to fully integrate these environments upfront, leading operators are taking a more pragmatic approach—introducing a knowledge layer that can sit across a limited set of systems and establish relationships incrementally.

Initial deployments do not require full API modernization or real-time integration across all domains. Instead, operators are focusing on selective data ingestion and mapping, building enough context to support specific operational use cases. This approach reduces complexity, accelerates time to value, and allows integration to evolve alongside demonstrated outcomes

Organizational Silos and Change Management

Telecom operations are structured across multiple domains—RAN, core, OSS/BSS, and increasingly AI/ML teams—each with its own processes and priorities. Rather than attempting broad organizational realignment, early Knowledge Plane deployments are being introduced within existing team structures, focused on shared operational problems.

By targeting cross-domain use cases such as incident resolution or service impact analysis, operators can create natural points of collaboration without requiring immediate changes to reporting lines or governance models. Over time, as value is demonstrated, alignment around shared semantics and data models emerges organically, reducing the need for large-scale change management programs.

Cost and Resource Requirements

Traditional approaches to knowledge modeling have often implied large, multi-year investments in data infrastructure, ontology design, and specialized tooling. In practice, operators are increasingly adopting incremental deployment models, where investment is aligned to specific use cases and scaled over time.

Rather than building a comprehensive knowledge base upfront, initial implementations focus on bounded domains, leveraging existing data infrastructure and extending it with semantic modeling capabilities. This allows operators to demonstrate measurable improvements—such as reduced time-to-resolution or improved service visibility—before expanding investment.

Building Skills for AI-Driven Operations

While specialized skills in data engineering, AI/ML, and semantic modeling are valuable, initial Knowledge Plane deployments do not require a fully developed internal capability from the outset. Operators are combining existing domain expertise within network operations teams with targeted external support and tooling that abstracts much of the underlying complexity.

Importantly, much of the knowledge required already exists within the organization—embedded in workflows, rules, and operational practices. The focus shifts from acquiring entirely new skills to capturing and structuring existing knowledge, making it accessible for automation and AI-driven use cases.

Legacy Systems

Legacy systems remain a critical part of telecom operations, containing not only data but decades of accumulated operational knowledge. Rather than viewing these systems as barriers, operators are increasingly treating them as valuable sources of insight that can be incorporated into a Knowledge Plane.

A key advantage of this approach is that it does not require system replacement. The Knowledge Plane can be introduced as a complementary layer, extracting and contextualizing knowledge from existing systems while preserving their operational role. This enables modernization without disruption and allows operators to build forward while leveraging the past.

STARTING NOW: INCREMENTAL DEPLOYMENT IN BROWNFIELD ENVIRONMENTS

Operators do not need to wait for perfect data, unified architectures, or organizational alignment to begin deploying a Semantic Knowledge Plane. In practice, the most effective approach is to start within constrained operational domains where value can be demonstrated quickly.

Typical entry points include:

- Cross-domain incident correlation (RAN to core to transport)
- Service impact analysis for high-value enterprise services
- Root cause analysis in complex multi-vendor environments

In these scenarios, the Knowledge Plane is introduced as a complementary layer, ingesting data from existing systems and applying semantic modeling to establish relationships and context.

Initial deployments typically:

- Integrate a limited number of data sources
- Focus on a single operational workflow
- Deliver measurable improvements in time-to-resolution or operational efficiency

Once validated, the Knowledge Plane can be extended incrementally - expanding coverage, refining ontology, and enabling more advanced AI-driven workflows.

This model aligns with how operators are adopting AI today from our research.

BUSINESS IMPACT

The business case for Knowledge Planes rests on several quantifiable benefits, many of which are already visible in early deployments.

Reduced Mean Time to Recovery and Improved Service Availability

AI-augmented incident detection and response can reduce mean time to recovery (MTTR) by 30 to 50 percent compared to manual processes. When agents can correlate symptoms across domains, identify root causes deterministically, and execute healing actions automatically, the time between detection and resolution collapses. For operators managing thousands of network elements, even modest improvements in MTTR compound into substantial availability gains and reduced SLA violations.

Operational Efficiency and Cost Reduction

Mature Knowledge Plane implementations can deliver 10 to 20 percent reductions in network operations center (NOC) operational costs. Automation of routine incident response, autonomous optimization of traffic steering and load balancing, predictive maintenance, and equipment utilization optimization all reduce human intervention. The knowledge captured in the Plane can be leveraged across multiple use cases, amplifying the return on the initial investment.

Improved Service Quality and Customer Experience

Real-time, AI-driven visibility across network and business domains enables proactive remediation before customer impact. Networks become self-optimizing and self-healing. Premium services—such as guaranteed SLA network slicing and outcome-based performance guarantees—become viable because operators have the observability and reasoning capability to honor them.

Faster Service Innovation and Reduced Time-to-Market

Well-defined APIs and a shared ontology reduce the friction of deploying new services and operational features. Service designers work against stable, semantically coherent abstractions rather than hunting down undocumented data models and negotiating with multiple teams. The result is faster design cycles and faster time to revenue.

Revenue Assurance and Fraud Detection

A unified view of customer transactions, network usage, and billing data enables faster detection of revenue leakage and fraud. Anomalies in usage patterns, billing inconsistencies, and unauthorized access attempts become visible. The cost of a single major fraud incident typically exceeds the annual operating budget of a small data governance program.

Regulatory Compliance and Auditability

Multi-domain knowledge graphs with lineage and audit trails support regulatory requirements—GDPR, network resilience, and security audits. When regulators ask how a decision was made or why an outage occurred, auditable Knowledge Plane logs provide answers that satisfy compliance teams.

VITRIA: OPERATIONALIZING THE KNOWLEDGE PLANE

Vitria Technology has spent over three decades building deep expertise in telecom domain knowledge, network operations, and automation. The company's VIA AIOps platform is deployed in some of the world's largest global networks, delivering measurable reductions in incident volume and MTTR, improved service availability, and accelerated time-to-value.

Vitria is now evolving the VIA platform by embedding a *Self-Evolving Knowledge Plane* directly into operational workflows. This is a critical distinction from point solutions. Rather than treating the Knowledge Plane as a separate software component that other systems query, Vitria is integrating the knowledge layer into the core of incident detection, analysis, and resolution. This allows operators to introduce semantic reasoning without disrupting existing OSS, BSS, or network systems.

Continuous Knowledge Acquisition

The VIA Knowledge Plane continuously discovers and refines network topology and service dependencies by mining CMDBs, real-time telemetry, and operational logs. It captures diagnostic and procedural knowledge from historical incident data—learning which symptoms indicate which failures, and which remediation sequences are most effective. It integrates structured configuration data with unstructured operational narratives, extracting causal relationships that would be invisible in either source alone. Unlike static knowledge models that require manual curation and quickly become stale, Vitria's approach autonomously evolves the ontology as the network changes.

Standards-Based and Transparent

Vitria’s Knowledge Plane uses W3C RDF standards for semantic representation and is aligned with TM Forum standards where applicable. This standards-based approach protects customers against vendor lock-in and enables interoperability with other systems.

Operational Proof Points

Vitria's VIA platform is currently deployed in environments managing millions of subscribers and tens of millions of dollars in daily revenue. Real-world deployments demonstrate 95 percent or higher incident resolution rates, with agents detecting and resolving issues before they impact end-user experience. In complex multi-hop network failure scenarios—such as when storage array latency cascades to impact a virtualized user plane function—Vitria's Agentic AI Resolution Planner does not simply alert a human operator. It executes a multi-step healing workflow: migrating the impacted service to healthy infrastructure, re-verifying performance metrics to confirm restoration, and updating the Knowledge Plane with the outcome so future similar incidents can be handled faster.

Key Differentiators

Vitria's approach is distinguished by several critical capabilities:

1. **Operational Embedding:** knowledge is applied directly within incident detection, analysis, and resolution workflows, not exposed as a separate query interface.
2. **Continuous Knowledge Acquisition**: the Knowledge Plane is automatically mined and evolved from real operational data, not manually built by consultants and frozen in place.
3. **Explainable AI**: all agent decisions and recommendations are traceable to Knowledge Plane relationships and rules, not opaque neural network weights.
4. **Brownfield Integration**: the platform deploys alongside existing OSS, BSS, and network systems without requiring replacement.
5. **Closed-Loop Execution:** insights immediately drive action, actions are validated, and outcomes feed back into the Knowledge Plane, continuously refining understanding.

As telecoms transition to 5G-Advanced and Network-as-a-Service business models, traditional data-driven approaches are proving context-poor. Vitria's approach addresses this by encoding three decades of deep telecom expertise directly into AI workflows, creating a System-of-Understanding layer that anchors AI agents in real-world network, service, and business logic. This virtually eliminates the hallucinations and inconsistencies that have previously stalled autonomous initiatives.

A Real-World Example

Business Problem

End-user experience issues on radio access network (RAN) devices are often symptoms of problems originating elsewhere—such as resource constraints on virtualized aggregation routers in the data center. However, correlating issues across RAN, transport, and application layers remains inherently difficult due to multi-hop separation and fragmented operational domains. As a result, diagnosis requires coordination across RAN, transport, SRE, and NOC teams, leading to prolonged MTTR, operational inefficiencies, and recurring customer-impacting issues.

Approach

A Self-Evolving Knowledge Plane addresses this challenge by continuously capturing, updating, and refining relationships between network elements, services, and dependencies across all layers. It combines deterministic knowledge (e.g., CMDB, topology) with inferred relationships derived from telemetry and behavioral patterns, creating a unified, continuously learning system of understanding that accurately connects symptoms to root causes across domains.

Solution Components

- **Telemetry Integration:** Real-time ingestion of network and application data (e.g., RAN metrics, transport signals, Kubernetes performance).
- **Relationship Inference:** Mining telemetry and operational data to uncover implicit dependencies and cross-domain interactions.
- **Topology & Dependency Modeling:** Integration of CMDB-enriched topology with auto-discovered dependencies and interactions across RAN, transport, and cloud layers.
- **Knowledge Graph / Knowledge Plane:** A semantic layer modeling entities and their relationships and causal dependencies.
- **Knowledge-Augmented Analysis:** Combines structured knowledge with AI-driven reasoning to correlate symptoms, identify the root cause (i.e., container memory exhaustion), and recommend remediation actions (i.e., increasing container memory allocation).
- **Agentic Remediation:** A remediation agent verifies key symptoms (i.e., container memory exhaustion) and generates a plan to increase container memory allocation.

Outcomes

Real-world deployments of Vitria's Self-Evolving Knowledge Plane in brownfield environments have achieved autonomous incident resolution rates of up to 95%+ for this and similar network issues, while dramatically reducing MTTR.

RECOMMENDATIONS FOR COMMUNICATIONS SERVICE PROVIDERS

1. **Start with a bounded operational use case.** Operators should begin by applying a Knowledge Plane to a specific, high-impact operational problem—such as cross-domain incident resolution or service impact analysis. The objective is not architectural completeness, but measurable improvement in operational outcomes within a defined scope.
2. **Develop a Multi-Year Knowledge Plane Roadmap**. CSPs should treat Knowledge Plane implementation as a strategic multi-year initiative, not a single project. Appledore recommends a phased approach: Phase 1 (Months 1–12) focuses on data unification and governance for a single high-value domain—for instance, 5G core network operations. Success here builds momentum and demonstrates value. Phase 2 (Months 12–24) extends the Knowledge Plane to a second domain—perhaps RAN operations—and begins integrating the two domains' semantics. Phase 3 (Months 24+) brings together all domains and introduces agentic AI workflows for closed-loop automation. This phased approach reduces risk, allows course correction, and spreads investment across budget cycles.

3. **Invest in Ontology Design and Governance Early.** Poor ontology design cascades into broken semantics and expensive rework. CSPs should align their ontologies with TM Forum SID standards and establish clear governance from the start. This includes deciding who owns the ontology, how changes are approved, how versions are managed, and how conflicts between domain experts are resolved. Undisciplined ontology design has derailed more companywide initiatives than technical constraints. The investment in governance and semantic design discipline at the outset pays dividends throughout the lifetime of the platform.
4. **Build or Acquire Data Mesh Capabilities in parallel – not as a pre-requisite.** Knowledge Planes sit on top of data infrastructure. But it is not necessary (may not even be possible) to have a complete data fabric in place before starting to develop the knowledge plane. Instead, operators should look for ways to develop the Knowledge Plane and data fabric in parallel, gradually expanding scope as additional data (such as real-time telemetry) becomes more accessible. The end goal remains the same: establishing a modern data foundation with data mesh patterns (federated data ownership with self-service analytics), data governance frameworks, deploying real-time streaming capabilities, and maintaining a data catalog that documents available data sources.
5. **Demand proof points from vendors before large commitments.** Vendors are increasingly claiming Knowledge Plane capabilities. CSPs should demand concrete evidence before committing resources. This includes reference architectures with customer logos and willingness to speak with references, published performance benchmarks and comparison studies, detailed technical roadmaps with committed timelines and milestones, and transparent licensing models that allow CSPs to understand the total cost of ownership. Brownfield deployments with controlled scope and clear success metrics are preferable to generic enterprise-wide claims of success.

A key takeaway for practitioners is that the Semantic Knowledge Plane can be introduced incrementally. Rather than a fully formed, immediately deployable capability, the Semantic Knowledge Plane is better understood as a viable path forward to deploy early use cases that deliver value.

About the Authors

Patrick has more than 25 years of experience in product management, business development, and technology consulting. He has advised executives and developed actionable business plans to help hundreds of technology companies profit in high growth software segments of the market. He and has published research in the areas of cloud economics, virtualization of the network, AIOps, network orchestration, analytics, service assurance, and customer experience management.

Patrick founded Appledore Research Group in 2014 to focus on the business impact of cloud and virtualization in the telecommunication market. Prior to Appledore, he was Research Director at Analysys Mason and held many technical roles in the field supporting both enterprise and service provider customers.

Contact: patrick.kelly@appledorerg.com

Robert Curran has 25 years of experience in the telecom software industry. Prior to Appledore, he was CEO at Aria Networks, a pioneer in the use of AI to design and optimize telecom networks.

Previous roles include Strategic, Corporate and Product Marketing at startups and blue-chip companies such as Ericsson, Amdocs, Cramer and Orange. Robert started his career as software engineer.

Robert has a B.Eng degree in Electrical and Electronic Engineering from Queen’s University, Belfast.

Contact: robert@appledorerg.com

About Appledore Research

Dedicated to the telecom industry, Appledore Research helps its clients navigate rapid change in technology, service innovation and operational practices to positively transform and grow their businesses. We think differently, challenge the status quo and identify the best ways to move forward. We want you to succeed.

Through our expert team of experienced telecom industry analysts and business practitioners, Appledore Research provides you with a unique blend of sector-specific market research reports, strategic advice and marketing support services, either as a part of a subscription package or as individual purchases. We're flexible.

Reach out to the authors for an initial discussion or follow us on LinkedIn or check out further research at appledoreresearch.com

Insight and analysis for telecom transformation.

@AppledoreVision

Appledore Research

www.appledoreresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

Appledore RESEARCH