Solution Overview

UnitySuite™ Software Platform

UnitySuite Software Platform

Solution Overview

Table of Contents

1. *Executive Summary* 2
2. *Platform Overview* 3
3. *UnitySuite AMI Manager (HES)* 4
4. *UnitySuite Network Manager (NMS)* 7
5. *UnitySuite Intelligence Hub* 9
6. *Security Architecture* 11
7. *Deployment Flexibility and Technology Foundation* 12
8. *Standards and Interoperability* 13
9. *Enterprise Integration* 13
10. *Why the Trilliant UnitySuite Software Platform* 14

1. Executive Summary

Utilities today are being asked to do far more than read meters. They are expected to operate resilient communication networks across service territories, manage device fleets sourced from multiple vendors across electric, water, and gas commodities, and turn the resulting flood of operational data into decisions that improve reliability, reduce losses, and serve customers better. Most metering platforms were designed to solve only part of this challenge. The Trilliant UnitySuite Software Platform is designed to bring all three together in one unified operating environment.

The UnitySuite platform brings together three tightly integrated pillars. The **UnitySuite AMI Manager (HES)** provides multi-commodity, meter and vendor agnostic device management with full multi-tenancy, scaling from pilot deployments to territories with millions of endpoints. It pairs this operational core with user-configurable reports, dashboards, and a built-in AI assistant that helps answer operational questions posed in plain language. The **UnitySuite Network Manager (NMS)** gives operations teams visibility and control over the communication infrastructure that carries metering traffic, exposing throughput, capacity, latency, and health metrics across field network technologies. The **UnitySuite Intelligence Hub** transforms the data these systems produce into actionable insight, providing an advanced analytics environment and the ability to deploy, run, and manage AI applications and agents both centrally and at the grid - edge.

Because the three pillars share a common architecture, a common security model, and a common data foundation, utilities avoid the integration seams, duplicated infrastructure, and inconsistent operational views that come from stitching together point products. The platform is vendor agnostic, meter agnostic, and network agnostic by design, and it supports open standards at every layer, which preserves procurement leverage, technology choice, and future flexibility firmly with the utility. The platform deploys anywhere the utility requires, whether on premises, in a private or public cloud, or in hybrid configurations, and it does so on a fully containerized foundation that keeps operational overhead low.

This document describes the solution architecture, the capabilities of each pillar, the security and interoperability posture of the platform, and the deployment and scalability characteristics that allow UnitySuite to grow with the utility over the life of the AMI program. Together, these capabilities position UnitySuite as a future-ready foundation for modern AMI operations, grid-edge intelligence, and long-term utility digitalization.

2. Platform Overview

The Trilliant UnitySuite Software Platform is a unified, modular software platform that manages the complete AMI value chain: the field devices that measure consumption, the communication networks that carry their data, and the analytical intelligence that converts that data into operational and business value. The platform is built as a set of containerized services orchestrated on a lightweight Kubernetes distribution, which gives every deployment the same characteristics regardless of where it runs: horizontal scalability, self-healing, easy upgrades without service interruption, and consistent operational tooling.

The three pillars are not separate products bolted together. They share core services for security, event handling, data processing, and platform management, and they exchange information over a common high performance messaging fabric. A meter alarm raised in the UnitySuite AMI Manager, a connectivity degradation detected by the UnitySuite Network Manager, and an anomaly flagged by an analytics model in the UnitySuite Intelligence Hub module can all be correlated in a single operational picture, because they flow through the same platform rather than through three vendor integrations.

At the northbound edge, the platform exposes its capabilities through standards-based APIs, including IEC CIM, MultiSpeak, etc., so that billing, CIS, MDMS, OMS, SCADA, GIS, and DERMS systems integrate through well understood industry interfaces rather than proprietary connectors. At the southbound edge, a driver-based architecture supports DLMS and ANSI metering protocols alongside legacy device families, and network drivers cover RF mesh, cellular, and multiple field area network generations. This dual layer of abstraction, standards at the top and drivers at the bottom, is what makes the platform genuinely meter, vendor, and network agnostic.

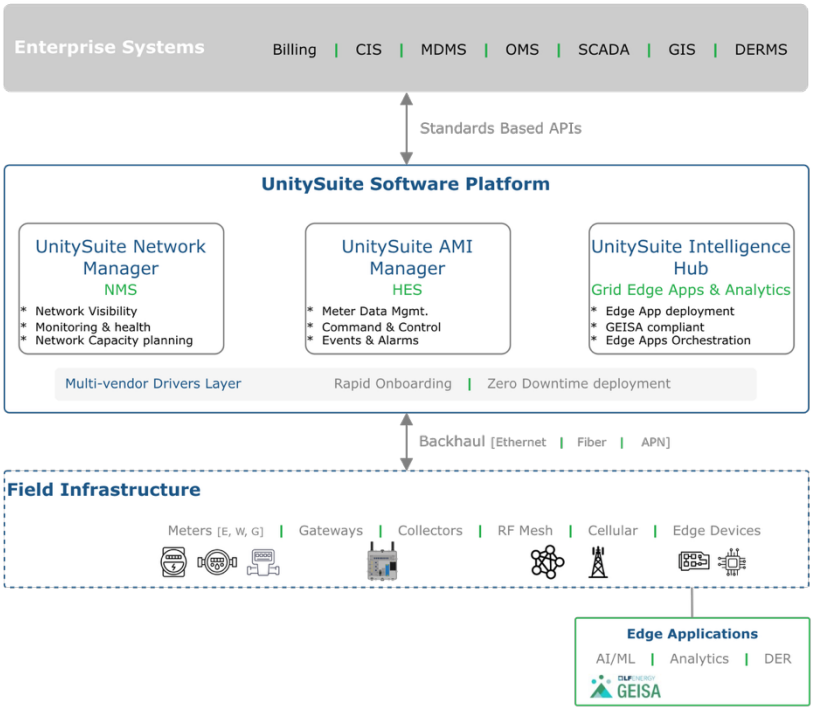

3. UnitySuite AMI Manager (HES)

The UnitySuite AMI Manager is the operational heart of the metering estate. It is responsible for acquiring interval and register data, executing on demand and scheduled reads, managing connect and disconnect operations, orchestrating firmware campaigns, processing events and alarms, and enforcing security across the device population. It performs these functions for electric, water, and gas meters concurrently, within a single deployment and a single operational model.

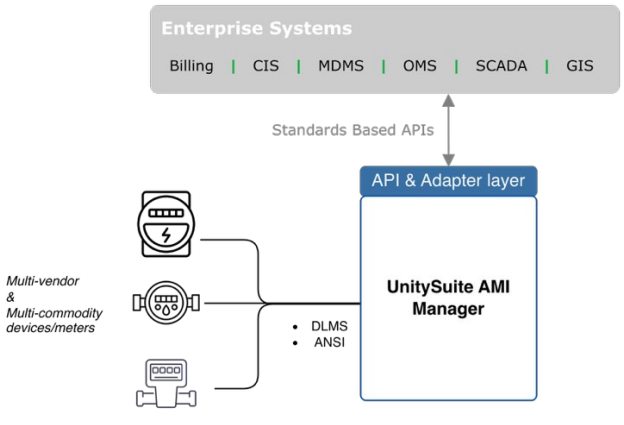

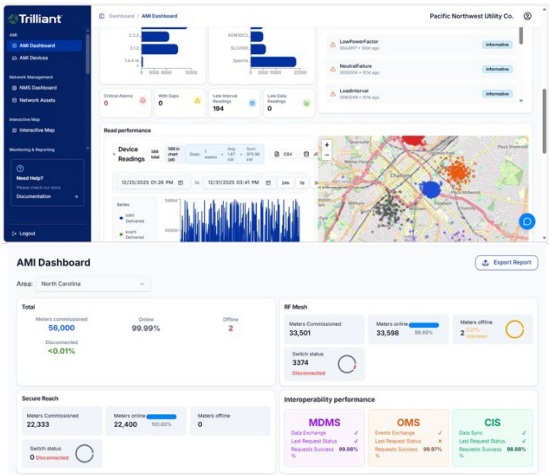

- ***Multi-Commodity and Vendor Agnostic Device Management***:
- Device connectivity is delivered through a layered architecture. Protocol drivers implement DLMS and ANSI standards as generic, configuration driven engines rather than as per meter implementations. In practice this means that support for a new standards compliant meter model follows a defined build, test, and integrate cycle and is deployed in weeks, with the work centered on configuration items such as object mappings, security suite selection, and communication profiles. This stands in contrast to the traditional model in which each new meter means a vendor product development cycle measured in quarters, and it removes the vendor lock that per meter driver development creates. Native meter data formats are preserved through the onboarding process, so nothing is lost in translation between the device and the enterprise.

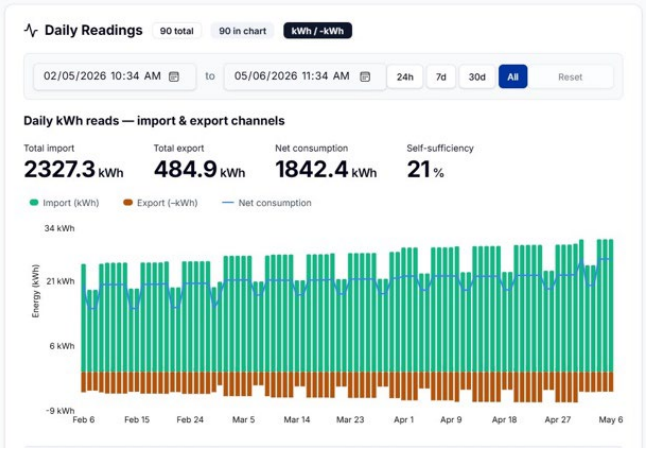

- ***Multi-Tenancy***: The HES is multi-tenant by design. A single platform instance can serve multiple utilities, operating companies, service territories, or organizational units, each with strict isolation of device populations, data, users, and configuration. Tenancy is enforced at the authentication and authorization layer and carried through data storage and messaging, so shared infrastructure never means shared visibility. This capability serves shared service organizations and metering as a service operators directly, and it gives single utilities a clean way to separate business units, pilots, or regulatory jurisdictions within one deployment.

- ***Core Services***: The HES core comprises a set of cooperating services for large scale device communication scheduling, firmware over the air orchestration for controlled and resumable upgrade campaigns, a device manager holding the authoritative registry of the device estate, configurable business logic for utility specific workflows, authentication and authorization services for both users and devices, event and alarm processing with configurable classification and routing, store and forward capabilities that preserve data integrity through communication interruptions, and a data processor that validates and normalizes incoming readings before they reach downstream consumers. A platform manager provides unified administration, health monitoring, and configuration across all services. These services communicate over an enterprise grade message broker supported by a high-performance in-memory cache, an architecture chosen specifically for sustained high throughput under the bursty traffic patterns that characterize AMI networks, such as morning read windows, storm driven event floods, and firmware campaign waves.

- ***User Customizable Reports and Dashboards***: Operational visibility in the HES is delivered through reports and dashboards that users shape themselves rather than through fixed vendor screens. Operations staff assemble dashboards from the measures that matter to their role, covering areas such as meters commissioned and online across each network type, read performance over selectable time windows, connect and disconnect status, event and alarm activity, and interoperability performance showing the health and success rates of data exchange with downstream systems such as MDMS, OMS, and CIS. Views can be filtered by service area, commodity, network, or device population, presented in tabular, chart, or map form. Reports built from the same foundation can be run on demand, scheduled for automatic distribution, or exported for regulatory and management purposes. The result is that each utility, and each team within it, works from an operational picture arranged around its own questions, without raising change requests against the vendor for every new view.

- ***AI Assistant for Operational Insight***: The HES incorporates an AI assistant that lets users ask operational questions in plain language and receive answers within minutes, collapsing a reporting cycle that traditionally involves raising a request, waiting for a report developer, and iterating over days or weeks. A field operations lead can ask for all meters that have not reported data in the last twenty-four hours grouped by neighborhood and immediately prioritize truck rolls against communication gaps. An operations manager can ask for a report on last week's outage suitable for sharing with leadership and receive a structured summary covering timeline, impact, and resolution. Analysts can probe for unusual consumption patterns in a region that might indicate transformer overloads, phase imbalance, or theft, before the first customer complaint arrives. The assistant operates strictly within the platform security model. It answers only from data the requesting user is authorized to see, respects tenant isolation in multi-tenant deployments and produces outputs that are traceable to the underlying operational data. It is a query and reporting capability for the metering estate, complementary to and distinct from the analytical model execution described in the Grid Intelligence and Analytics module.

- ***Scalability***: The HES scales horizontally by adding service instances and nodes rather than by replacing hardware. The platform has been engineered and validated against a tiered capacity model spanning deployments from one million to multi-million meters, with defined KPIs for read success rates, command latency, event processing throughput, and end to end data delivery at each tier. Performance validation is conducted against integrated, realistic traffic mixes sustained over continuous windows, not against isolated component peaks, so the published capacity figures reflect how the system behaves in production rather than in a laboratory corner case.

4. UnitySuite Network Manager (NMS)

AMI programs succeed or fail on the health of the communication network, yet in many vendor architectures the network is a black box visible only through failed reads. The Trilliant UnitySuite Network Manager treats the network as a first class managed asset. It provides operations teams with continuous visibility into throughput, capacity utilization, latency, packet delivery, signal quality, and topology across the entire communications estate, together with the tools to provision, configure, secure, and troubleshoot that estate.

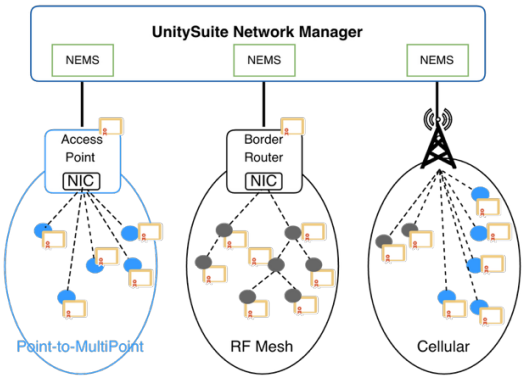

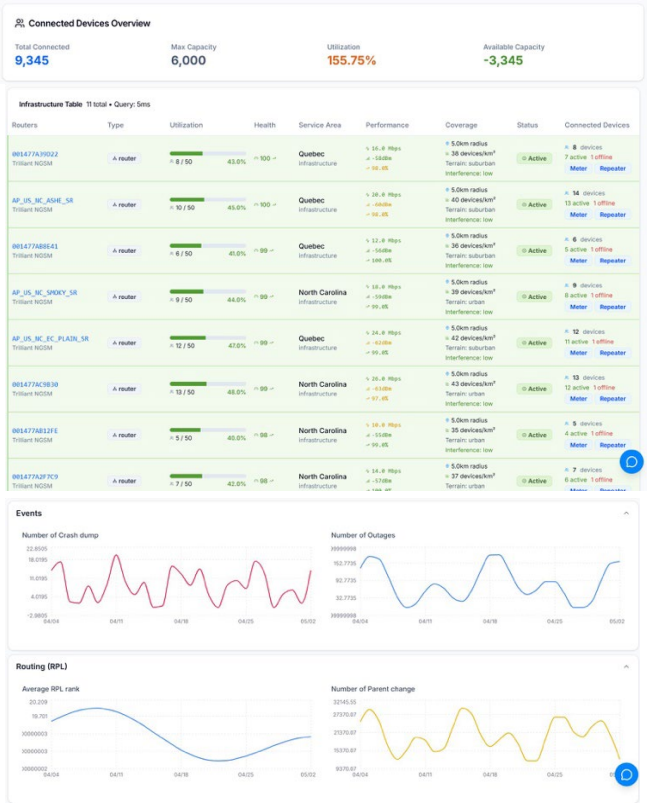

- ***Multi-Network Technology Support***: The NMS is network agnostic. It manages heterogeneous communication infrastructure with network element management provided per technology under a single management umbrella. Supported technologies span point to multipoint wireless networks, open standards based wireless mesh including Wi-SUN compliant mesh, and cellular connectivity, with the model designed to absorb future network types without architectural change. Gateways, access points, border routers, collectors, and edge devices are managed alongside the networks they participate in, and backhaul options including Ethernet, fiber, and cellular APN are supported. A utility can therefore operate mixed network technologies across its territory, chosen deployment by deployment on merit, and manage all of them through one system. Network agnosticism extends upward as well as downward. The NMS exposes its capabilities through an API layer that serves

not only the platform HES but also third-party head end systems, so a utility running another vendor's head end can still gain full visibility and control of its communication network. This decoupling of network management from meter management is a deliberate architectural position, and it protects the utility from having its network locked to any single metering vendor.

- ***Management Capabilities***: The management stack covers the full operational lifecycle. Network management provides real time and historical views of network performance and utilization. Provisioning management automates the secure enrollment of new devices and network elements. Firmware management delivers enhanced over the air update capability for network infrastructure, with campaigns that are targeted, paced against available network capacity, resumable after interruption, and tracked to completion, independently of meter firmware campaigns run from the HES. Security management governs credentials, certificates, and cryptographic material across the network layer. Topology management maintains an accurate, continuously updated map of how traffic flows through the mesh and backhaul. Reporting and SLA management gives both the utility and its service providers an objective, metrics-based view of whether the network is meeting its commitments, and fault and diagnostics capabilities support root cause analysis when it is not.

- ***Network and Mesh Analytics***: Beyond device level management, the NMS provides enhanced analytics over network behavior as a whole. Infrastructure views present each router and access point with its utilization against capacity, health score, throughput, coverage characteristics, and connected device population, so saturated or degrading elements stand out before they affect service. Mesh routing analytics expose how the mesh is actually organizing itself, including routing rank distributions and the frequency of parent changes, which are early indicators of instability. Availability analytics track node counts and average availability over time, and delivery analytics track transmission success rates at both the link and physical layers. Event analytics correlate crash and outage patterns across the estate. Capacity views aggregate connected device totals against maximum capacity so that planners can see utilization headroom, or the lack of it, across every part of the territory and plan expansion with data rather than estimates.

- ***Operational Value***: Because the NMS and HES share a common platform, network intelligence directly improves metering outcomes. Read schedules can respect network capacity. Firmware campaigns can be paced against actual throughput headroom. A cluster of missed reads can be immediately distinguished as a meter problem, a network problem, or a backhaul problem, which shortens resolution times from hours to minutes and reduces unnecessary truck rolls by putting remote diagnostics ahead of field dispatch. This is visibility that a traditional head end system is structurally blind to, because a head end sees the network only through the reads that fail, whereas the NMS sees the network itself.

5. UnitySuite Intelligence Hub

The UnitySuite Intelligence Hub is where the platform converts operational data into insight and automated action. It serves three complementary roles: an advanced analytics platform for exploring and reporting on data at scale, an AI and machine learning environment for building and operating models, and a management system for distributed intelligence applications and agents that run both centrally and on edge devices in the field.

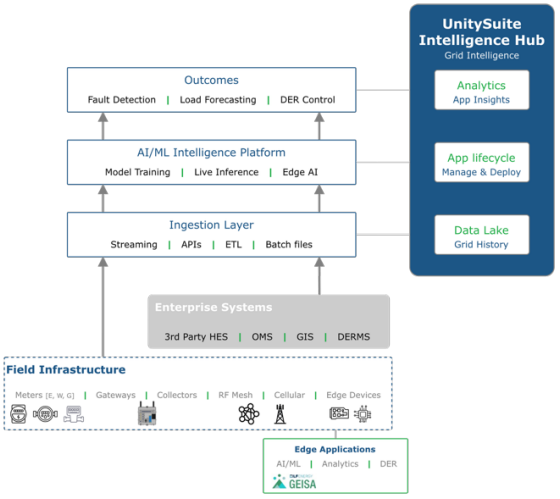

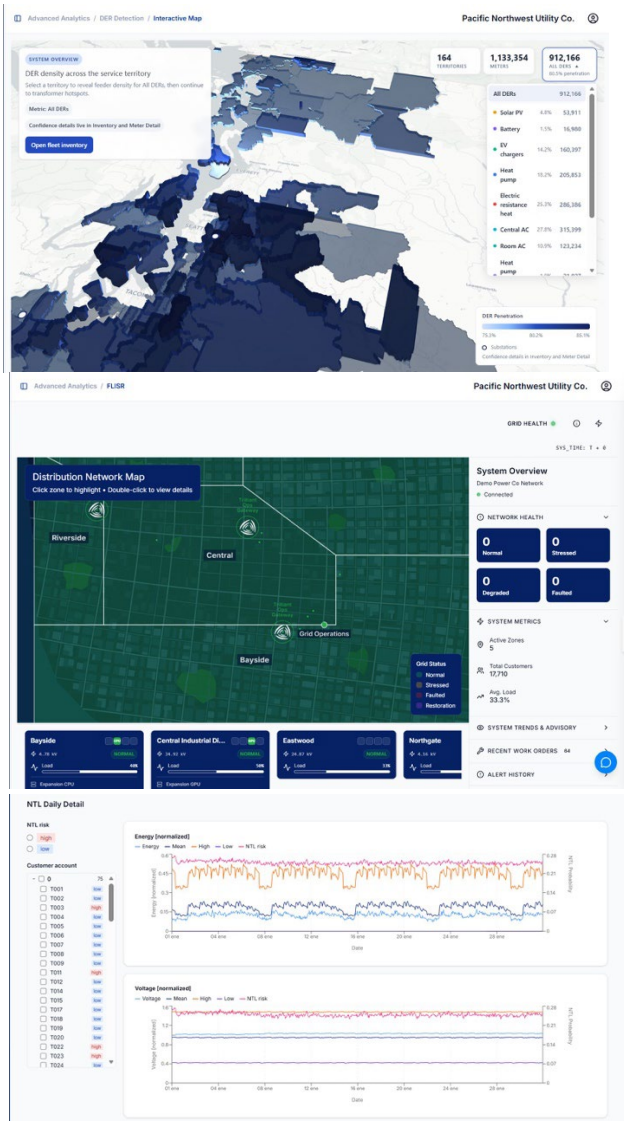

- ***Data Ingestion and the Analytics Foundation***: The module ingests data from across the platform and beyond it: interval and register reads, events and alarms, network performance metrics, device health telemetry, and external sources including third-party head end systems, OMS, GIS, and DERMS

platforms, as well as feeds such as weather data. The ingestion layer supports streaming for real time telemetry, API based exchange, ETL pipelines, and batch file loading, so data arrives through whichever channel each source naturally provides. This matters in practice because it means the module delivers value on the AMI data the utility already collects, including data originating from other vendors' systems, without requiring the metering or network estate to be replaced first. Ingested data lands in a governed analytical data foundation over scalable storage, organized in progressive refinement zones that take data from raw capture through validated and enriched form to business ready, curated datasets. This layered approach preserves full data lineage, supports reprocessing when business rules change, and keeps raw history available for future use cases that have not yet been imagined. The data lake is treated as a platform in its own right, a durable, open, query able grid history that outlives any individual application built on top of it. Analysts and business users interact with this foundation through an interactive analytics and visualization environment supporting dashboards, ad-hoc exploration, and scheduled reporting, with query access governed by the same role-based access model that protects the rest of the platform.

- ***AI and Machine Learning Platform***: The AI and ML capability operates behind the scenes on the data the module ingests, supporting the full model lifecycle: data preparation, training, validation, deployment, live inference, monitoring, and retraining. Its results reach users as finished analytics, dashboards, and reports rather than as raw model output, so business and operations teams consume insight without needing to interact with the machinery that produces it. This is deliberately distinct from the conversational AI assistant in the HES, which answers operational questions on demand; here, models run continuously against the analytical data foundation and surface their findings through the reporting layer. Typical outcomes include fault detection ahead of customer calls, load forecasting at feeder and transformer granularity, non-technical loss and theft detection, transformer loading and asset health analytics, outage prediction and restoration verification, voltage and power quality analysis, water and gas leak detection, and DER visibility, including territory wide views of DER density and penetration by type, from solar and battery storage to EV charging and electric heating, that inform both grid planning and DER control programs. Models can be operated centrally against the analytical data foundation.

- ***Distributed Intelligence and Edge Application Management***: The module includes a complete lifecycle management system for edge applications and AI agents. It handles application onboarding, signing and verification, targeted deployment to selected device populations, version management, health monitoring, messaging between edge applications and central systems, and controlled decommissioning. Edge execution follows the GEISA standard for grid-edge application environments, which keeps the utility free to source edge applications from multiple providers and to run them on compliant hardware from multiple vendors. This distributed intelligence capability

connects directly to grid programs. EV management, demand management, demand response programs, and fault location applications consume edge intelligence through the module, allowing the utility to act on conditions at the grid-edge in seconds rather than waiting for data to make a round trip to the head end.

6. Security Architecture

Security is implemented as a platform service rather than as a per module afterthought. All user and system access is governed by centralized identity and access management supporting single sign on (SSO), multi-factor authentication (MFA), fine grained role-based access control (RBAC), and full audit logging. Secrets, keys, and certificates are held in a dedicated secrets management service with controlled issuance and rotation. Transport security uses TLS 1.3 for platform communication, and device level security implements the cryptographic suites defined by the DLMS and ANSI standards, including high level security authentication mechanisms and secured firmware transfer.

The multi-tenant model extends the security architecture with tenant scoped roles, tenant isolated data domains, and tenant specific cryptographic material, so that shared infrastructure carries no shared risk. Database access follows a tiered model with per service accounts, strong password authentication, network scoped administrative access, and explicit default deny rules, with mutual TLS available where policy requires it. The platform architecture aligns with recognized utility security frameworks and supports the evidence gathering that regulatory and audit regimes demand.

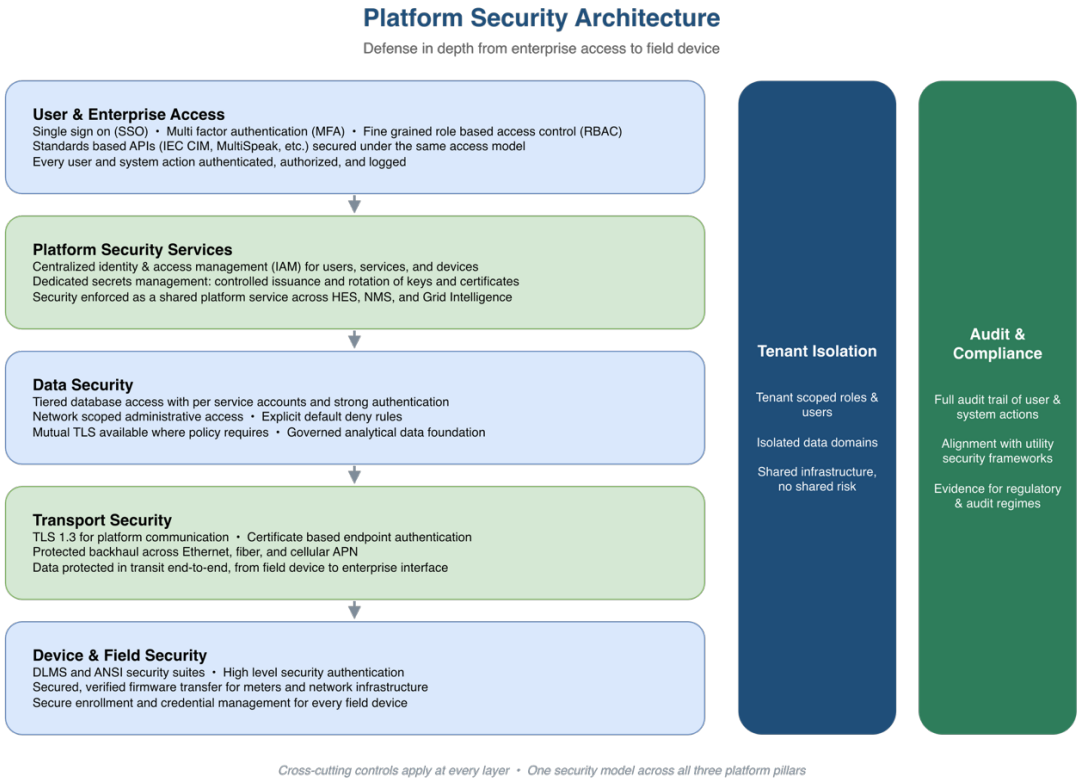

7. Deployment Flexibility and Technology Foundation

The platform is packaged entirely as containers and orchestrated on a lightweight, production grade Kubernetes distribution running on a hardened enterprise Linux operating system. This foundation is deliberately infrastructure neutral. The same software deploys on premises on utility owned hardware, in a private cloud, in any major public cloud, or in hybrid configurations where, for example, the AMI Manager and Network Manager run on premises for data sovereignty reasons while the Intelligence Hub environment runs in the cloud for elasticity. There is no functional difference between deployment models and no cloud provider lock in.

Data persistence uses a proven enterprise relational database as the primary store, with support for existing enterprise database investments where a utility requires it, including scenarios where a migration from an incumbent head end system. The analytical foundation adds scalable storage and open formats without introducing heavyweight new infrastructure, reusing the same orchestration, monitoring, and operational practices as the rest of the platform. Throughout the architecture, the guiding principle is operational simplicity: the minimum number of moving parts, reuse of common infrastructure, and a single consistent way to deploy, monitor, and upgrade everything.

The user experience layer is delivered as a modern, responsive web application requiring no client installation, with a consistent look and interaction model across the AMI Manager, Network Manager, and the Intelligence Hub modules.

8. Standards and Interoperability

The platform is built on the position that a utility should never have to choose between capability and openness. The table below summarizes the principal standards supported across the platform layers.

Layer	Standards	Purpose
Enterprise Integration	IEC CIM (IEC 61968/61970), MultiSpeak	Interfaces to MDMS, CIS, billing, OMS, GIS, SCADA, and other enterprise systems
Metering Protocols	DLMS, ANSI C12 family	Vendor agnostic communication with electric, water, and gas meters
Grid-edge applications	GEISA	Portable, multi-vendor edge application environment
Transport security	TLS 1.3, standards defined device security suites	Protection of data in transit, data at rest, and device level authentication

Standards based enterprise APIs mean that the platform participates in the utility ecosystem on the utility terms. Existing MDMS, billing, and operational systems integrate through interfaces their vendors already support, integration projects are estimated against known specifications rather than proprietary documentation, and the utility retains the freedom to change any surrounding system without disturbing the software platform.

9. Enterprise Integration

The platform integrates with the utility enterprise across the full application landscape. Billing and CIS systems receive validated consumption data through standards-based interfaces. MDMS platforms consume interval data, events, and device lifecycle information, and a dedicated integration adapter supports utilities migrating from incumbent head end systems while preserving existing MDMS investments and data flows. OMS platforms receive outage and restoration events with the timeliness that storm response requires. GIS integration keeps device and network locations synchronized. SCADA and DERMS integration connects AMI intelligence to grid operations, and demand side programs consume both meter data and edge application outputs. All integrations are delivered through the

standards-based API layer, monitored through the same operational tooling as the rest of the platform, and secured through the central identity and access framework

10. Why the Trilliant UnitySuite Software Platform

Four characteristics distinguish this platform in competitive evaluation. First, it is **unified by design**: metering, network management, and intelligence are pillars of one platform, with a common security model, shared operational experience, and consistent tooling across functions, rather than a collection of point products presented as a platform. Second, it is **agnostic and open by design**: vendor and meter agnostic at the device layer, network agnostic across mesh, cellular, and future technologies, standards based at the enterprise interface and grid-edge, and open even to service third-party head end systems through the NMS. Third, it provides **network visibility that traditional head-end systems cannot**: exposing network health, capacity, and failure trends directly rather than inferring them from failed reads, and shortening the path from operational question to actionable answer through its built-in AI assistant. Fourth, it **creates intelligence from the infrastructure the utility already owns**: applying analytics and AI to AMI data already being collected, including data from third-party systems, without requiring new field hardware and a forklift replacement to begin returning value.

Underpinning all of this is **proven scale**, with a published capacity model validated across tiers from one million to multi-million meters under realistic integrated load. The result is a platform that meets the requirements of the current tender while protecting the utility investment across the technology and regulatory changes of the coming decades. Trilliant welcomes the opportunity to demonstrate these capabilities in detail through workshops, reference discussions, and proof-of-concept engagements as part of the evaluation process.

DISCOVER THE *Power Of Choice*

The **Power of Choice** gives utilities the freedom to select the devices, communications technologies, applications, and deployment path that best fit their operational needs without being forced into a single technology approach.

Questions, comments: info@trilliant.com

©2026 Trilliant Holdings Inc., its subsidiaries, affiliates and/or licensors. All rights reserved. All trademarks are the property of their owners. This material is provided for informational purposes only; Trilliant Holdings Inc., its subsidiaries, affiliates and/or licensors assumes no liability related to its use and expressly disclaims any implied warranties or merchantability or fitness for a particular purpose. All specifications descriptions, and information contained herein are subject to change without prior notice.