Opti-Sphere: BubbleRAN Intelligent Optimization System

Version: 2026-02

Contents

1 General Description

Within the BubbleRAN Sphere product family, Opti-Sphere is an intelligent layered decision-support and optimization platform for 5G operations. It connects to existing network infrastructure (EMS/NMS/SMO/RIC/RAN/Core) environments through available APIs, transforms operational data into traceable insights, and produces governed recommendations and change plans. Execution is controlled by policy (recommend-only, human-approve, or optional closed-loop).

This datasheet describes the technical capabilities by layer, supported blueprints (e.g. agentic workflows), deployment options, interfaces, and indicative packaging/tiering.

Capability	Scope
Primary scope	5G performance, availability, configuration optimization, and operational assurance
Typical control loop	Sub-seconds to hours (operations and optimization loops)
User outputs	Dashboards, reports, APIs, recommendations, change plans, expected-impact estimates, audit trail
Integration	Adapters to APIs: telemetry/KPIs, alarms/events, configuration snapshots, topology/context, and (optional) change/execution interfaces
Automation levels	Recommend-only, Human-approve or Closed-loop (optional, bounded by policy and verification)

Opti-Sphere is organized as three optional capability layers (L1-L3) and an optional Digital Twin that spans all layers. Layers can be deployed incrementally and independently packaged.

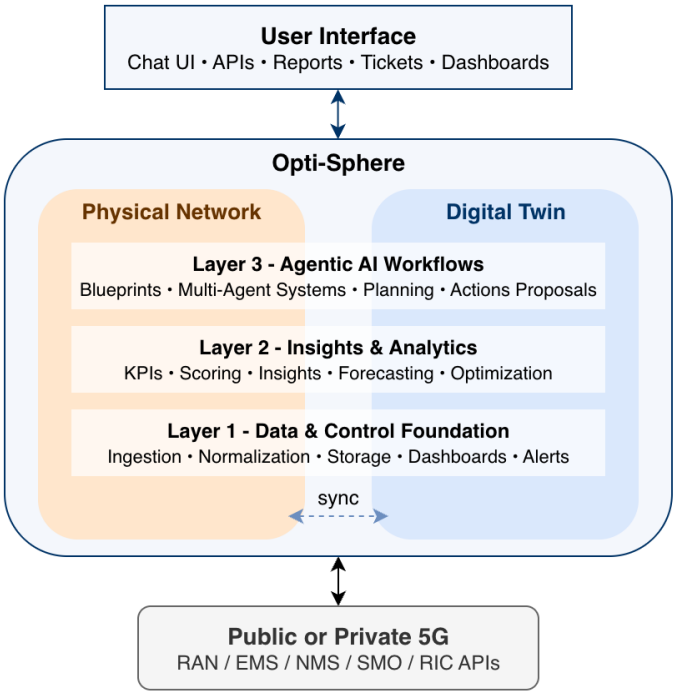

Figure 1: Opti-Sphere layered architecture: Intelligent decision-support and optimization platform for public and private 5G operations

2 Software Layers

2.1 Layer 1 (L1) - Data

The L1 telemetry stack is designed for high ingest throughput, low operational overhead, scalable storage/query, and cost-efficient retention for large-scale networks. It includes multi-source ingestion, normalization, time alignment, and queryable/auditable datasets.

Layer 1	Specification
Purpose	Make operational data consistent, queryable, and auditable across vendors and domains, including telemetry (metrics), logs, traces, events, configuration, and resource/energy signals.
Default capabilities	• **Multi-source ingestion** for metrics/stats, logs, traces, events, configuration snapshots, and infrastructure/resource telemetry (CPU/GPU, memory, storage, network). • **Normalization and schema management** (entity identifiers, units, time alignment, enrichment). • **Time-series optimized storage and querying** for high-cardinality operational telemetry, with retention policies and downsampling/rollups. • **Basic dashboards and plots** for KPIs/KQIs with threshold alerts. • **Dataset export APIs** (REST/gRPC where applicable) and role-based access controls. • **Data quality checks:** completeness, sampling drift indicators, and lineage/audit metadata. • **High-performance observability pipeline** designed for scalability, ease of operation, and cost efficiency (e.g., efficient storage, compression, and query performance). • **Energy & sustainability telemetry support**: power consumption and CO_2 footprint data ingestion, tagging, and reporting-ready exports.
Typical inputs	• **PM/KPIs and counters** (e.g., per-cell, per-slice/service where available). • **Alarms/events and logs** (service/system/application), including fault and audit logs. • **Distributed traces** (where available) and span metadata. • **Configuration management (CM)** snapshots and change history. • **Topology and context signals** (site, sector, time, route/environment tags). • **Infrastructure & sustainability telemetry**: compute/storage/network utilization; power, energy, emissions/CO_2 factors (site/device where available).
Typical outputs	• **Normalized datasets** (time-series & event/log indexes) accessible via APIs and/or exports. • **Dashboards and alerts** (threshold- and rule-based) integrated with ticketing/notification systems. • **Data feeds to Layer 2/3:** curated features/KPI sets for analytics, forecasting, optimization, and agent workflows.

	• **Audit artifacts:** lineage metadata, access logs, dataset versions, and config snapshot history.

2.2 Layer 2 - Insights & Analytics

L2 transforms L1 data into explainable insights, forecasts, and optimization primitives for proactive operations and downstream automation. It includes formulas, baselines, scoring, correlations, forecasting signals, and optimization primitives.

Layer 2	Specification
Purpose	Produce explainable insights and decision-ready signals, using formulas first and models second
Default capabilities	• **Baselines and scoring** (per KPI/KQI, per context) with confidence indicators. • **Correlation and root-cause** candidate ranking (evidence-backed hypotheses). • **Trend analysis, seasonality detection, and anomaly indicators** (statistical & ML, where enabled). • **Forecasting signals** (traffic, risk, degradation likelihood) for proactive operations. • **Optimization primitives:** objective definitions, constraint checks, and what-if scoring hooks.
Optional / modular add-ons	• **Advanced model training** (customer-specific supervised/unsupervised learning and feature engineering). • **Multi-objective optimization** extensions (e.g., energy vs QoE vs cost trade-offs). • **Synthetic data generation** pipelines (often coupled with the Digital Twin). • **Custom KPI/KQI libraries** and domain models.
Typical inputs	• **Curated datasets from L1:** normalized KPIs/KQIs, events/log features, traces, topology/context tags, CM snapshots and change history. • **Policy/intent signals and constraints:** SLO/SLA targets, thresholds, guardrails, cost/energy budgets. • (Optional) **Digital Twin outputs:** simulated KPIs, what-if scenarios, synthetic datasets.
Typical outputs	• **Scored KPIs/KQIs and health indices** with confidence levels and context. • **Explainable insights:** correlations, ranked hypotheses, and evidence summaries. • **Predictive signals:** forecasts, risk scores, degradation likelihood, capacity warnings. • **Optimization artifacts:** recommended parameter ranges, constraint feasibility checks, and what-if scorecards for downstream decisioning (Layer 3 or operators).

2.3 Layer 3 - Agentic AI & Decision Support

The L3 agentic workflow layer is designed to operationalize intent and constraints into governed, repeatable decision workflows, producing auditable change plans and structured recommendations. It includes blueprint-driven workflows to diagnose, propose, validate, and rollout changes under governance.

Figure 2 illustrates a sample optimization blueprint supported by Opti-Sphere for multi-agent cell reconfiguration. Functionally, a Cell Reconfiguration Planner (Reconfig) interacts with the UI and coordinates two agents via agent-to-agent (A2A) exchanges: (i) a Digital Twin (DT) Validator that validates candidate configurations against the digital twin, and (ii) a Physical Network (PHY) Optimization agent that refines and enforces the configuration on the physical network. Separately from these roles, the figure also emphasizes flexible LLM deployment: the LLM backbone can be provided via a cloud API (Reconfig agent) or hosted locally at the edge infra (DT agent) or on an available AI factory (PHY agent), depending on latency, privacy, and compute constraints.

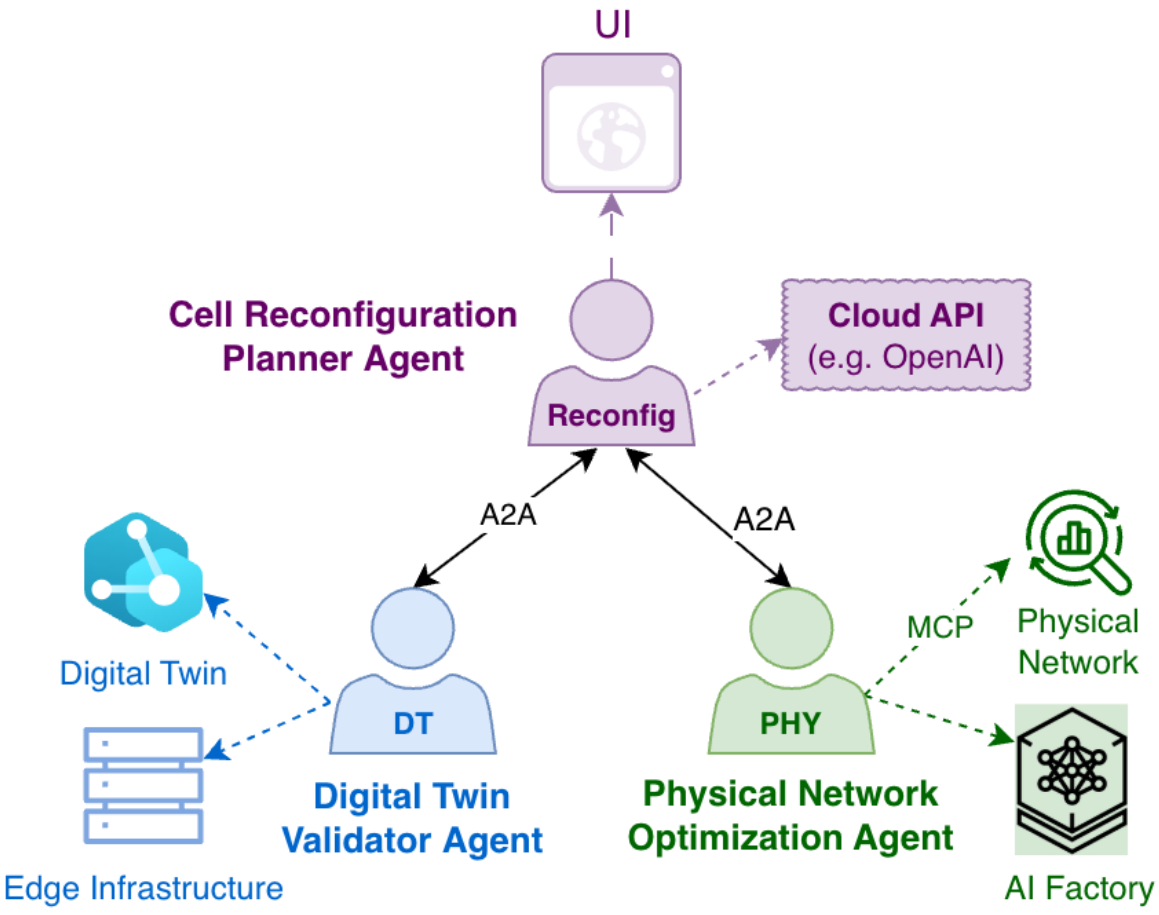

Figure 2: Sample Optimization Blueprint: multi-agent configuration management with cell reconfiguration planning, digital twin validation and physical network optimization with flexible cloud deployment.

Layer 3	Specification
Purpose	Execute repeatable decision workflows that turn intents/goals and constraints into governed change plans.
Default capabilities	• **Blueprint catalog** (configurable multi-agent workflows) with intent-to-task translation. • **Decision narratives** and structured outputs (recommended actions, expected impact, confidence, evidence links). • **Tooling integration:** calls to analytics/optimization engines, digital twin validation, and change-management systems. • **Governed execution modes:** recommend-only, human-approve, and (optional) closed-loop for bounded actions. • **Policy guardrails:** allowed action types, evidence thresholds, approval gates, and rollback hooks.
Typical inputs	• **Operator intent and goals** (natural language and/or structured policies), priorities, and constraints (SLA/SLE, cost, energy, risk, maintenance windows). • **L2 decision-ready signals:** scores, forecasts, anomaly indicators, ranked hypotheses, and optimization artifacts. • **L1 context as needed:** KPIs/KQIs, events/log summaries, CM snapshots and change history, topology and inventory metadata. • (Optional) **Digital Twin services:** what-if evaluation, safety checks, feasibility validation, and simulated impact estimates. • **Change governance inputs:** approval roles, action permissions, evidence thresholds, and rollback requirements.
Typical outputs	• **Recommended actions and bounded change plans:** configuration deltas, rollout steps, timing, and risk classification. • **Approval-ready artifacts:** justification narrative, evidence references, expected impact, and confidence. • **Execution payloads** (where enabled): tickets/changes, API calls, scripts/templates, and audit logs. • **Post-action tracking hooks:** monitoring objectives, success criteria, and rollback triggers.
Blueprint examples (included set varies by tier)	• **Config Planner:** propose and validate configuration deltas with approval-ready rollout plans. • **Observability Copilot:** guided exploration of KPIs, alarms, and context with traceable explanations. • **SLA/SLE Assurance:** detect degradations, quantify impact, and propose mitigations under constraints. • **Anomaly Detection & RCA:** early warning, evidence collection, RCA hypothesis generation, and next-step guidance.

2.4 Digital Twin - Validation & What-If

The Digital Twin capability is designed to reduce operational risk by providing a controllable, testable representation of the network for what-if analysis, validation, and rollout planning before production changes. It provides synced representation of network state and behavior for safe evaluation, data synthesization and trial planning.

Digital Twin	Specification
Purpose	Reduce risk by validating candidate actions and comparing scenarios before production changes.
Default capabilities	• **Digital representation of topology/config/state**, with synchronization hooks to the physical network where available. • **What-if evaluation of candidate changes** (scenario comparison, expected KPI deltas, risk scoring). • **Controlled trials and rollout planning** support (e.g., canary sites/cells, staged deployment). • **Synthetic scenarios** for rare faults or stress conditions (when enabled).
Typical inputs	• **Topology, inventory, and context metadata** (sites, sectors, cells, neighbors, service/slice context where available). • **Configuration baselines and change candidates** (deltas) from L3 or operators. • **Historical KPIs/KQIs and events** (for calibration and scenario baselining). • **Validation policies:** action bounds, safety constraints, success criteria, measurement windows.
Typical outputs	• **Scenario scorecards:** expected KPI/KQI deltas, risk classification, and constraint feasibility. • **Validation artifacts:** pass/fail results, evidence summaries, and recommended rollout strategy (canary/staged). • **Synthetic datasets** (when enabled) for analytics/model training and rare-event testing.
Integration requirements	• Access to relevant **configuration and topology data**. • **Defined action space** (which parameters/policies can be tested). • **Validation criteria** and measurement windows agreed with the customer.

3 Interfaces and Integration

Opti-Sphere is designed to integrate via northbound and southbound interfaces. Exact protocols depend on the connected 5G stack and customer environment.

Interface	Specification
Northbound (user outputs and control)	• **UI** for dashboards, recommendations, and workflow approvals. • **APIs** for queries, reports, and exporting structured recommendations/action plans (REST/gRPC; webhooks/SDK where applicable). • Integration to **ticketing/change management systems** for approvals and audit (optional).
Southbound (data in / actions out)	• **Data ingestion** via streaming or batch (e.g., Kafka/MQTT, Prometheus/OpenTelemetry, REST/gRPC, syslog/SFTP, NETCONF/YANG, O-RAN O1 when applicable). • **Action outputs** as configuration deltas, policy updates, workflow triggers, or change requests via EMS/NMS/SMO/RIC/vendor APIs (optional and governed).

3.1 Deployment

Opti-Sphere can be delivered as customer-hosted (on-prem / private cloud / edge) or as a managed service, depending on data sovereignty, latency, and integration constraints.

Deployment	Specification
Customer-hosted (on-prem / private cloud / edge)	• Runs in a **customer environment** (Kubernetes) with local data access. • Supports integration to customer observability stack and IAM. • Optional hardware bundle (GPU/CPU/storage) sized to data volume and analytics/AI requirements.
Managed service / API model	• Opti-Sphere **hosted by a provider** with API-based access for analytics and recommendations. • **Data ingestion** via secure connectors; output via APIs and webhooks. • The commercial model can combine base subscription & usage-based API units.

3.2 Operations

Operation	Specification
Recommend-only	Insights and ranked **recommendations**; no changes executed.
Human-approve	**Approval-ready change plans** with audit trail; changes executed only after explicit approval.
Closed-loop (optional)	**Autonomous execution** of bounded actions within policy and continuous monitoring/rollback rules.

3.3 Packaging

Packaging is organized into capability tiers aligned with the Opti-Sphere layers (L1–L3) and the optional Digital Twin. Exact commercial terms are customer-specific.

Package	Layers	Included (default)	Optional add-ons
Sense	L1	Ingestion, normalization, dashboards, threshold alerts, dataset export, data quality checks	Extra connectors, custom schemas, retention extensions
Infer	L1+L2	Baselines & scoring, trends, correlations, anomaly indicators, forecasting signals, optimization primitives	Advanced model training, KPI/KQI libraries, multi-objective optimization extensions
Automate	L1+L2+L3	Blueprint catalog, agentic workflows, governed execution modes, structured action plans, policy guardrails	Additional blueprints, custom tools, closed-loop enablement package
Twin (add-on)	Twin	What-if evaluation, risk scoring, rollout planning, optional sync hooks	Synthetic scenarios, higher-fidelity models, twin-to-EMS adapters

3.4 Deployment and Operations

Opti-Sphere can be delivered as customer-hosted (on-prem / private cloud / edge) or as a managed service, depending on data sovereignty, latency, and integration constraints.

Deployment	Specification
Customer-hosted (on-prem / private cloud / edge)	• Runs in a **customer environment** (Kubernetes) with local data access. • Supports integration to customer observability stack and IAM. • Optional hardware bundle (GPU/CPU/storage) sized to data volume and analytics/AI requirements.
Managed service / API model	• Opti-Sphere **hosted by a provider** with API-based access for analytics and recommendations. • **Data ingestion** via secure connectors; output via APIs and webhooks. • The commercial model can combine base subscription & usage-based API units.

4 What is typically included vs. scoped separately

Item	Specification
Included	• **Core platform software** for the purchased tier (Sense/Infer/Automate; Twin if purchased). • **Standard UI** (dashboards, recommendations view, approvals where applicable). • **Default blueprint set** for the purchased tier (e.g., Config Planner / Observability Copilot — exact set per tier). • **Standard documentation** (installation guide, admin guide, user guide) and baseline configuration templates. • **Standard integrations (where available)**: basic exports (CSV/JSON), standard API endpoints for reading outputs, and basic alerting/webhook patterns.
Scoped separately	• **Custom southbound adapters/connectors** to specific OSS/EMS/NMS/SMO/RIC or vendor APIs (e.g., proprietary vendor endpoints, custom data formats, customer-specific authentication/SSO constraints). • **Customer-specific model training** and feature engineering beyond defaults (e.g., bespoke KPI packs, supervised models tuned per region, advanced forecasting calibration). • **Complex workflow customization** (e.g., multi-stage approval chains, custom guardrail logic, integration with ITSM/ticketing processes, custom evidence thresholds). • **Deployment services**: on-site delivery, environment hardening, security reviews, data residency requirements, HA/DR architecture, and performance benchmarking.
Closed-loop automation	• Closed-loop execution is typically a **separately gated activity** enabled only after agreed prerequisites: • **Defined action space** (which parameters/actions are allowed). • **Policy guardrails** (limits, evidence thresholds, approval gates). • **Verification criteria** (success metrics, measurement windows). • **Rollback plan** (automatic/manual triggers). • **Example bounded actions:** limited parameter adjustment within safe ranges; staged rollout to a canary set; automatic rollback if KPI degradation exceeds threshold.

Sizing depends on number of sites/cells, KPI sampling rate, retention window, and enabled analytics/AI features. Common requirements are summarized below.

- Container runtime (Kubernetes) and persistent storage for datasets and metadata.
- Secure access to required data sources (KPIs, alarms/events, CM snapshots, topology).
- Optional GPU acceleration for L3 agentic workloads and advanced L2 model training (deployment-dependent).
- Connectivity to change-management and execution interfaces if human-approved or closed-loop modes are enabled

5 Frequently Asked Questions

Q1: Can Opti-Sphere run fully on-prem / at the edge?

Yes. Opti-Sphere supports customer-hosted deployments on Kubernetes (on-prem, private cloud, or edge). A managed service / API model is also available when data sovereignty and integration constraints allow.

Q2. Do you integrate with our EMS/NMS/SMO/RIC and multi-vendor stacks?

Opti-Sphere connects through available interfaces (REST/gRPC, Prometheus/OpenTelemetry, Kafka/MQTT, syslog/SFTP, NETCONF/YANG, and O-RAN O1 where applicable). Vendor-specific connectors may be included when available or scoped separately.

Q3. What do we need to get started?

A typical starting set includes KPI/PM counters, alarms/events (and relevant logs), configuration snapshots/change history (CM), and topology/context. Most deployments start with Sense (L1) to establish a consistent, auditable dataset.

Q4. Is the Digital Twin required?

No. The Digital Twin is an optional add-on. It is recommended for higher-risk changes where what-if validation, canary planning, and risk scoring materially reduce operational risk.

Q5. How do recommendations turn into execution?

Execution is policy-controlled: recommend-only (no execution), human-approve (execution after approval), or optional bounded closed-loop (autonomous execution within strict guardrails and rollback rules).

Q6. What makes closed-loop safe in production?

Closed-loop is bounded by an agreed action space, policy guardrails (limits and evidence thresholds), verification criteria (success metrics and measurement windows), staged rollout (canary), and rollback triggers. The Digital Twin can further validate candidates before production changes.

Q7. How is the expected impact estimated?

Opti-Sphere provides expected-impact estimates using baselines/scoring, what-if hooks, and (when enabled) Digital Twin scenario comparisons. Post-action tracking validates outcomes against success criteria and supports rollback when needed.

Q8. What hardware do we need?

Sizing depends on the number of sites/cells, KPI sampling rate, retention window, and enabled analytics/AI features. Sense/Infer can run CPU-only in many environments; optional GPUs help with advanced analytics/model training and agentic workflows in Automate.

Q9. How is data privacy handled?

Customer-hosted deployments keep data in your environment. Opti-Sphere supports RBAC, audit logging, encryption in transit (TLS) and at rest (environment-dependent), and configurable retention/export controls. Managed service deployments use secure connectors and customer-controlled policies.

Q10. How is Opti-Sphere packaged?

Opti-Sphere is offered as tiers aligned to the layers: Sense (L1), Infer (L1+L2), Automate (L1+L2+L3), plus Twin (Digital Twin add-on). Exact inclusions depend on your domain scope and integration requirements.

Ask your questions: Contact us with your target domains (RAN/Core/SMO/RIC), vendors, and governance needs for a tailored scope and rollout plan.

6 Change History

6.1 Denim 2026-02

- First Release of Opti-Sphere
- Released Layer 1 to Layer 3 software components
- Integrate Digital Twin layer.
- Released BubbleRAN Agentic Toolkits for Opti-Sphere.

7 License

This product includes software components from BubbleRAN.

- BubbleRAN software is © 2021–2026 BubbleRAN.

Redistribution of either software component, in whole or in part, without prior authorization from the respective copyright holders is prohibited.

The software is provided “as is”, without any express or implied warranties. In no event shall BubbleRAN be liable for any damages arising from the use of this software.

For detailed licensing terms and conditions, please refer to the respective license files provided with each software component.

8 Terminology

RAN	Radio Access Network
CN	Core Network
IMS	IP Multimedia Subsystem (3GPP Term)
BAT	BubbleRAN Agentic ToolKit
BRC	BubbleRAN Command Line Interface
DT	Digital Twin
CM	Configuration Management
EMS	Element Management System
KPI	Key Performance Indicator
KQI	Key Quality Indicator
NMS	Network Management System
OpenTelemetry	Standard for telemetry collection (metrics/logs/traces)
SLA	Service Level Agreement

BubbleRAN headquarters:
BubbleRAN
450 Route des Chappes
06410 Biot, France
contact@bubbleran.com
www.bubbleran.com

BREAKTHROUGH TECHNOLOGIES

France - US - Kenya