EMQX Neuron

PRODUCT DATASHEET

Industrial Connectivity Gateway for IIoT

EMQX Neuron is a high-performance industrial connectivity gateway designed for modern IIoT architectures. It bridges the factory floor to the cloud by collecting data from any industrial asset, processing it intelligently at the edge, and delivering unified, contextualized data to IT systems and AI platforms, enabling true OT/IT convergence.

With native support for 100+ industrial protocols, a built-in SQL stream engine, and Python/Go/JS-based AI inference, EMQX Neuron empowers industries to transform raw machine data into actionable intelligence at the edge in real-time.

100+	100+	≤100ms	<200MB
Industrial Protocols	Concurrent Drivers	Collection Frequency	Resource Footprint

Core Features & Technical Capabilities

Industrial Protocol Connectivity

- **100+ Native Drivers**: Full coverage — Modbus, OPC UA, OPC DA, Siemens S7, Mitsubishi, EtherNET/IP, Omron, BACnet, IEC 61850, SNMP, and more.
- **Parallel Collection**: Up to 100 simultaneous drivers, supporting 100,000 data tags with millisecond-level latency.
- **Bi-directional Control**: Device write-back via REST API and MQTT for closed-loop automation.
- **Extensible Architecture**: Modular driver design with SDK support for custom protocol extension.

Multi-Source Integration & Data Delivery

- **Enterprise System Integration**: Connects MES, WMS, ERP, ESB, databases, document systems, and video streams.
- **Flexible Northbound Output**: MQTT, SparkplugB, WebSocket, and HTTP to any IIoT platform or cloud.
- **Edge Data Storage**: MySQL, InfluxDB, Redis, Kafka, PostgreSQL, SQLite, SQL Server, and more.
- **OPC UA Server**: Expose collected data as OPC UA server for downstream SCADA and historian systems.

Edge Intelligence & Stream Processing

- **SQL-Based Rule Engine**: Built-in stream engine for real-time filtering, cleaning, aggregation, and multi-source fusion at the edge.
- **160+ Built-in Functions**: Event-driven rules for data fusion, anomaly detection, and automated alerting.
- **Natural Language to SQL**: LLM-integrated AI assistant converts plain language into stream processing rules.
- **Local Decision & Control**: Rule engine with reverse control for on-site alarms and device actuation, without cloud dependency.

AI/ML Algorithm Integration

- **Multi-Framework Support**: Integrate industrial mechanism models, machine learning, and deep learning for edge inference.
- **Low-Latency Inference**: Millisecond-level AI inference for predictive maintenance, anomaly detection, and quality inspection.
- **Flexible Plugin System**: Python, Go, and JavaScript plugins; external algorithm services callable via HTTP.
- **AI Operations Assistant**: Built-in knowledge base for natural language Q&A configuration and diagnostics.

Use Cases

Smart Manufacturing & Assembly

Collect data from mixed-brand PLCs, CNC machines, and robots across multi-line production environments. Apply real-time stream processing for production monitoring, quality traceability, and process parameter optimization, without ripping out existing equipment.

Energy & Power Management

Connect to energy storage systems, smart meters, power grids, and renewable energy assets using IEC 61850, DNP3, and other power-industry protocols. Aggregate data locally and trigger alerts or grid control actions at the edge with sub-second latency.

Process & Discrete Manufacturing

Integrate with DCS, SCADA, and historian systems in steel, pharmaceutical, and chemical plants. Enable predictive maintenance and defect detection through edge AI inference, reducing unplanned downtime and maintenance costs.

Multi-Site Enterprise Connectivity

Deploy standardized Neuron instances across hundreds of distributed sites using Docker/K8s for one-click replication. Combined with EMQX, build an enterprise-wide Unified Namespace (UNS) for consistent, real-time operational intelligence.

Why Choose EMQX Neuron?

Lightweight & Edge-First

<200MB memory footprint. Runs on x86, ARM, MIPS hardware — from small IoT gateways to industrial PCs — without heavy runtimes. Ideal for resource-constrained edge environments.

Cloud-Native Architecture

Native Docker and Kubernetes support. Seamless integration with Azure IoT Edge, AWS Greengrass, and any MQTT-compatible cloud platform. One-click multi-site replication via Rancher / K3s.

AI-Ready by Design

Built-in LLM-assisted rule generation, Python/Go/JS AI plugins, and millisecond-level edge inference. Transform raw OT data into actionable industrial intelligence, fully on-premise.

Technical Specifications

OS Support	Architectures	Memory Usage
Linux, Docker, K8s, K3s	x86_64, ARM 64-bit, MIPS	128MB – 4GB (by tag count)