Building Tomorrow's Healthcare Platform: Compliant and Built to Scale

ZEISS

Seeing beyond

A comprehensive guide for MedTech and Healthcare leaders based on state-of-the-art research insights

With funding from the:

CLUSTERS 4 FUTURE
Innovationsnetzwerke für unsere Zukunft

Federal Ministry of Research, Technology and Space

SEMECO.

zeiss.com/digital-innovation

Last updated: 2025-11

Executive Summary

Healthcare systems face sustained cost and staffing pressure. Digital platforms that connect devices, data, and clinical systems are now essential to cut errors, streamline work, and speed up product development. This report explains what such a platform requires, which design choices matter, and how these ideas are realized in the exemplary SEMECO project.

A modern edge-plus-cloud platform links edge and cloud to streamline clinical workflows, cut costs, reduce errors, and improve patient care. It also provides a governed foundation for research and faster development of innovative medical products and services. At the bedside, it supports remote monitoring, OR workflow orchestration, and alarm management, among other use cases; in the cloud it enables explainable decision support, population-scale analytics, and secure collaboration. Operationally, such a solution can power predictive maintenance, fleet and software lifecycle management, and command-center views that optimize throughput.

Composition of a medical platform

A modern medical platform can be separated into two parts. The edge runs inside the hospital network, close to devices, to handle low-latency tasks, local preprocessing, and retrofitting of legacy equipment. The cloud provides elastic storage and computing for cross-site coordination, analytics, and model training. Core building blocks across both parts include secure storage, compute and analytics, standardized data exchange (HL7/FHIR, SDC, DICOM), identity and access management, security monitoring, and automated compliance evidence.

Design considerations

Security and privacy obligations based on regulations such as MDR, FDA, HIPAA, GDPR, NIS2, and EU AI Act, need to be addressed from the start. Practical measures such as encryption, least-privilege access, software inventories, and audit logs, are built into everyday development and operations. Hosting balances capability with data sovereignty: A hybrid multi-cloud is typically best. Keep patient identifiers and keys in EU-resident clouds; use global clouds for de-identified analytics, while keeping workloads portable with containers, Kubernetes, and infrastructure-as-code (IaC). Investment follows a make-vs-buy rule: Buy mature, commodity services (compute, storage, observability, managed FHIR/IoT, consent) and build only what creates differentiation or does not address your organization's needs (e.g., interoperability adapters, privacy pipelines, and workflow components).

SEMECO as a concrete example

SEMECO is a German innovation cluster that brings together academia and industry to accelerate digitally enhanced medical devices. The SEMECO platform developed within the cluster applies the blueprint above: A microservices cloud, a modular edge, clear certification boundaries (only safety-critical services are treated as medical products), and security by design throughout. For SEMECO, we specifically focus on healthcare capabilities that create outsized value while reusing proven cloud and edge building blocks. Services, validated with partners, are grouped into three tiers:

- **Priority 1 – Core services:** Device interoperability based on Service-oriented Device Connectivity (SDC); anonymization for imaging and telemetry data (edge for residency, cloud for scale); workflow components that orchestrate clinical tasks and capture evidence.
- **Priority 2 – Supplementary services:** Consent management, device/fleet management, and edge retrofitting to bring legacy devices online.
- **Priority 3 – Additional services:** Audit logging across edge and cloud, scalable health-data processing, and pseudonymization where linkage is needed without full identification.

Two components already show our approach in practice: A **DICOM anonymization** module that reduces re-identification risk while preserving data quality for secondary use, and an **SDC-based interoperability** module that standardizes device communication and enables safer, more efficient workflows (e.g., alarm management, OR orchestration).

Deployment models and hosting choices

Selecting where to run the platform means trading off control, cost, and regulatory assurance. Most deployments pair an edge footprint with one or more cloud options:

- **Private cloud** maximizes customization and data sovereignty but requires high capital and specialist operations.
- **Colocation** offers professional facilities with lower capital cost, but less flexibility.
- **Public cloud** delivers elastic scale and rapid feature uptake, with careful design needed for EU-only residency.
- **Hybrid / multi-cloud** balances compliance in EU clouds with global scaling on hyperscalers, at the price of more complex governance.
- **Edge** provides low latency, on-prem processing, protocol bridging, and store-and-forward resilience, but adds site-level hardware and patching duties.

Provider comparison and legal context

Global hyperscalers (AWS, Azure, Google Cloud) offer unmatched scale and advanced services, but EU-only data residency requires careful design. As US providers, they may be subject to obligations under the CLOUD Act, which can compel disclosure of data under a valid US court order, even if the data is stored in the EU. European sovereign providers (e.g., StackIT, SysEleven, OVHcloud) simplify GDPR/NIS2 alignment and reduce extraterritorial risk, with smaller catalogs. A pragmatic path is **hybrid multi-cloud**: Place protected health information (PHI) and keys with EU providers; run de-identified analytics and AI on hyperscalers; keep everything portable and encrypted, and use anonymization/pseudonymization where appropriate.

Why this matters

This blueprint provides a clear path from pilot to production: A modular, compliant platform that improves clinical efficiency and patient care today, and a governed foundation for research and faster development of new medical products tomorrow.

Table of Contents

1. Introduction

Hospitals and MedTech companies operate under intense pressure: Staff shortages, rising costs, fragmented systems, and error-prone manual data entry slow care and increase risk. Data is locked in silos, devices are inconsistently connected, and compliance demands are growing. To fix this, healthcare providers and manufacturers need platforms that move data reliably from devices to clinical systems and research & data analysis environments without compromising safety, privacy, or regulatory obligations.

This report outlines the essential parts of such a platform and illustrates them with a concrete example: The SEMECO platform developed within the SEMECO research project. SEMECO is a German innovation cluster that unites universities, research centers, and industry to accelerate digitally enhanced medical devices. By combining AI, secure hardware, and modular software services, the initiative aims to shorten time-to-market while improving safety and efficiency.

Within SEMECO's "Secure & Trusted System Architectures" stream, we and our partners design a secure, compliant platform that spans on-premises and cloud environments. The work covers EU-resident compute and storage (including confidential computing), managed healthcare services such as FHIR APIs, gateways, and device management, end-to-end security (identity and access, encryption, monitoring), compliance automation (audit-ready logging, evidence bundles), Service-oriented Device Connectivity (SDC) for device interoperability, data anonymization for DICOM and FHIR datasets, and ready-made modules that digitalize clinical workflows.

This report provides a clear view of the regulatory mandates shaping cloud architectures (including MDR, FDA guidance, HIPAA, GDPR, ISO 27001, SOC 2, NIS2, and the forthcoming EU AI Act); a description of the core building blocks required for secure health-data workflows; a comparison of global hyperscalers and European sovereign clouds; guidance on when to buy standardized services and when to build differentiating capabilities; and a hybrid deployment strategy that delivers compliance, scalability, and cost efficiency together.

Taken as a whole, these insights help healthcare providers and MedTech companies assemble a robust, compliant platform that accelerates innovation while upholding the highest standards of data protection and patient safety.

semeco.info, zeiss.com

2. Regulatory Framework

Healthcare-focused platforms must embed compliance at every layer. The EU Medical Device Regulation (MDR) requires CE-certified software, unique device identification (UDI), and "state-of-the-art" IT security. In the United States, FDA regulations (21 CFR Part 820 and Part 11) mandate design controls, software validation, electronic records/signatures, and comprehensive cybersecurity plans with a Software Bill of Materials (SBOM). HIPAA enforces strict controls on protected health information under a Business Associate Agreement, while GDPR treats health data as a special category that demands lawful basis, purpose limitation, data minimization, pseudonymization, and tight cross-border transfer rules.

Whether a platform itself is a medical device under the MDR or other regulations depends on its intended use: If it only aggregates and distributes data, it may not be classified as a medical product; if it provides functions that can influence clinical decisions (e.g., decision support), it may fall under medical device regulation. For this reason, platforms are typically modular, so that only safety-critical components are classified and certified, while non-critical services remain outside the medical-device scope reducing regulatory burden without compromising safety.

To demonstrate trustworthy security, platform providers must establish a shield of trust and hold **ISO 27001** certifications (with healthcare guidance in ISO 27799), **SOC 2 Type II** reports, and, where relevant, **HITRUST CSF** attestations. Additional mandates such as **NIS2, EU AI Act** and Germany's **BSI C5** further raise the bar for network security, AI risk management and cloud-specific compliance.

Key point: Every component, from encryption at rest/in transit and key management through logging, incident response and vulnerability management, must map directly to one or more regulatory requirements.

EU MDR, EU GDPR, EU NIS 2, EU AI Act

Shield of Trust

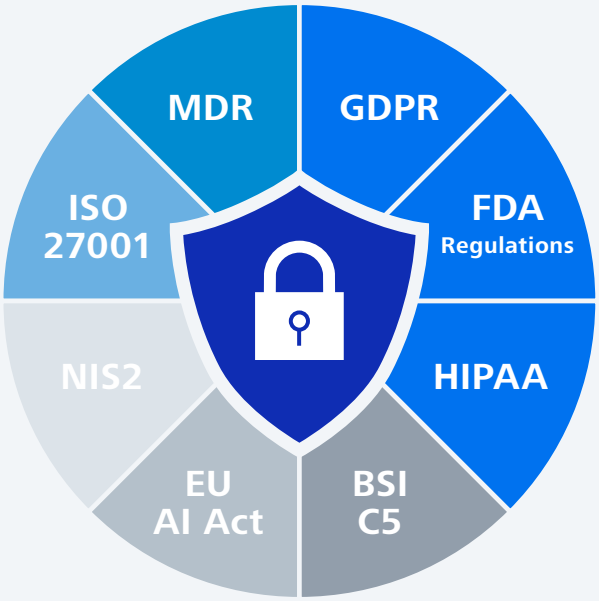

3. Core Platform Building Blocks

Modern medical platforms split responsibilities between Edge and Cloud to balance clinical constraints with scale. Each part is split into several layers, which are described in more detail in the following.

Edge (clinical side)
Runs on gateways or local servers next to devices. It enforces policy where data is created, handles on-prem processing for low-latency and offline-tolerant workflows, and respects data-residency constraints. Typical use cases: Device retrofit and protocol bridging (SDC/HL7/DICOM to FHIR), bedside alarm orchestration and early-warning scores, on-site imaging pre-processing, and anonymization/pseudonymization before data leaves the hospital.

Cloud (center side)
Provides elastic storage and compute, standardized interoperability services, and centralized identity, security operations, and compliance automation. It supports cross-site and longitudinal use cases: Population-scale analytics, European Health Data Space (EHDS)-ready secondary use, model training/serving, fleet management, consent enforcement, and enterprise reporting.

Interplay
Data moves from Edge to Cloud over secure, authenticated links (e.g., mTLS) and can be store-and-forwarded if the network is down. What goes up is already validated and privacy-protected data, plus basic telemetry about system health. In the other direction, from Cloud to Edge, the platform sends the things the Edge should apply at runtime, policies (who may send what, where), schemas (data formats), models (for analytics/alerts), and credentials. Both sides stay in sync through shared observability, automated CI/CD, and versioned contracts, so workflows remain resilient, auditable, and standards-compliant across clinical sites.

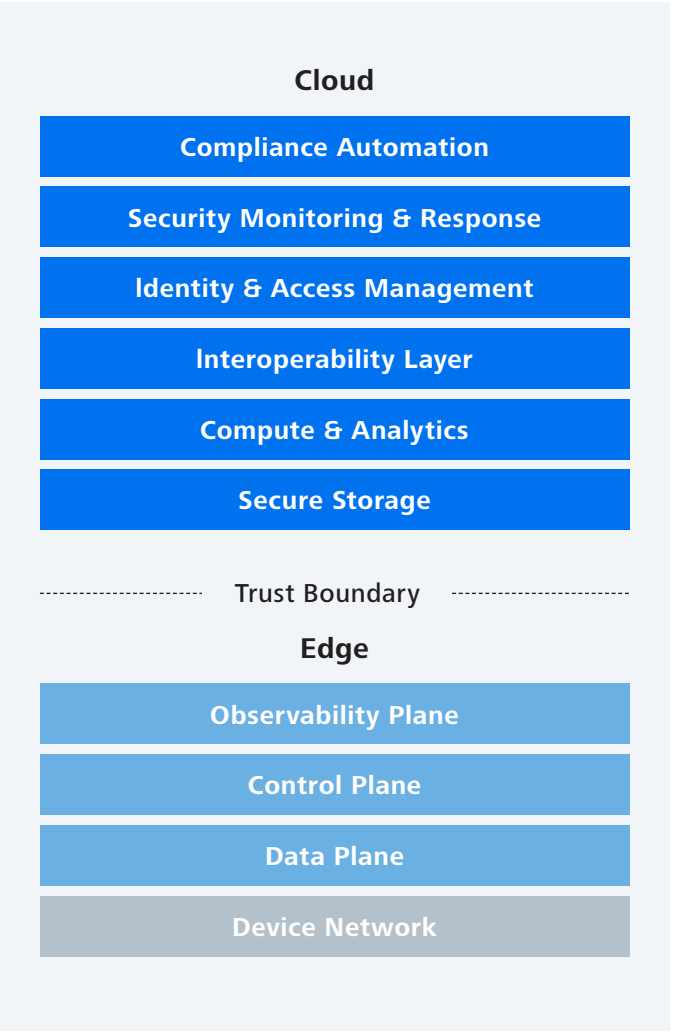

3.1 Cloud Layers

Purpose & scope
Cloud components provide elastic compute, durable storage, and standardized interfaces, e.g. for analytics. They complement edge deployments by handling longitudinal data, cross-site integration, and heavy processing while enforcing security, privacy, and regulatory controls at scale.

Architecture overview
A typical design separates data services (storage/streams), compute services (batch, real-time, ML), interoperability services (APIs/transformers), identity & access, security operations, and compliance automation. All run inside private networks (VPCs), with explicit egress paths, infrastructure-as-code, and end-to-end observability.

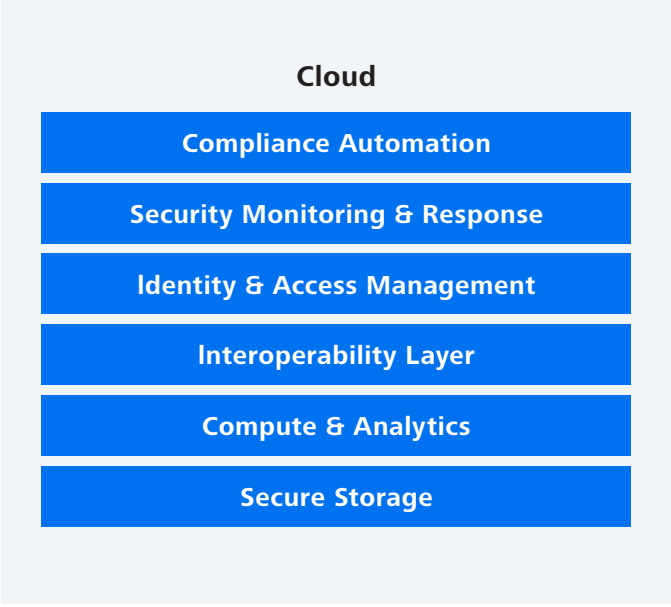

The cloud part consists of 6 layers, which provide the following functionality:

Secure storage
This layer provides the durable foundation for clinical data. Operational records live in relational or NoSQL databases, while images and documents reside in object storage, and long-term records are preserved in immutable archives. Encryption in transit and at rest, HSM-managed keys, granular access policies, and residency controls ensure that data remains protected and traceable through tagging and versioning.

Compute & analytics
Here the platform executes workloads reliably and at scale. Containers or virtual machines handle steady services, serverless functions respond to events, and distributed engines support population-level analytics. Private networking, zero-trust service access, and audited CI/CD pipelines maintain reproducibility; confidential computing is available when data must remain encrypted during processing.

Interoperability layer
This layer turns heterogeneous inputs into consistent, shareable clinical information. Legacy formats such as HL7 v2 and DICOM are transformed into normalized resources exposed through FHIR APIs, with semantic mappings (e.g., LOINC, ICD-10) preserving meaning across systems. Strong interface controls, authentication, rate management, and schema validation, protect patient safety and data quality.

Identity & access management
A central identity provider unifies authentication across cloud and on-prem environments. Least-privilege authorization, multi-factor authentication, and short-lived, just-in-time privileges reduce risk while keeping access aligned with role changes. Automated joiner–mover–leaver processes keep permissions accurate over time.

Security monitoring & response
Continuous posture management and centralized logging provide visibility across the platform. Network activity and workloads are monitored for threats, software artifacts are scanned and tracked via SBOMs, and defined playbooks guide containment and forensics. Integration with ticketing systems ensures incidents are documented and remediated end-to-end.

Compliance automation
Policies and controls are mapped to regulatory obligations and checked continuously for drift. Evidence, such as encryption attestations, access histories, deployment records, and model lineage, is generated from runtime telemetry, keeping audits lightweight and repeatable. When a control deviates, workflows trigger corrective actions and prevent unsafe changes from reaching production.

How it fits together
Together, these layers form a coherent cloud foundation. Storage provides durable, traceable data with clear residency, while compute and machine learning operate on governed datasets inside private networks. The interoperability layer standardizes data exchange across systems, and identity and access management ensures the right people and services have the right, least-privilege permissions. Security operations monitor, detect, and respond to threats, and compliance automation ties everything back to policies and auditable evidence. The result is a platform that scales analytics and innovation without compromising security or regulatory readiness.

enisa.europe.eu, microsoft.com, aws.amazon.com, cloud.google.com

3.2 Edge Layers

Purpose & scope
Edge runs locally inside clinical networks as a policy-enforcing trust boundary to device networks and the cloud. This design supports on-premise processing, data residency, and low-latency workflows; enables safe retrofitting of legacy devices; and aggregates, normalizes, and securely forwards data to cloud services for advanced analytics.

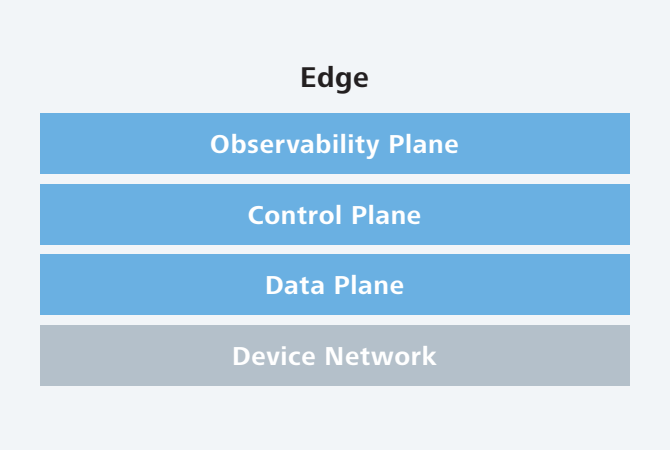

Architecture
The edge typically separates a data plane (ingest, normalize, enrich, buffer, egress), a control plane (configuration, policy, secrets, device registry, updates), and an observability plane (metrics, logs, traces, security telemetry).

These 3 layers provide the following functionality:

Data plane

The data plane is responsible for acquiring, transforming, and moving information at the edge. Connectivity adapters translate DICOM, HL7 v2, FHIR, SDC, and proprietary protocols into consistent schemas, while payloads are checked against profiles and business rules. Units and code systems are harmonized, duplicates are removed, and precise time stamps are maintained. Where privacy is required, de-identification and anonymization run on-premise; pseudonymization or tokenization with vault-backed keys can be applied to enable linkage without exposing identities. Low-latency use cases are supported through rule-based eventing and lightweight ML, and resilience is made possible. Data leaves the edge over authenticated, policy-controlled channels. Typical deployments range from bedside gateways to edge servers and small device-adjacent agents, with messaging standardized on MQTT/AMQP/HTTPS and governed by a shared schema registry.

Control plane

The control plane governs configuration, identity, and lifecycle. Devices enroll through zero-touch provisioning and authenticated bootstrap, establishing identities anchored in TPM-backed keys and short-lived credentials. Configuration is managed declaratively with drift detection and policy conformance so sites stay aligned with intended settings. Software supply-chain security is built in: Secure boot, disk encryption, signed artifacts, SBOM tracking, vulnerability scanning, and atomic A/B updates. Network segmentation and least-privilege access limit blast radius, while orchestration choices, from lightweight container runtimes to micro-Kubernetes, are selected to match site scale and operational maturity. Compliance follows recognized standards (ISO 27001, IEC 62443, IEC 62304 where applicable, ISO 27701, and the EU Cyber Resilience Act), with the classification of components determined by their intended use.

Observability plane

The observability plane provides continuous visibility and assurance. Metrics, logs, traces, and security telemetry are collected with clinical context such as device identifiers, encounter links, and location, giving operators a clear picture of health and performance across nodes. Incident response is supported by tamper-evident logging, remote snapshotting, and integration with the hospital security operations center, enabling fast investigation and recovery. Capacity profiles for CPU, memory, storage, and network throughput guide sizing and service-level objectives, while ongoing checks of policy adherence keep the edge auditable, resilient, and ready to interoperate with cloud analytics and coordination services.

How it fits together

Edge components collect data securely at the bedside, convert it into standardized formats, and apply privacy safeguards before any information leaves the clinical network. They then forward data reliably to upstream systems, even through outages, using authenticated and policy-controlled channels. Built with security by design and aligned to regulatory requirements, these components can start on a single gateway and grow to hospital-wide clusters. Because they follow established health data standards, they integrate cleanly with cloud services and existing clinical systems.

sciencedirect.com, sciencedirect.com, springer.com, mdpi.com, wiley.com

4. Bringing it together: SEMECO Platform – Modular, Hybrid, Regulatory-First

SEMECO is a German innovation cluster that accelerates digitally enhanced medical devices by combining AI, secure hardware, and modular software shortening time to market while improving security, safety and efficiency. Its "Secure & Trusted System Architectures" stream defines a secure, compliant platform across edge and cloud, including EU-resident compute and storage with confidential computing. The result is a modular, hybrid foundation that integrates with existing systems and adds healthcare-specific capabilities where they create the most value.

4.1 Vision, Scope and Implementation

The vision of SEMECO is to sit between diverse healthcare data sources, medical devices, hospital systems (HIS/PACS), and specialized research/industry platforms, and the user groups that need them. It shall provide a modular service layer with cloud-hosted capabilities (e.g., fleet management, consent management) and edge-hosted capabilities (e.g., SDC-based interoperability, data anonymization, pseudonymization). Clinics could streamline device connectivity and consent-aware workflows; research teams could access governed, de-identified datasets for studies; and industry could operate secure device fleets and post-market surveillance. Each capability could be adopted standalone or combined to form a cohesive, integrated platform. SEMECO will not reinvent the wheel. We deliberately assemble the platform from proven cloud ecosystems and standards, then add specialized healthcare components, e.g. interoperability, privacy services, and clinical workflow logic, that extend use cases and create differentiated value.

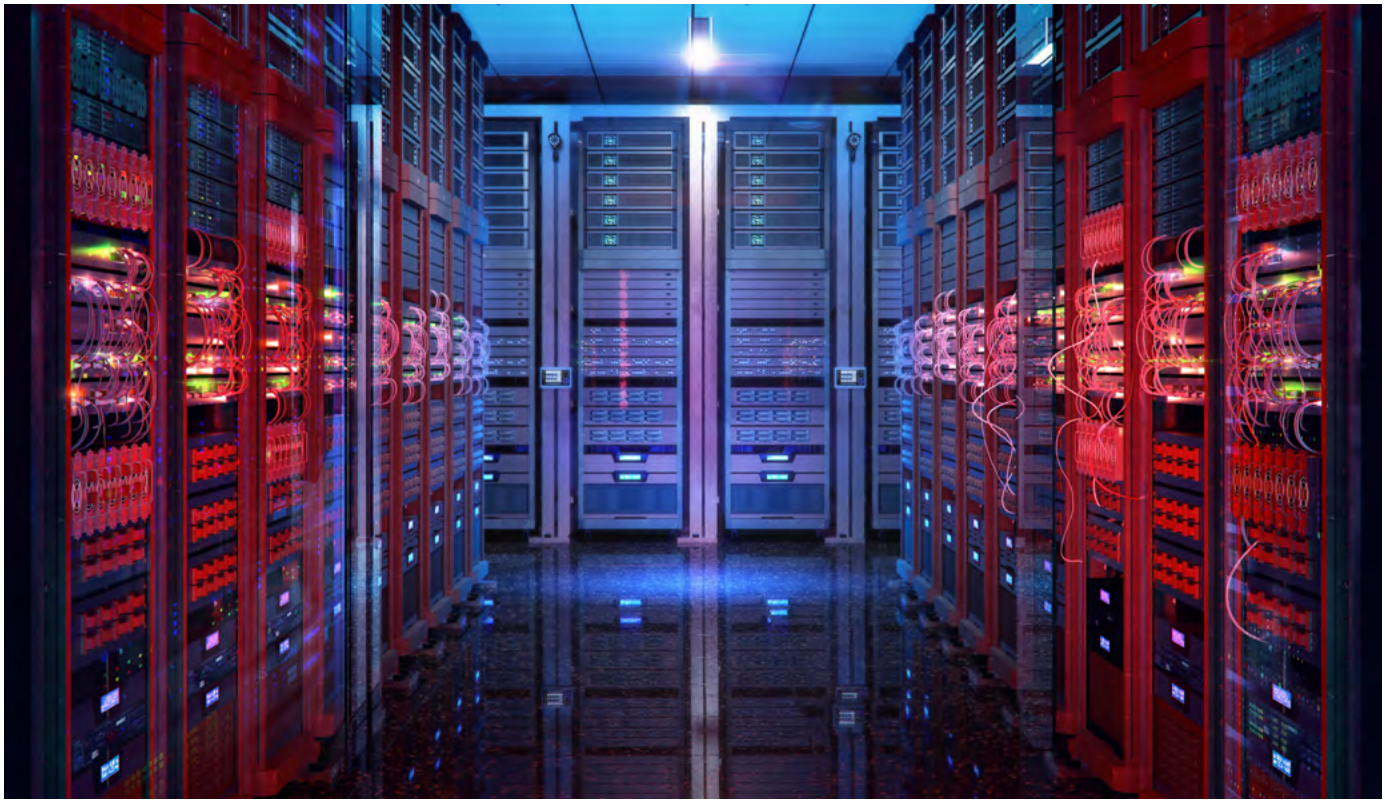

For example: Modern hospital IT spans several tiers: Medical devices at the bedside feed data into the edge network (OR/ICU gateways); higher up, clinical systems (HIS, PACS, EMR, VNA), vendor-specific data hubs, and IoMT integration platforms exchange information; at the top, a cloud & analytics layer delivers longitudinal insights. SEMECO shall support two deployment scenarios:

1. **Component overlay.** Hospitals and MedTech providers with mature infrastructure could pick targeted SEMECO modules such as SDC/FHIR/DICOM adapters at the device or edge, on-prem anonymization and policy gateways in the clinical network, or cloud-based analytics and compliance services, plugging them into existing stacks to close specific gaps.

2. **Modular data-platform build-out.** Organizations lacking an end-to-end architecture could adopt SEMECO as the backbone itself, using its edge, integration, and cloud components to assemble a full healthcare data platform that still connects cleanly to incumbent systems.

In both cases, SEMECO will help to enable a secure, standards-based data flow from device to cloud, letting organizations adopt only what they need today while keeping a path to a comprehensive platform tomorrow.

The figure below provides an overview for deployment scenario 1. Here, SEMECO potentially augments existing clinical platforms with the integration of healthcare-specific cloud or edge services into various infrastructures.

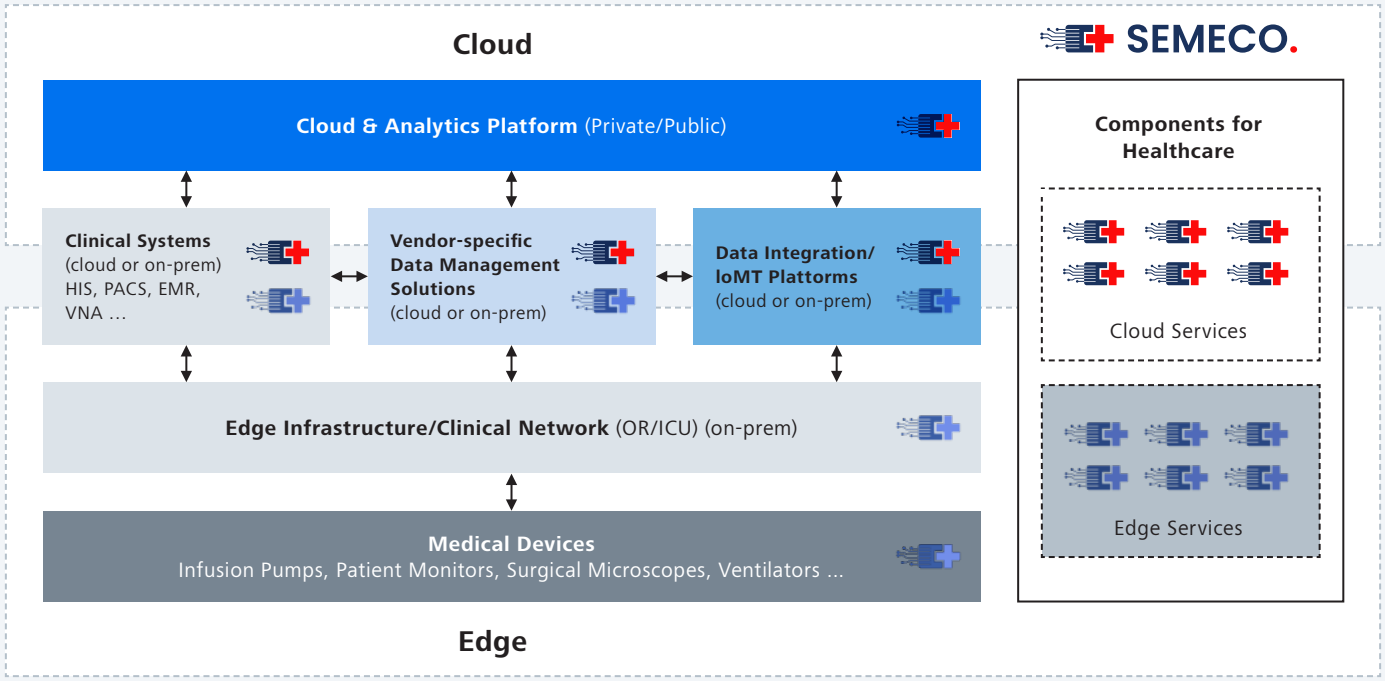

The implementation of the SEMECO platform is guided by modular, hybrid, and regulatory-first principles. It deliberately builds on proven ecosystems for generic capabilities while focusing engineering on healthcare-specific services. A microservices cloud, a modular edge, clear certification boundaries, and security by design ensure the platform can scale, remain compliant, and accelerate innovation without reinventing the wheel.

The following design principles are key for SEMECO:

Make-vs-buy
Generic capabilities such as compute, storage, networking, observability, and managed databases are sourced from hyperscalers and, where needed, regional providers for data residency. Engineering effort focuses on health-specific services like FHIR/SDC/DICOM bridging, device interoperability, anonymization and pseudonymization, and workflow components aligned to clinical processes.

Cloud architecture
The cloud tier follows a microservices approach: Containerized services (e.g., Docker) orchestrated by Kubernetes with API contracts and zero-trust networking. This enables independent scaling, fault isolation, and rapid iteration while preserving strong security, auditability, and cost governance.

Certification boundaries
Safety-critical functions, such as decision support, are isolated as certifiable medical products with dedicated risk management, verification, and evidence. Non-critical platform services remain uncompromised in security and quality but avoid medical device certification, allowing faster development and release cycles.

Edge architecture
At the edge, SEMECO uses a modular design deployed on secure gateways or local servers. Connectivity adapters normalize device data, on-prem privacy pipelines enforce policy before egress, and store-and-forward buffers provide resilience for low-latency and offline scenarios.

Interplay of cloud and edge
Policies, schemas, and models flow from cloud to edge; validated, privacy-governed data and telemetry flow back to the cloud over authenticated channels. Versioned contracts and automated tests keep both tiers in lockstep and reduce operational risk.

Regulatory and security posture
Compliance with MDR, GDPR, ISO 27001/27701, and related frameworks is built into CI/CD through policy gates, SBOMs, logging, and automated evidence generation. Security by design underpins the stack: Encryption, least-privilege access, signed updates, and continuous vulnerability management.

Why this design
This modular architecture isolates change and risk to small units, so teams can iterate quickly without destabilizing the whole platform. Building on mature cloud ecosystems for generic capabilities cuts time to value and operating cost, while our own effort targets healthcare-specific functions where differentiation matters. Clear certification boundaries confine regulatory scope to safety-critical services (e.g., decision support), allowing the rest to evolve on faster cycles. A hybrid edge-cloud split places policy and low-latency processing near devices and uses the cloud for scale, analytics, and coordination. By anchoring everything in open standards (FHIR, SDC, DICOM) and security by design, the platform scales across sites, remains portable across providers, and stays extensible through SEMECO's specialized components delivering compliance, performance, and innovation together.

zeiss.com

4.2 Advanced Healthcare-Specific Services

Focus and approach

As previously stated, SEMECO targets healthcare capabilities that create outsized value while reusing proven cloud and edge building blocks. We validated priorities with partners and conducted in-depth market research, then concentrated on components that accelerate digital workflows without rebuilding generic infrastructure. In the following, we discuss our results presenting the different service tiers and key implementation principles:

Priority 1 – Core services

SDC-based interoperability normalizes device data and enables a safe, standardized exchange with clinical devices and systems. Anonymization of patient data protects privacy across imaging and telemetry, running at the edge for data-residency needs and in the cloud for scale. Workflow components orchestrate clinical tasks, capture evidence, aggregate and forward data, and integrate with consent and audit trails to make processes repeatable and compliant.

Priority 2 – Supplementary services

Consent management records and enforces patient permissions, integrates with FHIR Consent, and drives allow/deny decisions at APIs and pipelines. Device/fleet management inventories unique device identifiers (UDIs), handles secure onboarding and signed over-the-air (OTA) updates, and aggregates telemetry to support post-market surveillance. Device retrofitting brings legacy devices online and makes them interoperable via edge adapters that translate proprietary outputs into standardized resources.

Priority 3 – Additional services

Audit logging provides tamper-evident trails across edge and cloud. Health-data processing offers scalable pipelines for transformation and enrichment. Pseudonymization enables controlled linkage across datasets where full anonymization would break utility.

Prioritized Platform Services

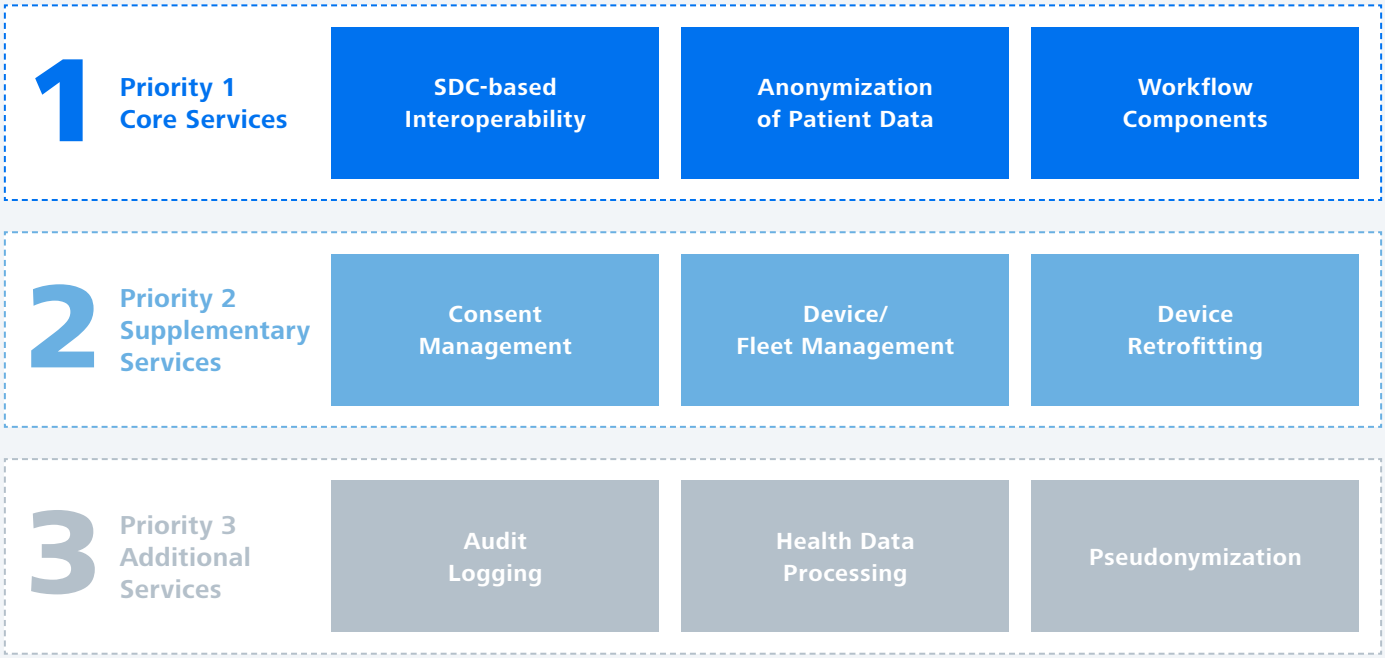

Our platform services are shaped by three practical decisions: Where they run (edge vs. cloud), how they inherit security and compliance from the platform, and what we build vs. buy to move fast without sacrificing quality:

Edge vs. cloud placement
Interoperability bridges, workflow components and retrofitting run primarily at the edge to handle local protocols, low-latency needs, and store-and-forward. Consent checks, fleet management, and longitudinal processing run mainly in the cloud for elasticity and cross-site coordination. Anonymization and pseudonymization are deployable in both tiers, chosen by residency and latency requirements.

Security and compliance by design
All components inherit platform controls, either from SEMECO or existing platforms, encryption, IAM with least privilege, SIEM, vulnerability management, and map to GDPR and relevant security standards. Evidence generation and configuration checks are automated so services remain audit-ready through their lifecycle.

Make-vs-buy
Commodity capabilities (compute, storage, observability, managed FHIR/DB, CI/CD) are sourced from established providers. SEMECO invests in the differentiated layer: Interoperability, privacy pipelines, workflow components, and retrofit tooling. Where mature solutions exist (e.g., consent engines, OTA frameworks), we integrate rather than build.

Outcome
By deciding where services run (edge for low-latency device integration and retrofitting; cloud for elastic consent, workflow, fleet, and longitudinal processing; anonymization/pseudonymization in either tier by residency needs), ensuring every service inherits security and compliance controls (encryption, least-privilege IAM, SIEM, vulnerability management, automated evidence), and applying a pragmatic make-vs-buy strategy (buy commodity cloud capabilities, build healthcare-specific differentiators), SEMECO delivers impact quickly without sacrificing quality.

This modular, hybrid approach brings the foundation consisting of interoperability, privacy, and workflow automation online first, then scales out with consent, fleet operations, and analytics as needs grow. The result is faster time to value, lower risk, and a clean path from pilot to production, all while staying audit-ready and portable across existing environments.

zeiss.com

4.3 Components at PoC Stage

We have initiated the development of two foundational components to demonstrate SEMECO's modular approach in practice. First, a DICOM anonymization module aimed at EHDS-ready data sharing, designed to run at the edge or in the cloud while preserving analytical utility. Second, an SDC-based interoperability module that standardizes device communication for clinical use cases (e.g., alarm orchestration and workflow support), usable either embedded in devices or as an edge service for retrofitting legacy equipment.

DICOM anonymization component
We are building a DICOM anonymization component to solve a core challenge: Anonymizing imaging data while preserving as much data quality as possible. This balance is essential for secondary-use scenarios under the European Health Data Space, such as research, clinical studies, and the development of new, innovative medical products, where privacy must be protected without compromising analytical value.

The core is an embeddable Java library exposing a configurable anonymization workflow (policies plus input/output streams). It couples a DICOM service (dcm4che3) for parsing and writing with an ARX-based privacy engine for k-anonymity, generalization, and suppression. The same library runs at the edge or in the cloud and can be utilized via CLI, desktop/GUI, web app via middleware, or cloud-native microservice, ensuring consistent processing across environments.

The tool classifies data into four categories: Identifying, insensitive, sensitive, and quasi-identifying attributes. For each dataset and its intended use, both qualitative and quantitative risk assessments are carried out following established guidelines and risk scoring methods. This step ensures a high level of data utility of the anonymized data while ensuring data privacy. Based on these assessments, data fields are then processed - either generalized, retained in full, or removed. To improve the quality of the anonymized output, some records may be fully suppressed.

Processed and anonymized data is written back to the original data or folder structure. To ensure security and compliance, the system validates data against schemas and business rules and adjusts privacy settings based on cohort size, while staying compatible with HL7/FHIR/DICOM standards.

zeiss.com, zeiss.com, zeiss.com, arx.deidentifier.org

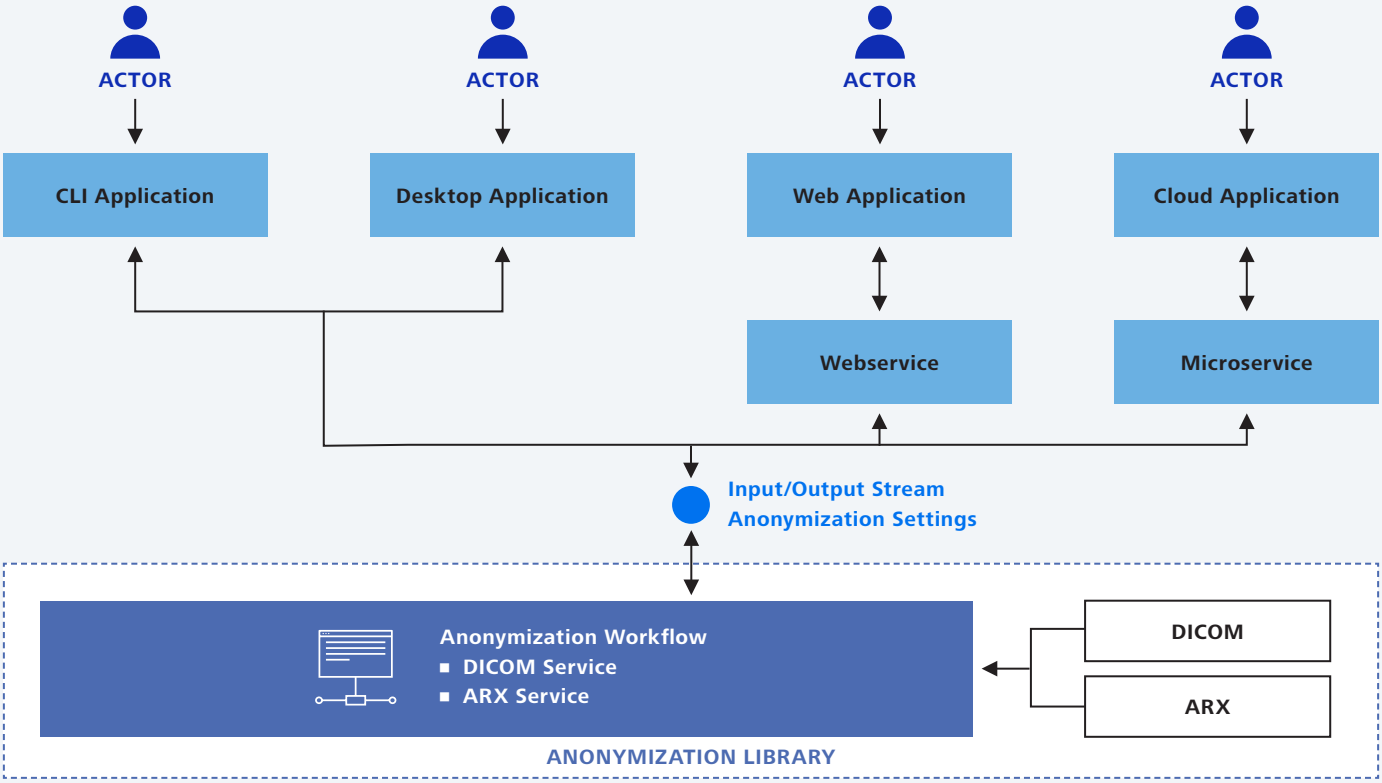

SDC-based interoperability component

Further, we are developing an SDC (IEEE 11073) interoperability component to standardize device-to-device communication and unlock clinical use cases such as automated alarm management, smoother staff workflows, and improved patient experience ultimately increasing efficiency and reducing errors in clinical settings. The component complements existing standards (e.g., HL7/FHIR, DICOM) rather than replacing them. The PoC provides a software library that can be embedded directly in device software and an edge-deployable service for retrofitting of legacy devices. It supports secure discovery, capability modeling, and streaming of metrics/events, with optional mappings to FHIR for downstream integration. Initial work includes use cases, architecture, and a reference library aligned with regulatory and security expectations.

The component is engineered under an Information Security Management System (ISMS) (ISO 27001) with Cyber Resilience Act readiness. Because it does not process real patient data, the focus is on integrity first, with strong confidentiality and availability requirements for the information it handles (e.g., device metadata, configurations). Auditability and traceability are built in.

Our SDC component can be integrated into medical device software or hosted as an edge service on gateways/servers, following secure interfaces and least-privilege access. While runtime apps and real patient data are out of scope for the PoC, the component is designed for clean handoff into regulated products and secure operations across the device lifecycle.

See the following architecture: A lightweight library sits between the product application and external SDC devices. It exposes an agnostic interface with Provider (publish capabilities/metrics) and Consumer (subscribe) services backed by MDIB models, while a Session Manager handles secure discovery, association, and streaming. A Protocol Adapter implements IEEE 11073 SDC transport; stage storage persists state; and centralized logging/error handling ensures traceability.

zeiss.com, zeiss.com

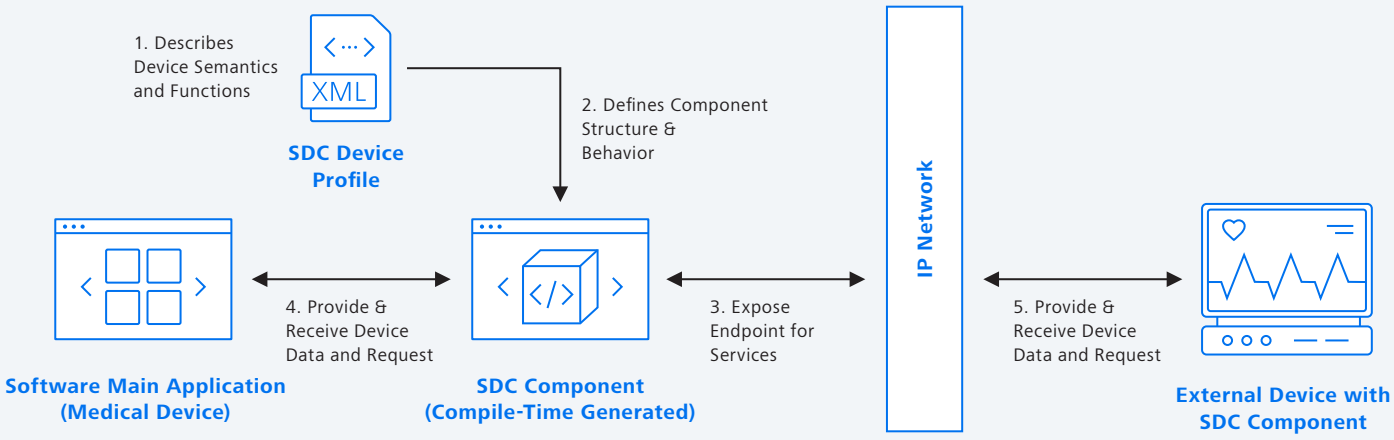

5. Deployment Models, Provider Comparison & Make vs Buy

5.1 Deployment Models

Selecting a hosting model for a medical platform, SEMECO included, means balancing three forces: How much control and data residency you require, the cost to build and operate, and the level of regulatory assurance you need. In practice, most healthcare deployments pair an edge footprint in clinical networks with one (or more) cloud options to meet latency, scalability, and compliance goals. The table below summarizes the advantages and drawbacks of each deployment scenario to help guide that choice.

As guidance, use the edge for bedside scenarios that demand low latency, local data processing, aggregation and forwarding, or retrofitting of legacy devices; let it buffer and normalize data before forwarding upstream. Pair this with a public or hybrid cloud for elastic analytics, AI/ML, cross-site coordination, and long-term data services. Choose colocation or private cloud when strict data sovereignty, dedicated networking, or predictable long-lived workloads justify higher ownership cost. Across all models, preserve portability with infrastructure-as-code, containers, and standard APIs to limit vendor lock-in and ease future migrations.

Model	Description	Advantages	Drawbacks
Private Cloud	Your own data centers; full control of hardware/software.	Maximum customization and data sovereignty; tailored security zones.	High CAPEX/OPEX; specialized ops staff; slower feature cadence.
Colocation	Your servers in a third-party data center.	Professional facilities, power, and connectivity; lower CAPEX than private.	Limited flexibility on racks/power; provider policy dependencies.
Public Cloud	IaaS/PaaS/SaaS (AWS, Azure, GCP, etc.).	Elastic scale; rapid feature uptake; pay-as-you-go; global reach.	Potential lock-in; careful design needed for EU-only residency.
Hybrid / Multi-Cloud	Mix of private/colo/public with workload placement by need.	Balance compliance (EU clouds) with burst compute/AI on hyperscalers; geo-resilience.	More complex networking, IAM federation, and cross-platform governance.
Edge	Gateways or local servers inside clinical networks.	Low latency and offline resilience; data stays on-prem when required; protocol bridging and legacy device retrofit; store-and-forward to cloud.	Site-by-site hardware management; scaling limited by footprint; patching/ physical security responsibilities.

Additional operational considerations

A production medical platform should be reliable, secure, affordable, and easy to audit. Plan for outages with clear recovery playbooks, a standby environment you can switch to, and regular failover drills. Keep connections fast and safe by using private links where possible and separating network zones so problems don't spread. Control costs with simple rules: Tag every resource, reserve capacity for steady workloads, and set budget alerts.

Build trust in AI by tracking models, checking them for bias, and explaining results. Reduce vendor risk by using standard security assessments, asking for certifications (e.g., ISO 27001 or SOC 2), defining service levels, and agreeing on exit terms before you start. Ease integration by placing an API layer in front of older systems and using common healthcare standards for exchange. Keep people ready with role-based training and regular security exercises. And design for sustainability by choosing efficient compute options when you can and measuring your carbon footprint.

5.2 Providers: Hyperscalers vs. European Sovereign Clouds

Selecting a hosting partner is a trade-off between scale and service breadth on one side and data sovereignty and legal clarity on the other. For medical platforms, both dimensions matter; the pragmatic answer is usually a hybrid multi-cloud. In the following, discover strengths and risks with mitigations as well as legal considerations for both scenarios.

Two main viable paths exist for hosting the cloud part of a medical platform. Global hyperscalers (AWS, Azure, Google Cloud) bring near-unlimited scale, global reach, and deep catalogs, managed FHIR/IoT, advanced data/AI, and mature security, backed by fast release cycles and large partner ecosystems. The trade-offs are tighter coupling and potential lock-in, careful engineering for EU-only residency, and possible exposure to the US CLOUD Act even for EU-hosted data.

European sovereign providers (e.g., StackIT, OVHcloud, SysEleven, Open Telekom Cloud, IONOS, Cloud&Heat, Aruba) offer EU-only operations, simpler GDPR/NIS2 alignment, and reduced extraterritorial risk; their use of open stacks (OpenStack/Kubernetes) improves portability and exit options, with local teams and local-language SLAs. The trade-offs are smaller service catalogs, fewer healthcare-specific managed offerings, and slower access to cutting-edge AI/ML or broad geographic capacity.

The table below compares these two groups across the dimensions that matter most, data sovereignty, service breadth, innovation pace, support, and portability/lock-in:

Aspect	Hyperscalers (AWS, Azure, GCP)	European Providers (e.g. StackIT, OVHcloud)
Data Sovereignty	Multi-region but requires careful data-residency configuration	EU-only by design; simplified GDPR/NIS2 compliance
Service Breadth	Extensive healthcare-specific modules (HealthLake, FHIR, IoT, AI)	Core IaaS/PaaS; fewer specialized managed healthcare services
Innovation Pace	Rapid release cycles; advanced AI/Big Data tools	Slower adoption of bleeding-edge technologies
Support & SLAs	24/7 global support; large partner network	Local teams, personalized SLAs; local-language support
Lock-in & Portability	High API lock-in potential; vendor-proprietary extensions	Based on OpenStack/K8s; easier migration

Legal considerations

The US CLOUD Act can compel US providers to disclose data under court order, even if stored in the EU. GDPR treats health data as highly sensitive and restricts transfers to jurisdictions without adequate protection. For healthcare, this creates a compliance risk when protected or personal health information (PHI) is hosted with a US provider. Non-compliance may trigger severe penalties, up to €20M or 4% of global turnover, plus potential breaches of medical confidentiality and civil claims. To mitigate risk, you can favor data minimization and anonymization/pseudonymization before data reaches hyperscalers; keep encryption keys in EU-controlled HSMs and enforce double-key or split-key models where available. Use EU-only regions, rigorous residency controls, and confidential computing for sensitive processing. Maintain containerized workloads and open APIs to preserve portability and negotiate exit clauses up front.

Decisions for SEMECO

When considering the modular data-platform build-out deployment scenario for SEMECO discussed previously, we aim to adopt a hybrid multi-cloud: Run elasticity-hungry analytics/AI and broad managed services on hyperscalers with de-identified data, while placing PHI stores, key management, and residency-critical workloads with EU-sovereign providers. Edge components handle local ingestion plus forwarding and privacy enforcement as well as promote interoperability; cloud microservices remain portable across vendors via Kubernetes, infrastructure-as-code (IaC), and standard interfaces.

5.3 Make-vs-Buy Decision Framework

While building a medical platform such as SEMECO, it is crucial to allocate resources between in-house development and third-party services. The following table summarizes key criteria, which we apply to all our intended services to decide what to buy and where we focus our internal development efforts:

Criterion	Recommendation: Buy	Recommendation: Make
Compliance assurance	Use certified vendor services for consent management, de-identification, SOC 2/ISO 27001 evidence collection, and MDR documentation to save audit effort.	Build only if regulations demand custom features not covered by existing solutions (e.g., special anonymization workflows for rare-disease research).
Time to market & cost	Buy standardized infrastructure (IaaS, managed FHIR/IoT, CI/CD tooling) to accelerate rollout and keep costs predictable via subscription/OPEX.	Make when you need long-term cost control and the functionality is central to your business, justifying CAPEX and specialized teams.
Feature fit & flexibility	Buy when 80–90% fit is acceptable and vendor roadmaps align with your needs; configure instead of customizing.	Make for highly specific workflows, UX, or integrations that define your differentiation and cannot be achieved by configuring vendor tools.
Lock-in & portability	Buy when portability is maintained via open standards, containerization, and exit clauses in contracts.	Make if vendor lock-in is unavoidable and long-term control over APIs, data, or algorithms is critical.
Strategic differentiation	Buy commodity plumbing like compute, storage, observability, and generic FHIR stores. These do not differentiate your solution.	Make core intellectual property: Device interoperability (SDC/legacy adapters), advanced analytics engines, privacy-preserving data pipelines, and clinical workflow components.
Expertise & staffing	Buy non-core services where vendor expertise exceeds what you can efficiently build (e.g., security operations, managed databases).	Make where your in-house domain knowledge gives you an edge (e.g., MedTech-specific clinical processes, regulatory-aligned innovation).

For SEMECO, the guiding principle is to buy where maturity exists and make where differentiation matters. This means relying on established providers for infrastructure-adjacent services such as IaaS, managed databases, observability, FHIR/IoT backends, and consent management. These areas are compliance-heavy, commoditized, and benefit from vendor certifications and economies of scale. Outsourcing them accelerates time to market and reduces operational risk.

Internal development efforts, in contrast, should focus on modules that define SEMECO's unique intellectual property. This includes SDC-based device interoperability and retrofit adapters that bring legacy devices into modern ecosystems, advanced analytics engines that generate actionable insights from medical data, innovative clinical workflow components, and advanced data anonymization tailored towards the EHDS. These components represent the true differentiators of the platform. They drive efficiency in hospitals, enable new clinical use cases, and provide strategic value that off-the-shelf tools cannot.

By drawing this boundary deliberately, SEMECO avoids wasting effort on commodity "plumbing" while ensuring full ownership of the features that will shape its competitive edge. The result is a balanced strategy: Fast rollout and regulatory confidence from vendor services, combined with focused innovation where SEMECO adds the most value.

6. Impact of a Medical Platform

Once a holistic medical platform is in place, such as a future SEMECO, you gain an edge–cloud infrastructure that streamlines clinical workflows, reduces costs and errors, and improves patient care. It also creates a governed foundation for research and faster development of new medical products and services. Examples include:

Care delivery

- Remote patient monitoring: Edge gateways collect vitals; policy-checked data streams to cloud dashboards with real-time alerts and full audit trails.
- OR workflow & alarm orchestration: SDC-enabled devices share state in real time; alarms are deduplicated and routed, reducing fatigue and improving team coordination.
- Imaging pre-processing & teleradiology: On-prem anonymization/compression; secure cloud review and collaboration without exposing PHI.
- Clinical decision support: Validated models run at the edge or in the cloud with explainability, versioned models, and evidence capture.

Operations & safety

- Predictive maintenance: Device telemetry feeds anomaly detection to schedule service proactively and reduce downtime.
- Fleet & software lifecycle management: Secure onboarding, signed OTA updates, SBOM tracking, and vulnerability remediation at scale.
- Hospital command center: Unified signals (beds, equipment, staffing) inform throughput optimization and escalation pathways.

Data governance & compliance

- Consent and access control: FHIR consent-driven allow/deny decisions enforced at APIs and data pipelines.
- Automated compliance & traceability: Continuous control checks, tamper-evident logs, and one-click evidence for MDR/FDA/GDPR audits.

Research & innovation

- Secondary use / EHDS participation: Anonymization and pseudonymization pipelines enable GDPR-compliant data sharing and cohort building.
- Real-world evidence & post-market surveillance: Aggregated device and outcomes data supports vigilance and iterative product improvement.
- Clinical studies: eConsent, patient-reported outcomes, and sensor data collection with governed access for investigators.

Together, these capabilities create a learning health platform delivering safer daily operations, measurable efficiency for stretched care teams, and a compliant engine for continuous R&D so your organization is ready for the future.

7. Summary & Conclusion

This report presented a practical blueprint for building a compliant, scalable medical platform that blends edge and cloud. The core idea is simple: Enforce policy and privacy where data is born, and use the cloud for scale, analytics, and cross-site coordination. By grounding the architecture in open standards (FHIR, SDC, DICOM), strong security controls, and automated compliance evidence, organizations can meet MDR, FDA, HIPAA, GDPR, NIS2, and EU AI Act expectations without slowing innovation.

A deliberate make-vs-buy strategy keeps teams focused on what differentiates them. Mature, commodity capabilities, compute, storage, observability, managed FHIR/IoT, consent, are best sourced from proven providers. Engineering effort concentrates on healthcare-specific value: Device interoperability, privacy-preserving data pipelines, and workflow components. In hosting, a hybrid multi-cloud approach balances sovereignty and capability by placing PHI and keys with EU providers while using hyperscalers for de-identified analytics and AI kept portable via containers, Kubernetes, IaC, and open interfaces.

The result is a platform that moves quickly from pilot to production: Safer, more efficient clinical workflows; better patient experiences; and a governed foundation for research and new product development. The SEMECO implementation illustrates how this modular, regulatory-first design can be realized today leveraging existing ecosystems, adding targeted healthcare components, and delivering measurable impact with confidence.

What matters most

- **Modular architecture:** Microservices in the cloud and pluggable edge components enable independent scaling, faster releases, and clear certification boundaries for safety-critical functions.
- **Security & compliance by design:** Encryption, IAM least-privilege, SBOMs, audit trails, and automated evidence generation are embedded in CI/CD.
- **Data governance & interoperability:** Standardized APIs (FHIR, SDC, DICOM) and privacy pipelines (anonymization/ pseudonymization) unlock safe data exchange and EHDS-ready secondary use.
- **Hybrid multi-cloud:** Use EU-sovereign clouds for PHI/keys and hyperscalers for de-identified analytics/AI maintaining portability via containers, IaC, and open interfaces.
- **Real impact:** From OR alarm orchestration to predictive maintenance and research cohorts, such a platform reduces errors, lowers costs, and improves patient outcomes while accelerating product innovation.

Bottom line: A holistic, standards-based platform is now a strategic necessity. By leveraging existing ecosystems and adding targeted healthcare components, such as DICOM anonymization and SDC interoperability, organizations can move from pilot to production with confidence: Safer workflows for clinicians, better experiences for patients, and a faster path to innovative medical products.

Ready to build your compliant healthcare platform?

If you're modernizing clinical workflows, integrating devices, or preparing for EHDS-ready data sharing, let's partner. We bring a modular Edge-Cloud reference architecture, accelerators like **SDC interoperability** and **DICOM anonymization**, and a clear **make-vs-buy** playbook, so you move from pilot to production faster, with confidence on safety and compliance.

Let's co-design a roadmap that reduces errors, lowers cost, and improves patient care while accelerating product innovation.

Get in touch to schedule a discovery session.

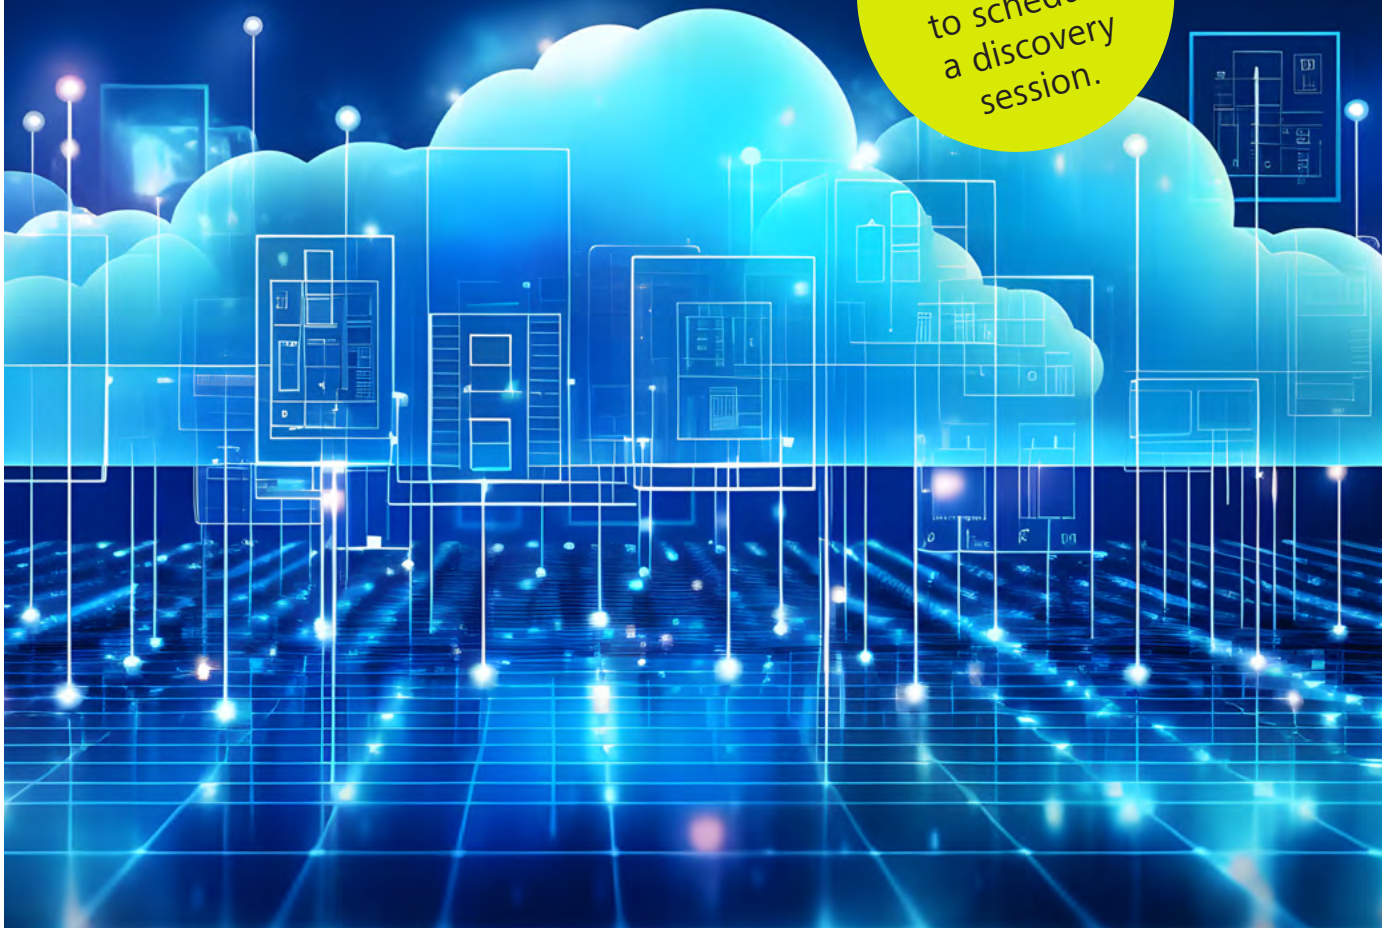

Authors

- **Dr. Andreas T. Bachmeier, Solutions Enablement Team Lead | SEMECO Research Lead, Berlin**
- **Dirk Asmus, Senior Solution Specialist, Berlin**
- **Richard Bieck, Senior Solution Specialist, Dresden**
- **Dr. Max Rockstroh, Senior Business Development Manager, Leipzig**

Contact

Dr. Andreas T. Bachmeier
Andreas.bachmeier@zeiss.com
LinkedIn

Carl Zeiss Digital Innovation GmbH
Fritz-Foerster-Platz 2
01069 Dresden
Germany

Phone: +49 351 49701 – 500
contact.digitalinnovation.de@zeiss.com
zeiss.de/digital-innovation