



Developing Connected Medical Diagnostic and Monitoring Products: An End-to-End Engineering Blueprint

Executive Summary

Industry trends toward the management of chronic diseases, aging, and Post-acute Care (PAC) outside of the hospital are creating powerful tailwinds for connected diagnostics, remote monitoring, and AI augmented decision support. Service delivery systems and payers need longitudinal, high-fidelity data streams rather than episodic point in time measurements. Clinicians are seeking technologies that remove cognitive load from their workflows, not just shoveling more raw data into their already noisy information spaces.



Commercially Shipping Solution Archetypes Already Include

- Connected point of care analyzers with secure cloud connectivity and EMR (Electronic Medical Records) integration
- Wearables and patches for cardiology, sleep, and activity monitoring with multi-day autonomy
- Remote monitoring hubs aggregating data from BLE (Bluetooth Low Energy) peripherals, medical grade peripherals, and telehealth kiosks
- AI assisted diagnostic applications that triage, annotate, and prioritize measurements and images for clinicians

Designing, implementing, and certifying these systems for the market requires skills and competencies across a spectrum of technical domains. These include analog and digital hardware, embedded firmware and RTOS, connectivity stacks and protocols, cloud native data platforms, AI/ML pipelines, cybersecurity, and regulated lifecycle management under standards such as ISO 14971, IEC 60601 1, IEC 62304, and IEC 62366. A comprehensive end-to-end engineering partner helps reduce program risk with reusable reference architectures, pre-validated building blocks, integrated verification environments, and regulatory grade documentation frameworks supporting design controls, risk management, and cybersecurity evidence.

Market and Technology Landscape

Diagnostic platforms are evolving from closed, standalone instruments with native UIs and offline data export toward integrated systems that are software defined and analytics driven. API exposure, time-series cloud data streaming, and Over the Air (OTA) firmware upgrades are becoming table stakes, allowing vendors to continuously improve and remediate issues in the field.

Applications include ambulatory cardiology (Holter monitors and ECG patches) for secure, AI assisted arrhythmia detection, imaging adjacent devices and home-based vital sign monitors, smart auto injectors for chronic therapies, and telehealth kiosks that expand diagnostic reach into retail clinics, workplaces, and patient homes. This requires robustness in more uncontrolled environments and reliable operation by non-expert, everyday users.



Several technological vectors are enabling these shifts

- Miniaturized sensing elements and integrated analog frontends tuned for biopotential, optical, pressure, and biochemical measurements
- Low power System on Chips (SoCs) and wireless radios enabling multi-day wearables with always on sensing and intermittent data connectivity
- Edge AI accelerators built into MCUs/MPUs for on device inference, required for low latency alarms and data reduction
- Secure connectivity stacks and protocols with mutual authentication, certificate-based identities, and encrypted channels aligned with HIPAA and GDPR expectations
- Cloud native architectures that decouple ingestion, storage, analytics, and integration via microservices and standard healthcare APIs like FHIR and DICOM

System Reference Architecture



Sensing and Data Acquisition Layer

Medical equipment that measures physiological parameters requires specialized sensors that combine multiple high-level modalities to accurately capture physiological parameters. The types are:

- ECG, EEG, and EMG using Ag/AgCl electrodes or dry electrodes, demanding high common-mode rejection and low-noise differential inputs
- PPG/SpO2 sensors with multi-wavelength LEDs and photodiodes for oxygen saturation, heart rate, and perfusion indices
- Pressure sensors for blood pressure cuffs, catheters, and infusion lines; as well as load cells and strain gauges
- Contact and non-contact temperature sensors for core and skin measurements
- IMUs (3 axis accelerometers and gyroscopes) for posture/fall detection, tremor analysis, and motion artifact suppression
- Point of Care (PoC) analyzers for blood gases, glucose, electrolytes, and biomarkers, often integrating microfluidics and optical/electrochemical transduction

Analog front-end design plays a critical role in managing the data signals generated by these sensors. It is important to consider some of the design criteria—such as input protection and isolation barriers—to ensure compliance with patient leakage current requirements. Differential amplifier topologies with high CMRR (Common-Mode Rejection Ratio), and programmable gain stages that dynamically adjust to avoid saturation with wide input dynamic ranges are also considered for the design. Noise and distortion can be controlled through guarding and shielding, with optimized PCB stack up and layer allocation. EMC and EMI testing are key for filtering, transient suppression, and compliance (IEC 60601 1 2) that includes conducted and radiated emissions/immunity, electrostatic discharge, and surge withstand.

After sensor section and design consideration, sampling strategy and data acquisition strategy for physiology (bandwidth, redundancy, critical features) aspects need to be focused on while considering design constraints, such as form factor and battery size. Multi-channel and synchronized sampling acquisition is necessary for multi-lead ECG, multi-site PPG, and simultaneous EEG montages. You can also use multiple ADCs (Analog-to-Digital Converter) or multiplexed converters with high precision timing control. Another key architectural decision point is what to do on the device. Whether to execute feature extraction such as QRS detection, HRV metrics, apnea events, or step counts or stream near raw time-series data to the cloud for more advanced model training and inference, at the cost of additional bandwidth and privacy concerns.



Embedded Compute and Device Software

MCU (Microcontroller Unit) and MPU (Microprocessor Unit) selection is challenging as it requires dealing with the constant tradeoff among key capabilities that include compute, memory, power, connectivity integration, and chip security. Wearable solutions and patches and monitoring systems that operate on battery, need ultra-low power MCUs with integrated BLE (Bluetooth Low Energy) radios and trust zones like security partitions. The larger legacy diagnostic systems and therapeutic devices leverage Linux capable MPUs with secure boot and hardware cryptographic accelerators. The selection criteria for an MCU/MPU should include support for DSP instructions to enable for efficient real-time signal processing, FPU/NEON availability for AI inference, sensor interfaces (SPI, I2C, UART, ADC), ADC resolution/frequency, and safety documentation. The safety manual or pre-certified components under IEC 61508/IEC 62304 are required for documentation purposes.

RTOS or embedded Linux are used for device software development in compliance with IEC 62304. The device software architecture is layered into HAL, middleware (protocol stacks, file systems, security libraries), application services (device logic/algorithms/features, alarms), and UI components. As per the regulatory bodies, safety classes (A/B/C) are allocated based on the software's ability to impact patient safety. Software requirements, architecture, and detailed design down to code are defined by the high-level clinical use case. Traceability and justifications are maintained in a Design History File (DHF) that has traces of all the design versions and their justifications. Regulatory bodies depend on this extensive technical document for deciding certification and classes of the device.

Analytics can also be executed at the edge, enabling safety-critical alarms (e.g., asystole detection, lead-off alerts), data reduction through computing features that compress data before uplink, and greater system resilience when cloud connectivity is impaired. Another important design decision is to define fail safe behavior for the device. If sensor operation is impaired, for example with lead-off, loss of confidence in sensor contacts, or RF interference, the system should have a defined mode that ideally reverts to alarm. Or it logs the fault condition for diagnosis and alerts the user in a meaningful way. Watchdog timers, memory protection, and other hardening mechanisms ensure that the device's firmware does not project uncontrolled behavior. When integrating third-party libraries, services, or AI engines on resource-constrained hardware, this becomes important.



Connectivity and Interoperability

Connected medical devices need communication channels that connect with hospitals, labs, private practitioners, ambulatory home care and payer companies. Short range wireless includes BLE for wearables and accessory sensors, Wi-Fi for fixed devices and hubs, and sometimes proprietary sub-GHz links for home range. Long range cellular (4G/5G, LTE M, NB IoT) is used either embedded in the device or via mobile/gateway companions. Clinical QoS (Quality of Service) and latency requirements drive architecture decisions (e.g., near real-time arrhythmia notifications vs. delayed batch upload of trend data).

EMRs and HIS (Hospital Information System) are key data management and storage systems. Hence, for connected devices, it is required that data is mapped to standards to support storage systems. HL7 v2 is very common for legacy integrations, used to feed vitals and lab-like results into hospital information systems, while FHIR APIs are increasingly being used for modern, RESTful integration of observations, diagnostic reports, and device metadata. Imaging adjacent devices leverage DICOM to associate structured data and waveforms with imaging studies, enabling worklist coordination and PACS/VNA archival. Gateway devices often perform protocol translation, local buffering, and security mediation between low-power devices and cloud.

Secure data transport is required under HIPAA and GDPR, including strong encryption in transit and robust authentication. All IP based communication uses TLS with modern cipher suites, and application layer security may also include token or certificate based mutual authentication. Device identities, keys, and certificates must be securely provisioned at manufacturing time and lifecycle managed using secure elements or trusted platform modules, including rotation and revocation in case of compromise or ownership changes.



Cloud and Data Platform

Cloud native architectures for diagnostic and monitoring systems use streaming ingestion pipelines that handle high-volume time-series data, device telemetry, and event logs. Message brokers, streaming services, or MQTT bridges are used to terminate device connections and forward normalized messages to microservices responsible for parsing, validation, and routing. Time-series databases handle the high-frequency physiological data, while object stores, blob stores, and relational databases store waveforms, images, metadata, and configuration artifacts with tiered storage policies for short-term operational access and longterm archival with lifecycle rules aligned with clinical and regulatory retention needs.

Microservices-based development helps manage independent functionalities such as patient enrollment, consent and data sharing preferences, device and patient fleet management, firmware update sharing, alert rules engines, and clinician notification and engagement workflows. These services expose APIs for integration with EMRs, care management platforms, and third-party analytics applications, while role-based access control ensures least privileged access to PHI.(Protected Health Information) The observability stack of logs, metrics, and traces is required for programmatic operational monitoring, cybersecurity incident detection, and audit trail generation for HIPAA's security rule.

AI and Analytics Enabled Diagnostics



AI Use Cases and Patterns

For connected diagnostics, time-series signal analytics works as a fundamental AI workload. Automated arrhythmia detection from multi-lead ECGs, ischemia detection, QT-interval monitoring, and heart rate variability analysis for assessing autonomic function are some key cardiovascular use cases. Sleep monitoring systems analyze multi-channel ECG/PPG, respiratory belts, SpO2, and accelerometry to derive sleep staging, apnea detection, and periodic limb movement features. Respiratory diagnostics cover detection of hypopneas, apneas, and abnormal breathing patterns. Hemodynamic trend analysis correlates blood pressure, SpO2, and cardiac output surrogates to flag early deterioration.

Diagnostic adjacent analytics often perform triage, quality control, and workflow optimization, rather than full end-to-end interpretation. Portable imaging tasks include detecting incomplete views or motion artifacts and flagging for repeat acquisition. This automatically identifies regions of interest in ultrasound or radiographs to speed radiologist review, and protocol optimization to select optimal scan parameters based on patient settings and previous studies. For POC analyzers, algorithms can triage samples or cartridges to detect anomalies and recommend repeat measurements or alternative workflows.

At the population level, analytics distill digital biomarkers and risk scores from the longitudinal data to enable patient segmentation for proactive outreach, therapy optimization, and capitated risk-based care management. Wearable adherence monitoring uses behavioral patterns inferred from device usage logs, while clustering and survival models can identify high-risk patient subgroups with differentiated risks of hospitalization or mortality. Risk scores and event probabilities are reported to care management platforms or payer systems to drive value-based care interventions.



AI/ML Lifecycle in Regulated Environments

In regulated environments, data curation workstreams are as important as model development and test strategies. Pipelines for curated datasets need to handle acquisition of labels from clinical experts, adjudication processes for ambiguous or disputed cases, as well as capture of metadata describing devices, firmware versions, acquisition settings, and more. Addressing class imbalance (e.g., rare arrhythmias), multi-site variability (different hospitals, sensors, and care practices) and device drift over time and deployment requires stratified sampling strategies for training and test sets, domain adaptation and fine-tuning approaches, and regular dataset refresh cycles.

Model development and tuning is focused on clinically meaningful performance metrics and plots beyond raw accuracy. Sensitivity, specificity, Positive Predictive Value (PPV), and Negative Predictive Value (NPV) along with calibration curves, decision likelihood ratios, and net reclassification statistics are important to select meaningful clinical decision thresholds. Cross validation must be done using patient-level and site-level splits to avoid leakage. External validation on independent cohorts is often necessary for regulatory submissions. For continuous monitoring systems, per event-based performance (such as false alarms per 24 hours, time to detect) may be more relevant than sample level accuracy.

MLOps and version control for medical devices manage datasets, models, model serving configurations as regulated configuration items. Versioning systems allow tracking relationships between raw data, curated datasets, feature extraction code and parameters, training job configurations, model binaries, and mapping between software modules or firmware versions and deployed model. This can be built on open-source systems like DVC, modelDB, GitOps for ML with added security controls. But in any case, it enables exact reconstruction of any released model or derivation tree from primitives. Rollback strategies should allow rapid reversion to prior models if field performance is found to degrade, with either A/B testing or phased rollouts to mitigate risk. Pipeline monitoring to track model drift, case mix changes, and safety signals (missed critical events, unexpected alarm behavior) feeds back to a post market surveillance process required under ISO 14971 and regulator expectations.



Responsible and Trustworthy AI

Bias and fairness must be assessed systematically in responsible AI products across demographics, comorbidities, and care settings, particularly for cardiovascular and sleep related measures where there are important differences in disease prevalence and presentation in different populations. Analyses stratify performance by age, sex, ethnicity, and by device or site variations, to detect systematic over or under performance that could lead to different levels of care. Remediation strategies include balanced sampling, reweighting, and model calibration adjustments, along with clearly labeling the intended population and limitations in the instructions for use.

Explainability mechanisms are also key to clinician trust. For waveform analytics, this can take the form of explainable features used in the model (e.g., intervals, amplitudes, variability metrics) in addition to model outputs, as well as saliency methods or other data viz techniques that highlight contributing segments of the waveform. For imaging adjacent tasks, highlighting the regions of interest using heatmaps or bounding boxes is common, while confidence scores and guardrails (such as no output zones when input quality is low) can reduce overreliance on uncertain predictions. Regulatory bodies are also starting to issue guidance on adaptive AI in medical devices and expected processes to support this, such as predetermined change control plans that define how a model can evolve in post-market conditions under controlled processes. For example, CLSI GL34 and EMA MDSUB also address post market changes.

Safety, Risk Management, and Cybersecurity by Design



Risk Management Under ISO 14971

ISO 14971 specifies a process for applying risk management throughout the lifecycle of medical devices. Hazard identification includes hardware risks (electrical, mechanical, thermal), software risks (logic faults, user interface errors), usability risks (misuse, confusion), data risks (loss, corruption, latency) and any dependencies on external systems (mobile apps, cloud services, dependencies on third-party hardware or software) or environments (daily use, clinic vs home, ECG leads). Tools include Preliminary Hazard Analysis (PHA), task analysis, and review of similar devices or databases of prior incidents to capture foreseeable misuse and environmental stresses.

Quantitative and qualitative approaches to risk analysis include FMEA/FMECA and Fault Tree Analysis (FTA). FMEA examines failure modes of components, subsystems, or processes, listing causes and effects and current detection controls, and is well suited for electronics, embedded software modules, manufacturing processes and changes to process flow. FTA is a deductive, top-down approach where a top-level hazard (failure to alarm life threatening arrhythmia) is modeled as a logical combination of lower-level failures. This enables exploration of the effects of complex interactions and redundancy architectures. Outputs from these activities are fed into risk control decisions and residual risk evaluations, documented in a risk management file that is expected to align with regulatory expectations.

Risk controls in the design phase may include hardware redundancy for critical measurements, self-tests for sensors and memories, diagnostic coverage via watchdogs, plausibility checks and alerts, software interlocks, and alarm management strategies. Safe states and graceful degradation modes are defined for loss of connectivity, sensor failure, battery depletion, or cloud outages, ensuring that the patient and clinicians are clearly informed when the device is in a limited state.



Functional and Patient Safety

Functional safety and patient safety are directly linked to standards such as IEC 60601 1 for basic safety and essential performance, IEC 60601 1 2 for EMC, and IEC 60601 1 8 for alarm systems. Special standards (IEC 60601 2 xx) for different device categories (ECG, infusion pumps, ventilators) guide test plans and safety architectures for basic safety and EMC, in alignment with the general requirements. Usability engineering under IEC 62366 requires formal human factors to ensure that user interfaces for both clinicians and patients support safe and effective operation in clinical and home environments.

Alarm safety is a major focus, given extensive research on alarm fatigue and its impact on both patient safety and user engagement. Alarm design requires consideration of alarm prioritization (life threatening vs advisory), suppression of non-actionable alarms, context-aware alarm thresholds, and multimodal presentation (audible, visual, tactile) that is discriminable and identifiable amongst other devices and in noisier clinical environments. User studies in realistic use environments assess response times, error rates, and cognitive load to refine alarm logic and UI representation.

Verification of measurement accuracy is achieved through traceability to reference instruments and standards supported by calibration strategies at different points across the device lifecycle. Environmental robustness testing extends performance across temperature, humidity, shock, vibration, and electromagnetic disturbances, as well as aging and wear that is relevant to the long-term home use lifecycle or expected wash and re-use of accessories.



Cybersecurity for Connected Diagnostics

Threat modeling frameworks such as STRIDE or attack trees are applied to devices, mobile applications, gateways, and cloud levels to identify and prioritize cybersecurity risks. Assets include PHI, cryptographic keys, firmware images, and control commands, while threats are eavesdropping, tampering, injection, denial of service, ransomware or other activity that could impact healthcare environments.

Security controls for endpoint devices include secure boot (verify signed firmware images), secure firmware updates with encryption and signature checks, and secure storage of PHI and keys using a hardware security module or secure enclave where available. Implementing least privilege, defense-in-depth, and secure coding practices minimize the attack surface, while security event logging enables the detection of anomalous behaviors. OTA mechanisms also have guidelines for verifying integrity before installing, rollback capabilities if installs fail or produce unexpected behavior, and monitoring for failed or suspicious update patterns.

Cloud and integration architectures must align with HIPAA and GDPR requirements around encryption in transit and at rest, robust authentication and authorization, detailed logging and audit trails, and documented incident response workflows. Audit logs must capture user access events, data modifications, and administrative actions with sufficient detail to support forensic analysis and regulatory audits. Data minimization, where possible and subject to limitations on clinical data, pseudonymization, as well as documented handling of consent and data subject rights are key to privacy compliance.

Regulatory Strategy and Quality Engineering



Regulatory Pathways and Classifications

A regulatory strategy for diagnostic and monitoring products should start with accurate product classification, a clear definition of intended use, and thorough documentation of the state of technology and design inputs. In the U.S. regulatory context, most devices are Class II devices following the FDA 510(k) premarket notification pathway, based on demonstrating substantial equivalence to predicate devices. Novel technologies or AI heavy products, without appropriate predicates, may need De Novo classification as novel and without a predicate. EU MDR classification rules consider invasiveness, active monitoring, and diagnostic impact and are often Class IIa/IIb or higher for systems that measure or monitor vital physiological parameters, or that inform therapy decisions.

Software as a Medical Device (SaMD) vs Software in a Medical Device (SiMD) is a distinction that influences some aspects of architecture and documentation. SaMD components—such as a cloud-based AI triage service—must be treated as a standalone regulated product with its own regulated configuration items, configuration management and versioning, cybersecurity controls and post market surveillance, even if it is part of a broader ecosystem. Evidence strategies combine bench testing of hardware and software verification and validation, along with clinical performance studies or evaluations in many cases to demonstrate safety and effectiveness in the intended population and use cases.



Design Controls and Documentation

A structured approach to managing requirements, design, verification, validation, and changes are made mandatory in the design control requirements. Details such as how the product has evolved and what design changes need to be documented in the Design History File (DHF). A Device Master Record (DMR) defines the as-built configuration, and the technical documentation dossier is the assembly of evidence for the regulatory review process. Requirements engineering covers user needs, system requirements, and user specifications. System requirements and software requirements with unique identifiers are bidirectionally traceable to design elements, test cases, and risk controls.

For any changes, including hardware revisions, software updates, and AI model changes, the change management process needs to be established. Configuration management tools are needed to maintain baselines for released software/hardware versions. It also ensures that the deployed devices, documentation, and manufacturing processes remain in sync with the configuration management system. This is critical where CI/CD pipelines interface with the regulated constraints on changes to deployed systems.



Verification, Validation, and Test Automation

Verification and Validation (V&V) spans unit tests and static analysis for embedded code, hardware in the loop setups for protocol and driver logic, and system tests that bring in sensors, UI, connectivity, and cloud components for end-to-end behaviors. Interoperability testing with EMRs, PACS, and hospital networks covers compliance with HL7, FHIR, DICOM profiles, as well as robustness to real-world integration patterns and interactions with installed EMRs or hospital networks.

Performance, reliability, and environmental testing include accelerated life tests, drop and shock tests, Ingress Protection (IP) rating for home and wearable devices, and comprehensive EMC testing against all collateral and standards as applicable. Statistical reliability methods are applied to both field data and simulated stress tests to estimate failure rates and maintenance intervals. Automated test frameworks for embedded firmware, mobile apps, and cloud services can be built on open-source systems with the appropriate security controls; these with CI/CD pipelines to orchestrate builds, regression suites, and static and dynamic security scans help maintain quality while managing change in regulated environments.

Human Centric Design for Clinical and Home Environments



Clinical Workflow Integration

Clinical adoption and usability depend on matching device behavior to real-world workflows in hospitals, clinics, EMS, and telehealth use settings. Workflow analysis and contextual inquiry to identify when and how measurements are taken, who interacts with the device and where, and what information is actionable at each step, is the basis for engineering UI states, default configurations, and endpoint integration.

Role-based UX designs account for clinicians, technicians, patients, and caregivers with differing cognitive load, training levels, and access rights. Reducing cognitive load often requires good visualization design, where we present trends, risk scores, and alerts prioritized by severity rather than overwhelming users with unfiltered waveforms and numeric streams. Interactive dashboards and device UIs should support quick recognition of deterioration and relative status across multiple timepoints, as well as drill down to raw data, when necessary, while minimizing alarm proliferation and information clutter. Careful integration with existing systems (EMR, pagers, secure messaging) also ensures that insights show up where clinicians already work.



Home and Wearable Use Considerations

Home care and wearable devices need to be easily accessible to people with no or limited clinical training, variable dexterity, sensory impairments, and differing levels of digital literacy. For better industrial design, ergonomic aspects need to be considered along with weight distribution, skin friendly materials, cable management, and ingress protection suitable for daily life. Battery life should support the intended wear time of the device without being so short that it limits usability, while also enabling a longer overall device lifecycle. Clear status indicators and limited physical controls also reduce misuse and dependence on manuals for operation.

Setup and onboarding flows can use guided instructions, animations, or a companion app to walk the user through steps for initial setup, such as electrode placement and positioning. Edge automation can detect improper sensor contact, motion artifacts, or other missing peripherals and coach the user toward corrective action before measurements are accepted or captured. Architectures must also tolerate intermittent connectivity—buffering and delayed upload, delayed or skipped sync in error cases—and have conflict resilient synchronization to ensure clinical data is consistent and not lost even with network outages, roaming between locations, or concurrent use of multiple devices.

About the Author

Purva Shah - Product Marketing Manager at eInfochips (an Arrow Company)

Purva Shah is a product marketing manager for the medical device practice of eInfochips, an Arrow Electronics company that provides end-to-end product engineering services for numerous industries, including performance reliability, embedded systems, AI/ML integration and regulatory compliance. With a foundation in both engineering and marketing, her expertise helps translate complex technologies and industry demand into successful healthcare solutions.

FOLLOW US



/einfochips



/einfochipsltd



/einfochips



/einfochipsindia