MEDICAL

Designing Medical Devices for 10+ Year Lifecycles

A Practical Framework for Managing Component Obsolescence

Executive Summary

Medical device design is advancing quickly, shaped by progress in embedded computing, display technology and system integration. Yet medical platforms are held to a different standard to most other electronics applications: they are expected to deliver consistent, dependable performance over lifecycles that frequently exceed ten years.

This creates a genuine engineering tension. Design teams want to integrate modern, high-performance technology, while also guaranteeing long-term availability, stability and regulatory compliance. Component obsolescence sits at the center of that tension. When a processor, display or power supply reaches end of life, the consequences rarely stop at sourcing a replacement part. Engineers are often drawn into system-level redesign, software rework and regulatory revalidation, each of which introduces cost, delay and risk into a programme that may already be years into deployment.

This white paper sets out why obsolescence should be treated as a design consideration from the outset, rather than a supply chain problem to be managed later. It examines where lifecycle risk actually originates, what it costs when it is left unmanaged, and how a structured, engineering-led approach to component selection, system architecture and integration can materially reduce that risk.

The objective is not to eliminate change. Technology will continue to move forward, and medical devices must be able to draw on it. The objective is to design systems that can absorb component-level change without triggering disproportionate redesigning, revalidation or delay. Done well, this approach protects the programme's timeline and budget, and it protects the long-term reliability that patients, clinicians and regulators depend on.

This paper is written for the engineers, R&D leads and procurement managers who carry that responsibility: the people specifying, architecting and validating platforms that will still need to perform reliably a decade or more after they leave the design phase.

The Lifecycle Challenge in Medical Device Design

Medical device development operates under constraints that most embedded and industrial applications do not share. Technology continues to evolve at pace, but medical systems are

expected to remain stable, supported and compliant over deployment periods that commonly span seven to fifteen years, and sometimes longer. This creates a structural mismatch between the lifecycle of the components a device is built from and the lifecycle the finished system is expected to deliver.

Diverging lifecycles

Many of the core technologies used in medical devices, including embedded processors, TFT displays and power architectures, are developed primarily for commercial and consumer markets. Those markets move quickly, and components can reach end of life within three to five years of launch. A medical device specified against consumer-grade or commercially-driven components is, in effect, built on a foundation with a shorter working life than the product itself. Bridging that gap is one of the central challenges of long-lifecycle design.

Regulatory and validation constraints

Medical devices cannot be modified in the field with the freedom afforded to many other electronic products. Even a minor component substitution can trigger design verification and validation updates, a risk assessment in line with ISO 14971, and revisions to the technical file. Depending on the scale of the change, re-certification may follow. A component change is therefore rarely just an engineering task. It becomes a compliance and quality exercise, typically involving engineering, regulatory affairs and manufacturing teams working in parallel.

System-level dependency

Modern medical devices are highly integrated systems, with hardware, software, mechanical design and the user interface all interdependent. Replacing a single obsolete component can have effects that extend well beyond the part itself. A new processor may demand changes to firmware and power architecture. A substitute display can affect optical performance, touch integration or enclosure design. A power supply change may shift thermal behavior and require re-certification in its own right. These dependencies mean the true cost of an obsolescence event is rarely confined to the component being replaced.

Left unaddressed, these pressures tend to surface at the worst possible moment: mid-production, mid-deployment, or during a design freeze ahead of a regulatory submission. Unplanned redesign cycles, extended timelines, additional compliance cost and disrupted supply are the practical result, and in critical care or high-volume applications, they can also affect product availability for clinicians who are already relying on the device in the field.

Where Obsolescence Risk Really Originates

Obsolescence is often treated as an unavoidable supply chain event. In practice, the root causes are usually set much earlier, in decisions made during initial component selection and

system design. Understanding where that risk originates is what allows engineers to design it out, rather than manage it after the fact.

Commercial technology cycles

Processors, graphics chipsets and interface ICs used in medical platforms are frequently derived from high-volume commercial markets, where the priority is rapid iteration and performance improvement rather than long-term availability. Product generations are refreshed often, legacy parts are phased out to streamline manufacturing, and support windows are set by commercial logic that has little regard for a fifteen-year medical deployment. Specifying against this kind of roadmap, without adjustment, embeds risk into the design from day one.

Single-source dependency

A significant proportion of obsolescence risk comes from reliance on single-source components with no viable alternative. This is particularly common with specialised displays that have unique optical or mechanical characteristics, embedded platforms tied to a specific chipset or form factor, and power solutions built to a custom specification. When a component like this reaches end of life, there is no drop-in substitute to fall back on. The result is a full redesign rather than a controlled, planned transition.

Interface and integration constraints

Even where an alternative component technically exists, integration constraints can prevent a straightforward swap. Fixed interface standards such as LVDS or eDP, software written against a specific hardware architecture, mechanical tolerances built around one component's footprint, and thermal or power characteristics unique to the original part can all stand in the way of substitution. In these cases, replacing one component quietly becomes a broader system change.

Display and HMI variability

Displays carry a disproportionate share of obsolescence risk. Panel roadmaps vary between manufacturers, and even a like-for-like replacement panel can differ in brightness, colour consistency or optical stack. Touchscreen technology and controller compatibility add a further layer of variability. In applications where visual accuracy and reliable user interaction are safety-relevant, even small differences of this kind can trigger a redesign and revalidation cycle that a design team did not budget for.

Supply chain visibility

Finally, limited visibility into supplier roadmaps leaves many design teams exposed to end-of-life announcements they did not see coming. Without clear lifecycle data, proactive notification from suppliers, and a structured last-time-buy process, engineering teams are pushed into reactive decisions, usually under time and cost pressure that makes those decisions harder to

get right. Fragmented sourcing compounds the problem further: when compute, display and power are bought independently, with no single view of how their lifecycles align, responsibility for spotting an approaching end-of-life event tends to sit with no one in particular, and it is often noticed only once a supplier has already announced it.

Taken together, these five factors show that obsolescence is not random. It is the predictable outcome of component selection, system architecture and integration choices made early in a programme. Recognising this is what makes lifecycle risk manageable rather than merely unfortunate.

The Cost of Getting It Wrong

When obsolescence is left unmanaged, its impact is rarely confined to the cost of the component itself. In medical device design, the consequences are systemic, touching engineering resource, regulatory compliance and overall programme delivery.

An obsolete component rarely permits a simple substitution. Because of the system-level dependencies described above, even a modest change can require schematic and PCB layout rework, firmware modification, mechanical adjustment and thermal re-evaluation. This work consumes engineering capacity that would otherwise go towards new product development, and on complex systems, the cost of redesign can escalate quickly once more than one subsystem is affected.

The regulatory burden compounds this. A component change may trigger design verification and validation updates, a fresh risk management review under ISO 14971, and revisions to technical documentation, even where functional equivalence has been maintained. Evidence still has to be generated to demonstrate that compliance continues to hold, and that work draws in teams well beyond engineering.

Timelines suffer in ways that are rarely linear. A single obsolescence event can trigger delays while a replacement component is identified and qualified, extended testing and validation phases, and rework that ripples across hardware and software teams simultaneously. For products already in production, the picture is different again: last-time-buy decisions taken under deadline pressure, higher procurement costs for parts nearing end of life, and disruption to manufacturing and service continuity for devices already in the field.

Viewed only at component level, these costs are easy to underestimate. Once engineering effort, compliance work, delayed revenue and supply chain overhead are accounted for, the total cost of ownership of an unmanaged obsolescence event is substantially higher than the price of the part it started with. That gap between the visible cost of a component and the true cost of replacing it is precisely why lifecycle risk is worth quantifying at the outset of a programme, rather than discovering it in the middle of one. It is the case for addressing lifecycle risk early, when mitigation is a design decision rather than a recovery exercise.

A Framework for Designing for Longevity

Designing a medical device for a ten-year-plus lifecycle is not achieved through component choice in isolation. It is the outcome of coordinated decisions across three areas: component selection, system architecture and integration strategy. Each is addressed below.

1. Component selection

Component selection is the first, and arguably the most consequential, point at which lifecycle risk is either built in or designed out.

Lifecycle availability should be treated as a specification requirement in its own right, not an afterthought. Components with published long-term availability, typically five to fifteen years or more, backed by a clear manufacturer roadmap, give a design team a genuine planning horizon rather than a guess.

Industrial and medical-grade components are generally engineered for extended operation, with more stable specifications and more controlled revision cycles than their consumer-grade equivalents. The specification premium is usually modest set against the redesign cost it avoids.

Supplier transparency matters as much as the component itself. Working with suppliers who provide honest lifecycle data, proactive end-of-life notification and structured last-time-buy planning gives engineering teams the lead time to plan a transition, rather than react to one.

Where possible, second-source options should be identified at the point of selection, not searched for in a crisis. Components with compatible electrical, mechanical and software characteristics turn an obsolescence event into a controlled substitution.

Finally, performance margin deserves more attention than it typically receives. Specifying to the absolute limit of a component's rated performance leaves no headroom to accommodate its eventual replacement. Deliberate margin in thermal, electrical and optical performance keeps future substitution options open.

Approached this way, component selection becomes a risk management decision as much as a specification exercise.

2. System architecture

A well-defined architecture is what allows a system to absorb component-level change without a ground-up redesign.

Modular design, with compute, display and power treated as separable subsystems, allows one to be updated without disturbing the others. Standardised interfaces, such as eDP, HDMI, USB and COM Express, simplify future component replacement because the surrounding system does not need to be redesigned around a proprietary connection. Scalable platforms,

chosen with headroom for performance upgrades, extend a design's working life without requiring it to be re-engineered from scratch.

Thermal and power headroom, built in from the start, gives a system the flexibility to accommodate a future component with a different power or thermal profile. Mechanical flexibility, in mounting points, enclosure tolerances and connector placement, does the same for the physical design. None of this needs to be over-engineered; it needs to be deliberate.

The cumulative effect of these choices is a system that depends less on any single component, and one that can evolve under control rather than under pressure.

3. Integration strategy

Integration is where lifecycle risk is either mitigated or amplified, and it is frequently the point at which a well-selected, well-architected design still runs into difficulty.

Treating compute, display and power as an integrated subsystem, rather than as components sourced and assembled independently, ensures compatibility is designed in from the start rather than discovered during testing. This is particularly true of the display assembly: engineering optical bonding, touch technology and interface electronics together as a stable, long-term subsystem materially reduces the variability and revalidation risk described earlier in this paper.

Electrical and signal interfaces should be designed with tolerance for variation, so that a future component change does not immediately fall outside specification. Validation processes should be aligned across subsystems, so that a change in one area does not force a disproportionate amount of rework elsewhere. And component and subsystem lifecycles should, wherever practical, be aligned with one another, so that a device does not end up with a display, processor and power supply reaching end of life on three different timelines.

A cohesive integration strategy reduces fragmentation, simplifies future updates and materially improves a system's overall resilience.

Designing Beyond the Component Level

Longevity in medical device design is not the product of a single well-chosen part. It is the result of consistent, engineering-led decisions carried across selection, architecture and integration, made early enough that they shape the design rather than react to it. Applied together, these decisions reduce exposure to obsolescence risk, limit the redesign and revalidation burden when change does occur, and improve system stability over an extended operational life.

Designing for longevity is, ultimately, a question of engineering foresight: making decisions today that protect performance, compliance and availability years from now.

Talk to FORTEC UK

FORTEC UK works with medical device design engineers and OEM teams across the specification, integration and long-term support of TFT displays, embedded computing and power supplies, drawing on more than 25 years of experience in displays and embedded computing and 30 years in magnetics manufacturing.

If you are approaching the design of your next medical device, or reviewing the lifecycle exposure of an existing platform, we would welcome the opportunity to discuss your specification. Contact FORTEC UK to explore the right solution for your application.