

Scaling digital pathology in the petabyte era



Digital pathology is now embedded in routine practice across many UK centres. Whole slide imaging is no longer exceptional, and petabyte-scale datasets are increasingly common as organisations expand their digital workflows. As data volumes grow, managing how these images are stored, accessed and retained is emerging as one of the most significant operational challenges in digital pathology.

GE HealthCare hosted a **Digital Pathology Roundtable** in late 2025 to explore how UK pathology organisations are addressing this issue. The event brought together clinical pathologists, service managers and digital and IT specialists from NHS Trusts, Health Boards, pathology networks and private providers. The event created space for open discussion about the barriers, opportunities and future direction of digital pathology in the UK, and reinforced a growing recognition across the pathology community: storage is no longer simply a technical consideration, but a strategic issue that will shape the pace and sustainability of digital pathology adoption.

The pace of adoption of digital pathology continues to accelerate and laboratories are increasingly generating datasets at the petabyte scale. But how should these images be managed? Insights from a UK roundtable highlight the challenges of image storage, and discuss strategic approaches that will be essential for success.

The retention question

One of the most immediate questions facing organisations adopting digital pathology is how long whole slide images should be retained. As scanning volumes increase, storage decisions quickly become operational and financial concerns, but there is still little consensus on appropriate retention periods.

In practice, policies vary widely across organisations. Some centres retain digital slides for only a few weeks, while others store them for many years. Glass slides and paraffin blocks are preserved for decades, but the potential value of long-term digital storage is still being hotly debated. At the same time, the legal status of digital slides as evidential material has yet to be definitively tested in UK courts. In the absence of clear precedent, many organisations are adopting cautious policies to mitigate potential medico-legal risk and this has significant consequences. Storing images for extended periods rapidly increases data volumes, placing pressure on both infrastructure and budgets.



The cost of scale

As datasets expand, the financial implications of storage become increasingly visible. Many departments now face monthly costs running into the thousands, with annual expenditure exceeding £100,000 in some cases. On-premise servers, while offering a sense of control, are expensive to maintain and rapidly become obsolete. Furthermore, limited storage capacity can restrict slide scanning, slowing digitisation programmes just as adoption begins to accelerate.

In contrast, cloud-based models provide the scalability that expanding datasets need. Elastic storage reduces reliance on local hardware and supports growth without the need for repeated capital investment. However, scalability introduces its own questions around cost predictability, vendor dependency and data sovereignty. A coordinated NHS-scale cloud solution could mitigate these concerns by combining elasticity with stronger governance and greater purchasing power.



Building resilient infrastructure

Cost and capacity are only the start of the storage challenge though. As digital pathology becomes embedded in routine workflows, reliability and resilience become equally important, as system failures carry immediate clinical consequences. A generation of pathologists trained primarily on digital platforms may have limited recent experience reverting to glass-based workflows, raising the stakes of downtime.

Resilience therefore requires more than storage capacity alone. Mirrored, multi-site back-ups, robust disaster recovery planning, and sustained investment in network performance are all essential. Whole slide images are data intensive, and insufficient bandwidth or high latency can erode workflow efficiency and clinician confidence. Storage strategy must be considered alongside connectivity and business continuity to ensure a reliable service.



The need for national direction

Underlying many of these challenges is the absence of firm national guidance. Existing recommendations on retention – often cited as eight years or two UK Accreditation Service conformity assessment (UCA) cycles – are interpreted inconsistently and lack a clearly articulated clinical rationale. As a result, there's huge variation between trusts reflecting different perceptions of risk rather than differences in patient care requirements.

Without a nationally agreed, clinically grounded and legally robust framework, organisations are left to determine retention policies independently. This fragmentation complicates financial planning, increases duplication and creates uneven medico-legal exposure. Clear, binding standards would not only harmonise practice but also provide a stable foundation for long-term infrastructure decisions.



From bolt-on to built-in

These challenges are compounded when storage is treated as an afterthought in procurement. Too often, scanners and image management systems are purchased first, with storage solutions assembled later. This piecemeal approach leads to fragmented architectures and limits opportunities for standardisation.

Embedding storage from the outset results in integrated, interoperable solutions that reduce complexity and support flexibility. Scanner-agnostic systems designed with AI-readiness in mind can accommodate future innovation without costly retrofitting. Aligning procurement cycles with long-term storage strategy also improves accountability and cost control.





Smarter storage strategies

Storage decisions now sit within a broader context of environmental responsibility. Data centres are energy intensive, and large-scale digital archives carry a measurable carbon footprint. Procurement frameworks increasingly require suppliers to demonstrate emissions reductions and deliver social value, placing sustainability alongside performance and price in decision-making criteria. At the same time, innovation offers opportunities to reduce this burden; intelligent compression at the point of image capture lowers file sizes while maintaining diagnostic integrity, provided legal defensibility is preserved.

Tiered storage models – with frequently accessed cases held in high-performance environments and older material moved to lower-cost archival layers – align cost more closely with clinical utility.



Looking ahead

Discussions at the roundtable hosted by GE HealthCare highlighted the value of bringing together clinicians, service leaders and technology specialists to share experiences and explore practical solutions to common challenges. As digital pathology continues to grow, this kind of collaboration will be essential to shaping approaches that are both sustainable and resilient.

Ultimately, digital pathology has transformed how images are generated and reviewed. Its continued evolution depends on strategic thinking about how those images are stored. Nationally-defined standards, scalable infrastructure, built-in resilience, integrated procurement and sustainability by design will determine whether storage becomes a limiting factor or an enabler of the next phase of digital pathology.

The perspectives presented in this article reflect insights shared during the roundtable discussion and represent current experiences and expert opinion from participants, rather than formal positions of any individual organisation.