




INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

DICOM Whole Slide Imaging Integration: Scanner to AI Guide

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-10-20 · dicom whole slide imaging integration · dicom wsi · whole slide imaging · digital pathology pacs integration · dicomweb · ihe dpia · vendor neutral archive digital pathology · digital pathology ai integration

Original article: <https://intuitionlabs.ai/articles/dicom-whole-slide-imaging-integration-architecture>



In this article

1. Executive Summary
 2. Introduction and Background
 3. Key Changes in the DICOM WSI Standard and Its Ecosystem
 4. Native, Dual-Personality or Proprietary: Choosing the Path
 5. Implementation Considerations and Process Changes
 6. Data Analysis and Evidence
 7. Implications and Future Directions
 8. Frequently Asked Questions (FAQs)
 9. Conclusion
-

Executive Summary

DICOM (Digital Imaging and Communications in Medicine) whole slide imaging (WSI) integration is a systems-engineering problem more than a file-format choice. The standard stores each pyramid level of a scanned slide as a tiled, multi-frame image object, with tiles as frames (dicom.nema.org). It carries specimen and container identity inside the image object (dicom.nema.org). Retrieval runs over **DICOMweb**, and annotations and AI outputs travel as separate objects. The practical question for pathology IT teams is therefore how to get slide identity from the laboratory information system (LIS) to the scanner, and how to get tiles, annotations and model results back to viewers, without losing either.

Three paths exist. A **native DICOM** path has the scanner emit WSI instances directly; vendors document this as an option or a default on some current devices (leicabiosystems.com) (www.usa.philips.com). A **dual-personality or conversion** path rewrites proprietary files, and Microsoft's Azure documentation states plainly that a format conversion is required for WSI in other formats (learn.microsoft.com). A **proprietary** path keeps vendor formats and viewers, and the literature observes that manufacturers, developers and users continue to depend on such formats (pmc.ncbi.nlm.nih.gov). Most enterprises will run a mixture during migration, so DICOM should be the system of record and conversion a testable stage.

The identity chain is the least forgiving part. The IHE (Integrating the Healthcare Enterprise) Digital Pathology Image Acquisition (**DPIA**) supplement, published for trial implementation on March 6, 2026, uses three HL7 v2 transactions between an **Acquisition Manager** and the scanner (www.ihe.net). It assumes the slide barcode uniquely identifies the physical glass slide (www.ihe.net). As of October 2026 DPIA is not final text, so procurement should ask for conformance statements and test evidence rather than assume the profile is universally implemented.

On scale, NEMA's worked example is an 80,000 x 60,000 pixel image of about **15 GB** uncompressed (dicom.nema.org), and published institutional figures run from a projected **4.3 petabytes per year** for a multisite network (pmc.ncbi.nlm.nih.gov) to **9 PB** allocated at one cancer center (pmc.ncbi.nlm.nih.gov). This report supplies a capacity worksheet that uses slides per day, source size, conversion multiplier and retention copies, plus a concurrent-tile model, and it insists on local measurement rather than borrowed benchmarks. It also provides an identifier lineage table, a platform support matrix, a failure and retry list, and an eight-step pilot checklist. Annotations at scale belong in **Microscopy Bulk Simple Annotations** rather than **Structured**

Reports, according to a 2026 preprint from the Imaging Data Commons team (arxiv.org). Vendor conformance varies by scanner and release, so every claim here should be re-checked against the specific conformance statement in your RFP.

Introduction and Background

This report examines how to connect scanners, laboratory information systems, picture archiving and communication systems (PACS), vendor neutral archives (VNA), viewers and machine-learning pipelines around DICOM whole slide images. It is written for digital pathology program owners, pathology IT architects, imaging core laboratories, medical-device product teams and AI-platform engineers. It is not a scanner ranking, a PACS comparison or a de-identification guide. The guidance is vendor-neutral: identity, retrieval and reconciliation testing deserve more design time than file-format debates.

History matters only as lineage. DICOM Working Group 26 was established in 2005 to extend the standard to pathology (digitalpathologyassociation.org). **Supplement 145** introduced the whole slide IOD (information object definition), and a later peer-reviewed account notes that adoption lagged despite its publication in 2010 (pmc.ncbi.nlm.nih.gov). Supplement 145 is now folded into the maintained standard, so implementers should read the current edition of PS3.3, PS3.5 and PS3.18 rather than the historical supplement. The NEMA implementation page was last updated on January 27, 2026, and the standard text cited below is current as of October 2026.

The rationale for DICOM is infrastructure reuse. NEMA states that PACS capabilities for storing, archiving, retrieving, searching and managing images can be leveraged (dicom.nema.org). A multi-vendor implementation study concluded that DICOM lets healthcare enterprises leverage a wealth of existing medical imaging infrastructure (pmc.ncbi.nlm.nih.gov). A companion review argued that pathology laboratories going fully digital need to include WSI in an enterprise-level VNA, and that digital pathology storage needs dwarf those of radiology (pmc.ncbi.nlm.nih.gov).

Interoperability is also being tested in public. A 2025 Journal of Pathology Informatics report describes a **Connectathon** with over 30 digital pathology vendors and healthcare providers from around the world (pmc.ncbi.nlm.nih.gov). Existing explainers mostly answer PACS, viewer or vendor questions, so this report concentrates on the integration path.

Key Changes in the DICOM WSI Standard and Its Ecosystem

The phrase "key changes" here means the structural shifts a team adopts when moving from proprietary slide files to DICOM: a tiled object model, identity inside the object, web retrieval, separate annotation objects and public conformance testing.

Change 1: A tiled multi-frame object replaces vendor tile stores

The current PS3.3 defines the **VL Whole Slide Microscopy Image IOD** as a multi-frame visible light image divided into a set of tiles (dicom.nema.org). Each frame encodes a single tile within a three-dimensional imaged volume, and the storage SOP class is named VL Whole Slide Microscopy Image Storage. NEMA describes the organization as a pyramid of image data. The basic mechanism stores the tiles of a single

pyramid level as individual frames in a multi-frame object (dicom.nema.org), and each level is stored separately in the series (dicom.nema.org). A series may also hold ancillary images such as a slide label or macro image.

Two attributes decide how a reader finds a tile. The total pixel matrix describes the size of the entire imaged volume as one extent across all frames. The dimension organization is either **TILED_FULL**, where tile positions are implicitly defined (dicom.nema.org), or **TILED_SPARSE**, where positions are explicitly defined by per-frame functional group entries (dicom.nema.org). **TILED_FULL** is simpler for servers to index and cache; **TILED_SPARSE** accommodates scanners that skip blank tiles. A conformance test suite should include both. A layer's tiles may also span several instances of a concatenation or of a series, so a reader must not assume one file per level.

Frame arithmetic (illustrative calculation). For the NEMA example of 80,000 x 60,000 pixels, a 256 x 256 tile grid gives 313 x 235, or **73,555** base-level frames. A 512 x 512 grid gives 157 x 118, or **18,526** frames. NEMA gives typical tile sizes from 240 x 240 pixels (172 KB uncompressed) to 4,096 x 4,096 pixels (50 MB uncompressed) (dicom.nema.org). Frame count drives request count, so tile size is an architecture decision.

Change 2: Compression and transfer syntax become an interoperability variable

NEMA names JPEG and JPEG 2000 as the most frequently used lossy techniques. Its worked example shows JPEG 2000 reducing the roughly 15 GB image to about 300 MB, an approximately **50:1** ratio (dicom.nema.org). A 2022 review reports commonly used ratios of **10:1 to 20:1** for JPEG and **30:1 to 50:1** for JPEG 2000 (pmc.ncbi.nlm.nih.gov). One observer study on breast biopsy regions found performance nearly constant up to a ratio of **32:1** and a significant decrease at 64:1 and 128:1 (pmc.ncbi.nlm.nih.gov). Those are ranges from specific studies, not validated settings for a given scanner or diagnostic use.

PS3.5 explicitly allows RGB with lossy compression for whole slide images, to allow conversion from proprietary formats without loss due to color space transformation (dicom.nema.org). The change proposal behind this notes that compressed RGB components rather than YBR really are used by some WSI vendors. The consequence: a converter that passes compressed tiles through can preserve source pixels, but only if the server and viewer accept the resulting photometric interpretation and transfer syntax pair. In a multi-vendor study, baseline JPEG produced the smallest data size (pmc.ncbi.nlm.nih.gov). Platform support differs, and the **Table 3** matrix below shows how the major services treat transfer syntaxes.

Change 3: DICOMweb is the retrieval contract

PS3.18 defines DICOMweb as the REST web services of the standard. Bulk data resources retrieve large data elements such as Pixel Data from instances, and pixel data is a subset of bulk data. The standard lets a service constrain the transfer syntax parameter on media types, which is why two compliant servers can return different bytes for the same frame request. Rendered images returned as image/jpeg are baseline lossy 8-bit (dicom.nema.org), so a rendered tile is a display derivative, not diagnostic source data.

Annotations and AI results are likewise conveyed in objects separate from the image, as NEMA states (dicom.nema.org); they are treated in the AI section below. As of October 2026 the IHE DPIA profile that governs the LIS-to-scanner link is still a trial-implementation document, discussed under the identity chain.

Native, Dual-Personality or Proprietary: Choosing the Path

The first architectural decision is whether to normalize scanner output into DICOM at the source, convert it downstream, or coexist with proprietary formats. *Table 1* below separates the three paths by what each one demands from the scanner, the integration layer and the archive. The table does not rank scanners; vendor statements are cited only to show that each path exists in documented form.

Path	How it works	Documented examples and limits	Main test burden
Native DICOM	The scanner or its bundled software writes VL Whole Slide Microscopy Image instances and sends them with DICOM or DICOMweb.	Leica's Aperio GT 450 DX brochure lists a native DICOM and SVS saving option (leicabiosystems.com), and a press release says the scanner produces DICOM WSI natively (www.prnewswire.co.uk). Philips announced configurable native DICOM JPEG and JPEG XL output on its SGI scanner, which it describes as a world first (vendor claim, September 2025) (www.usa.philips.com). 3DHISTECH lists optional DICOM output for PACS interoperability (www.3dhistech.com).	Confirm which transfer syntaxes, tile sizes and Specimen module fields the scanner writes, per release and per country configuration.
Conversion or dual-personality	A converter reads a proprietary file and writes DICOM, or writes a file that is both a TIFF and a DICOM Part 10 file.	wsidicomizer opens proprietary formats and optionally converts them to DICOM (github.com); lossless conversion is offered for files supported by opentile (github.com), and z-stacks or multiple focal paths are documented as not supported (github.com). Philips documents a DICOM export that converts iSyntax to lossless JPEG 2000 (www.usa.philips.com). Bio-Formats can write dual-personality DICOM/TIFF files (bio-formats.readthedocs.io).	Pixel fidelity, metadata completeness, pyramid construction and the cost of keeping both the original and the DICOM copy.
Proprietary with viewer lock-in	The archive and viewer consume the vendor format. DICOM is produced only on export, if at all.	A 2019 paper notes that manufacturers, developers and users continue to depend on proprietary file formats (pmc.ncbi.nlm.nih.gov). One enterprise imaging vendor states that not all scanner vendors produce images in a standard file format and that it supports several proprietary formats (medical.sectra.com).	Exit cost, multi-vendor viewing, AI access to pixels and long-term retention readability.

The interpretive point is that "native" is a property of a scanner model, software release and licensed module, not of a brand. Philips states that it intends to support both DICOM and iSyntax formats for its IntelliSite Pathology Solution (www.usa.philips.com). Evidence in the table is vendor-authored, so it indicates availability, not behavior in your environment; only a conformance statement plus your own acceptance tests are dependable.

Conversion fidelity: what is documented and what is not

Lossless pass-through of compressed tiles needs a library that reads tiles from the source container without recompression; opentile describes exactly that for WSI TIFF files (github.com). The dual-personality proposal permits transcoding without further loss than present in the original source (pmc.ncbi.nlm.nih.gov). By contrast, wsidicom documents that levels must have a close to factor-2 scale and the same tile size (github.com), and Bio-Formats supports uncompressed, baseline JPEG, JPEG 2000 and RLE lossless transfer syntaxes for DICOM (bio-formats.readthedocs.io), narrower than what some scanners emit. OpenSlide lists DICOM as a supported multi-file format (openslide.org) but is a reader.

No public benchmark of conversion fidelity across scanner families is known. A 2018 multi-vendor study showed that whole slide image data can be encoded with patient and specimen metadata as DICOM objects ([pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)), which establishes feasibility, not a fidelity guarantee. Define local acceptance criteria: byte-identical compressed tiles where lossless pass-through is claimed, a pixel-difference tolerance where re-encoding occurs, and a metadata completeness check against the LIS record.

The decision rule that follows from the table: choose native DICOM where a scanner model documents it and your archive accepts its transfer syntax; choose conversion for mixed or legacy fleets, budgeting a converter service and a second data copy during transition; accept coexistence only where a diagnostic viewer is bound to a proprietary format, and even then require DICOM export so exit remains possible.

Implementation Considerations and Process Changes

Whole slide imaging scanner to LIS integration: specimen-container-slide identifiers

The workflow standards sit in the IHE Pathology and Laboratory Medicine (PaLM) domain. The identity chain has four hops: LIS case and specimen, container and slide label, scanner-read barcode, and the DICOM Specimen module. NEMA states that workflow outside the image object uses HL7 V2 transactions standardized in the IHE DPIA profile (dicom.nema.org), with patient, specimen and slide metadata retrievable using the identifier from the scanned barcode on the slide label (dicom.nema.org).

DPIA defines an Imaging Work Order Step (**IWOS**) as a service request that defines the detail for image acquisition (www.ihe.net). The work order uses an HL7 v2 **OML^O33** message that contains zero or more IWOSs for one or more specimens (www.ihe.net). DPIA states that the barcode uniquely identifies the physical asset and can be used to link case data to slide data across all actors (www.ihe.net). Its Study Instance UID maps one to one with the case accession number (www.ihe.net), and a pathology accession equates to one study (www.ihe.net). That mapping has a design consequence: in an illustrative layout, a case with 40 slides becomes one study with about 40 series, which affects QIDO paging and viewer worklists.

DICOM also supports a worklist route. PS3.17 describes the Modality Worklist use for specimens, in which the contents of the Specimen Module may be conveyed in the Scheduled Specimen Sequence of the worklist query (dicom.nema.org). In practice, DPIA acknowledges that the scope of the acquisition modality generally includes the scanner and its vendor middleware (www.ihe.net), so the LIS-to-scanner link often hides inside one vendor's software.

Table 2 below traces each identifier from source to image object. Mapping rows come from the DICOM Specimen module text and the DPIA appendix; the reconciliation column is analytical guidance and should be validated locally.

Identifier	Created by	Where it lands in DICOM	Reconciliation check
Accession number (case)	LIS	Study-level accession; maps to Study Instance UID one to one in DPIA	One study per accession; no study split across archives.
Specimen identifier	LIS	Specimen Identifier, a departmental information system identifier; DPIA maps it from HL7 SPM-2 (www.ihe.net)	Every series resolves to a specimen in the LIS; no orphan specimens.

Identifier	Created by	Where it lands in DICOM	Reconciliation check
Container (slide) identifier	LIS, printed as slide barcode	Container Identifier, the container that contains the specimens being imaged (dicom.nema.org); DPIA maps it from SAC-3	Barcode read equals LIS record; duplicates flagged before archive commit.
Issuer of identifiers	LIS site policy	Issuer of the Container Identifier Sequence, carried in the same Specimen module as the container identifier	Multi-site: issuer namespaces are unique, or identifiers collide after merger.
Specimen UID	LIS or Acquisition Manager	Specimen UID, carried inside the Specimen Description Sequence (dicom.nema.org)	Stable across rescans of the same slide.
Slide label image	Scanner	Ancillary image in the series	Label OCR or barcode matches metadata; label image retention follows privacy policy.

The table makes visible that the same logical slide appears under at least four names. Where a pathology system also exposes HL7 FHIR, a **FHIR ImagingStudy** provides information on a DICOM imaging study (hl7.org), and a FHIR Specimen carries a separate accession identifier assigned by the lab when accessioning specimens (hl7.org); the two should reference one another by identifier, never by free-text label.

Conversion pipeline design

Where conversion is needed, treat it as a service with its own queue, idempotency key and audit trail.

Microsoft's documentation recommends uploading each converted WSI as a DICOM series with multiple instances (learn.microsoft.com). A defensible design has three parts:

- **Landing zone and converter workers:** the scanner writes proprietary output to a quarantined share; stateless workers run a documented tool (for example `wsidicomizer` or a vendor export) with pinned versions recorded in each output.
- **Metadata join:** the worker fetches case, specimen and container data from the LIS by slide barcode, and refuses to write if a required field is empty.
- **Verification and commit:** compare frame counts, total pixel matrix dimensions and a sampled tile hash against the source before STOW, then write a manifest of every SOP Instance UID; keep the source file until the manifest reconciles.

Open-source tools are good test baselines, but production use needs a validation plan appropriate to your regulatory context.

DICOMweb whole slide images: STOW, QIDO and WADO data flow

The data flow is: scanner or converter sends instances with **STOW-RS** (store); viewers find studies and series with **QIDO-RS** (query); viewers and AI workers fetch metadata and frames with **WADO-RS** (retrieve). A 2018 enterprise-WSI briefing lists these archive requirements: take care at 2 GB and 4 GB boundaries such as frame byte offsets (dclunie.com), and support QIDO-RS and WADO-RS for viewer query, metadata and frame-level retrieval (dclunie.com). Those are still good test cases.

Service limits differ and should be written into the RFP. Azure's DICOM service limits requests to 4 GB (learn.microsoft.com) and accepts only explicit-VR transfer syntaxes (learn.microsoft.com); **AWS HealthImaging** announced import of instances up to 4 GB (aws.amazon.com); Google lists a limit of 4 GB per tag on ingested files (docs.cloud.google.com). A 15 GB uncompressed instance would exceed these, so compressed storage or concatenated instances must be settled before purchase. For frame retrieval, Microsoft recommends parallel single-part GET requests (learn.microsoft.com), while Google notes that retrieving an entire instance is generally faster per frame than batch retrieval (docs.cloud.google.com). These apply to different access patterns; only local measurement shows which fits your viewers and AI jobs.

Failure and retry matrix

Define each failure mode, the expected system behavior, and the reconciliation action before go-live:

- **Barcode unreadable or LIS order not found:** hold the slide in the landing zone, retry the IWOS query on a schedule, and escalate after a defined number of attempts; never create DICOM without a resolvable specimen.
- **STOW timeout or partial series:** retry with the same SOP Instance UID so the second attempt is idempotent, then compare QIDO counts with the manifest and re-send only missing instances.
- **Unsupported transfer syntax:** the archive rejects or transcodes; Azure documents transcoding to the requested transfer syntax on retrieval (learn.microsoft.com), which helps viewers but hides what is stored, so record the stored syntax in the manifest.
- **Rescan:** produce a new series under the same study with an auditable supersession flag, not an overwrite.

Digital pathology PACS integration: VNA, object store and cache architecture

A PACS or VNA is acceptable for WSI only if it handles multi-GB instances, frame-level access and large studies. The 2018 VNA review argues that fully digital laboratories need WSI in an enterprise-level archive, and the 2018 briefing lists frame-level retrieval and query support as archive requirements. Memorial Sloan Kettering's published design keeps about six months of scanned slides on a fast tier for viewer access and moves older slides to cheaper disk (pmc.ncbi.nlm.nih.gov). Google documents a digital pathology proxy that performs just-in-time frame caching to serve frames from an in-memory Redis cache (docs.cloud.google.com).

For storage planning, three rules follow. First, keep a hot tier sized to recent cases plus conference and tumor-board reuse, not to the whole archive. Second, put a frame cache in front of the archive, sized from measured viewport patterns. Third, define lifecycle policy per copy so that DICOM, the original proprietary file and any derived results can be retired independently.

Viewer and annotation interoperability

Viewers differ in what they consume. **Slim** runs fully client side and communicates with servers via standard DICOMweb (pmc.ncbi.nlm.nih.gov); its authors report interoperability with five commercially available enterprise-grade VNAs (pmc.ncbi.nlm.nih.gov). **OHIF** advertises out-of-the-box support for image archives which support DICOMweb (github.com), and **QuPath** reads DICOM WSI through Bio-Formats

(qupath.readthedocs.io), so an analyst tool may read files rather than use DICOMweb. Viewer selection should therefore be a test matrix of image, annotation, segmentation and measurement objects, not a feature checklist.

AI inference, bulk annotations, segmentation and structured results

Digital pathology AI integration architecture has four parts: tile access, inference orchestration, result encoding and result display. For tile access, The dicomslide library, for example, dynamically retrieves only the image frames (tiles) needed for a requested region (dicomslide.readthedocs.io). A pipeline that copies whole slides to local disk first gives up that benefit and creates an unmanaged extra copy. An AWS reference architecture places machine-learning models on SageMaker AI endpoints that perform inference on whole slide images (aws.amazon.com); it is vendor material, not an independent evaluation.

For orchestration, the IHE AI Workflow for Imaging (**AIW-I**) profile specifies workflow-management transactions based on the DICOM UPS-RS worklist service (www.ihe.net), and tasks may be managed with pull, triggered-pull or push workflows (www.ihe.net). The companion **AI Results** profile addresses capture, distribution and display of medical imaging analysis results (www.ihe.net). Both are radiology profiles, and the AI Results profile text contains no pathology content, so using them for slides is an adaptation that must be tested. Platform support for the worklist half is uneven: Azure lists a worklist service with UPS push and pull (learn.microsoft.com), while Google states it does not support the Worklist Service (docs.cloud.google.com).

For results, the choice among objects is driven by cardinality and geometry:

- **Microscopy Bulk Simple Annotations (MBSA):** a collection of points, closed polygons, open polylines and simple geometric shapes, organized into annotation groups that share common characteristics. **Supplement 222** says the annotations share a 3D Frame of Reference with a whole slide image instance and introduce no new image encoding mechanism (dicom.nema.org).
- **Segmentation:** binary, fractional or label map segmentations, suited to masks that behave like images.
- **Structured Reports (SR):** the Python library **highdicom** supports Structured Report documents containing numerical results, qualitative evaluations and vector graphic annotations (github.com). SR suits a manageable number of regions with rich measurements.
- **Presentation states:** the Advanced Blending Presentation State blends two or more sets of images referenced from within the IOD, which suits channel combinations and display recipes, not diagnostic findings.

The 2026 Imaging Data Commons preprint gives practical guidance: MBSA should be preferred when there are many regions, with or without measurements (arxiv.org), because automatically generated cell boundaries in a pathology image can easily number in the tens of thousands (arxiv.org). It reports that MBSAs can contain hundreds of thousands of regions (arxiv.org). It also identifies a gap: DICOMweb does not yet define an API for spatial querying of MBSAs or SRs (arxiv.org). The architectural implication is that a viewer must download an entire annotation object, so result objects should be partitioned (for example per tissue region or per model output group) if viewer responsiveness matters. **Supplement 255**, a draft supplement at the time of review (check the DICOM supplements page for final-text status), proposes enhanced annotations for quantized heatmap regions, holes and voids, and nested mutually exclusive concepts (www.dicomstandard.org); do not build to it before it is final.

Whichever object you pick, store the producing algorithm identity, model version and input series UIDs in a way the archive can query, and verify that this information round-trips through your archive and viewer. If the information is held in a private convention, document it in your own conformance statement so that a later migration does not lose provenance.

Performance, integrity, retry and reconciliation testing

Performance tests should use your slides, your network and your viewers. Published numbers are specific to their setups: a 2018 multi-vendor study reported that baseline JPEG gave the smallest data size, and a viewer-server paper reports interoperability results rather than a universal latency. A vendor-published throughput figure for one cloud service is not a basis for sizing a different deployment, so this report does not reuse any. A minimum test plan:

- **Tile latency distribution:** record p50, p95 and p99 from the viewer for pan and zoom at 20x and 40x on a representative slide set, cold and warm cache.
- **Concurrency ramp:** raise concurrent readers until p95 exceeds the target; record the knee.
- **Integrity:** hash every frame at the source and after retrieval for a sample of slides; compare SOP Instance UID lists to the manifest.
- **Transfer syntax matrix:** store and retrieve each syntax your scanners emit through each service, and check whether the service transcodes or preserves bytes (docs.aws.amazon.com).
- **Failure injection:** kill the converter mid-STOW and drop the LIS interface for an hour; confirm reconciliation recovers without manual DICOM edits.
- **Rendered versus native frames:** rendered image/jpeg output, which is baseline lossy 8-bit, must never serve as AI input or archive copy.

Table 3 below is a platform support matrix drawn from vendor documentation. It lists what each documents for DICOMweb and large-image handling and is not a ranking; every cell should be re-verified against the current release before purchase.

Platform (documentation)	DICOMweb services documented	Large-image and transfer-syntax notes	Pathology-relevant notes
Azure DICOM service	STOW-RS, WADO-RS, QIDO-RS and a worklist service with UPS push and pull	4 GB request limit; explicit-VR transfer syntaxes only	States a format conversion is required for WSI in other formats (learn.microsoft.com); designed for multi-GB instances (learn.microsoft.com).
Google Cloud Healthcare API	Studies service; no URI service, worklist service or capabilities transactions	4 GB per tag limit on ingested files	Pyramid levels stored as frames (docs.cloud.google.com); import via DICOMweb, DIMSE or Cloud Storage (docs.cloud.google.com).
AWS HealthImaging	DICOMweb representations not available through AWS CLI and SDKs (docs.aws.amazon.com)	Any SOP class UID importable (docs.aws.amazon.com); HTJ2K default storage, JPEG 2000 Lossless configurable (docs.aws.amazon.com)	Faster import of large instances such as whole slide imaging announced May 2025 (aws.amazon.com).

Platform (documentation)	DICOMweb services documented	Large-image and transfer-syntax notes	Pathology-relevant notes
Orthanc (open source)	DICOMweb plugin; the full standard is not implemented yet (orthanc.uclouvain.be)	WSI plugin re-encodes some tiles to PNG for display (orthanc.uclouvain.be)	Includes a converter and web viewer for whole-slide images (orthanc.uclouvain.be).

The matrix shows why a product name is not a conformance claim: two services can both support DICOMweb while differing on worklists, size limits and transcoding. Ask each supplier for a current conformance statement and run the same acceptance suite against each.

Migration, pilot plan and RFP questions

A pilot sized for evidence, not demonstration, has eight steps:

1. **Inventory and success metrics:** list scanners, firmware, formats, LIS interfaces, viewers and AI tools; set targets for identity match rate, tile latency, integrity pass rate and reconciliation time.
2. **Select a slide set** with special cases: multi-section slides, immunohistochemistry, large blocks and rescans.
3. **Build the identity chain first:** LIS order to barcode to Specimen module, tested before any pixels move.
4. **Run native and converted paths side by side** for the same slides, comparing metadata and tiles.
5. **Load test** the archive and cache using the capacity worksheet below.
6. **Add AI results** as DICOM objects and confirm the viewer displays them without custom code.
7. **Validate for diagnostic use** under your governing guidance. The College of American Pathologists validation guideline update recommends at least **60 cases** for one application or use case (pmc.ncbi.nlm.nih.gov) and a washout period of at least 2 weeks between viewing digital and glass slides (pmc.ncbi.nlm.nih.gov), as quoted in a secondary paper. Consult the guideline and your regulator for binding requirements.
8. **Plan exit:** prove that you can export every slide and result as DICOM without vendor tools.

For the RFP, ask each supplier these questions:

- Which SOP classes, transfer syntaxes and tile sizes are written and read, per release, and is the LIS-to-scanner interface (DPIA role) documented?
- Which DICOMweb transactions are supported, with what size, paging and timeout limits, and can stored bytes be retrieved unmodified?
- How are annotations, segmentations and AI results stored and queried, and what export path exists for the full archive in standard DICOM?
- What evidence from public connectathons or independent tests exists for the exact version offered?

Data Analysis and Evidence

The quantitative core combines published reference points with a worksheet that readers fill with local measurements.

Published reference points

NEMA's example of a 20 mm by 15 mm sample at 0.25 micrometers per pixel gives 80,000 x 60,000 pixels, about 4.8 gigapixels (dicom.nema.org), or about 15 GB at 24-bit color (dicom.nema.org). Its extreme example, 125 gigapixels per plane across ten planes, reaches 3.75 TB (dicom.nema.org). Generating pyramid levels a factor of 2 apart adds **about 32%** to the data set (dicom.nema.org). Institutional figures vary by scanner, magnification, compression and policy: a 2022 review states large pathology centers may produce over 1 million digital slides and around 1 petabyte per year (pmc.ncbi.nlm.nih.gov); a 2018 multisite estimate projected about 4.3 petabytes per year (pmc.ncbi.nlm.nih.gov); and Memorial Sloan Kettering reported allocating 9 PB for slide storage (pmc.ncbi.nlm.nih.gov). They differ in year, compression and scope, and none is a forecast for your laboratory.

Capacity worksheet

The storage formula is: annual stored volume = slides per day x operating days x average source size x conversion multiplier x retention copies. The concurrency formula uses Little's law: required throughput (requests per second) = concurrent tile requests in flight divided by target response time in seconds.

The worksheet inputs, each to be replaced with a local measurement, are:

- **Slides per day and operating days:** from LIS accession logs.
- **Average source size:** the mean file size on your scanners at your magnification and compression settings.
- **Conversion multiplier:** 1.0 for native DICOM only, 2.0 if the proprietary original is kept alongside a DICOM copy.
- **Retention copies:** primary, disaster recovery and any research replica.
- **Concurrent tile requests and target response time:** from viewer logs and a reader-experience target agreed with pathologists.

(Hypothetical Example) A laboratory scans 800 slides per day on 250 operating days. Local measurement of recent scans gives an average source size of 1.0 GB. The laboratory keeps the proprietary original and a DICOM copy during transition (conversion multiplier 2.0) and holds two retention copies (primary and disaster recovery). Annual volume is $800 \times 250 \times 1.0 \text{ GB} = \mathbf{200 \text{ TB}}$ of source data, times 2.0 for the dual representation = **400 TB**, times 2 copies = **800 TB per year**. After migration, with native DICOM only (multiplier 1.0), the same laboratory would store 400 TB per year. If instead the conversion regenerates pyramid levels a factor of 2 apart that the source did not store, the NEMA overhead of about 32% would raise the native-only figure to about **528 TB**.

For viewing, suppose a 1920 x 1080 viewport over 256-pixel tiles needs about 40 tile requests per view refresh (8 columns x 5 rows). With 30 concurrent readers panning at once, 1,200 requests are in flight. With a target response time of 250 ms, the origin must sustain $1,200 / 0.25 = \mathbf{4,800 \text{ requests per second}}$ at peak. A cache hit ratio of 70% would reduce origin load to **1,440 requests per second**. Every input in this example is an assumption to be replaced by measurement, and the numbers illustrate the method, not a recommendation.

The first sensitivity worth running is the conversion multiplier: whether the proprietary original is retained for the full retention period or only until reconciliation completes can change volume by a factor of two.

What the evidence does and does not support

Strong evidence supports the structural claims (tiled frames, pyramid levels as separate instances, specimen identity in the Specimen module, DICOMweb retrieval), because they come from the standard itself. Moderate evidence supports compression ranges, which come from a review and an observer study with specific test conditions. Cross-vendor conversion fidelity, enterprise-scale DICOMweb tile performance and independent archive comparisons lack public benchmarks, which is a reason to budget for local testing, not to defer standards-based design.

Implications and Future Directions

Three shifts are likely to shape procurement over the next planning cycle. First, native DICOM output is becoming a documented option across more scanner families, with at least one vendor announcing native DICOM output in both JPEG and JPEG XL as of September 2025, and AWS announcing JPEG XL support for its imaging service in early 2026, naming whole slide imaging systems as an example consumer (aws.amazon.com). Codec choice is therefore a joint scanner and archive decision.

Second, the identity chain is moving from vendor middleware toward a published profile. DPIA is at trial-implementation status (www.ihe.net), and the draft IHE Digital Pathology Ordering and Workflow work says adherence to DPIA plays an important part in a successful implementation (build.fhir.org). Programs that capture identifiers cleanly now will find later profile adoption cheaper.

Third, AI results are becoming first-class DICOM objects: the Slim viewer is described as the slide microscopy viewer of the NCI Imaging Data Commons (pmc.ncbi.nlm.nih.gov), but the absence of a spatial query API for bulk annotations means viewers will keep downloading whole annotation objects for some time.

For organizations deciding whether to normalize to DICOM, the implications are practical. A mixed fleet will keep a conversion tier for years, so budget it as production infrastructure. Choose archives on documented frame-level access, size limits and transfer-syntax behavior, and keep an exit path: exporting every slide, annotation and AI result in standard DICOM is the strongest hedge against format and vendor change. Independent expertise can help structure the RFP, acceptance suite and governance; it does not replace vendor conformance statements, a clinical validation plan or regulator guidance, and quality and compliance teams should be involved before any diagnostic deployment.

Frequently Asked Questions (FAQs)

What is DICOM whole slide imaging?

DICOM whole slide imaging is the use of the DICOM standard to store and exchange digitized microscope slides as tiled, multi-frame images with a multi-resolution pyramid. The IOD specifies a multi-frame visible light image divided into a set of tiles, and each pyramid level is stored separately in a series.

Is DICOM Supplement 145 still the reference?

Supplement 145 is the historical origin of the whole slide IOD, published in 2010. For normative details use the current standard, for example PS3.3 section A.32.8 and PS3.18 for DICOMweb.

How does a scanner connect to the LIS?

Usually through HL7 v2 transactions, standardized in the IHE DPIA profile, with the slide barcode as the link to case data. Many deployments hide this link inside vendor middleware, so request interface documentation.

DICOM WSI versus proprietary scan formats: which is better?

Neither choice is universally better. DICOM enables reuse of enterprise imaging infrastructure, while proprietary formats persist in scanners and software. The decision depends on exit cost, multi-vendor viewing and AI access, covered in **Table 1**.

Conclusion

DICOM whole slide imaging integration succeeds or fails on the chain from LIS identifier to scanner barcode to Specimen module, and on whether the archive can serve frames at the speed readers and models need. The standard supplies tiled multi-frame objects, specimen metadata, DICOMweb retrieval and separate annotation objects; the vendor ecosystem supplies native, converted and proprietary paths, each documented to a different degree and changing by release.

The recommended stance is to treat DICOM as the system of record, conversion as a tested service, and conformance statements as the start of local acceptance testing. Size storage and delivery with the worksheet, using your own slide sizes and tile latency measurements, and insist on an export path in standard DICOM. Because DPIA is still at trial-implementation status as of October 2026 and vendor support differs by scanner and release, revalidate every decision here against the specific products and versions in your procurement.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://dicom.nema.org/dicom/dicomwsi>
2. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.7.6.22.html
3. https://leicabiosystems.com/sites/default/files/media_product-download/2024-04/Brochure_Aperio_GT_450_DX__240094_Rev_A.pdf
4. <https://www.usa.philips.com/healthcare/article/dicom-in-digital-pathology>
5. <https://learn.microsoft.com/en-us/azure/healthcare-apis/dicom/dicom-digital-pathology>
6. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6489422>
7. https://www.ihe.net/uploadedFiles/Documents/PaLM/IHE_PaLM_Suppl_DPIA_Rev1-4_TI_2026-03-06.pdf
8. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6236926>
9. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10550754>
10. <https://arxiv.org/pdf/2609.14530>
11. <https://digitalpathologyassociation.org/blog/dicom-for-digital-pathology-interoperability>
12. <https://pmc.ncbi.nlm.nih.gov/articles/PMC5869966>
13. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6289005>

14. <https://pmc.ncbi.nlm.nih.gov/articles/PMC12616093>
15. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_A.32.8.html
16. https://dicom.nema.org/medical/dicom/current/output/chtml/part03/sect_C.7.6.17.html
17. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9577055>
18. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3352607>
19. https://dicom.nema.org/medical/dicom/current/output/chtml/part05/sect_8.2.html
20. https://dicom.nema.org/medical/dicom/current/output/chtml/part18/sect_8.7.4.html
21. <https://www.prnewswire.co.uk/news-releases/leica-biosystems-enhances-integration-of-digital-pathology-into-the-diagnostic-pathway-with-dicom-imaging-853351848>
22. <https://www.usa.philips.com/about/news-and-insights/news/philips-announces-digital-pathology-scanner-with-native-configurable-dicom-jpeg-and-jpeg-xl-output-in-world-first>
23. <https://www.3dhistech.com/software/pannoramic-diagnostic-scanner-software>
24. <https://github.com/imi-bigpicture/wsidicomizer>
25. <https://bio-formats.readthedocs.io/en/stable/formats/options.html>
26. <https://medical.sectra.com/news-press-releases/news-item/51F95E435A285238>
27. <https://github.com/imi-bigpicture/opentile>
28. <https://github.com/imi-bigpicture/wsidicom>
29. <https://bio-formats.readthedocs.io/en/stable/formats/dicom.html>
30. <https://openslide.org/formats>
31. https://dicom.nema.org/medical/dicom/current/output/chtml/part17/sect_nn.7.html
32. <https://hl7.org/fhir/imagingstudy.html>
33. <https://hl7.org/fhir/specimen.html>
34. https://dclunie.com/papers/ECDP_NDP_2018_DICOMEnterpriseWSIPreSeminar_Clunie.pdf
35. <https://learn.microsoft.com/en-us/azure/healthcare-apis/dicom/dicom-services-conformance-statement-v2>
36. <https://aws.amazon.com/about-aws/whats-new/2024/03/aws-healthimaging-import-dicom-objects-hij2k-transfer-syntaxes>
37. <https://docs.cloud.google.com/healthcare-api/quotas>
38. <https://docs.cloud.google.com/healthcare-api/docs/how-tos/dicom-digital-pathology>
39. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10033920>
40. <https://github.com/OHIF/Viewers>
41. <https://qupath.readthedocs.io/en/stable/docs/intro/formats.html>
42. <https://dicomslide.readthedocs.io/en/stable/introduction.html>
43. <https://aws.amazon.com/blogs/industries/revolutionizing-healthcare-with-ai-driven-digital-pathology>
44. https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_Suppl_AIW-I.pdf
45. https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_Suppl_AIR.pdf
46. <https://docs.cloud.google.com/healthcare-api/docs/dicom>
47. https://dicom.nema.org/medical/dicom/Supps/PC/sup222_pc_WSIAnnotations.pdf
48. <https://github.com/ImagingDataCommons/highdicom>

49. <https://www.dicomstandard.org/news-dir/current/docs/sups/sup255.pdf>
50. <https://docs.aws.amazon.com/healthimaging/latest/devguide/get-image-frame.html>
51. <https://docs.aws.amazon.com/healthimaging/latest/devguide/using-dicomweb.html>
52. <https://docs.aws.amazon.com/healthimaging/latest/devguide/supported-sop-classes.html>
53. <https://docs.aws.amazon.com/healthimaging/latest/devguide/supported-transfer-syntaxes.html>
54. <https://aws.amazon.com/about-aws/whats-new/2025/05/aws-healthimaging-dicomweb-qido-rs-search-enhanced-data-management>
55. <https://orthanc.uclouvain.be/book/plugins/dicomweb.html>
56. <https://orthanc.uclouvain.be/book/plugins/wsi.html>
57. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10087327>
58. <https://aws.amazon.com/about-aws/whats-new/2026/01/aws-healthimaging-adds-jpeg-xl>
59. <https://build.fhir.org/ig/IHE/PaLM.DPOW>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.