MACH7 TECHNOLOGIES · FIELD GUIDE

Imaging Informatics 101

The Complete Guide

A plain-language field guide to how medical images are created, stored, moved, displayed, and shared across the modern health enterprise.

For everyone who touches medical imaging, not just radiologists.

MACH7

mach7t.com

Contents

Why this guide exists

Medical imaging is the largest data type in healthcare, and one of the least understood outside the reading room. A single hospital produces studies from dozens of devices, stores them for decades, moves them between facilities, and puts them in front of clinicians who are rarely radiologists. The systems and standards that make all of this work quietly run underneath almost every care journey, yet most of the people who depend on them have never been taught how they fit together.

> **Enterprise imaging data can account for up to roughly 90% of a health system's total storage consumption, and hospitals generate on the order of 50 petabytes of data per year.** *HealthTech Magazine, 2024*

This guide is the map. It explains imaging informatics from first principles, in language a new engineer, a clinical champion, an IT administrator, or a curious executive can follow. It is vendor neutral by design. The goal is understanding, not a sales pitch.

Who this is for

- Healthcare IT and imaging leaders evaluating or running enterprise imaging systems.
- Clinical and informatics professionals who work alongside imaging every day and want the full picture.
- Anyone new to the field who needs a fast, accurate grounding in how it all connects.

How it is organized

The guide moves in five parts: the foundations of the field, the standards that hold it together, the systems that store and display images, the infrastructure and trust that keep it safe, and where artificial intelligence and regulation are taking it next. Each part stands on its own, so you can read straight through or jump to what you need.

Part 1: Foundations

What imaging informatics is

Imaging informatics is the discipline concerned with how medical images and their associated data are created, stored, moved, displayed, interpreted, and shared across the healthcare enterprise. It sits where clinical medicine, computer science, and imaging science meet.

It spans a wide territory: the standards and interoperability that let systems talk to each other, the systems themselves from scanners to archives to viewers, the science of displaying and visualizing images, and the security and privacy practices that protect patient data through its whole life. It is a team sport. Radiologists, IT staff, engineers, physicists, and vendors all live here.

Why it matters

The point of all this machinery is simple to state and hard to deliver: the right image, on the right patient, in front of the right clinician, quickly. Imaging drives diagnosis and treatment decisions across nearly every specialty. Web based viewers now put prior studies and new results in front of clinicians wherever they are, whether in the reading room, the clinic, the operating room, or at home.

It is also a safety and trust discipline. Correct patient identity, calibrated displays, complete studies, and auditable access are patient safety concerns, not just IT housekeeping. And because imaging is the largest data type in healthcare by volume, moving and rendering it efficiently is an engineering challenge in its own right.

A brief history: from film to cloud

Every era of imaging expanded who could see an image, how fast, and from where. That single thread runs from the first X-ray to today's cloud services.

- The film era (1895 to 1970s): Roentgen's X-ray produced images on film, read on lightboxes, stored in film libraries, and couriered between sites.
- Digital is born (1970s to 1980s): CT in 1971 and MRI in the 1980s created inherently digital images, and the ACR-NEMA standard in 1985 began standardizing formats.
- PACS and DICOM 3.0 (1990s): DICOM 3.0 in 1993 enabled multi-vendor networking, and filmless radiology departments became real.
- Web and enterprise (2000s to 2010s): web and zero-footprint viewers, vendor neutral archives, and enterprise imaging brought images beyond radiology to the whole health system.
- Cloud and AI (2020s): DICOMweb, cloud-native imaging services, server-side rendering, and AI moved imaging onto internet-era systems.

The clinical context: who creates and consumes images

Radiology is the historical home of medical imaging, covering X-ray, CT, MRI, ultrasound, and nuclear medicine. But imaging now reaches far beyond it. Cardiology produces echo, angiography, cardiac CT and MR, and ECG waveforms. Ophthalmology captures fundus photography, OCT, and visual fields. Pathology scans whole slides of tissue. Dermatology and wound care rely on clinical photography over time. Point-of-care and surgical settings add ultrasound, endoscopy, and intra-operative imaging.

Just as important are the consumers: referring physicians, surgeons, oncologists, emergency teams, tumor boards, and increasingly patients themselves through portals. The key insight for anyone building or running imaging systems is that most viewer users are not radiologists.

Imaging modalities

Imaging devices fall into five broad families. The two-letter codes in parentheses are DICOM modality codes, and you will see them everywhere.

- Projection radiography (CR, DX, MG): X-ray passes through the body onto a detector.
- Cross-sectional (CT, MR): the body reconstructed as slices or volumes.
- Real-time and dynamic (US, XA, RF): live, moving images during an exam or procedure.
- Functional and molecular (PET, SPECT, NM): physiology and metabolism, not just anatomy.
- Visible light and specialty (OCT, fundus, whole-slide, endoscopy): cameras and optical scanners.

The main modalities in brief

X-ray (CR, DX). A short exposure through the body onto a digital detector. It produces one or a few high-resolution 2D projections per study. The images are small, but they are produced in enormous numbers, so throughput and time-to-first-image matter more than raw study size. X-ray is usually the majority of a hospital's exam count.

Computed tomography (CT). An X-ray tube and detector spin around the patient while the table advances, and a computer reconstructs the body as cross-sectional slices. Studies are stacks of hundreds to thousands of thin slices, often in several series, that recombine into multi-planar and 3D renderings. Thin-slice CT demands fast scrolling and server-side reconstruction, and radiation dose is captured in structured dose reports.

Magnetic resonance imaging (MR). A strong magnet and radio-frequency pulses produce images with no ionizing radiation. Each exam contains many sequences, often more than a dozen series, each highlighting different tissue. With this many series, hanging protocols that lay sequences out consistently are the difference between a fast read and a frustrating one.

Ultrasound (US). A handheld transducer builds images from returning sound echoes in real time, captured as still frames and cine loops, with Doppler adding blood-flow color and waveforms. Quality is operator dependent, and echo measurements travel as structured reports that viewers must parse and present.

Angiography and fluoroscopy (XA, RF). A C-arm produces continuous low-dose X-ray video while contrast is injected, often at 7.5 to 30 frames per second. These are working images in a live procedure, so cine must play back instantly and smoothly, and dose management is a first-class concern.

Beyond radiology, PET and SPECT overlay function onto anatomy and must fuse two datasets, ophthalmic OCT correlates laser cross-sections back to a fundus image, and whole-slide imaging turns a glass slide into a multi-gigapixel image you pan and zoom like a map.

Modalities at a glance

Modality	Ionizing	Form	Typical study size	What that means for software
X-ray (CR/DX)	Yes	2D projection	10 to 50 MB	Small and fast, but enormous volume
Mammo (MG/DBT)	Yes	2D plus pseudo-3D	50 MB to 3 GB	High-res display, priors side by side
CT	Yes	3D slice stack	100 MB to 1 GB+	Fast scroll, MPR and 3D, many slices
MRI	No	3D, multi-sequence	100 MB to 1 GB	Many series; smart layout matters
Ultrasound / Echo	No	2D plus cine loops	10 to 500 MB	Smooth video playback
PET/CT	Yes	Functional plus anatomic	Hundreds of MB	Fusion overlay of two datasets
OCT (ophthalmic)	No	3D B-scan volume	50 to 500 MB	Frame-to-fundus correlation
Whole-slide (SM)	No	Gigapixel 2D pyramid	1 to 5+ GB	Tiled, progressive streaming

Sizes are approximate and highly variable. The point is the range: roughly four orders of magnitude between a chest X-ray and a pathology slide.

The imaging workflow

Every study follows the same chain, and every link in it is a standards conversation.

- Order: a clinician orders a study in the EHR, and the radiology information system schedules it.
- Acquire: the modality pulls patient details from the worklist and scans.
- Store: images land in the PACS or VNA as DICOM objects.
- Interpret: a radiologist reads on a diagnostic viewer with prior studies alongside.
- Report: findings become a report linked to the study.
- Distribute: results reach referrers, specialists, and patients.

Orders travel as HL7. The worklist and the images are DICOM. Reports flow back as HL7 or FHIR. Part 2 unpacks each one.

A day in the life of one CT scan

At 07:42 a patient arrives in the emergency department with chest pain, and the physician orders a CT angiogram in the EHR. By 07:55 the scanner has pulled the order from the modality worklist, so the patient name, ID, and study details arrive pre-populated and nothing is typed by hand. At 08:03 the scan completes and roughly 1,400 thin slices are sent to the archive within seconds. At 08:10 a radiologist opens the study, prior exams load automatically for comparison, and 3D reconstructions help rule out a dissection. By 08:31 the report is signed and flows back to the EHR, and the emergency physician reads it minutes later. From then on the study is part of the patient's longitudinal record, retrievable years later, at any facility, from any authorized viewer.

Part 2: The standards that hold it together

DICOM: the language of medical imaging

DICOM is the standard for medical images. Its data model is a hierarchy worth learning once and using forever: a Patient has one or more Studies, each Study has one or more Series, and each Series contains Instances, the individual objects such as an image, a structured report, or an encapsulated PDF.

Metadata and pixels travel together. Every DICOM object carries tags describing the patient, exam, and acquisition alongside the pixel data itself, for example Patient Name, Study Date, the Study Instance identifier, and the pixel data. SOP classes define what an object is, such as CT Image Storage or Ultrasound Multi-frame, and unique identifiers glue the whole hierarchy together. Every query, worklist, prefetch rule, and viewer layout is built on this Patient, Study, Series, Instance model.

Transfer syntaxes and compression

A transfer syntax is the agreed encoding of a DICOM object, negotiated between sender and receiver. Compression is where storage cost, network load, and image fidelity meet, and the choice between lossy and lossless is a clinical and regulatory decision, not just a technical one.

Transfer syntax	Type	Where you will see it
Explicit VR Little Endian	Uncompressed	The default for interchange, simple but large
JPEG Baseline	Lossy	Ultrasound, XA cine, visible-light imaging
JPEG Lossless / JPEG-LS	Lossless	Classic archive compression
JPEG 2000	Lossless or lossy	Archives; progressive, resolution-scalable decode
HTJ2K (High-Throughput JPEG 2000)	Lossless or lossy	Modern cloud archives and streaming pipelines
RLE	Lossless	Older modalities and systems

DICOM networking: classic and web

Classic DICOM networking, known as DIMSE, still runs most of the world's imaging. A handful of services do the work: C-ECHO verifies that two systems can talk, C-STORE pushes objects to a destination, C-FIND queries by patient, study, or series, and C-MOVE or C-GET asks an archive to send studies. The modality worklist lets a scanner pull scheduled exam and patient details. Systems identify themselves with application entity titles and talk over point-to-point connections on set ports.

DICOMweb brings imaging onto the same web technology that runs the rest of the internet. QIDO-RS searches with an HTTP query, WADO-RS retrieves studies down to single frames and rendered tiles, STOW-RS stores objects with an HTTP POST, and UPS-RS offers worklists as a web service. The same transport, security tokens, load balancers, and content delivery networks that run modern web services now apply to imaging, and frame-level and tile-level retrieval is what makes progressive rendering of huge studies possible in a browser.

Beyond DICOM: HL7, FHIR, and IHE

Images and their metadata are DICOM. Almost everything around them, the orders, the results, and patient identity, is HL7 or FHIR, choreographed by IHE.

HL7 v2 is the veteran messaging standard, a stream of pipe-delimited events flowing between the EHR, the RIS, and the PACS: ADT messages for patient admit, move, and merge; order messages when an exam is ordered; and result messages when reports flow back.

FHIR is the modern successor, offering web-friendly resources such as Patient, ImagingStudy, and DiagnosticReport. It powers portals and connected apps and is increasingly the door into the EHR.

IHE is not a standard but a playbook for using standards together, tested publicly at events called Connectathons. Profiles such as Scheduled Workflow describe the imaging chain, and others cover cross-enterprise image sharing and the correction of mistakes.

Part 3: The systems that store and display images

The enterprise systems map

Four groups of systems, connected by standards, do the work. The EHR and RIS handle orders, scheduling, results, and the patient record, and talk in HL7 and FHIR. Modalities pull the worklist and acquire images, sending them as DICOM. The PACS or VNA is the archive that stores, routes, and serves studies. Viewers provide diagnostic and clinical viewing for every user, increasingly over DICOMweb. Around all of it sits image exchange, the sharing of studies with other organizations and with patients. The archive is the hub, and everything else connects through standards.

PACS, VNA, and the viewing layer

PACS (Picture Archiving and Communication System) is the departmental workhorse: reading worklists, routing, and the working archive, historically a single vendor's closed world.

VNA (Vendor Neutral Archive) is one standards-based store for all departments, letting an organization own its data independent of any single viewer or PACS vendor.

The diagnostic viewer is radiologist-grade reading software with hanging protocols, high-fidelity display, and advanced tools, where the formal interpretation happens.

The clinical or zero-footprint viewer is browser-based viewing for everyone else, launched from the EHR with no installation on any device. It is the widest-reaching system in the stack.

> **The vendor neutral archive market is projected to grow from about US$2.94 billion in 2024 to roughly US$8.89 billion by 2033, a compound annual growth rate near 13%.** *Grand View Research*

The clear industry direction is to separate the archive from the viewing layer and serve the whole enterprise, not just radiology.

How pixels reach the screen

Modern viewers combine two rendering approaches. Client-side rendering displays 2D images, such as X-ray, CT and MR slices, and ultrasound cine, right on the device, so scrolling, window and level, and measurement feel instant with no round trip. Intelligent streaming and prefetch move the pixels so that even large studies display quickly. Server-side rendering handles advanced visualization: multi-planar reconstructions, maximum-intensity projections, and 3D volume renderings are created in real time on GPU-backed servers, and the device simply streams the result, so a thin-slice volume never has to download in full to be explored in any plane. The best systems use both, 2D on the client for responsiveness and advanced visualization on the server for power, engineered so both stay fast regardless of study size and all delivered in the browser.

Tools of the reading room

- Hanging protocols auto-arrange series and priors the way each reader expects, by modality, body part, and user.
- Window and level, plus lookup tables, map deep grayscale to the screen using presets such as lung, bone, and brain.

- Measurements and annotations record lengths, angles, regions of interest, and standardized uptake values, with correct calibration behind every number.
- Presentation states save exactly how an image was displayed, including zoom, window and level, and markup, as a DICOM object others can recall.
- Key image selection flags the images that matter so the next reader goes straight to them.
- Multi-planar and 3D tools rebuild slice stacks into any plane and into volume renderings on demand.

Specialty deep dives

Whole-slide imaging. A scanned pathology slide is a multi-gigapixel image organized as a pyramid, from a low-resolution overview down to full detail. Viewers fetch only the tiles in view, at the zoom level needed, the same idea behind online maps. Progressive, tile-level retrieval is not optional at this scale.

Ophthalmic imaging. An OCT volume is a stack of laser cross-sections of the retina. Viewers draw reference lines on the fundus photo showing exactly where each cross-section sits, and true-scale display matters because measurements are in micrometers.

Cardiology and enhanced objects. An echo study is images plus dozens of measurements carried as structured reports. ECGs and hemodynamic traces are stored as DICOM waveform objects and rendered as scrolling signals. Cath-lab cine runs at up to 30 frames per second, and modern scanners pack a whole acquisition into a single enhanced multi-frame object. The further you move from single-frame radiology images, the more the viewer’s job becomes interpreting structure, not just drawing pixels.

Part 4: Infrastructure and trust

Where imaging systems run

Imaging systems run on-premises, in the cloud, or in a hybrid of the two. On-premises means hospital-owned data centers with full control and data residency, but capacity has to be bought years ahead and scaling is slow. Cloud offers elastic compute and managed services paid as you go, though imaging's sheer size makes network egress and data residency choices critical. Hybrid is the common reality: a cloud archive with on-prem edge caching and rendering, or the reverse. Because server-side rendering makes infrastructure part of the product, GPU capacity, instance selection, and network latency directly shape what the user experiences.

Cloud-native imaging services

The major cloud providers now offer managed imaging services, and the common thread is that all of them speak DICOMweb. AWS HealthImaging is a managed, HIPAA-eligible DICOM store built on HTJ2K with frame-level retrieval. Azure Health Data Services offers a DICOM service alongside FHIR under one umbrella. Google Cloud Healthcare API provides DICOM stores with native DICOMweb, tied to analytics and AI tooling. Because they are standards-based, a standards-based viewer can sit on top of any of them, which is exactly why the DICOM and DICOMweb foundations matter so much.

Security, privacy, and compliance

Images and their metadata are protected health information, among the most personal data there is, governed by HIPAA in the United States, PIPEDA and provincial law in Canada, GDPR in Europe, and more. Sound practice rests on a few pillars: access control through single sign-on, least-privilege roles, and break-glass procedures for genuine emergencies; auditing of every view, export, and share, using standard, reviewable trails; and de-identification for research, teaching, and AI training, which means scrubbing both metadata tags and text burned into pixels, which is harder than it sounds. Encryption in transit and at rest is the baseline everywhere, on the assumption that every study belongs to a real person.

Interoperability and image exchange

Exchange is where every standard gets stress-tested against the messy real world. Cross-enterprise sharing profiles and commercial networks move studies between organizations for referrals, transfers, and second opinions. Getting outside prior studies in front of a reader avoids repeat scans, repeat radiation, and repeat cost. Patient portals and share links put imaging in patients' own hands, which is both a regulatory expectation and a product opportunity. And discs still arrive daily, so import, patient reconciliation, and correction of mismatched records remain essential, unglamorous work.

Part 5: Where imaging is going

AI in the imaging workflow

Artificial intelligence shows up at every step of the imaging workflow. Before the read, it triages and prioritizes the worklist so that suspected bleeds and clots jump the queue. During the read, detection and quantification overlays highlight nodules, fractures, and volumes. In reporting, it drafts text and auto-populates measurements for the radiologist to verify. In operations, it supports protocol selection, scheduling, and dose and quality monitoring. And at population scale, it powers screening programs and opportunistic findings. The constant through all of it is that a human clinician remains responsible for the diagnosis. AI reorders, highlights, and measures. It does not sign the report.

AI results are just more DICOM

AI output travels through the same standards as everything else. Structured reports carry findings and measurements in machine-readable form. Segmentation objects are pixel-level masks of organs and lesions, overlaid on the source images. Presentation states and secondary-capture images carry simple, universally displayable markup. And most diagnostic AI is regulated software, cleared and version-controlled like any other medical device. The viewer is where AI meets the clinician, so results must render faithfully alongside the source images and be clearly marked as machine-generated. Display is a safety function.

Regulatory and quality: imaging software is a medical device

Diagnostic imaging software is a regulated medical device, and that shapes how it is built. Market clearance is required before use: an FDA 510(k) in the United States, a licence from Health Canada, or CE marking under the EU Medical Device Regulation. A software lifecycle standard governs documented requirements, design, verification, and maintenance, scaled to the safety class. Quality management and risk management standards prove that the software is built safely on every release, and change control means every change is assessed for clinical, safety, and regulatory impact before it ships. For anyone new to the field, the mindset shift is that process here is not bureaucracy. It is how a viewer earns the right to be used in patient care.

The whole field in three ideas

One mental model. Patient, Study, Series, Instance, moving along the imaging chain from order to distribution.

One rule of thumb. Images are DICOM. Everything around them, orders, results, and identity, is HL7 or FHIR, choreographed by IHE.

One mindset. Every study belongs to a real person. Safety, privacy, and process are not overhead. They are the product.

A note from Mach7 Technologies

Mach7 Technologies builds enterprise imaging software that helps health systems own and move their imaging data on open standards, independent of any single vendor. Our work centers on the interoperability, data independence, and practical modernization this guide describes: standards-based storage and exchange, and enterprise viewing that reaches every clinician, not just radiology. If the ideas here are relevant to your organization, we would welcome the conversation.

Want to go deeper? Talk with our enterprise imaging team about openness, data independence, and interoperability at mach7t.com.

Glossary: acronyms you will hear in week one

Term	Meaning
DICOM	The imaging standard: objects plus networking
PACS	Picture Archiving and Communication System
VNA	Vendor Neutral Archive
RIS	Radiology Information System
EHR / EMR	Electronic Health / Medical Record
HL7 / FHIR	Health data messaging standards
IHE	Integration profiles and Connectathons
MWL	Modality Worklist
AE title	A DICOM node's network identity
SOP class	A DICOM object-type definition
UID	Unique identifier gluing objects together
GSPS	Grayscale Softcopy Presentation State
SR	Structured Report
MPR / MIP	Multi-planar / maximum-intensity reconstructions
KOS	Key Object Selection (key images)
WSI	Whole-Slide Imaging (pathology)
OCT	Optical Coherence Tomography (retina)
SaMD	Software as a Medical Device

Imaging Informatics 101: The Complete Guide | Mach7 Technologies | For educational use.