
CERF as an Enterprise Document Management System for Regulated Scientific Discovery

Why 21 CFR Part 11 Compliant Records Management Matters — and Why General-Purpose Tools Such as SharePoint Are Inadequate in Stringent Discovery Settings

Lab-Ally LLC · Columbus, Ohio · White Paper, 2026

Abstract

Scientific organizations operating under 21 CFR Part 11, Good Laboratory Practice (GLP), Good Manufacturing Practice (GMP), and analogous regulatory frameworks face an enduring tension between the practical convenience of general-purpose collaboration platforms — most prominently Microsoft SharePoint, shared network drives, and consumer cloud file services — and the strict evidentiary requirements of regulated record-keeping. This article examines the regulatory and scientific rationale for adopting a purpose-built, compliant Enterprise Document Management System (EDMS), surveys the structural deficiencies that render general-purpose tools unsuitable for stringent discovery settings, and describes the architecture and compliance features of CERF (Collaborative Electronic Research Framework), an integrated Electronic Laboratory Notebook and EDMS designed from inception to meet the requirements of 21 CFR Part 11 and the ALCOA+ data integrity principles. Particular attention is paid to CERF's immutable audit trail, cryptographically anchored digital signatures, permanent version retention, side-by-side file differencing (DIFF) and granular role-based access control, as well as content indexing for search and inline display of a broad range of common and specialized file formats.

Keywords: 21 CFR Part 11; ALCOA+; electronic laboratory notebook; enterprise document management; scientific data integrity; GLP; GxP; electronic signatures; audit trail; version control; SharePoint alternative.

1. Introduction

The defensibility of a scientific result rests not only on the quality of the underlying experiment but also on the integrity, completeness, and long-term accessibility of the records that document it. In commercial drug discovery, pre-clinical research, contract research, histopathology, medical device development, or indeed, any environment regulated by the United States Food and Drug Administration (FDA) or comparable agencies, those records must satisfy a specific and rigorous set of statutory and regulatory criteria. Among the most consequential of these is Title 21 of the Code of Federal Regulations, Part 11 (21 CFR Part 11), which governs the use of electronic records and electronic signatures when used in lieu of paper. Comparable expectations are codified in EudraLex Annex 11 in the European Union, in the OECD Principles of Good Laboratory Practice, and in the data integrity guidance published by the United Kingdom Medicines and Healthcare products Regulatory Agency (MHRA), the World Health Organization (WHO), and the Pharmaceutical Inspection Co-operation Scheme (PIC/S).

Despite the maturity and consistency of these expectations, a substantial fraction of regulated laboratories continue to manage their scientific and quality documentation using software that was never

engineered for evidentiary recordkeeping. Microsoft SharePoint, in particular, has become the de-facto document repository in many life-sciences organizations not because it is well-suited to the task, but because it is bundled with infrastructure the organization already owns. Shared network drives, Dropbox, Google Drive, OneDrive, and a long tail of departmental tools complete the picture. The aggregate result is an ecosystem in which critical experimental data are dispersed across systems that lack a coherent audit trail, do not enforce permanent version retention, cannot index the content of or display specialized file formats, and cannot produce the structured, inspection-ready record packages that an auditor will request.

This article makes the case that the appropriate response is not to layer ever more elaborate procedural controls and bespoke add-ons on top of general-purpose tools, but rather to adopt a purpose-built compliant EDMS. Using CERF (Collaborative Electronic Research Framework) — a 21 CFR Part 11 compliant ELN and EDMS developed and supported in the United States by Lab-Ally LLC — as the worked example, this article describes what such a system must do, how it should do it, and what is gained by doing so. The discussion is intended for the practicing scientist, the laboratory manager, and the quality-assurance professional rather than for the software engineer; technical detail is provided only where it is needed to make the regulatory or scientific argument intelligible.

2. The Regulatory Imperative: 21 CFR Part 11 and ALCOA+

2.1 Origins and Scope

Promulgated in March 1997 by the United States FDA, 21 CFR Part 11 was the first comprehensive regulatory framework to recognize electronic records and electronic signatures as legally equivalent to their paper and ink counterparts. Its provisions apply to any record required to be maintained under FDA regulation, and to any signature required by such regulation, when those records and signatures are created, modified, maintained, archived, retrieved, or transmitted electronically. In practical terms, virtually every preclinical study record, GMP manufacturing batch record, design history file, validation record, training record, and standard operating procedure (SOP) acknowledgment generated in a regulated environment falls within its scope.

The rule distinguishes between *closed systems* — those under the control of the organization that uses them — and *open systems* — those in which the persons responsible for the records do not control system access. The technical and procedural controls required for open systems are more demanding, but the underlying objective is identical in both cases: the records must be trustworthy, reliable, and equivalent to paper records as evidence.

2.2 Core Technical Requirements

Several of Part 11's provisions translate directly into hard requirements for any system that aspires to host regulated records. They are summarized below in the language of the rule itself, with explanatory commentary.

§ 11.10(a) — **Validation.** Systems must be validated to ensure accuracy, reliability, consistent intended performance, and the ability to discern invalid or altered records. Validation is not a one-off activity; it is a continuous discipline that includes Installation Qualification (IQ), Operational Qualification (OQ), and Performance Qualification (PQ), each documented and traceable.

§ 11.10(b) — **Accurate and complete copies.** The system must be able to generate accurate and complete copies of records in both human-readable and electronic form suitable for inspection, review, and copying by the agency. In practice this means that an inspector must be able to obtain a self-contained, intelligible record package without depending on the vendor’s proprietary tooling.

§ 11.10(c) — **Protection of records.** Records must be protected to enable their accurate and ready retrieval throughout the records-retention period. This implies durable storage, defensible backup, and the absence of silent data loss across software upgrades.

§ 11.10(d) — **Limited access.** System access must be limited to authorized individuals. This is more than password protection: it requires granular, role-based, and auditable access enforcement at the level of the individual record.

§ 11.10(e) — **Audit trail.** The system must employ secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records. Record changes must not obscure previously recorded information. Crucially, audit trail documentation must be retained for at least as long as the underlying record and must be available for agency review and copying.

§ 11.50 and § 11.70 — **Signature manifestation and binding.** Signed electronic records must contain information that clearly indicates the printed name of the signer, the date and time of signing, and the meaning of the signature. The signature itself must be linked to its respective record to ensure that it cannot be excised, copied, or otherwise transferred to falsify a record by ordinary means.

§ 11.200 — **Signature uniqueness and non-repudiation.** Electronic signatures must be unique to one individual and must not be reused by, or reassigned to, anyone else. Signatures based on biometrics must be designed to ensure that they cannot be used by anyone other than their genuine owners.

2.3 ALCOA and ALCOA+

Regulators have increasingly framed Part 11 expectations in terms of a mnemonic originally introduced by Stan Woollen of the FDA in the 1990s and since extended by the MHRA, WHO, and PIC/S. Data — whether on paper or electronic — must be:

Attributable to the individual who generated them; Legible and permanently readable; Contemporaneous with the act they describe; Original, or a verified true copy; and Accurate. The four extensions known as ALCOA+ add that data must additionally be Complete, Consistent, Enduring, and Available. These nine attributes are the operational vocabulary in which contemporary regulators evaluate

data integrity, and the gap between general-purpose document management and a compliant EDMS is most easily appreciated by considering how each tool satisfies — or fails to satisfy — each criterion.

3. The Failure of General-Purpose Tools in Stringent Discovery Settings

3.1 Microsoft SharePoint and the Illusion of Compliance

Microsoft SharePoint is, in the broad sense, an extraordinarily capable collaboration platform. It is also the single most common platform on which regulated laboratories attempt — and fail — to host compliant records. The reasons are structural rather than incidental, and they cannot be remedied by configuration alone.

First, SharePoint’s audit log is administratively mutable. A tenant administrator may disable auditing, alter the retention period, or purge log entries; the actions of administrators are themselves logged, but the log is not cryptographically anchored to the records it describes, and a sufficiently privileged user can suppress entries before they are exported. This stands in direct opposition to § 11.10(e), which requires that the audit trail be secure, computer-generated, and that record changes do not obscure previously recorded information. A laboratory whose audit trail can be silently truncated by an IT administrator does not, in any meaningful sense, have an audit trail.

Second, the “approve” action in a SharePoint workflow is not an electronic signature in the sense intended by § 11.50 or § 11.70. It is a Boolean field on a list item. There is no cryptographic binding between the approver’s identity and the bytes of the approved document; if the document is subsequently modified — by means as innocuous as a content type migration — the approval persists without any indication that the underlying content has changed. True compliance requires that the signature be cryptographically computed over the content of the signed record, so that any alteration of the content invalidates the signature mathematically, not merely procedurally.

Third, SharePoint’s version history can be configured to retain only a limited number of versions, and even when configured to retain all versions, those versions can be purged by an administrator. Some Microsoft 365 retention features place legal-hold-like constraints on deletion, but their scope and behaviour have changed repeatedly across Microsoft service updates. A regulated record must remain available unchanged for the full retention period; building that guarantee on top of a platform whose behaviour shifts at the vendor’s discretion is a brittle proposition.

Fourth, SharePoint exposes only a coarse-grained file preview. It can render Microsoft Office documents and PDFs respectably; it cannot display the file formats that scientists actually generate. DICOM medical images, Nikon ND2 microscopy stacks, OME-TIFF whole-slide images, ChemDraw structures (CDX, CDXML), Chemical Markup Language (CML), MOL/SDF/RXN, GenBank and FASTA biological sequences, ABI trace files (.ab1), MATLAB .mat files, SAS .sas7bdat datasets, Photoshop PSD images, and the many instrument-specific formats produced by mass spectrometers, plate readers, and sequencers are either unreadable or shown as opaque binary attachments. Scientists end up downloading files to local

workstations in order to inspect them, which fragments the audit trail, multiplies the risk of uncontrolled copies, and renders centralized search effectively meaningless.

Fifth — and this is the criticism most frequently raised by quality professionals who have lived through the experience — SharePoint has no native, defensible mechanism for comparing any two selected versions of a scientific file. A reviewer who needs to verify that the only changes between SOP version 3.2 and version 3.3 are the changes that were approved by the change-control board has, in SharePoint, no recourse beyond a manual visual comparison. For images — microscopy images, gels, blots, histology slides, instrument screenshots — there is no mechanism at all. The result is that version control in SharePoint provides traceability of *that* a change occurred but not of *what* changed.

Sixth, granular access control in SharePoint is implemented by means of permission inheritance trees that are notoriously difficult to maintain across reorganizations, project boundaries, and personnel changes. The specific requirement that an unauthorized user not see, in search results, evidence that a restricted record exists is met in SharePoint only with considerable configuration effort, and is easily broken by routine administrative actions.

Seventh and finally, validating SharePoint for 21 CFR Part 11 use is a famously expensive undertaking that must be repeated, at least in part, with each Microsoft 365 service update. Many organizations have done this work; many have done it more than once; few would describe the outcome as stable. The combined cost of initial validation, ongoing re-validation, third-party compliance add-ons, and the bespoke procedural controls required to compensate for residual gaps frequently exceeds the cost of deploying a purpose-built EDMS in the first place.

3.2 Shared Drives, Consumer Cloud, and Paper

The failure modes of shared network drives, Dropbox, Google Drive, OneDrive, and analogous consumer-grade tools are essentially those of SharePoint, only more severe. None offers a computer-generated audit trail; none offers cryptographic signature binding; none enforces permanent version retention; none indexes specialized file formats; none can produce an inspection-ready record package. Paper notebooks, although still in use in many academic and small-biotech settings, satisfy the original ALCOA criteria reasonably well but are increasingly impractical: they cannot be searched, cannot be shared in real time across multiple sites, cannot host the data files that instruments produce, and are vulnerable to physical loss in ways that a properly backed-up electronic system is not.

4. The Anatomy of a Compliant Enterprise DMS: CERF as a Worked Example

The deficiencies just enumerated define, by negation, the feature set that a compliant enterprise DMS for scientific use must offer. The remainder of this article describes how CERF realizes each of those features. CERF was first developed in the early 2000s and has been continuously developed and supported by Lab-Ally LLC since 2016. The system has been deployed in pharmaceutical research organizations, contract research organizations, histopathology laboratories, medical device manufacturers, government

laboratories, and academic groups for nearly two decades. During that time many other systems have come and gone, leaving CERF as the prima facie most sustainable and long-lived solution of its kind. CERF 6, released in 2026, is built on Java 25, Apache Tomcat 11, MySQL 9 (or Microsoft SQL Server or Oracle), and a curated stack of open-source scientific libraries; this is mentioned not as a software-engineering boast but because the use of mainstream, open-source, long-lived components is itself a feature of a system intended to outlast the organization that purchased it.

4.1 An Immutable, Computer-Generated Audit Trail

Every action that occurs within CERF — the creation of a record, an edit, a check-out, a check-in, a metadata change, a successful or failed login attempt, the granting or revoking of a permission, the application of a signature, the export of a record — is captured by the server in a time-stamped log that cannot be modified or deleted from within the application. The audit trail is computer-generated, not user-generated; there is no facility, even for the system administrator, to alter a log entry after the fact. When a regulatory inspector requests the complete history of an SOP or a notebook page, CERF produces it as a structured report with a single command. This is § 11.10(e) realized as a system invariant rather than an operational policy.

4.2 True PKI Digital Signatures Anchored to Record Content

CERF is, to our knowledge, the only commercially available ELN and EDMS that implements electronic signatures using genuine Public Key Infrastructure (PKI) cryptography by means of the U.S. federal Digital Signature Algorithm (DSA). When a CERF user account is created, the server generates a unique cryptographic key pair that is bound to that user. When the user signs a record, the act of signing computes an MD5 content digest and applies the user’s private key to produce a signature that is mathematically bound to the bytes of the signed content. Any subsequent alteration of the content — by any user, by any means, including direct manipulation of the underlying file store — invalidates the signature in a manner that is immediately detectable. The signature cannot be excised, copied to another record, or reassigned to another user; the cryptographic binding is the assurance, and that assurance is mathematical rather than procedural.

The CERF signature subsystem is configurable in a manner that reflects the real organizational realities of regulated science. Signing actions — *Submit*, *Co-author Submit*, *Witness*, *Peer Review*, *Manager Approval*, *Legal/Regulatory Sign-off*, and others — are defined declaratively in an ontology file and may be extended or amended without altering CERF source code. Each signing action carries a configurable endorsement statement (for example, “*As part of my submission of this document, I am signing the document in compliance with 21 CFR Part 11 regulations*”) that is recorded together with the signature itself. Signature workflows — for example, a four-stage Submitter → Peer Reviewer → Manager → Legal sequence — may be defined per workgroup, and CERF refuses to allow a notebook page to be signed unless all required SOP placeholders within that page have been satisfied with the currently effective controlled-

document version. Non-compliant submissions are not flagged as exceptions; they are structurally impossible.

4.3 Permanent Version Retention and Side-by-Side Differencing

CERF treats every check-in of a version-controlled resource as the creation of a new, numbered version (1.0, 1.1, 1.2, 2.0, and so on). No version is ever deleted, neither by the originating user, by a workgroup administrator, nor by the CERF system administrator. The platform supports linear histories and explicit branching for cases in which an older version must form the basis of an alternative line of development; in such cases, the branch is itself recorded in the audit trail. Reverting to an earlier version does not delete intervening versions; it creates a new version whose content matches the earlier one, preserving the complete history of edits between the two states. These behaviors collectively satisfy the ALCOA+ requirements of *completeness*, *consistency*, *enduringness*, and *availability* while remaining intuitive to scientific users who are not specialists in document control.

What distinguishes CERF most sharply from SharePoint and similar systems is its built-in differencing tool, known to CERF users simply as the DIFF viewer. Any two versions of a textual document can be compared in a side-by-side display in which inserted, deleted, and modified text is visually highlighted, with navigation arrows that step through detected changes one at a time. The same DIFF viewer operates on images. Two versions of a histology slide, a Western blot, a gel image, an instrument screenshot, or any other raster image can be compared with a sensitivity slider that allows the reviewer to suppress trivial differences (such as those introduced by lossy recompression) and surface substantive differences (such as the addition or removal of a region of interest). For a regulated laboratory, the DIFF tool transforms version review from a manual, error-prone exercise into a routine, defensible activity. We are not aware of any general-purpose document management platform that offers comparable functionality for scientific image data.

CERF additionally supports three distinct edit states for any resource: *Version Controlled* (the default, in which every check-in creates a new permanent version), *Versionable* (in which the user may choose to overwrite the current version, intended for active drafts and explicitly disabled for records subject to regulated retention), and *Final* (in which further edits are prohibited entirely). The Final state is distinct from a digital signature; a record may be final without being signed, or signed without being final, and the two concepts compose cleanly to support the diverse review and lock-down conventions of different regulated environments.

4.4 Granular Role-Based Access Control and Workgroup Isolation

Access in CERF is organized around the concept of the Workgroup, which represents a project, a department, a client engagement, or any other operationally meaningful grouping of users and resources. Each Workgroup is the owner of its files and notebooks, and every user is assigned a role within each Workgroup of which they are a member. CERF defines nine distinct access roles — including Read Only, Annotator, Editor, Notebook Editor, Manager, and Workgroup Administrator — each conferring a precise

and well-documented bundle of permissions. The same user may hold different roles in different Workgroups, which is essential for contract research organizations in which a scientist may be a Manager on one sponsor's project and a Read-Only collaborator on another. Permissions are enforced at the record level rather than at the folder level: an unauthorized user cannot see, even in search results, that a restricted record exists.

Capabilities — distinct from access roles — provide a finer-grained delegation mechanism for specific administrative privileges such as the ability to create new Workgroups, to manage controlled vocabularies, or to manage signature groups. Capabilities are granted and revoked by the system administrator and are themselves recorded in the audit trail. Session security is reinforced by single-session enforcement (a user cannot be logged in concurrently from multiple workstations), configurable session timeout, account lockout policies, and optional Time-based One-Time Password (TOTP) multi-factor authentication using any RFC 6238-compliant authenticator application.

4.5 SOP and Controlled Document Lifecycle Management

CERF includes a dedicated module for the lifecycle management of controlled documents — SOPs, protocols, assay procedures, training materials, non-disclosure agreements, and any other documents that must be approved before release, distributed to a defined audience, acknowledged by that audience, periodically reviewed, and ultimately retired. When a controlled document is released, CERF enforces a configurable acknowledgment workflow: users to whom the document is assigned receive a notification, must wait a configurable minimum review period, and must explicitly acknowledge the document before they are permitted to sign notebook pages that reference it. SOP expiry dates are monitored automatically, with configurable warning notifications delivered to managers and assigned users in advance of the expiry. The result is that a notebook page cannot be signed by a scientist who has not acknowledged the current version of every SOP it depends upon — another instance of structural rather than procedural compliance.

The training-record function deserves particular emphasis. A complete, time-stamped log of every user's acknowledgment of every version of every controlled document is maintained automatically and can be exported at any time for inspection. The single most common deficiency in regulated laboratories — the absence of contemporaneous evidence that a particular scientist had read the version of a particular SOP that was current at the time they performed a particular experiment — is addressed by CERF as a routine system operation rather than as a quality-assurance project.

4.6 Formal Metadata, Annotations, and Semantic Search

CERF distinguishes carefully between two categories of descriptive information that are commonly conflated in less rigorous systems. *Formal metadata* — added through the Edit Metadata dialog by users with the Metadata Editor role or higher — is part of the official, version-controlled, auditable record of the resource. Every change to formal metadata creates a new version, and an inspector may examine exactly what metadata were associated with any version of the resource at any point in time. *Annotations* — informal tags, notes, star ratings, and informal relations — are deliberately lightweight: they may be created

or removed by any Annotator without creating a new version, and they are intended for working notes, search aids, and exploratory categorization rather than for documentation of scientifically significant assertions. The distinction is enforced by the system rather than left to user discipline, and it preserves the integrity of the compliance record without depriving users of the convenience of informal tagging.

CERF's search engine indexes the full text of MS Office documents, PDFs, plain text, biological sequence files, chemistry structure files, and many scientific data formats, in addition to file properties, formal metadata, and annotations. Searches may combine any number of parameters — contributor, date range, resource type, tag, custom metadata field, signature status, controlled-vocabulary term — using Boolean AND/OR logic. Queries may be saved as reusable smart cohorts; results may be sorted and filtered by any column. A distinctive *Find Experts* facility identifies colleagues who have contributed to resources matching the user's query, an unusually useful capability for large, distributed organizations seeking subject-matter expertise that has not been formally catalogued.

4.7 Native Display and Indexing of Specialized Scientific File Formats

A central and frequently underappreciated advantage of CERF over general-purpose tools is the breadth of file formats it can display and index without recourse to external applications. The following summary is illustrative rather than exhaustive.

File category	Representative formats displayed and indexed in CERF
Microsoft Office	Word (.doc, .docx), Excel (.xls, .xlsx), PowerPoint (.ppt, .pptx); Indexed and rendered as inline previews.
Apple iWork	Pages, Numbers, Keynote (.pages, .numbers, .key); Indexed and rendered as inline previews.
Adobe and general documents	PDF, RTF, plain text (.txt), HTML, Markdown. Indexed and rendered as inline previews.
Raster and microscopy images	JPEG, PNG, TIFF, GIF, BMP, PSD (Photoshop), JP2 (JPEG 2000), OME-TIFF, displayed with EXIF metadata extracted and indexed. Highly specialized images can be displayed using a feature called 'Official Print Copy'.
Vector graphics	SVG rendered natively at any zoom level.
Chemistry structures	CDX and CDXML (ChemDraw), MOL/MOL2, SDF, RXN, RDF, CML, PDB, MRV, XYZ, SMILES, SMARTS, InChI, and others — over twenty-four chemistry formats in total.
Biological sequences	GenBank (.gb), FASTA (.fa, .fasta), ABI trace files (.ab1), CLUSTAL alignments, Phylip (.phy). Most are indexed, all of these are displayable by CERF.
Statistical and engineering data	MATLAB .mat, SAS .sas7bdat, Microsoft Access .mdb and .accdb, Outlook .pst archives. Can be displayed using a feature called 'Official Print Copy'.

File category	Representative formats displayed and indexed in CERF
Technical/source files	Plain-text viewable: .bat, .css, .ini, .js, .md, .php, .py, .sh, .sql, and many others — relevant for bioinformatics, statistics, and software-adjacent laboratories. Indexed and rendered as inline previews.

Table 1. Selected categories of scientific and office file formats supported by CERF for native inline display and full-text or metadata indexing. For files whose format is too specialized to render natively (certain proprietary instrument outputs, for example), CERF provides the *Official Print Copy* mechanism, which permits the user to designate a separate, viewable, printable surrogate (typically a PDF or image) that is bound to the original file and that serves as the human-readable version of record for inspection and signature purposes. Original files are preserved bit-for-bit in their native form.

The practical consequence is profound. A reviewer who must examine a ChemDraw structure, a chromatogram, a sequence alignment, and a histology image, all referenced from the same notebook page, can do so within a single CERF window without installing or licensing any specialized viewer. The full-text index encompasses the textual content of Word and PDF files, the residue sequence of FASTA and GenBank files, the SMILES string of a chemistry file, and many other format-specific content channels, so that a single search query reliably surfaces every record in which a specific gene, protein, compound, instrument identifier, study code, or other term of interest appears, regardless of the file format in which that term is encoded.

4.8 Round-Trip Editing and Automated Chain of Custody

CERF supports the familiar check-out / edit / check-in workflow: a user checks out a file, edits it in their preferred desktop application (Word, Excel, ChemDraw, ImageJ, or any other), and checks it back in. CERF stores the new version, retains the previous version permanently, and records the entire transaction in the audit trail. Check-out may be performed offline — for instance, before travelling to a location without network access — and the system reconciles the resulting versions automatically on reconnection.

For automated data capture, CERF includes a component called the CERF Automaton, which monitors designated network locations and ingests files into CERF as they are produced by instruments, scripts, or upstream systems. Ingestion is configured through a graphical interface; no programming is required for typical use cases. The CERF Email-to-CERF facility permits any user to forward files and attachments into CERF by email, with automatic routing to the correct workgroup and automatic attribution. Both mechanisms preserve the chain of custody from the moment of file generation, which is precisely the requirement that general-purpose tools find most difficult to honor.

4.9 Inspection-Ready Export and Data Sovereignty

When an inspector or sponsor requests a complete record package, CERF produces it as a single export operation. The export comprises the original native files (not transcoded or repackaged), the full version history of each file, all formal metadata, all annotations, the complete audit trail of every transaction touching the exported records, the signature manifest, and a structured XML manifest describing the

relationships among the exported items. The CERF Exporter utility supports both single-record and bulk exports, and the resulting package is intelligible without CERF: a sponsor or regulator can read the native files in any platform-appropriate application and the XML manifest in any text editor or XML viewer.

The architectural commitment to native file storage — files are stored in their original format on a secure file store, not compressed into a proprietary container — is the foundation of CERF’s data-sovereignty story. The relational database (MySQL, Microsoft SQL Server, or Oracle, at the customer’s election) holds metadata, audit-trail records, signatures, and relationship information, but the underlying scientific files are always recoverable in their native form. This eliminates the vendor-lock-in failure mode that has bedeviled many earlier scientific data systems: the customer’s data are never hostage to the continued operation of the vendor.

4.10 Flexible Deployment and Long-Term Sustainability

CERF can be deployed on a cloud platform of the customer’s choice (AWS, Azure, Google Cloud, or any other), on a private cloud instance managed by Lab-Ally, on the customer’s own on-premise infrastructure, or on a completely sealed local-area network without any internet connectivity. The last of these configurations is increasingly relevant for laboratories operating under defense, intelligence, or export-controlled-research constraints. CERF is database-agnostic — MySQL is included as the default but Microsoft SQL Server and Oracle are fully supported — and the entire stack is built on open-source components selected for technological longevity. A perpetual-license option is offered for customers who wish to ensure that their system continues to operate indefinitely, irrespective of any future change in the customer’s relationship with Lab-Ally.

Lab-Ally and its qualified validation partners provide Installation Qualification (IQ) and Operational Qualification (OQ) services as part of standard deployment for GLP, GMP, and 21 CFR Part 11 environments. These services are tailored to the customer’s specific configuration, in accordance with the FDA expectation that validation reflects the system as actually deployed rather than the vendor’s generic baseline.

5. Summary Comparison: CERF and General-Purpose File Management

The discussion of the preceding section may be condensed into the comparison shown in Table 2. The contrast does not arise from any particular failing on Microsoft’s part; SharePoint was simply not designed to host regulated scientific records, and adapting it to that purpose by configuration alone is not feasible.

Requirement	CERF ELN	SharePoint / shared drive / consumer cloud
Computer-generated, immutable audit trail (§ 11.10(e))	Built in; cannot be altered from within the application.	Mutable by administrator; logs may be truncated or disabled.

Requirement	CERF ELN	SharePoint / shared drive / consumer cloud
Cryptographic binding of signature to record content (§ 11.50, § 11.70)	PKI/DSA signatures with per-user key pairs; tampering invalidates the signature mathematically.	“Approve” field has no cryptographic binding to file content.
Permanent version retention	No version is ever deleted; reverting creates a new version.	Configurable, mutable, frequently capped and subject to administrative purge.
Side-by-side DIFF for documents and images	Built-in; navigable diff with sensitivity control for raster images.	Not provided natively.
Granular access enforcement in search results	Record-level enforcement; unauthorized users do not see that a record exists.	Achievable only with elaborate permission configuration; easily broken.
Native display of specialized scientific formats	DICOM, ND2, OME-TIFF, ChemDraw, CML, GenBank, FASTA, AB1, MATLAB, SAS, and many others.	Microsoft Office and PDF only; scientific files appear as opaque attachments.
SOP/controlled-document lifecycle with enforced acknowledgment	Built-in module with review periods, training records, and expiry monitoring.	Requires third-party add-on; behaviour shifts with vendor service updates.
Inspection-ready record export (native files plus XML manifest)	Single-command export of records, history, audit trail, and metadata.	Manual assembly; no integrated manifest.
Sealed-LAN deployment	Supported and routinely used for restricted environments.	Not feasible for cloud-tier services.
Long-term data sovereignty	Native file storage; perpetual license option; open-source component stack.	Data and access dependent on continued vendor service.

Table 2. Selected compliance and functional requirements of a regulated scientific DMS, contrasted between CERF and a representative general-purpose tool such as Microsoft SharePoint. The intent is illustrative; a full requirements matrix mapped to specific 21 CFR Part 11 sub-sections is maintained separately by Lab-Ally and is available on request.

6. Practical Implications for Researchers and Organizations

6.1 Intellectual-Property Defensibility and Patent Support

Even outside formally regulated environments, the discipline imposed by a compliant EDMS confers material benefits on academic and early-stage-biotech laboratories. The single most common point of failure in patent prosecution and inventorship disputes is the absence of contemporaneous, attributable, immutable records of conception and reduction to practice. A CERF notebook page, digitally signed at the time of the work and witnessed by a peer, constitutes precisely the kind of evidence that patent counsel and licensing partners require. Many of the small academic groups that have adopted CERF report that the

system pays for itself the first time it is used to defend a publication priority claim or to support a patent filing.

6.2 Beyond Pharma and Histopathology: Adjacent Use Cases

The features that make CERF suitable for FDA-regulated laboratories also make it suitable for any organization that must maintain auditable, long-lived, defensible documentary records under stringent access controls. Legal departments managing matter files, deal teams handling due-diligence packages, project-management offices retaining design rationale across multi-year programs, and government laboratories operating under classified or controlled-unclassified-information rules have all been served by deployments of CERF. The system is discipline-agnostic in its evidentiary mechanics; what distinguishes its scientific deployments is principally the breadth of scientific file formats it understands.

6.3 Deployment Effort and Total Cost of Ownership

A common objection to dedicated EDMS adoption is the perceived implementation burden. In practice, a basic CERF deployment is operational within days; a fully validated enterprise deployment with customized workgroup configuration, signature workflows, and IQ/OQ documentation typically requires between four and twelve weeks, depending on organizational complexity. Per-user pricing is comparable to or lower than the per-user cost of SharePoint with the third-party add-ons required to approximate compliance, and the perpetual-licence option permits the total cost of ownership to be known in advance rather than escalating with each renewal cycle.

7. Conclusion

21 CFR Part 11 and the ALCOA+ data-integrity principles articulate, in mature and operationally well-understood language, what a regulated scientific record must be: attributable, legible, contemporaneous, original, accurate, complete, consistent, enduring, and available, with a computer-generated audit trail, cryptographically bound electronic signatures, permanent version retention, and granular access control. These requirements are not aspirational; they are the minimum threshold for evidentiary admissibility in a regulatory inspection.

General-purpose collaboration tools — Microsoft SharePoint chief among them — were not engineered to satisfy this threshold, and the cost of attempting to make them do so by configuration and procedural overlay is high, the result fragile. A purpose-built compliant EDMS, by contrast, satisfies the threshold by construction. CERF — through its immutable audit trail, PKI-anchored signatures, permanent version history, side-by-side DIFF tool for both documents and images, granular workgroup access controls, dedicated SOP lifecycle module, native indexing and inline display of specialized file formats, automated capture and chain-of-custody facilities, inspection-ready exports, and flexible deployment options — illustrates what such a system looks like in practice. For any laboratory whose data must withstand scrutiny by a regulator, a sponsor, a patent examiner, or a future scientist attempting to reproduce a result a decade hence, the case for a compliant EDMS is not a marketing claim; it is a precondition for doing the work.

Further Reading and References

- [1] U.S. Food and Drug Administration. *Title 21 Code of Federal Regulations Part 11 — Electronic Records; Electronic Signatures*. Final rule, 1997; current edition.
- [2] U.S. Food and Drug Administration. *Data Integrity and Compliance With Drug CGMP: Questions and Answers — Guidance for Industry*. December 2018.
- [3] European Medicines Agency. *EudraLex Volume 4, Annex 11: Computerized Systems*. 2011.
- [4] Medicines and Healthcare products Regulatory Agency (UK). *'GXP' Data Integrity Guidance and Definitions*. March 2018.
- [5] World Health Organization. *Guideline on Data Integrity*. WHO Technical Report Series, Annex 4. 2021.
- [6] PIC/S. *Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments (PI 041-1)*. 2021.
- [7] Organization for Economic Co-operation and Development. *OECD Principles of Good Laboratory Practice*. ENV/MC/CHEM(98)17, as revised.
- [8] Lab-Ally LLC. *CERF ELN as a Document Management System*. Available at <https://cerf-notebook.com/document-management.html>
- [9] Lab-Ally LLC. *What is CERF?* Available at <https://cerf-notebook.com>

Correspondence and product enquiries: Lab-Ally LLC, 247 E. 9th Avenue, Columbus, Ohio 43201, USA. Telephone +1 (614) 407-4547. Web: lab-ally.com and cerf-notebook.com. *To request a live demonstration of CERF or a free evaluation deployment, please contact info@lab-ally.com.*