Data Integrity as Enterprise Risk: What the 2026 Warning Letter Surge Means for the Boardroom

A cGMP Consulting White Paper on Governance, Detection, and Remediation in FDA-Regulated Manufacturing – September 3, 2026

Pharmaceutical and Biologics Manufacturing | Quality Systems | Data Governance

An educational resource for executive and technical leaders in FDA-regulated life sciences.

Executive Summary

Data integrity used to be something the quality laboratory worried about. It isn't anymore. It now reaches the chief executive, the board, and the balance sheet. In fiscal year 2025 the FDA issued a sharply increased number of drug quality warning letters, and the first quarter of 2026 brought a run of cases citing deliberate falsification: backdated documents, discarded batch records, and in at least one instance records destroyed while investigators were standing in the building. Call these what they are. Not documentation errors. They are the kind of failure the agency now treats as a reason to distrust every quality decision a manufacturing system produces.

The logic behind that is simple, and it is brutal. Once the FDA decides a facility cannot demonstrate its records are reliable, it can classify the products made there as adulterated under the Federal Food, Drug, and Cosmetic Act, even when nobody has found a speck of contamination. The finding rests on lost confidence in the data, not on a proven defect in the drug. For a commercial-stage company that reasoning turns one integrity lapse into an enterprise event: import alerts, consent decrees, distribution holds, investors heading for the exits.

Here is the structural problem almost everyone shares. Internal monitoring programs tend to confirm that audit trails exist. They rarely confirm anyone reads them. And the quality signals that would warn you almost never reach the executives who actually own enterprise risk until an inspection has already pried the gap open. So leadership learns about a systemic failure at the exact moment it becomes a public, enforceable liability. The failures that hurt most are not the ones you catch. They are the ones the FDA catches first.

This paper walks through the regulatory framework behind these expectations, the enforcement trends shaping inspections in 2026, the operating conditions that quietly build exposure, and what poor implementation actually costs. Then it lays out a program that works: real audit trail review, an executive reporting line for integrity risk, investigations and corrective actions that mean something, and data governance sized to the risk each record carries. Treat data integrity as an enterprise governance function instead of a quality checklist and you get two things at once, a defensible posture under inspection and an operation that makes better decisions.

Introduction

Pharmaceutical and biologics operations generate staggering volumes of data across research, development, manufacturing, laboratory testing, and distribution. Every record backs a decision about safety, identity, strength, quality, or purity. The reliability of those decisions rests entirely on the reliability of the data underneath them. When the data is trustworthy, the quality system does its job. When it isn't, every downstream conclusion is suspect and the whole system loses its evidentiary footing.

For years the industry filed data integrity under good documentation practice, handed it to quality assurance, and treated it as a technical compliance matter. That framing is out of step with where enforcement has gone. The FDA has made data integrity a sustained priority, and the language in recent warning letters gives it away. The agency isn't citing a stray batch record error anymore. It is increasingly asserting that a firm's quality system does not adequately ensure the accuracy and integrity of the data supporting the drugs it makes. That phrasing is chosen on purpose. It tells you the agency wants evidence of an examined and restructured quality culture, not a tidy fix to the one item it named.

Several forces pushed the stakes up together. Global supply chains multiplied the number of contract manufacturers and testing labs whose data a sponsor has to trust. Digitalization raised both the volume of electronic records and the sophistication needed to govern them. Inspections got more data-driven, with investigators picking through metadata, audit trails, and the timing of when records were created. And enforcement outcomes got more severe, with consent decrees historically running into the hundreds of millions in remediation. So the question for executives is not whether data integrity matters. It is

whether your organization can prove it under inspection conditions, on a bad day, with an investigator in the room.

This paper is written for the executives and technical leaders who have to answer that question. It assumes you know good manufacturing practice and concentrates on the strategic, organizational, and operational calls that decide whether a data integrity program actually holds.

Industry Background

The legal foundation here is old. Requirements for accurate manufacturing and laboratory records have existed in one form or another since the Federal Food, Drug and Cosmetic Act took shape in the mid twentieth century. What changed is the intensity of the attention and the technical complexity of the systems that now create and store the data.

The modern data integrity conversation took off after a string of inspections exposed patterns of misconduct that gutted confidence in submitted and manufacturing data. Investigators found audit trails switched off right when they showed up. Unofficial trial testing of samples. Failing results quietly discarded. Testing repeated until something passed. Records signed by people who never did the work. Those findings pushed regulators worldwide to publish guidance spelling out how data should be created, reviewed, and retained.

Out of that came a shared vocabulary, the ALCOA framework: data should be Attributable, Legible, Contemporaneous, Original, and Accurate. Regulators later leaned on the extended set, ALCOA+, adding Complete, Consistent, Enduring, and Available. These attributes run across the whole data lifecycle, from creation and modification through processing, review, retention, retrieval, and eventual disposition. They apply to paper and electronic records alike, and they are the practical yardstick an inspector uses to judge whether a record can be believed.

The technical side has matured too, and this is where many programs fall behind. Today's manufacturing and laboratory environments run on computerized systems that produce dynamic records, generate metadata, and keep audit trails. Governing them means understanding user access controls, electronic signatures, time and date controls, backup and archival practice, and the difference between a static record like a printed report and a dynamic record that can be reprocessed or reinterpreted. The maturity of a data integrity program is now mostly a question of how well it governs those computerized systems, not how neatly it controls paper forms.

Regulatory Landscape

Data integrity expectations sit on binding regulation and get elaborated through guidance from several authorities. The distinction matters, because the FDA draws its enforcement authority from the regulations and the underlying statute, while the guidance documents only explain how the agency currently reads them.

United States: 21 CFR Parts 210, 211, and 11

In the United States, current good manufacturing practice for finished pharmaceuticals lives in Title 21 of the Code of Federal Regulations, Parts 210 and 211. Plenty of provisions establish data integrity obligations without ever using the phrase. Records must be reviewed for accuracy, completeness, and compliance with established standards. Production and control records must be reviewed before a batch goes out. Laboratory records must document the complete data from all tests. The review and verification provisions are scattered across sections 211.22, 211.68, 211.100, 211.160, 211.180, 211.182, 211.186, 211.188, 211.192, and 211.194, among others. Electronic records and electronic signatures fall under 21 CFR Part 11, which governs records kept in electronic form that are otherwise required by the predicate rules in Parts 210, 211, and 212.

The FDA laid out its interpretation in the guidance for industry titled Data Integrity and Compliance With Drug CGMP: Questions and Answers, issued in December 2018. It defines the key terms, data integrity, metadata, audit trail, and answers the practical questions: how to restrict access to computerized systems, when workflows have to be validated, how often audit trails should be reviewed. One point deserves underlining because so many firms miss it. Audit trail review frequency should track the review frequency of the underlying data. Since production and laboratory records must be reviewed before batch release, the guidance indicates the relevant audit trails have to be reviewed before that batch is released. The same document also puts responsibility for a healthy quality culture squarely on management with executive responsibility, the people meant to make sure employees understand why integrity matters and feel able to raise concerns.

International Guidance: MHRA, PIC/S, WHO, and EMA

The United Kingdom's Medicines and Healthcare products Regulatory Agency published its GXP Data Integrity Guidance and Definitions in March 2018. It runs across the good practice disciplines, pushes a risk-based approach that weighs data criticality against inherent integrity risk, and expands the ALCOA attributes with real emphasis on record permanence. The Pharmaceutical Inspection Co-operation Scheme issued PI 041, Good Practices for Data Management and Integrity in Regulated GMP/GDP Environments, in 2021, to harmonize what inspectors expect across member authorities. The World Health Organization covered the ground in Technical Report Series No. 996, Annex 5, Guidance on Good Data and Records Management Practices, published in 2016. The European Medicines Agency has its own questions and answers on the subject. The documents differ in scope and detail, but they line up around ALCOA and ALCOA+ and a risk-based, lifecycle view of data governance.

The landscape keeps moving. The European Union has advanced a draft revision of EU GMP Chapter 4 on documentation that would write data governance and computerized system controls more explicitly into

binding regulation, and the EU GMP Annex 11 revision on computerized systems has moved toward finalization with expanded audit trail scope and multi-factor authentication expectations. If you operate across jurisdictions, build one harmonized control set. Treating each regulator as a separate project is how you end up doing the same work three times and still failing an inspection.

Enforcement Trends and Inspection Focus

The enforcement data backs up the priority. Analyses of recent warning letters keep finding data integrity deficiencies as a primary or contributing factor in a substantial majority of drug good manufacturing practice letters. In 2026 the frequent offenders are inadequate out-of-specification investigations, falsified or manipulated chromatographic data, shared or uncontrolled login credentials, and missing or unreviewed audit trails. Foreign sites, especially in certain regions with concentrated production, draw a lot of the attention, though domestic facilities are far from exempt. One thing stands out this year: the agency is willing to escalate fast, in some cases refusing to accept data from contract laboratories where it found integrity failures, which effectively pulls those labs out of every supply chain that leaned on them.

> *Inspection focus in 2026 has moved from confirming that audit trails are generated to confirming that they are reviewed, by whom, and how often. Generating the audit trail is now the baseline. Reviewing it before batch release is the control that actually protects product quality.*

Current Industry Challenges

Even firms that mean well struggle to put data integrity controls in place that survive an inspection. The trouble spreads across technology, process, people, and the partners you depend on.

The Audit Trail Review Gap

This is the big one, and the most common. There is a world of difference between generating audit trails and reviewing them. Plenty of organizations configured their systems to record audit trails years ago and mentally checked the box. Inspectors in 2026 ask a sharper question: who reviews it, how often, against what criteria, and what did they find. When a lab can produce an audit trail but cannot show that a qualified reviewer looked at it before releasing the result, the control has failed in substance even though it exists on paper. Reviewing audit trails at the frequency the underlying data demands, and documenting that you did, is a grind. Most organizations have not fully operationalized it, and they know it.

Computerized System Controls and Access Management

Shared login credentials keep showing up. When several people work under one account, you cannot pin an action on an individual, which knocks out a foundational ALCOA attribute and breaks the requirement that second-person review be attributable. The neighbors of that problem are just as familiar: administrative privileges broad enough to let an analyst alter or delete data, weak control over date and time settings, systems that let data be reprocessed without keeping the original. Legacy instruments and standalone systems that never had modern access controls are the hardest case, because retrofitting them is expensive and sometimes flat-out impractical.

Investigation Quality and the Out-of-Specification Problem

Out-of-specification investigation inadequacy is one of the most cited issues in current enforcement, and the failure mode repeats itself. A failing result gets pinned on operator error with no scientifically sound, documented justification, the result gets invalidated, and testing continues until something passes. Regulators treat invalidation as something that requires a valid, documented, scientifically sound justification, and they expect the full batch record handed to the quality unit to include the original, invalidated data next to the investigation. When investigations run as paperwork exercises instead of genuine root cause work, they manufacture data integrity exposure and flag a deeper weakness in the quality system.

Supply Chain and Contract Partner Oversight

Your release decisions are only as reliable as the data your contract manufacturers and testing labs generate. A contract laboratory can move through an inspection observation, an inadequate response, a warning letter, and an import alert over several months without ever picking up the phone to tell you. Firms that rely on annual audits and periodic questionnaires are structurally blind to the failures that happen and escalate between those touchpoints. And when a partner's release data later turns out to be falsified, the products you released on it are resting on unreliable evidence, which drags the exposure into material already in distribution.

Culture, Staffing, and Workload

Most data integrity failures trace back to conditions that make honest, contemporaneous record-keeping hard. Crushing workload. Too few people for quality oversight and investigations. Production pressure that raises both the odds of an error and the temptation to shade the data. Regulators expect management to allocate resources so the pressure on the people generating data does not become an opening for integrity problems. A culture where employees can raise concerns without fear is a control, not a slogan, and its absence is a well-documented precursor to systemic failure.

Risks of Poor Implementation

Weak data integrity bites across regulatory, financial, operational, and reputational lines. And because the agency treats lost data reliability as grounds to distrust the whole quality system, one serious finding can do damage wildly out of proportion to how narrow it looked at first.

Risk Category	Potential Consequence
Regulatory action	Form 483 observations, warning letters, import alerts, and in severe cases consent decrees that impose external oversight and remediation obligations.
Product disposition	Classification of product as adulterated based on unreliable records, leading to distribution holds, recalls, and rejected lots even without confirmed contamination.
Financial impact	Remediation costs historically reaching into the hundreds of millions of dollars for consent decrees, plus retrospective data reviews, third-party audits, and lost production.
Market and supply	Loss of market access through import refusal, supply interruptions to patients, and the inability to accept data from a blacklisted contract laboratory.
Investor confidence	Public warning letters and enforcement actions that erode investor confidence, hit valuation, and complicate financing, partnering, and M&A activity.
Patient safety	The underlying risk that unreliable data lets a substandard product reach patients, which is the harm all of these controls exist to prevent.

One feature of the current environment surprises people: the agency now grades your response. Firms that answer a warning letter with narrow actions aimed only at the cited observations increasingly get inadequate response letters back, which stretches the enforcement cycle and signals that the agency wanted evidence of systemic examination. A thin, narrow response can turn a manageable situation into a long and expensive one. The remediation is part of what gets evaluated, so treat it that way.

The retrospective dimension is its own headache. When investigators find a data integrity failure, they usually expect you to establish the scope across time and across products, which can mean a comprehensive review of historical data and, in some cases, written certification of the integrity of every batch made during the relevant period. That review routinely costs more, in money and disruption, than the proactive controls that would have prevented the finding in the first place.

Best Practices

A defensible program comes down to a handful of principles applied without letting up: govern data across its full lifecycle, review audit trails in substance rather than in form, mature investigations and corrective action, push oversight out to your contract partners, and lift integrity risk to the part of the organization that actually owns enterprise risk. Here is what that looks like in practice.

Govern the Full Data Lifecycle

Good programs treat data governance as a lifecycle discipline covering creation, modification, processing, review, retention, retrieval, and disposition. Start with a data inventory and a criticality assessment that identifies which records support quality, safety, and efficacy decisions and therefore carry the most integrity risk. Then apply controls in proportion. High-criticality dynamic records from computerized systems get the strongest controls over access, audit trail review, and change management. Lower-criticality records get proportionate rigor. Aim the resources where they matter and stop spreading equal effort across records of unequal importance.

Make Audit Trail Review Real

Write it into procedure: who reviews audit trails, how often, against what criteria, and with what documented outcome. For records that support batch release, that review happens before release, consistent with regulatory expectations. Train reviewers to spot the signals that actually matter, altered or deleted results, changes to integration parameters, unexplained reprocessing, aborted runs, edits to date and time settings. You are not reviewing every system event. You are reviewing the events that could change the reliability of a quality decision. Document what you looked at and what you found, and an audit trail that was just sitting there becomes an operating control.

Strengthen Computerized System Controls

Govern access with unique user credentials and privileges assigned by role, so every action ties back to a person. Restrict and monitor the administrative privileges that allow data to be altered or deleted. Configure systems so data cannot be reprocessed or overwritten without preserving the original, and lock down date and time controls against manipulation. Where a legacy system simply cannot support these controls, put compensating controls and a documented remediation plan in place, assess the risk, and communicate it. What you do not do is pretend the gap isn't there.

Mature Investigations and Corrective Action

Investigations should establish root cause through evidence, not assertion, and corrective actions should be verified for effectiveness before they close, not closed because the calendar says so. Handle out-of-specification results through a documented, scientifically sound process, with invalidation backed by valid justification and the original data retained in the record given to the quality unit. When you find a data integrity problem, run it through a formal investigation that establishes scope and root cause, weighs the effect on patient safety and product quality, and drives a global corrective action, not a quiet local patch.

Extend Oversight to Contract Partners

Since your decisions ride on their data, oversight of contract manufacturers and laboratories has to move from periodic audits toward continuous awareness of where each partner stands. Quality agreements should spell out data integrity expectations, audit rights, and notification obligations, so a partner is contractually bound to tell you about significant integrity events and regulatory actions. Watch the public enforcement status of your critical suppliers to cover the space between audits, and risk-tier suppliers by exposure so oversight lands where failure would hurt most.

Elevate Integrity Risk to Enterprise Governance

The most important practice is organizational, and it is the one most often skipped. Integrity risk should be visible to the executives, and where appropriate the board members, who own enterprise risk, with a defined reporting line and metrics that mean something. Management with executive responsibility should visibly back a culture where concerns get raised and handled. Treat integrity risk as an enterprise governance matter instead of a quality department chore and you are far more likely to find the problem before an inspection does.

Implementation Roadmap

The roadmap below moves from a current-state assessment to sustained governance. Each phase feeds the next, and the sequence is built so early findings shape later spending instead of the other way around.

FIGURE 1

Data Integrity Program Implementation Lifecycle

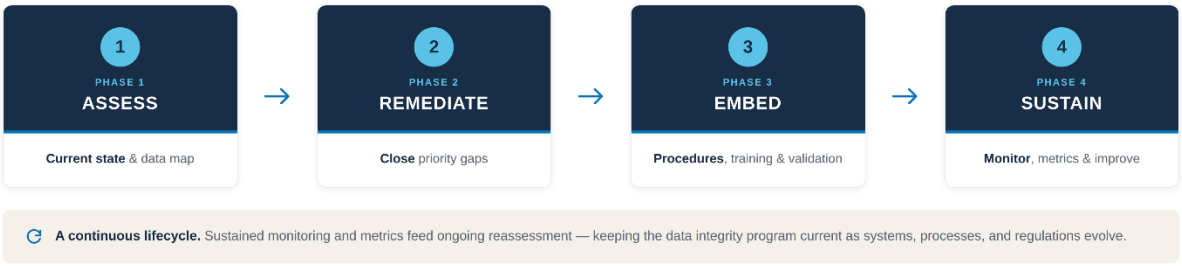

Step 1: Assess Current State

Start with an honest look at where you are: a data inventory, a criticality assessment, and a real evaluation of computerized system controls, audit trail review, investigation quality, and contract partner oversight. Be blunt about the gaps and document them so they can be prioritized.

Step 2: Perform a Gap Assessment

Compare the current state against regulatory expectations and recognized guidance, and find the gaps in access controls, audit trail review, investigations, and governance. Rank them by risk to product quality and by how likely they are to be cited, so remediation follows risk instead of convenience.

Step 3: Develop a Strategy

Build a prioritized remediation strategy with clear ownership, timelines, and resource commitments. Cover the quick wins that cut immediate exposure and the longer plays like system upgrades or replacements. And define the governance structure that will carry integrity risk up to leadership.

Step 4: Update Procedures

Revise procedures to match the strategy, with specific expectations for audit trail review, access management, investigation conduct, and lifecycle governance. Make them clear enough that a trained employee knows exactly what to do, and keep them aligned with controls you can actually sustain rather than ones that look good on paper.

Step 5: Train Personnel

Train people on the procedures and on the reasons behind them, so they understand why contemporaneous, attributable, accurate record-keeping protects patients and the company. Teach them how to recognize and report an integrity concern, and make sure the training reaches management, not just the bench.

Step 6: Validate Systems and Controls

Validate the computerized systems and the controls that depend on them, using a risk-based approach that concentrates assurance where a failure would most affect data reliability. Confirm that access controls, audit trails, and data protection actually work as intended.

Step 7: Monitor Performance

Give leadership metrics that show the current state of integrity health, audit trail review completion, investigation timeliness and effectiveness, access control exceptions, supplier status. Make them leading indicators where you can, so they surface risk before it becomes a finding.

Step 8: Drive Continuous Improvement

Use monitoring results, internal audit findings, and lessons from other firms' enforcement to keep improving. Data integrity is not a project with an end date. It is an ongoing property of the organization that has to be maintained as systems, products, and people change.

Technology Considerations

Technology cuts both ways. It is where a lot of the risk comes from, and it is also the most powerful way to control that risk. The same computerized systems that spin up dynamic records and metadata can, when they are configured and governed well, enforce controls that paper never could.

Electronic quality management systems, laboratory information management systems, and manufacturing execution systems can enforce contemporaneous record-keeping, attributable actions, and secure audit trails, provided they are configured with real access controls and validated for their intended use. Electronic batch records take out the transcription errors and blank-form control problems that paper invites. Automated audit trail review tools help you meet the review-before-release

expectation at scale by surfacing the events that deserve a human look, which is exactly the resource crunch that makes manual review so painful.

Cloud brings its own governance questions, starting with shared-responsibility models where you and the provider each control different pieces of the environment. Clear agreements, defined responsibilities, and genuine oversight of the provider are what keep you in control of regulated records. Artificial intelligence and machine learning tools are showing up more in quality operations, and their outputs need authorized human review before they become controlled records, in line with emerging regulatory expectations. Anything that generates or influences a quality decision falls inside the scope of data governance and has to be validated and controlled like it.

The finalized Computer Software Assurance approach supports a risk-based way of building confidence in these systems, putting the assurance effort where a failure would most affect product quality, patient safety, and data integrity. Used well, it makes it feasible to validate the growing pile of computerized systems in a modern operation without watering down the rigor where rigor counts.

Future Outlook

A few trends will shape what data integrity looks like over the next several years. Regulatory frameworks are modernizing, with the European Union advancing revisions to documentation and computerized system requirements that write data governance more explicitly into the rules, and international authorities continuing to harmonize around lifecycle governance and the full ALCOA+ set. If you operate globally, one harmonized control set will serve you far better than chasing each jurisdiction on its own.

The arrival of AI in manufacturing and quality will sharpen the focus on governing algorithmic systems. As models influence more decisions, intended use, human review, model validation, and control over updates move to the center of the data integrity conversation. The first enforcement actions touching artificial intelligence in good manufacturing practice documentation are a clear signal that regulators will hold firms accountable for the reliability of records their AI systems help produce.

Continuous manufacturing, heavier automation, and more connected systems will pump out larger data volumes and more tangled data flows, which raises the value of automated monitoring and of governance that can keep pace. Meanwhile inspections keep getting more data-driven, with investigators using metadata and audit trails to reconstruct what actually happened rather than what the summary report claims. The organizations that come out ahead will be the ones that treat data integrity as a designed property of their systems and their culture, not a control bolted on after the fact.

How cGMP Consulting Can Help

Building a data integrity program that holds up under inspection takes regulatory judgment, technical depth, and the kind of operational experience you only get from having lived through it. cGMP Consulting works with pharmaceutical, biologics, and device organizations to assess where they stand, close the priority gaps, and stand up governance that reaches the level of the organization where enterprise risk is actually owned.

Our consultants run data integrity assessments that inventory records, evaluate computerized system controls, and test whether audit trail review works in substance rather than on paper. We help design and implement remediation strategies, revise procedures, and validate systems with a risk-based approach that concentrates effort where it counts. We support inspection readiness and the preparation of credible, evidence-based responses to observations, and we help organizations facing enforcement build the systemic, well-documented corrective action programs regulators now expect. For firms that depend on contract manufacturers and laboratories, we help design the supplier oversight and quality agreement provisions that close the gap between audits.

Our approach is consultative and grounded in how inspections actually unfold. We aim to leave you with capabilities you can sustain, not binders you cannot maintain. Whether you are strengthening a mature program, preparing for a critical inspection, or remediating a serious finding, the goal is a defensible posture and a program that protects both patients and the business.

Conclusion

Data integrity is an enterprise governance issue now, full stop. The regulatory logic that treats lost data reliability as grounds to distrust an entire quality system means one serious finding can set off consequences far bigger than its apparent scope, from adulteration findings and distribution holds to consent decrees and an investor exodus. The 2026 environment rewards organizations that can demonstrate genuine control and punishes the ones whose controls exist on paper but not in the building.

The path to a defensible posture is not a mystery. Govern data across its full lifecycle. Review audit trails in substance and before release. Mature investigations and corrective actions. Extend oversight to contract partners. Lift integrity risk to the level of the organization that owns enterprise risk. It is demanding work, but it is achievable, and it pays off operationally as well as in compliance, because a system that produces reliable data is a system that makes better decisions.

The firms that win here are the ones that quit treating data integrity as a quality checklist and start treating it as a designed property of their systems and their culture. That is a leadership decision. Make it deliberately and you hold a durable advantage over everyone still waiting for an inspection to make it for them.

> *If your organization is evaluating its data integrity posture or preparing for future regulatory expectations, the experienced consultants at cGMP Consulting can help assess your current state, identify compliance gaps, and develop practical solutions tailored to your business objectives. To start a confidential conversation, complete the Contact Us form at* cgmpconsulting.com/contact-us*.*

References

European Medicines Agency. (2016). Data integrity: Questions and answers. European Medicines Agency.

Medicines and Healthcare products Regulatory Agency. (2018). 'GXP' data integrity guidance and definitions (Revision 1). MHRA.

Pharmaceutical Inspection Co-operation Scheme. (2021). Good practices for data management and integrity in regulated GMP/GDP environments (PI 041-1). PIC/S.

U.S. Food and Drug Administration. (2016). Data integrity and compliance with CGMP: Draft guidance for industry. U.S. Department of Health and Human Services.

U.S. Food and Drug Administration. (2018). Data integrity and compliance with drug CGMP: Questions and answers; Guidance for industry. U.S. Department of Health and Human Services.

U.S. Government Publishing Office. (n.d.). Title 21 Code of Federal Regulations Part 210: Current good manufacturing practice in manufacturing, processing, packing, or holding of drugs; general. Electronic Code of Federal Regulations.

U.S. Government Publishing Office. (n.d.). Title 21 Code of Federal Regulations Part 211: Current good manufacturing practice for finished pharmaceuticals. Electronic Code of Federal Regulations.

U.S. Government Publishing Office. (n.d.). Title 21 Code of Federal Regulations Part 11: Electronic records; electronic signatures. Electronic Code of Federal Regulations.

World Health Organization. (2016). Guidance on good data and records management practices (WHO Technical Report Series No. 996, Annex 5). World Health Organization.

European Commission. (2011). EudraLex Volume 4: EU guidelines for good manufacturing practice, Annex 11: Computerised systems. European Commission.