The Fire That Never Started

Quality, and Why It Only Gets Noticed When It Fails

Raphael Jorand

Cosmael.com

THE

FIRE

THAT NEVER STARTED

Quality, and Why It Only Gets Noticed When It Fails

Raphael Jorand

Cosmael.com

INTRODUCTION

The Least Glamorous Word in Business

I think that nobody grows up dreaming of working in quality.

Kids want to be ... well a lot of things, but nobody dreams of writing a non-conformance report or auditing a supplier's records. Quality is seen as the department that fills out forms, likes to say no, and slows everyone down.

I want to argue that this view is wrong.

Because the thing with quality is that when it works, you do not see it. There are no recalls, no angry customers, no batch thrown in the bin. Nothing happens, and “nothing happened” is the sign of success. Quality is the fire that did not start, but no one gives a medal to the person who prevented a fire that did not happen.

So the people who do it well are invisible, and its bad reputation comes from the moment it becomes visible: when it slows a project down, and when it fails. In between, in the vast stretch of time where it quietly holds everything together, it is ignored. That is the tragedy of the field, and it is also the reason so many companies underinvest in it right up until the day it would have saved them.

I have spent a good part of my working life inside this world. Among other things, I helped implement final inspection of finished goods before shipping, what we called FIQA, along with equipment qualification (IQ, OQ, and PQ), and preventive maintenance programs. I have written and argued over corrective actions, chased suppliers for real answers instead of polite ones, and done other work to retain existing customers and bring in new ones. I am not going to pretend that all of it was thrilling. Some of it was genuinely tedious. But that feeling definitely grew on me: everything is good and nothing is happening. I came out of it convinced that quality is one of the few places in a company where you can actually see, in advance, the shape of a future disaster, if you are willing to look.

What this primer is

This primer has two parts.

The first is about quality, as a working reality: what it is, the systems that make it up, and the strange jargon that surrounds it. I will talk about what happens at the two ends of a production line, incoming and final inspection, and why both matter. And I will decode the alphabet soup that quality people speak in, the NCRs and NCMRs and CARs and SCARs and CAPAs, because behind every one of those acronyms is a simple idea that got dressed up in initials.

The second part is about ISO certification, which is where quality meets the outside world. What these standards actually are. The difference between the general one, the one for medical devices, and the one for laboratories, which behave quite differently from each other. How to get certified. And a distinction that trips up almost everyone, including people who should know better: the difference between saying you meet a standard, being certified that you meet it, and being accredited, which are three very different claims with three very different levels of trust behind them.

The idea underneath

Before we start, let me give you the one idea that runs under everything, because if you take nothing else from this book, take this.

Quality is meeting requirements consistently, and being able to prove it.

Anyone can make one good product. A talented person, on a good day, with enough care, can make almost anything once. That is not quality. Quality is making the same thing, to the same standard, the ten-thousandth time, on a Tuesday, when the person who normally does it is off sick and their replacement has never seen the job before. And it is being able to open your records and show, to a stranger who does not trust you, that this is true.

That second half, the proving, is what is hard to understand for outsiders, what a quality department lives by. In this world, work that is not documented did not happen. A test you ran but did not record is a test you cannot rely on. A process that works but that you cannot demonstrate works is a process you do not really control, you are simply getting away with it. This obsession with evidence looks like bureaucracy from the outside. From the inside, it is the only thing standing between a company and the slow accumulation of small unnoticed failures that eventually become one large, very noticeable one.

So this is about consistency, evidence, and the discipline of caring about things before they go wrong.

PART I

Building Quality Inside the Company

CHAPTER 1
What Quality Actually Means

Ask ten people what quality means and you will get ten answers, most of them about luxury. A quality watch, a quality hotel, a quality bottle of wine. In everyday language, quality means expensive, refined, better than the ordinary thing.

That is not what the word means in this world, and the confusion causes real damage.

In quality management, quality is the degree to which something fulfills its requirements. A product has quality if it does what it is supposed to do, within the limits that were defined for it. By this definition, a cheap plastic fork can be a quality product if it is safe for its intended use, performs as specified, and arrives consistently as promised. And a luxury car with “quality materials” can be a quality failure if the electronics behave differently in every unit and half of them need a warranty repair in the first year.

Quality is not about how impressive the thing is. It is about the gap between what you promised and what you delivered, and whether that gap is under control. A humble product that always meets its modest promise has quality. A magnificent product that meets its grand promise only sometimes does not.

Two words that sound the same and are not

Now the distinction that every newcomer gets corrected on within their first week: quality control and quality assurance are related, but they are not the same thing.

As a practical shorthand, quality control, QC, is mainly about checking the product. It is the checking, the testing, the inspecting, the measuring. Did this specific item meet the specification? QC catches the bad unit before it reaches the customer. It is reactive, because by the time you inspect something, it has already been made. QC is the goalkeeper: essential, but by the time the goalkeeper is involved, something has already gone wrong upfield.

Quality assurance, QA, focuses more on the system and process that provide confidence the requirements will be met. It is the set of activities designed so that the product comes out right in the first place, so that you rely less and less on catching problems at the end. QA asks not "is this unit good?" but "is our way of making things capable of producing good units consistently?" QA is proactive. It is the coaching, the training, the tactics, the whole system designed so that the ball rarely reaches your goalkeeper in a dangerous way at all.

Here is why the distinction matters. A company that relies mainly on end-of-line inspection is playing defense forever. It inspects, it finds defects, it throws them away or reworks them, and it congratulates itself on catching them, without ever noticing that it is paying to produce defects and then paying again to remove them. Every defect found at final inspection was paid for twice: once to make it, once to catch it, and if it slipped through, a third time in the customer's hands. QA is the attempt to stop paying that tax by fixing the process instead of endlessly filtering its output. A mature quality organization is slowly, deliberately trying to make its own QC less necessary.

The cost nobody puts on the invoice

This leads to an idea that quality people talk about constantly and that finance people rarely see, called the cost of quality, and specifically the cost of poor quality.

Everyone can see the cost of doing quality work: the salaries, the testing equipment, the time a batch spends waiting for release. It sits right there on the budget, visible, and in a bad year it is one of the first things someone proposes to cut.

What almost nobody adds up is the cost of not doing it, because that cost is scattered across the whole company under other names. Scrap. Rework. Warranty claims. Returns. Expedited shipping to replace a rejected batch. The engineer pulled off new work to investigate an old failure. The customer who did not complain and simply never ordered again, whose absence appears in no report at all. When you total these honestly, the cost of poor quality is frequently much larger than the visible quality budget.

This is the argument quality professionals can never quite win, because their budget is a visible line and the disasters they prevent are invisible. The best of them learn to make the invisible visible, to put a number on the scrap and the rework and the returns and lay it next to the cost of the inspection, so that the trade becomes obvious even to someone who has never set foot on the production floor.

Prevention, appraisal, and two kinds of failure

There is a useful four-part model for where quality money goes.

You can spend on prevention: designing requirements and processes well, training people, maintaining equipment, managing risk, and qualifying suppliers. Money spent here reduces the chance that defects will exist.

You can spend on appraisal: inspecting, testing, auditing, and checking. Money spent here detects problems that already exist before they travel farther.

You can pay for internal failure: scrap, rework, retesting, downtime, and investigations before the problem reaches the customer.

Or you can pay for external failure: complaints, warranty, field service, recalls, legal exposure, lost trust, and customer harm after the problem has escaped. This is usually the most dangerous category because the company no longer controls where the defect goes or what it costs.

The whole art of quality management is shifting effort from external failure to internal detection, from detection to prevention, and eventually into better design. Companies under pressure often do the reverse. They stop training, defer maintenance, thin out supplier oversight, and reduce preventive work because its absence takes time to hurt. Nothing bad happens for a while, which feels like vindication. Then several things fail at once, which feels like bad luck. It was not luck. It was the delayed bill for prevention that was cancelled months or years earlier.

With that foundation, we can go to the factory floor and start at the beginning of the line, before you have made anything at all: the moment other people's materials arrive at your door.

CHAPTER 2

The Two Ends of the Line

Every production line has two doors. Materials come in one end, and finished product leaves the other. It turns out that the smartest place to put your attention is at those two doors, because both of them are borders, and borders are where responsibility changes hands.

The door where things come in

Almost nothing you sell is made entirely by you. You buy raw materials, components, parts, ingredients, packaging, and you build on top of them. Which means that a large part of your product's quality was decided before it ever reached you, by suppliers you do not control, in factories you have never seen.

Incoming inspection is the check you perform on those materials as they arrive, before you accept them into your process. It exists because of a simple and unforgiving piece of arithmetic: a defect that enters your line at the beginning gets more expensive with every step it survives.

A faulty component caught at the receiving dock costs you almost nothing. You reject it, you send it back, you move on. That same faulty component, if it slips past the door and gets built into your product, is now wrapped inside your labor, your other materials, and your time. Catch it halfway through assembly and you scrap everything you have added to it. Catch it at final inspection and you scrap a finished unit. Miss it entirely and it reaches the customer, where it costs you a return, a warranty claim, a damaged reputation, and possibly a safety incident. The same defect, depending only on how far it traveled before you noticed, can cost you a dollar or a hundred thousand. Incoming inspection is your chance to catch it while it is still worth a dollar.

Now, you cannot inspect everything to the same depth, and you should not try. This is where quality becomes a matter of judgment rather than paperwork. The right level of incoming inspection depends on how critical the part is and how much you trust the supplier. A trivial component from a supplier with a perfect ten-year record might need only a quick count and a glance. A safety-critical part from a new supplier might need every unit tested. Most things sit in between, and the professional's job is to put the scrutiny where the risk is, not to spread it everywhere.

There is even a formal discipline for deciding how many items to check out of a large batch, called acceptance sampling, which uses statistics to let you inspect a defined sample and make a controlled accept-or-reject decision. It does not prove that every item is good, and it always carries producer's and consumer's risk. The details are beyond this primer, but the underlying idea is worth keeping: you are trying to collect enough evidence for a decision at a cost and confidence level that make sense for the risk involved.

The door where things leave

At the other end of the line sits the last check before your product goes to the customer. I call it final inspection: the last set of eyes, the last measurement, the last chance to stop a bad unit from becoming a customer's bad experience.

Final inspection is the last and most consequential checkpoint before the product reaches the customer, because everything that passes it becomes real. Up to this point, a defect is an internal matter, embarrassing but private. Past this point, a defect is a customer's experience, and customers do not see your process, they see the one unit they received. To them, that unit is your entire company. A single bad product in the hands of the wrong customer at the wrong moment can undo the reputation built by a hundred thousand good ones.

But here is the trap of final inspection. It is the most visible quality activity, so it is the one companies lean on most, and leaning on it is a sign of weakness, not strength. If final inspection is regularly catching lots of defects, that is not good news that your net is working. It is bad news that your process is producing defects for the net to catch. Every unit stopped at final inspection was fully built and then thrown away or reworked, which means you paid for it in full before deciding it was worthless. A company that is proud of how much its final inspection catches is like a person proud of how often their smoke alarm goes off. The alarm working is good, but the number of fires is the actual problem.

So the two doors tell you different things, and you need to check both. A rise in incoming rejections tells you something about your suppliers. A rise in final inspection rejections tells you something about your own process. And a healthy factory, over time, wants both numbers to fall, not because it is inspecting less, but because less is going wrong to find.

The border principle

Notice what these two doors have in common. Both sit exactly where responsibility changes hands: from your supplier to you, and from you to your customer. Quality problems love borders, because a border is where one party's job ends and another's begins, and things fall into the gap between them.

The supplier assumes you will check, so they relax. You assume the supplier delivered good material, so you relax. The gap between those two assumptions is where the bad part slips through. The same happens at the customer border: the customer assumes you tested thoroughly.

This is why the two ends of the line deserve special attention. They also help to identify the problems that lie in the middle. Which brings us to the border that causes more quality trouble than any other, and that deserves a chapter to itself: the relationship with your suppliers.

CHAPTER 3
The People You Depend On

Your quality is only as good as the quality of the people you buy from. You can run the most disciplined factory in your industry, and if a supplier ships you a bad batch of raw material and you do not catch it, your customer gets a bad product with your name on it. Not your supplier's name. Yours. The customer never sees the supplier. As far as the world is concerned, you made it, all of it.

This is one of the hardest truths in quality, because it means you are held responsible for work you did not do, performed by people you do not manage, in places you rarely visit. And the only way to live with that responsibility is to take the supplier relationship as seriously as you take your own production, which most companies do not, until a supplier teaches them why they should.

Qualifying a supplier before you need them

The time to find out whether a supplier is any good is before you depend on them, not after. This sounds obvious and is routinely ignored, because qualifying a supplier properly is slow and unglamorous work, and there is always pressure to just place the order and get moving.

Supplier qualification is the process of establishing, with evidence, that a supplier is actually capable of consistently meeting your requirements. Not that they seem nice. Not that their sales representative was impressive and their brochure was glossy. That their factory, their process, and their quality system can genuinely deliver what you need, again and again.

Doing it properly means looking at real things. Their own quality certifications, and whether they are real and current. Their track record with other customers. Sample parts, tested against your actual specification rather than their claims. Sometimes an audit, where you or someone you trust physically visits and looks at how they work, because a factory floor tells you in an hour what a year of phone calls will not. For critical suppliers, you may run a formal qualification where they must demonstrate, on your parts, that their process produces conforming product consistently before you approve them for regular supply.

This front-loaded effort feels expensive, and it is. It is also far cheaper than the alternative, which is qualifying your supplier by accident, in production, using your customers as the test. Companies that skip qualification are not avoiding the cost of finding out whether a supplier is good. They are only delaying it, and moving it to the worst possible moment, when the bad batch is already inside a shipped product.

Proving the process, not just the sample

There is a subtle point here that separates serious quality operations from casual ones, and it applies to your own processes as much as to your suppliers'.

When you evaluate a supplier by testing a few sample parts, all you have learned is that they can make a few good parts. That is not the same as proving they can make good parts reliably, at volume, over time. A supplier can hand-pick perfect samples the way a job applicant picks their best references. What you actually need to know is whether their normal, everyday, unsupervised process produces good parts consistently.

This is the idea behind process qualification, terminology we will return to in the chapter on proving that things work. The principle is the same wherever it appears: the goal is not to prove that a good outcome is possible, it is to prove that the good outcome is what the process reliably produces. Anyone can achieve a good result once, under ideal conditions, with everyone watching. Quality is about what happens on an ordinary day when nobody is watching, and proving that requires looking at the process running normally, over enough time and enough units to trust it.

Confirming that something can work and concluding that it will work are different claims, and the gap between them is very important.

When a supplier lets you down

Even good suppliers fail sometimes. What separates a manageable problem from a chronic one is what you do about it, and here quality has built a specific tool, with a specific acronym, for exactly this situation.

When a supplier ships you something that does not conform, you first record it. That record is a non-conformance report, and when it is specifically about physical material or product, it often has its own name, a non-conforming material report, which is where the acronym NCMR comes from. The report is not an accusation, it is a fact: this delivery, this batch, this quantity, failed this requirement, and here is the evidence.

But recording the problem is not fixing it. If you simply log the bad batch, send it back, and move on, the same problem will arrive again next month, because nothing has changed at the supplier's end. To force a real fix, you issue what is called a supplier corrective action request, universally shortened to SCAR.

A SCAR is a formal demand, and the word demand is deliberate. It is not a polite email hoping they will do better. It requires the supplier to investigate the problem, find its actual root cause, put a fix in place, and demonstrate that the fix works. A good SCAR process refuses to accept vague reassurances. “We have spoken to the operator and it will not happen again” is not an answer. A real answer identifies why the failure was possible in the first place, what specifically has been changed so that it cannot happen again, and how the supplier will prove, in the coming deliveries, that the change actually worked. Until that evidence exists, the problem is not closed.

The difference between a company that issues SCARs well and one that does not is visible over a couple of years. The first has suppliers who improve, because every failure is turned into a permanent fix. The second has suppliers who apologize, repeatedly, for the same failure, because nothing was ever forced to change. Apologies are free, which is exactly why they are worthless as a quality tool. A SCAR costs the supplier real effort, and that effort is the point.

The relationship, not the transaction

I want to end this chapter on the human note, because supplier quality is often treated as an adversarial game and it should not be.

The goal is not to catch your suppliers out, or to squeeze them until they cut corners somewhere you cannot see. The goal is a relationship where their quality genuinely improves because of how you work with them, where your requirements are clear enough

to actually meet, where your feedback is specific enough to act on, and where a failure produces a fix rather than a fight. The best supplier relationships look less like a customer policing a vendor and more like two companies who both understand that a defect anywhere is a problem everywhere.

But, and this is the balance, none of that warmth substitutes for verification. You can have an excellent, trusting, collaborative relationship with a supplier and still inspect what they send you, still issue the SCAR when something fails, still require the evidence. Trust and verification are not opposites. In quality, trust is what lets you work together efficiently, and verification is what keeps the trust honest. A relationship with one and not the other is not a relationship, it is either a fight or a gamble, and both end badly.

CHAPTER 4

The Alphabet Soup

If you walk into a quality meeting for the first time, you will get confused very quickly. NCR, NCMR, CAR, SCAR, CAPA, OOS, RCA. The room is thick with initials, and it can feel deliberately exclusionary, a private language that keeps outsiders out.

It is not, or at least it does not have to be. Behind every acronym is a simple idea. The names vary between industries and companies, a standard usually cares more about the function than the initials printed on the form. Once you see the functions, the alphabet stops being intimidating and becomes useful shorthand.

The system underneath the jargon answers a sequence of questions. Something went wrong: what exactly happened? What must we do about the item or situation in front of us? Why did it happen? What should change so that it does not happen again? And what evidence will show that the change worked?

Naming and controlling the problem

It starts when something does not conform to a requirement. A part is out of tolerance, a step was skipped, a document is wrong, a test result falls outside its limits. The first thing you do is write it down honestly and specifically. Many companies call that record a non-conformance report, or NCR.

The NCR states the fact: what was found, where, when, against which requirement, and what may be affected. It is not yet a complete explanation and not yet a permanent fix. It is the honest naming of a problem, and everything else is built on it. A company that cannot write honest non-conformances cannot do quality, because it has broken the first step: admitting on paper that reality differed from the plan.

When the non-conformance is specifically about physical material or product, a batch of raw material, a lot of finished goods, a shipment of parts, it often gets a more specific name, the non-conforming material report, or NCMR. The distinction is practical rather than philosophical: material non-conformances need extra information that process ones do not, such as the exact quantity affected, the batch or lot number, and where the suspect material is currently sitting so it can be quarantined and not accidentally used. If you hear NCR and NCMR used almost interchangeably, that is because they are close cousins. NCR is the general case, NCMR is the version for physical stuff you can put your hands on and, crucially, lock away.

That locking away has its own importance. The moment you identify non-conforming material, you have to make sure nobody uses it by mistake while you decide what to do with it. That decision, scrap it, rework it, send it back, or in some cases use it as-is with proper approval, is called disposition, and controlling the suspect material until disposition is decided is one of those unglamorous disciplines that prevents genuine disasters. More than one company has shipped a defective batch it had already identified as defective, simply because the bad material was never physically separated from the good.

Demanding a deeper fix

Naming a problem does not fix it. To force a real solution, you issue a corrective action request, a CAR. This is the formal way of saying: this is serious enough that I need you to investigate it properly and stop it from recurring, not just patch this one instance.

The CAR is where the tone changes from recording to demanding. When that request is aimed specifically at a supplier, it becomes the supplier corrective action request, the SCAR we met in the previous chapter. Same tool, pointed outward. Internally you raise a CAR, externally you issue a SCAR, but the demand is identical: investigate, find the root cause, fix it, and prove the fix works.

Not every non-conformance deserves a CAR, and this is a judgment that separates mature quality systems from bureaucratic ones. A one-off, trivial, easily corrected problem may just need to be recorded and fixed. Raising a full formal investigation for every tiny thing buries the organization in paperwork and, worse, trains people to dread reporting problems at all. The art is to escalate to a corrective action when a problem is serious, or systemic, or keeps coming back, and to simply fix and record the rest. Over-using the heavy machinery is its own failure, because a quality system that punishes people with paperwork for honesty will, very quickly, stop receiving honesty.

Correction, corrective action, preventive action

These three phrases are often blurred together, but they are very different.

If the NCR names the problem and the CAR demands a fix, CAPA is what actually produces the fix.

A correction fixes or contains the detected problem. Corrective action removes or controls causes of an actual non-conformity so it is less likely to recur. Preventive action addresses causes of a potential non-conformity before it happens.

The heart of it all: understanding causes

Underneath corrective action sits one of the most important skills in quality: root cause analysis (RCA).

Root cause analysis is the discipline of not stopping at the first answer. Because the first explanation is often only a symptom. The part failed. Why? The material was weak. Why was weak material used? It passed the incoming inspection. Why did it pass? The method or limit was unsuitable. Why was that unsuitable condition approved? Now the investigation is approaching decisions, assumptions, and controls that can actually be changed.

The Five Whys is a well-known technique because it resists the urge to stop at the first blameable answer. Five is not a magical number and complex failures often have several interacting causes: a weak design, an unclear procedure, inadequate training, supplier variation, production pressure, and a detection control that did not work.

"The operator made a mistake" is where lazy investigations end and useful ones begin. People will make mistakes. The harder questions are why the task made the mistake likely, why the system failed to detect it, and why the consequence was allowed to reach the

product. A process that depends on no human ever making an error is not a proper process.

This is where quality stops being about forms and becomes about honesty and courage, because real root cause analysis often leads somewhere uncomfortable. It leads to a decision management made, a corner the company chose to cut, a warning someone raised and someone else ignored. The temptation to stop the analysis one layer above the uncomfortable truth is enormous, and giving in to it is how companies run elaborate corrective-action systems for years while the same disasters keep happening, because every investigation was quietly halted just before it reached the thing nobody wanted to change.

Closing the loop

A corrective action is not complete when someone has written an action plan. It is complete when the action was implemented, the relevant documents and training were updated, and evidence shows that the action was effective without creating a new problem elsewhere.

So here is the alphabet soup in plain language. Something went wrong, and we recorded exactly what. We controlled the immediate risk. We decided whether the event needed a deeper investigation. We followed evidence past the easy explanation, changed the system where the causes lived, and checked that the change worked. Then we kept the record so that a stranger can see that the problem was not merely hidden until the next audit.

The initials are shorthand for a grown-up discipline: admit the mistake, understand it honestly, control its consequences, and learn enough that it does not have to teach you the same lesson twice.

CHAPTER 5

You Cannot Inspect Quality Into a Bad Design

There is a sentence attributed in different forms to several quality thinkers: you cannot inspect quality into a product. Whatever its exact origin, the idea is right, and it reaches farther than inspection.

A factory can follow every procedure correctly and still produce a bad product if the product was badly conceived. Inspection can confirm the material matches the drawing. Final inspection can confirm the finished unit matches the specification. Neither can tell you that the drawing expressed what the customer needed, that the specification made sense, or that the product is safe in the hands of the person who will actually use it.

A perfect inspection plan can prove that you made the wrong thing perfectly.

Requirements before specifications

Quality begins before the first component is ordered, with the question that sounds obvious and is routinely answered badly: what must this product actually do?

A requirement is not a wish such as "easy to use," "robust," or "high performance." Those phrases may express a need, but they are not yet requirements you can design against or test. A usable requirement has enough clarity that two competent people can agree whether it was met. How much load? Under which temperature? For which users? With what accuracy? For how long? In which environment?

The purpose is not to reduce everything human to a number. It is to remove hidden assumptions before they get into hardware, software, tooling, and contracts. Every vague sentence at the beginning becomes an argument later, when changing it is expensive.

Intended use and the real user

The most important requirements usually begin with intended use: who will use the product, for what purpose, in what setting, and with what consequences if it fails.

Engineers naturally design around the product they understand. Users operate the product they encounter. And those are not always the same thing. A control that is obvious to the designer may be confusing for the user. A connector that cannot be inserted incorrectly on the bench may be forced incorrectly while someone is wearing gloves. A warning that is technically present may be invisible in a specific scenario.

This is why serious design work includes real users and foreseeable misuse, not only ideal use. "They should not do that" is not a risk control when ordinary people predictably do it. Quality does not require protecting against every imaginable act, but it does require honesty about how the product will behave outside the calm demonstration where everyone already knows the answer.

Inputs, outputs, and traceability

Design inputs translate needs, intended use, regulations, risks, and interfaces into requirements the design must satisfy. Design outputs are the drawings, specifications,

software, bills of material, test methods, manufacturing instructions, and other information that define what will actually be built.

The quality discipline is to keep a traceable line between them. Each important input should lead to an output, and each output should be verifiable against the input it is meant to satisfy. Traceability is not paperwork for its own sake. It prevents requirements from disappearing during the journey from a customer problem to a technical drawing.

Without that line, products accumulate orphan features nobody requested and missing functions everyone assumed were included. At the end, the team discovers that it tested the things it built rather than the things it needed to prove.

Verification and validation

Two words carry much of this work: verification and validation. Industries use the terms with slightly different emphasis.

Verification asks whether the design outputs meet the design inputs. Did we build the thing according to the requirements and specifications? It is the world of calculations, inspections, code reviews, bench tests, dimensional checks, and objective comparison.

Validation asks whether the resulting product fulfills its intended use and user needs under the conditions in which it will actually be used. Did we build the right thing? A product can pass verification and fail validation. Every component may meet its drawing while the whole product remains too difficult, too slow, too fragile, or simply unsuitable for the real job.

The familiar summary is useful because it is almost impossible to forget: verification asks whether we built the product right, validation asks whether we built the right product.

Risk is part of the design, not a form at the end

Risk management is often performed too late, as a document written after the design is essentially fixed. At that point it becomes an exercise in explaining why the chosen design is acceptable rather than using risk to improve it.

Done properly, risk work begins early and changes decisions. What can go wrong? How severe could the harm or business consequence be? How likely is the chain of events? Can the hazard be removed by design rather than covered by a warning? If it cannot be removed, what protective control reduces it? How will we verify that control?

The hierarchy matters. A warning label is weaker than a design that makes the hazardous action impossible. An inspection is weaker than a process that cannot create the defect. Training is weaker than an interface that does not invite the error. Good quality moves the control as close as possible to the source of the risk.

Design transfer: where good ideas go to die

Eventually the design has to leave the small group that created it and become something other people can manufacture, test, service, and support. This is design transfer, and it is a border, which means quality problems love it.

The prototype may have been assembled by the engineer who invented it, using judgment that was never written down. Production receives drawings that are technically complete but practically ambiguous. Test methods depend on knowing which result "looks wrong." Service discovers that a critical component cannot be reached without destroying three others. The design worked because its creators were part of the process, and then the company removed the creators and expected the result to stay the same.

A good transfer proves that ordinary trained people, using controlled documents, tools, materials, and methods, can build and test the product consistently without depending on the designer standing beside them. If that proof does not exist, the prototype is not ready for production.

Change control: the design is never truly finished

Products change. Suppliers discontinue parts. Software is updated. A complaint reveals a weakness. Manufacturing finds a more efficient method. Change is normal, uncontrolled change is the danger.

A disciplined change process asks what the change touches before approving it. Which requirements, risks, tests, documents, suppliers, inventories, regulatory commitments, and customer configurations are affected? What must be reverified or revalidated? How will old and new versions be identified?

The smallest-looking changes are often the most treacherous. A component described as equivalent behaves differently at the edge of temperature. A software library update changes timing. A new adhesive cures faster but ages worse. The individual change looks local, the product is a system, and systems transmit consequences.

This is the deeper reason quality cannot begin at incoming inspection. By the time the material reaches the door, most of the important decisions have already been made. The requirements defined what success meant. The design decided which failures were possible. Risk controls decided how visible and containable those failures would be. Inspection can confirm the result but it cannot rewrite the history that produced it.

PART II

Proving It to the Outside World

CHAPTER 6
What ISO Actually Is

Everything so far has been about quality as an internal discipline, the work you do inside your own walls. Now we turn to the moment quality meets the outside world, where you have to give customers, regulators, partners, and strangers a reason to trust how you work. This is the world of standards and conformity assessment, and it starts with three letters that appear on company walls and marketing materials everywhere, usually badly understood: ISO.

ISO is the International Organization for Standardization, an independent, non-governmental international organization that develops and publishes standards. "ISO" is not an English acronym, the short name was inspired by the Greek isos, meaning equal. More importantly, ISO itself does not certify companies or accredit laboratories.

ISO standards are not laws, but a standard can become mandatory when a regulation incorporates it, a customer writes it into a contract, an industry requires it for market access, or an organization voluntarily commits to it.

What a standard actually is

A standard is an agreed way of defining or doing something so that different people in different places can work consistently and trust each other's work.

ISO has published standards covering everything from screw threads to information security. Most are invisible to the public and quietly essential. They allow parts, services, and expectations to travel across organizations and borders without being reinvented every time.

The standards this primer cares about concern management systems and technical competence. They do not tell you the exact design of your product. They define what a trustworthy system should control: requirements, responsibilities, documents, risk, resources, competence, suppliers, non-conformities, audits, improvement, and evidence.

That is the key idea. A standard converts a claim "trust us, we work properly" into a shared set of requirements that another party can examine.

Why bother

Companies pursue standards for a mix of reasons, and it is worth being honest about all of them.

The best reason is that building a real quality system genuinely makes you better. Defining processes, controlling documents, assigning responsibility, tracking failures, reviewing risk, and closing actions force the company to replace memory and improvisation with a real system.

The most common reason is that customers or regulators demand it. In many industries, a relevant certificate or accreditation is a passport. Without it, you may not be considered for a contract, regardless of how persuasive your sales presentation is.

And the least noble reason is the badge. Some organizations want the logo, the tender checkbox, or the sentence on the website, and treat the system as a performance for the

auditor. That is possible, and it is why experienced people do not mistake a certificate for proof that every product is good or every employee follows the process on the night shift.

A management-system certificate provides evidence that a defined system was assessed against a standard. It does not certify product perfection, guarantee culture, or transfer responsibility away from the organization.

CHAPTER 7

Three Standards, Three Worlds

There are thousands of ISO standards, but for anyone working around manufacturing, medical devices, and laboratories, three are especially useful to compare. They answer related but different questions about trust.

ISO 9001: the general quality-management standard

ISO 9001 is the best-known quality-management-system standard. It can be applied to a factory, a software company, a bakery, a law firm, a logistics operation, or almost any other organization because it is deliberately general. It does not care what you make as much as it cares whether the system used to deliver it is defined, controlled, evaluated, and improved.

What ISO 9001 asks is, in essence, the codified version of good management. Do you understand your customers' requirements? Have you defined your processes and who is responsible for them? Do you control your documents so people work from the current version rather than an old one? Do you track things that go wrong and actually fix them? Do you check whether your system is working and improve it over time? It is built around a small number of durable ideas: focus on the customer, involve leadership genuinely rather than decoratively, take a risk-based approach to what could go wrong, and commit to continual improvement rather than treating quality as a box ticked once.

Because it is so general, ISO 9001 is both the most useful and the most criticized of the standards. Useful, because its principles genuinely apply everywhere and building a real 9001 system makes almost any organization more deliberate. Criticized, because its very generality makes it possible to satisfy on paper without much substance, to produce the documents and pass the audit while changing little about how the work actually happens. The standard is a powerful tool and a hollow badge, depending entirely on the sincerity of the company holding it, and it cannot tell the difference between the two. That is not a flaw in the standard so much as a limit on what any general standard can do.

ISO 13485: the regulatory-focused standard for medical devices

ISO 13485 looks, at first glance, like ISO 9001 with a narrower audience. It is the quality management standard for medical devices, and it shares much of 9001's structure. But the resemblance hides a fundamental difference, and the difference is instructive about how quality changes when the stakes change.

ISO 9001 is built around customer satisfaction and continual improvement. Its goal is a happy customer and an ever-better organization. ISO 13485 is built around something more sober: safety, effectiveness, and regulatory compliance. When your product is a heart valve or an insulin pump, the goal is not primarily to delight the customer or to improve year over year. It is to ensure the device is safe and does what it claims, every single time, because the cost of failure is not a refund, it is a patient.

This shift in purpose changes the whole texture of the standard. ISO 13485 is far more demanding about documentation, traceability, and risk management. It wants to know the full history of a device: what went into it, who made it, how it was tested, so that if a problem emerges in the field, every affected unit can be traced and recalled. It ties tightly

into formal risk management, the disciplined analysis of everything that could go wrong and what has been done about each possibility. And notably, it is more cautious about the very "continual improvement" that 9001 celebrates, because in a regulated medical world, you cannot simply change your process because you found a nicer way, changes themselves must be controlled, assessed, and often approved, since an uncontrolled improvement is just an unvalidated change wearing a friendly name.

The gap between 9001 and 13485 is quite interesting: as the consequences of failure rise, quality shifts from satisfying the customer toward protecting them, from improvement toward control, from "are they happy" toward "can we prove it was safe." The two standards are cousins, but they come from different concepts.

Certified does not mean regulator-approved. ISO 13485 certification does not by itself establish compliance with every law, authorize a medical device for sale, or replace a regulatory inspection. In the United States, the FDA Quality Management System Regulation incorporates ISO 13485:2016 by reference but retains FDA-specific statutory and regulatory requirements. Certification, FDA registration, premarket authorization, reporting obligations, and FDA inspection are different things.

ISO/IEC 17025: competence and valid laboratory results

ISO 17025 is the standard for testing and calibration laboratories. It is not about making products at all, it is about the competence of a laboratory to produce results you can actually trust. When a lab tests whether a drug is pure, or calibrates an instrument that other people's quality depends on, the entire value of its work rests on one thing: can you believe the result it gives you?

Because a laboratory's product is a result rather than an object, the standard cares about different things. It is deeply concerned with technical competence: are the people qualified, are the methods scientifically valid, is the equipment properly calibrated, is the required metrological traceability established through a documented chain to an appropriate reference? It cares intensely about the integrity of the result and about impartiality: does the lab have any hidden reason to report one answer over another? A calibration laboratory whose numbers cannot be trusted is worse than useless, because everyone downstream is building on a false foundation, unknowingly.

But the most important thing about ISO 17025, for our purposes, is not what it measures. It is that it works through a fundamentally different mechanism from the others, a mechanism called accreditation rather than certification, and that difference, which sounds like mere vocabulary, is the subject of the next chapter and one of the most practically useful things in this book. Because it turns out that "we meet ISO 9001" and "we are accredited to ISO 17025" are not just different claims about different standards. They carry different levels of assurance.

CHAPTER 8

Who Says So? Assessment, Certification, and Accreditation

Here is a question that seems trivial and is not: when an organization says it meets a standard, who checked?

The answer does not automatically tell you whether the assessment was good, but it tells you the relationship between the assessor and the organization being assessed. In conformity assessment, the conventional language is first party, second party, and third party.

First party: we checked ourselves

A first-party assessment is performed by the organization on itself. Internal audits, internal inspections, and a self-declaration of conformity are first-party activities.

Self-declaration can be legitimate. An organization may build a system around ISO 9001, assess itself honestly, and state that it conforms without purchasing third-party certification, provided the wording is accurate and no law or contract requires certification. The weakness is that it provides no independent assurance to an outsider.

A serious internal audit may be more demanding than a poor external one. But a stranger cannot know that from the claim alone, because the organization making the claim also benefits from it. First-party evidence is essential for running a system, its value as a public signal depends heavily on existing trust and the consequence of being wrong.

Second party: the customer checks the supplier

A second-party assessment is performed by a party with an interest in the relationship, usually a customer assessing a current or potential supplier.

Large companies conduct supplier audits because a certificate may not answer the specific questions they care about. A medical-device manufacturer may need to inspect a critical contract manufacturer's process, records, contamination controls, validation, or capacity. A customer audit can be extremely rigorous because the customer understands the exact risk it is buying.

The limitation is that the assessment is usually designed for one customer's needs and may not travel as a general credential to the rest of the market.

Third party: an independent body checks

A third-party assessment is performed by an organization independent of the supplier-customer relationship. Management-system certification is a familiar example.

A certification body audits the organization against a defined standard and scope. If the requirements are met, it issues written assurance in the form of a certificate, normally followed by surveillance and recertification activities. This is what most people mean when they say a company is ISO 9001 certified or ISO 13485 certified.

Third-party certification is more portable than a private customer audit. Multiple customers can rely on the same independent assessment instead of each repeating the entire exercise. But the certificate is only meaningful when you read it carefully: which

legal entity, which sites, which activities, which standard, which edition, which scope, and which certification body?

Accreditation: confidence in the assessors

Then comes accreditation, which is easiest to understand as checking the checker.

A company can be certified to a standard such as ISO 9001 or ISO 13485 by a certification body. But who checks that the certification body itself is competent and impartial? That is where accreditation comes in. An accreditation body assesses the certification body and formally recognizes that it is competent to perform specific certification activities.

But certification bodies are not the only organizations that can be accredited. Testing and calibration laboratories, inspection bodies, and other organizations whose work people rely on can also be accredited.

For example, a testing or calibration laboratory is normally accredited to ISO/IEC 17025 for a defined scope. In this case, the laboratory is not certifying another company. It is producing measurements or test results that others need to trust. Accreditation provides independent evidence that the laboratory has the people, methods, equipment, traceability, and technical competence needed to produce those results reliably.

A manufacturing company can therefore own an accredited laboratory, but the accreditation applies to that laboratory and its defined scope, not automatically to the rest of the company.

So the basic idea is simple: certification asks whether an organization meets a standard, accreditation asks whether the organization doing the assessment or producing trusted evidence is competent to do so.

How to read the claim

When you see a quality claim, do not stop at the logo.

Is the organization claiming conformity, certification, accreditation? Who issued the claim? Is the certification body accredited by a recognized accreditation body? What exactly is in scope? Is the certificate current, and can it be verified?

"ISO certified" by itself tells you very little. ISO does not certify organizations. "ISO/IEC 17025 certified laboratory" is usually the wrong phrase, the normal claim is accredited, and the scope matters. "FDA registered and ISO certified" does not mean FDA approved. The nouns and verbs do matter. They tell you what was assessed and what was not.

None of these claims proves that the organization is excellent every day. They tell you how much structured and independent scrutiny stands behind a defined statement. That is genuinely useful, as long as you read the whole statement properly.

CHAPTER 9

How You Actually Get Certified, or Accredited

Enough theory. What actually happens when an organization decides it wants ISO 9001 or ISO 13485 certification, or when a laboratory seeks ISO/IEC 17025 accreditation? The routes share some foundations, but they are not the same process and should not be described as if they were.

First decide what you need, and why

It begins with a decision that should be more deliberate than it usually is: which standard, and why. This sounds obvious, but companies routinely pursue certification for the wrong reason, chasing a logo a customer mentioned without asking whether it is the right standard for what they actually do, and the answer shapes everything that follows. A manufacturer of general goods might pursue ISO 9001. A medical device maker may need ISO 13485 certification. A testing laboratory may seek ISO/IEC 17025 accreditation, which as we now know is a different animal entirely. Choosing wrong here means either a certificate that does not open the doors you needed, or a far harder road than your situation required.

The honest question underneath this decision is why you are doing it at all. If the answer is purely "a customer requires the logo," you will be tempted to treat the whole thing as a box-ticking exercise, and you will get a hollow certificate that mostly wastes everyone's time. If the answer is "we want to genuinely run a better, more consistent operation, and the certificate is the proof," you will get something valuable, and the badge will be the smallest part of it.

The common foundation: understand the gap

The real work begins by comparing where you are against where the standard requires you to be. This is called a gap analysis, and it is exactly what it sounds like: an honest audit of the distance between your current reality and the standard's requirements.

It demands honesty about your own shortcomings. Where are your processes undocumented? Where do people work from memory or habit rather than defined procedures? Where do things go wrong without being formally recorded and fixed? Where is responsibility vague? The gap analysis maps all of it, and its value is proportional to how brutally honest you are willing to be. A gap analysis that flatters you produces a certification effort that fixes nothing, and a wall decoration at the end. A gap analysis that genuinely hurts is the beginning of an operation that actually improves.

Building and running the system

Then you close the gaps, and this is the bulk of the effort and the bulk of the time. You document processes that lived only in people's heads. You define who is responsible for what. You put in place the machinery from Part One of this book: the way you record non-conformances, the way you run corrective actions, the way you control which version of a document people use, the way you manage your suppliers, the way you qualify your equipment.

Crucially, this is not only writing documents. The single most common failure of certification efforts is to write beautiful procedures that describe a company that does not exist, an idealized version of how work supposedly happens, disconnected from how it actually happens on the floor. An auditor's favorite move, and the one that catches this instantly, is to read your documented process and then walk out to watch the real one and see if they match. If your documents describe a fantasy, they will find out within the hour, because the people doing the work will describe what they really do, not what the binder says. The goal is not documentation that looks good. It is documentation that is true, describing the actual system you actually run, improved to meet the standard. Writing down a lie and improving nothing is the worst of both worlds: all the cost, none of the benefit, and a permanent gap between the paperwork and the reality that every future audit will probe.

The certification route: ISO 9001 and ISO 13485

For management-system certification, the organization selects a certification body and agrees the legal entity, sites, activities, and scope to be assessed. Where recognized accreditation matters, the organization should verify that the certification body is accredited for the relevant management-system standard and sector.

The initial certification audit is commonly performed in two stages. Stage 1 reviews key documentation, scope, context, readiness, and whether the system is sufficiently implemented for the main assessment. Stage 2 examines implementation in practice: records, interviews, operational controls, internal audits, management review, corrective action, and the consistency between what the organization says and what it does.

Non-conformities found during the audit must be addressed according to their significance and the certification body's rules. Passing does not mean the assessor found nothing. It means the organization demonstrated conformity and resolved any issues required before the certification decision.

The certificate is issued for a defined scope and cycle, with surveillance audits between full recertification assessments. The organization must also manage changes that could affect the certification, such as sites, ownership, major processes, or scope.

The accreditation route: ISO/IEC 17025

Laboratory accreditation begins with defining the scope being requested. Which tests, calibrations, sampling activities, methods, materials, ranges, or measurement capabilities does the laboratory want recognized? This decision is central because the accreditation will say exactly what competence was assessed.

The laboratory then has to demonstrate both management control and technical competence. That may include method selection and validation or verification, equipment suitability and calibration, metrological traceability, measurement uncertainty, environmental conditions, sampling, handling of items, quality control of results, proficiency testing or interlaboratory comparison, impartiality, staff authorization, data integrity, and reporting.

An accreditation body reviews the application and documents, then performs an assessment using assessors and technical experts appropriate to the scope. They may

examine records, interview staff, witness tests or calibrations, and evaluate whether the laboratory can competently perform the activities it wants listed.

Any non-conformities must be corrected and accepted before the accreditation decision. When granted, the laboratory receives a certificate and, more importantly, a detailed scope. The scope is the useful document. It tells the customer which results are covered by the independent recognition.

Accreditation also continues through surveillance, reassessment, scope changes, proficiency activities, and ongoing demonstration of competence. A laboratory does not become permanently competent on the day the symbol is issued.

What certification and accreditation do not replace

External recognition does not replace the law, product approval, or the organization's own responsibility.

For example, ISO 13485 certification can support a medical-device quality system, but in the United States it does not replace FDA establishment registration, device listing, applicable premarket authorization, medical-device reporting, corrections and removals obligations, or FDA inspection. The FDA Quality Management System Regulation incorporates ISO 13485:2016 by reference while retaining FDA-specific requirements.

Likewise, ISO/IEC 17025 accreditation does not make every result from the laboratory accredited. Only work performed within the accredited scope and under the applicable system can be represented that way.

And then it does not stop

The certificate or accreditation is not the finish line. Organizations drift. People leave. Processes change. Documents age. Shortcuts become habits. Surveillance and reassessment exist because a system demonstrated once is not automatically a system maintained forever.

The companies and laboratories that gain the most are the ones that would have wanted the underlying control even without the badge. For them, the audit is an external challenge to a living system. The ones that gain the least are those that wanted only the logo and are surprised to learn that the logo comes attached to a permanent obligation to mean what it says.

CONCLUSION

The Fire That Never Started

I said at the beginning that nobody grows up wanting to work in quality, and that when it works, you see nothing. I want to end there, because that invisibility is the key to understanding why quality is so consistently undervalued right up until the moment it is desperately missed.

Think about what we have covered. Defining the right requirements before the design hardens around the wrong ones. Inspecting materials as they arrive, so a cheap defect never becomes an expensive one. Checking products and processes before a customer becomes the person who discovers the flaw. Qualifying suppliers before you depend on them. Verifying designs, validating intended use, qualifying equipment, and proving processes instead of relying on demonstrations. Recording problems honestly, understanding root causes, and closing actions with evidence. Building a system and technical competence strong enough that an independent stranger can examine them.

Every one of these activities is defined partly by an absence. The requirement that was clarified before it became a redesign. The defect that never reached the customer. The breakdown that never happened. The recall that was never needed. The batch that was never scrapped. Quality's product is often the disaster that did not occur, and a disaster that does not occur is, by its nature, invisible.

This is the structural tragedy of the discipline. Quality spending is visible and many of its benefits are invisible, which means that under pressure, prevention is always tempting to cut. For a while, nothing bad happens. The savings appear immediately. The costs are delayed, scattered, and deniable, and when they finally arrive they wear the disguise of bad luck: a run of breakdowns, a batch of returns, a supplier who suddenly failed, a recall out of nowhere.

It is often not bad luck. It is the bill for risk that was accepted without being named, maintenance that was deferred, evidence that was never collected, requirements that were never clarified, and problems that were corrected but not understood. The organization did not eliminate the cost. It borrowed it from its future at an interest rate it did not calculate.

The one idea, again

Quality is meeting requirements consistently, and being able to prove it.

Meeting the requirement once is evidence of one event. The discipline of quality asks a harder question: will it still be met the ten-thousandth time, on an ordinary day, when the usual person is away and nobody is paying special attention? And can you show, to someone with every reason to doubt you, that this is genuinely true?

The proving matters because memory is unreliable, people leave, systems drift, and "trust me" is not a control. But proof is not the same as paperwork. A document that describes work nobody performs is not evidence of quality. It is evidence of a gap. The record has value only when it accurately represents a controlled reality.

Where the two halves meet

The two parts of this primer, the internal discipline and the external recognition, are the same idea viewed from two sides.

Inside the organization, quality builds a controlled, evidence-based way of working. Outside it, standards, certification, accreditation, customer audits, and regulatory oversight provide structured ways for other people to examine defined parts of that claim.

And the deepest lesson connecting them is that the proving is worthless without the being. A certificate on the wall of a company that does not live its system is not merely useless, it can manufacture trust that has nothing real behind it. An accredited symbol used outside its scope does the same. The entire structure exists to transmit justified confidence from people who do controlled work to people who need to rely on it.

So do the real work first. Define the right requirement. Design with risk in view. Build the process so good outcomes are normal rather than heroic. Record the truth. Learn from failures while they are still cheap. Be genuinely, boringly, provably consistent, and let the certificate or accreditation be what it should be: not the goal, but the honest confirmation of something real.

Selected Official References

The following sources are useful starting points for checking terminology, current editions, and regulatory expectations. Standards and guidance should always be read in the edition applicable to the organization and market.

Quality management and standards

International Organization for Standardization, ISO 9000:2026, Quality management - Fundamentals and vocabulary.

International Organization for Standardization, ISO 9001:2015, Quality management systems - Requirements. ISO 9001:2026 was under publication at the time of writing.

International Organization for Standardization, ISO 13485:2016, Medical devices - Quality management systems - Requirements for regulatory purposes.

International Organization for Standardization, ISO/IEC 17025:2017, General requirements for the competence of testing and calibration laboratories.

Conformity assessment

International Organization for Standardization, Conformity assessment and certification resources from ISO/CASCO.

United States regulatory guidance

U.S. Food and Drug Administration, Process Validation: General Principles and Practices.

U.S. Food and Drug Administration, Quality Management System Regulation, 21 CFR Part 820, effective February 2, 2026.