



InduX Manufacturing Risk Handbook

A director's framework for understanding change, dependency and resilience

FIRST EDITION · SEPTEMBER 2026

*Every successful manufacturing business changes.
The question isn't whether your business has changed.
It's whether your understanding of its risk has changed with it.*

Contents

Chapter One	The Shape of Manufacturing Risk	3
Chapter Two	Building a Risk-Aware Operation	9
Chapter Three	Growth & Change	14
Chapter Four	Operational Resilience	19
Chapter Five	Financial Exposure	24
Chapter Six	People & Workforce	29
Chapter Seven	Claims & Defensibility	35
Chapter Eight	Emerging Risk	40
Chapter Nine	Working With Advisers and the Insurance Market	45
Chapter Ten	Building a Risk-Resilient Manufacturing Business	51
	Glossary	57
	Source and verification notes	58
	About InduX	59

InduX Manufacturing Risk Handbook

A director's framework for understanding change, dependency and resilience

First edition - September 2026

Every successful manufacturing business changes. The question isn't whether your business has changed. It's whether your understanding of its risk has changed with it.

About this handbook

This handbook is written for directors and leadership teams in UK manufacturing businesses. It offers a practical way to examine how business change alters exposure, creates dependencies and affects the organisation's ability to withstand disruption, fund consequences and demonstrate what it did.

It is manufacturer-first. Insurance is considered where it is relevant as one possible way to finance or transfer defined loss, but it is not the architecture of the framework and it is not a substitute for operational control.

Evidence labels used in this edition

- **Real anonymised case:** drawn from a genuine case with identifying details removed.
- **Composite scenario:** combines recurring patterns from more than one situation and is not a real business.
- **Illustrative scenario:** created to explain a principle, calculation or decision pathway.

Every named business in the chapters that follow is a composite scenario.

General disclaimer

This handbook provides general information about manufacturing risk and governance. It is not legal, regulatory, health and safety, engineering, cyber, tax, accounting or regulated insurance advice. Requirements and policy outcomes depend on the activity, jurisdiction, contracts, wording and facts. Obtain appropriately authorised or qualified advice before acting where specialist judgement is required.

Regulatory and guidance references were checked against official primary sources for this first edition on 9 September 2026.

CHAPTER ONE

The Shape of Manufacturing Risk

Why the risk picture moves every time the business does

Calder Precision Components is a composite scenario. It is not a real business, and no single manufacturer or client sits behind it. It combines patterns that recur across growing manufacturing SMEs to show how otherwise sensible decisions can interact.

Eighteen months ago Calder was a sixty-person subcontract machining business, steady turnover, a broad customer base, nothing on the risk register that would keep a director awake. Then three things happened, each of them a good decision on its own.

A major automotive tier-one awarded Calder a three-year supply contract. It was the kind of win the business had chased for years, and it took turnover from a spread of forty customers to one account representing over a third of it. To deliver the volume, Calder bought a five-axis machining centre, a serious piece of capital equipment that became the only machine in the building capable of running the new parts to tolerance. And because the new centre needed programming and maintenance nobody on the existing team could do, Calder brought in one specialist engineer who now holds, largely in his own head, the knowledge that keeps the contract running.

None of that was reckless. It was ambitious, well-financed, and exactly what growth is meant to look like. But look at what it produced. One customer now represents over a third of revenue. One machine now represents the entire capacity for the product line that customer buys. One person now represents the only route to keeping that machine running. Three separate decisions, each sound in isolation, have created a concentrated chain of dependencies that did not exist eighteen months earlier.

Calder's insurance schedule has not changed in that time. Neither has its risk assessment. The business that exists today and the business described on last year's renewal are no longer quite the same business, and nobody made a decision to let that happen. It simply accumulated, one contract, one machine and one hire at a time.

That is the shape of manufacturing risk. This chapter sets out why it behaves the way it does, and introduces the framework the rest of this handbook uses to examine it.

1. Why manufacturing risk behaves differently

Manufacturing risk is not a list of hazards waiting to be priced. It is the product of a system, and systems behave differently to lists.

A manufacturing business links production processes, physical assets, contractual commitments, and specialist knowledge into a single operating chain. Pull on any one link and the others move with it. A new customer contract changes liability exposure and delivery pressure at the same time. New machinery changes both capacity and the skills needed to run it. A single supplier decision can alter

your exposure to a whole region you had never previously thought about. Nothing on a manufacturing shop floor sits in isolation, which means nothing in a manufacturing risk profile does either.

The physical assets involved tend to carry more weight than they do elsewhere too. A specialist machine is rarely a commodity item that can be replaced from a catalogue within days. It may be built to order, imported, or dependent on an installation and commissioning process that takes months before it produces a single usable part. The capital tied up in plant, and the lead time behind replacing it, both sit well above what most other sectors carry as a matter of course.

For a manufacturer, recovery may involve replacing specialist machinery, restoring a certified process, recommissioning equipment and proving output to a customer's required standard. The physical asset may be only the beginning of the recovery journey. The gap between disruption and recovery is measured in production days, customer commitments and cash - not only in the cost of replacing property.

Supply chains in manufacturing tend to run tight rather than resilient, particularly where just-in-time delivery or single-source materials are part of the model. A disruption two tiers upstream, at a supplier you may not even hold a direct relationship with, can stop your own production line just as effectively as a fire in your own building.

And some of the consequences of what happens on a shop floor today do not surface for years. An exposure created now, a process change, a material substitution, a piece of equipment run without the right guarding, can produce a claim or a liability that lands on a desk long after the decision that caused it has been forgotten. Manufacturing risk does not always announce itself at the point it is created.

None of this makes manufacturing uniquely dangerous. It makes it interconnected, and interconnected systems need to be understood as systems, not audited as checklists.

2. Change is the trigger

Where risk review is tied mainly to insurance renewal or an annual management cycle, the formal picture may be updated only once a year. Manufacturing risk does not wait for that calendar. It moves when the business does.

Every one of the following is an ordinary commercial decision, made for good reasons, that also changes what the business is exposed to the day it happens, not the day the policy comes up for renewal:

Buying new machinery, particularly anything that becomes the sole means of producing a product line. Winning or losing a major contract, especially one that concentrates revenue or introduces liability terms nobody has read closely. Automating a process or connecting equipment to a network for the first time. Moving premises, adding a site, or extending an existing one. Exporting for the first time, or entering a market with a different legal and product liability environment. Acquiring another business, and inheriting whatever risk history comes with it. Growing headcount quickly, or coming to rely more heavily on agency and contract labour. Changing a critical supplier, bringing a process in-house, or sending one out.

Calder's three decisions all sit on that list. None of them would have prompted a phone call to a broker under most renewal cycles, because none of them looked, on their own, like an insurance event. Together, they were the most significant change to the business's risk profile in years, and the business carried on trading against a risk picture that had already moved without anyone deciding it should.

The practical implication is straightforward. Waiting for renewal to review what has changed means running for up to twelve months on a risk picture that no longer describes the business. A machine bought in January and a contract signed in June sit unreviewed until the following March, by which point both have been operating, unexamined, for the best part of a year. A stronger operating discipline is to treat significant change itself as the trigger for review.

3. The six pillars

To make a system this interconnected reviewable, it helps to break it into parts that leadership can actually interrogate, without pretending those parts operate independently. That is what the six pillars are for. Each one asks a single leadership question. None of them is taught in full here. Chapters Three to Eight take the six pillars in turn. Chapter Nine examines how the manufacturer works with advisers, and Chapter Ten turns the framework into a practical operating rhythm. This is the map, not the territory.

Growth & Change. Has the business's understanding of risk moved with its decisions, growth and operating model?

Operational Resilience. What could stop production or delivery, how quickly would the effect spread, and how would the business recover?

Financial Exposure. Which losses, commitments or interruptions could the business not absorb, and what assumptions support the protection it believes it has?

People & Workforce. Which capabilities, decisions, relationships or authorisations depend on particular people, and can workforce change be absorbed safely?

Claims & Defensibility. Could the business demonstrate what it knew, decided, controlled and did if externally scrutinised?

Emerging Risk. What is changing around or through the business, how could it reach operations, and what would trigger action?

Calder's situation touches at least four of the six before a single insurance question is asked: Growth & Change, because the contract and the machine both moved the business on from what its risk arrangements assume. Operational Resilience, because one machine now carries a product line. People & Workforce, because one engineer now carries the machine. Financial Exposure, because a third of turnover now sits with a single customer. That overlap is not a coincidence. It is the normal shape of a real exposure, and it is the reason the next section exists.

4. Dependencies and concentration

A pillar tells you where to look. A dependency tells you what you would actually lose.

Across the six pillars, nine categories of dependency turn up again and again in manufacturing businesses: a major customer or connected customer group, a contract or manufacturing programme, a site used for production, storage or testing, a technology, including a critical machine, production line or process, a digital system or piece of operational technology, a specialist person holding skill or authority that sits nowhere else, a supplier or subcontractor, a particular material, component or approved specification, and a country, region or transport route the business depends on to move goods or materials.

Most businesses can point to one or two of these without much thought. A director will usually know, roughly, who their biggest customer is or which machine would hurt most to lose. Fewer have ever mapped where several of the nine sit on top of one another, because that picture only appears when someone deliberately draws the connections rather than considering each dependency on its own.

The more serious exposures tend to sit exactly where several dependencies overlap on the same point of failure, as they do at Calder. The customer, the machine and the specialist engineer are three separate dependency categories, but they concentrate around one contract. Lose any one of the three and the other two become far less valuable, because none of them functions without the others. That is concentration, and it is very easy to build without ever intending to, one sensible decision at a time.

Identifying a dependency is not, on its own, a problem to fix. Every manufacturer depends on people, machines and customers. The question worth asking is whether the dependency has been identified at all, and whether anyone has tested what happens if it is removed.

5. How InduX examines risk

Every chapter in this handbook, and the InduX Risk360 review itself, works through the same sequence:

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Start with what changed. Work out what exposure that change created or shifted. Identify which dependencies sit behind that exposure. Establish what the impact would actually be under a credible disruption, including 30 days without the dependency. Look at what controls currently exist against it. Check whether the business could produce evidence that those controls are real, not assumed. Only then decide on a response, which may or may not involve insurance at all.

Run Calder's contract win through it. The change is the new three-year agreement. The exposure includes the liability and delivery terms Calder signed, and the capacity now committed to one customer. The dependencies are the customer, the machining centre, and the engineer, sitting on top of each other. The impact of losing any one of the three, tomorrow, is a production stoppage on the highest-value line in the business, with no fallback capacity and no second person who can run the machine. The controls worth checking are whether the contract terms were reviewed before signature, whether the machine has a maintenance and breakdown plan, and whether the engineer's knowledge

has been documented anywhere. The defensibility question is whether any of that could be evidenced to a customer, an insurer or a court if it mattered.

And the response, once all of that is understood, is rarely a single action. For Calder it might mean cross-training a second operator on the new machine, having the contract's liability clauses reviewed before the next renewal comes round, documenting the engineer's maintenance and programming knowledge properly for the first time, and only after that, if a genuine financial gap remains, adjusting insurance arrangements to reflect a business that looks nothing like the one on last year's schedule.

Insurance sits at the end of that sequence, as one possible response among several. It is rarely the first thing to fix, and it is never the only thing that needs fixing.

6. Why insured does not automatically mean protected

Holding an insurance policy confirms that certain risks may be covered subject to its wording, limits, exclusions, conditions and the circumstances of a loss. It does not establish that declared values remain accurate, that the period allowed for recovery is sufficient, that the evidence required for a claim exists, or that every material exposure facing the business is insured.

Chapter Five examines financial exposure, including underinsurance, recovery assumptions and the distinction between funding a loss and preventing one. Chapter Seven examines the evidence and response capability required when an incident, claim or other challenge occurs.

7. How to use this handbook

This is a director's working framework, not a legal manual, an insurance policy guide, or a compliance checklist. Nothing in it replaces regulated advice, and nothing in it should be read as a complete statement of the law in any area it touches.

What it is built to do is give a manufacturing leadership team a shared way of talking about risk that starts from the business itself, not from an insurance product. Each chapter from here takes one pillar and works through it properly. You do not need to read them in order. The most useful way to use this handbook is to start with whatever has actually changed in your business recently, find the chapter that covers it, and work through the sequence above against your own operation.

8. Five board questions

What has changed in this business in the last twelve months that nobody has formally reviewed?

Where do two or more dependencies, a customer, a machine, a site, a person, a supplier, currently sit on top of each other?

What assumption about this business's resilience has never actually been tested?

If we had to prove tomorrow that a specific risk had been properly managed, what evidence would we produce, and does it exist?

Which of our current arrangements were sized for the business as it was, rather than the business as it is now?

9. One immediate exercise

Pick the single most consequential change your business has made in the last twelve months. It might be a contract, a machine, an acquisition, a site move, or a person. Work it through the sequence in this chapter: what changed, what exposure it created, which dependencies sit behind it, what the credible impact would be across the first 30 days if it became unavailable, what controls exist, whether you could evidence them, and what response actually follows.

That one exercise, done properly on one change, will tell you more about where to focus than a general review of everything at once. Risk360 applies the same sequence to help a manufacturer identify where further review may be most valuable.

CHAPTER TWO

Building a Risk-Aware Operation

Compliance, controls and evidence

Brennock Sheet Metal is a composite scenario. It is not a real business and no single manufacturer sits behind it. It combines patterns that recur across manufacturing SMEs to show how gaps develop between changing operations, working controls and the evidence intended to support them.

Brennock is a forty-five person fabricator, laser cutting and folding for construction and shopfitting customers. The risk assessment pack in the site office is genuinely thorough, professionally written, and three years old. In that time Brennock installed a new fibre laser cutter, brought two agency operators on for a busy quarter, and moved a travelling crane into daily use on the new line to handle heavier sheet stock. None of it was hidden. None of it was reckless. It simply happened at the pace the business happened, while the paperwork stayed where it was.

Then a near miss. An agency operator overrides an interlock guard on the new laser to clear a jam, gets away with it, and a supervisor reports it rather than letting it go. The investigation that follows, led by Brennock's own team with support from a competent health and safety adviser, asks three questions in quick succession. Does the risk assessment for that machine actually describe the machine now on the floor? Had the agency operator received adequate machine-specific information, instruction, training and supervision - or only a generic induction that predated the laser? And when was the crane last thoroughly examined, and where is the certificate?

Brennock could answer none of the three quickly. Not because nothing had been done, but because what had been done was not written down anywhere a stranger could find it under pressure.

That is the gap this chapter is about. Not the absence of good intentions or good practice, which Brennock had plenty of, but the absence of a system that keeps assessment, competence and evidence moving at the same pace as the business itself.

1. Why controls and evidence behave differently in manufacturing

A risk assessment is useful only while it remains connected to the work it describes. Machinery, materials, guarding, software, layout, staffing and supervision can all change the way a process is actually performed. The relevant question is therefore not how old the document is, but whether it remains suitable and sufficient for the current activity and whether a material change should have triggered review.

Manufacturing evidence is often inseparable from the physical process. Maintenance logs, calibration certificates, inspection reports, competence records, permits, batch records and photographs may be needed to show what equipment, material and control existed at a particular time. If the operating system does not create and retain that evidence as work is performed, it may be impossible to reconstruct reliably after an incident.

Competence is equally specific. A general induction does not by itself demonstrate that a person is trained, assessed and authorised to use a particular machine or carry out a critical process. The evidence should match the task, the equipment, the individual and the level of supervision actually provided.

2. Change is the trigger, again

Chapter One set out the list of ordinary commercial decisions that move a manufacturer's risk position: new machinery, new contracts, new premises, new people, new suppliers. The same list applies here, pointed at a narrower question. Every one of those changes also creates, or invalidates, a specific control, a specific training requirement, or a specific piece of evidence.

A new machine needs its own risk assessment and, where it involves lifting, pressure or particular categories of equipment, its own statutory inspection regime from the day it arrives, not from the next scheduled review. A new material needs a COSHH assessment before it is used, not after an exposure. A new process needs a method statement that describes what actually happens, not what the previous process happened to involve. A new starter, agency or permanent, needs training signed off against the specific equipment they will use, not a generic site induction. A new customer contract often comes with an audit clause that assumes all of the above already exists in a form someone else can read.

Even a quieter change counts. A shift pattern extended to cover a busy period changes who is supervising a process and when. A supplier substitution changes the material behind a COSHH assessment written for a different one. None of these announce themselves as compliance events, which is exactly why they get missed.

Brennock's laser cutter, its agency operators and its crane are three instances of the same pattern. Each one needed something updating the day it arrived. None of them got it, because nothing in the business's routine treated the arrival itself as the trigger.

3. The six building blocks of a risk-aware operation

Six things, done consistently, are what separate a business that manages this well from one that discovers the gap during an investigation.

Risk assessments that match reality. Written to describe the process as it is actually run today, reviewed when the process changes rather than on a fixed annual cycle regardless of what has happened in between. A risk assessment that has not changed in three years is not evidence of stability. It is usually evidence that nobody has checked.

Control ownership. Every control, a guard, a permit system, an inspection schedule, belongs to a named person who knows it is theirs, not to "health and safety" as an abstract department that exists somewhere else in the business. A control with no named owner tends to be the one nobody notices has failed.

Competence. Training and certification that map directly onto the equipment and processes actually in use, refreshed when either changes, and checked rather than assumed for agency and temporary labour. A signature on an induction form from an agency is not the same thing as evidence of machine-specific competence.

Documentation. Proportionate to the risk, not bureaucratic for its own sake. A one-page method statement that gets read is worth more than a comprehensive manual that does not, and a business that over-documents low-risk activity often has less capacity left to document the activity that actually matters.

Escalation. A route for a near miss, a concern or a broken control to reach someone with the authority to act on it, used often enough that using it feels normal rather than exceptional. Brennock's supervisor reporting the interlock override, rather than letting it pass, is escalation working as intended. The system around it had not kept up.

Evidence that survives scrutiny. The output of the previous five, kept somewhere it can actually be found and produced, on the day someone outside the business asks to see it, rather than reconstructed under pressure once they already have.

Brennock had elements of most of these. What it lacked was the connective tissue between them, the habit of updating all six together whenever something changed, rather than treating each one as a separate exercise on its own schedule.

4. Where the law sets the floor

The controls in this chapter are an operating discipline. Underneath them sit legal duties that depend on the activity, equipment, substances, product and market involved. This section is an orientation, not a complete statement of law.

Risk assessment and safe work. Employers must make a suitable and sufficient assessment of risks to employees and others affected by the undertaking. Where machinery, material, layout, staffing or the way work is organised changes, the assessment and any safe system of work should be reviewed in proportion to the risk and complexity of the activity.

Information, instruction, training and supervision. Permanent, temporary and agency workers need adequate information, instruction, training and supervision for the work they are assigned. A process-specific competence record should show who has been trained, assessed and authorised rather than relying on a generic induction or an untested assurance from a labour supplier.

Plant and equipment. LOLER requires qualifying lifting equipment to undergo thorough examination at statutory intervals or in accordance with an examination scheme. PUWER inspection intervals are determined by the equipment, risks and circumstances rather than one universal timetable. Before qualifying pressure equipment is used, a written scheme of examination must be in place and an examination undertaken. Applicability and frequency should be established with a competent person.

Hazardous substances and occupational health. COSHH assessments, exposure controls and health-surveillance arrangements should match the substances and work actually present. Some health-surveillance records have long statutory retention periods, including 40 years in specified COSHH circumstances.

Environmental responsibilities. Changes to processes, materials, waste streams, emissions, storage or site capacity may alter permit, exemption, consent and waste-duty-of-care requirements. Establish the position before the change is implemented and retain the evidence supporting it.

Product compliance. The applicable conformity route depends on the product, the manufacturer's role and the market in which it is placed. CE and UKCA recognition has changed repeatedly; establish and verify the current requirements for each product and destination instead of assuming one mark applies universally.

Data protection. Training, health, incident and employment records may contain personal data. Keep what is necessary for a defined purpose, restrict access and apply a documented retention schedule; long retention should be based on a specific legal or evidential need.

Leadership responsibility. Directors and senior officers can face personal consequences where the statutory tests for consent, connivance or neglect are met. Criminal fines are not generally insurable; any cover for defence costs depends on the policy and circumstances. Seek qualified advice after a serious incident, material near miss, enforcement contact or when the required competence is not available internally.

5. How InduX examines this

The same sequence from Chapter One applies here.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Brennock's laser cutter through it. The change is the new machine, plus the agency labour brought in around the same time. The exposure is a machine operating without a matching risk assessment, and operators without machine-specific training on record. The dependency is the single new production line, and the specific interlock control that got bypassed. The event was a near miss, not a claim. The investigation should establish whether the engineering control was adequate, whether foreseeable misuse had been considered, how the procedure, competence and supervision worked in practice, what production pressure existed and why the override occurred. It should not assume in advance that the interlock was sound or that training alone explains the event. The defensibility question, the one Brennock could not answer quickly, was whether any of that could be evidenced on request.

The response is not a single fix. It is a new risk assessment for the actual machine in use, machine-specific sign-off for every operator including agency staff, a confirmed inspection date for the crane, and a habit of treating the next piece of new equipment as the trigger rather than waiting for the next scheduled review. None of this begins as an insurance exercise. It begins with a competent review before equipment is put into use, followed by controls, training and evidence proportionate to the actual risk.

6. From evidence to defensibility

Everything in this chapter produces the same output: a record that something was properly managed. That record has no value sitting in a folder nobody opens. Its value is realised only when it is actually needed, in an investigation, an audit or a claim, and can be produced quickly and stand up to scrutiny. Chapter Seven, Claims & Defensibility, takes that idea further and sets out what "stands up to scrutiny" actually requires when something has already gone wrong.

7. Five board questions

Which of our current risk assessments describes a version of the process that no longer exists?

Who personally owns each of our significant controls, by name, not by department?

Could we produce, today, the training record for every agency or temporary worker currently on site?

When were our lifting equipment, pressure systems and other statutory inspections last checked against what is actually due?

If a regulator or an insurer asked to see our evidence on any one control tomorrow, how long would it take us to find it?

8. One immediate exercise

Pick one piece of equipment, one process or one control that has changed in the last twelve months. Find its risk assessment. Check the date against the date of the change. Find the training records for everyone who uses it, including anyone temporary or agency. Find any statutory inspection certificate that applies to it, and check it is current rather than merely on file. If any one of those four is missing, out of date, or hard to locate, that is this chapter's exercise completed. Fix that one gap before moving on to the next.

CHAPTER THREE

Growth & Change

When the business moves faster than the risk picture

Fenwick Componentry is a composite scenario. It is not a real business, and no single manufacturer or client sits behind it. It combines recurring patterns to show how several reasonable growth decisions can interact when an underlying weakness is exposed.

Fenwick is a ninety-person precision plastics and electronics manufacturer. Over fourteen months it did three things, each one a clear commercial win. It signed a large new UK contract with a tier-one customer, on terms drafted by the customer's own legal team and signed quickly to secure the order. It acquired a smaller regional competitor, gaining capacity, a customer book, and a set of tooling nobody at Fenwick had used before. And it shipped its first export order, to a US medical device assembler, its first exposure to a product liability environment built on an entirely different legal foundation to the one at home.

Then a batch of components supplied under the new UK contract came back with a dimensional defect. The cause traced to a legacy mould inherited in the acquisition, one that had gone straight into production without being re-validated against Fenwick's own tolerances. The UK customer invoked the indemnity clause in the contract it had signed, a clause that assigned liability well beyond what ordinary negligence law would have imposed. Looking further, two of the acquired business's older product lines were still in the field with no liability history ever collected on them. And because a portion of the same production run had gone to the US customer, the same defect now had to be answered against a different legal system entirely.

None of Fenwick's three decisions was reckless. Each was exactly the kind of move a growing manufacturer is meant to make. Together, at the moment a single defect surfaced, they produced a problem that touched contract law, inherited liability and cross-border product regulation all at once, and none of it had been reviewed as one picture at any point along the way. Fenwick's turnover had grown by almost half in the same fourteen months. Nobody had asked, in that time, whether anything else needed to grow alongside it.

Growth & Change asks: How has the business changed, and has its understanding, control and financing of risk changed with it?

1. Why growth changes risk differently in manufacturing

Risk instruments update once a year. Growth decisions do not wait for that date, and in a manufacturing business they rarely arrive one at a time.

Manufacturing growth can change several commitments at once because capacity, equipment, people, materials, contracts, quality and working capital are connected. A new contract may require machinery, recruitment, supplier capacity, revised controls and cash before the first customer payment arrives. A new contract needs machinery to fulfil it. New machinery needs people to run it. New volume

needs suppliers who can deliver at the new scale. A single growth decision drags several others behind it, and each one carries its own exposure.

Manufacturing growth decisions are also harder to undo than most. A services business that overextends can usually scale back. A manufacturer that has signed a three-year supply contract, bought a machine to fulfil it, or absorbed another company's product range cannot simply reverse the decision if the risk picture turns out to be wrong. The commitments are structural, not seasonal, and they tend to stay structural for years after the decision that created them.

The interconnection described in Chapter One is what makes this compound. A contract does not sit next to a machine and a workforce as three separate items. It sits on top of them, and growing the contract without checking whether the machine and the workforce underneath it have kept pace is how a business ends up structurally exposed without a single reckless decision anywhere in the chain.

Recovery capacity moves with growth too, and rarely in the direction anyone notices. A business running at sixty per cent of its old capacity had headroom if a machine failed or an order was lost. The same business running flat out on new volume, on the equipment it had before growth arrived, often has none left at all.

2. Review before commitment, not after implementation

The trap is familiar and easy to fall into. Something changes, it gets noted mentally as worth mentioning at the next renewal, and the diary moves on. Nine months of growth is nine months of the business being materially different from the one your risk arrangements describe, and the gap does not pause because the renewal date has not arrived yet.

What you disclose to an insurer, and when, is its own subject and gets the attention it deserves in Chapter Seven. The point here is narrower and comes before any insurance conversation. A growth decision changes the operational and contractual shape of the business the day it is made. Fenwick's contract, acquisition and export order were each worth a conversation with someone about risk on the day they happened, not nine months later when a defect forced the conversation regardless.

3. The numbers move before the protection does

Growth changes more than turnover. It can alter:

- working-capital requirements and the timing gap between cost and payment;
- peak stock, work in progress and customer-owned property;
- customer and programme concentration;
- debt, lease payments and lender covenants;
- fixed costs and the minimum viable level of utilisation;
- reliance on subcontractors and suppliers;
- replacement values for buildings, machinery and stock;

- the time and cash required to restore output after disruption.

These numbers should be revalidated when the decision is made. Where business interruption insurance forms part of the response, its financial assumptions and maximum indemnity period should be checked against the current recovery plan rather than carried forward automatically.

4. Growth creates new exposures, not just bigger ones

The part that gets missed most often is that growth does not simply enlarge the risks a business already carries. It introduces categories of exposure that were not there before.

New products carry new liability. A product line the business has never made before is a design and specification risk before it becomes anything else, and the decisions behind it are usually made quickly, under commercial pressure, by people not thinking about liability at the time.

New markets carry new legal systems. Fenwick's US order is the clearest example. Exporting into a different jurisdiction is not an incremental change to an existing exposure. It is a different liability environment, different procedural rules, and often a different cost of defending a claim, arrived at the moment the first shipment leaves the building.

New customers carry their contracts. A larger customer brings its own terms, and those terms allocate risk in ways a standard supply relationship would not. Section six of this chapter looks at exactly that.

New machinery carries new duties. Beyond the capacity it adds, new equipment brings statutory inspection obligations, training requirements and single points of failure that Chapter Two's controls and evidence discipline needs to catch on day one, not at the next scheduled review.

New people carry new gaps. Rapid hiring compresses training, increases reliance on agency and contract labour, and puts the newest people closest to hazards they understand least. Chapter Six takes this further.

Subcontracting carries someone else's process. When demand outpaces capacity, work gets placed out. The customer still holds the original business responsible for quality, even though it no longer controls how the work is actually done, and that responsibility rarely travels back down to the subcontractor on equal terms. A subcontractor's quality failure becomes the manufacturer's problem to explain to the end customer, even when the manufacturer never touched the part.

5. Acquisition brings history

An acquisition can bring inherited liabilities, records, people, products, contracts, tooling, environmental obligations and claims history. What transfers depends on whether the transaction is an asset purchase, share purchase or another structure, the contracts involved and the applicable law. Due diligence should therefore identify which obligations and evidence move with the business and which remain elsewhere.

Financial diligence alone is not a complete assessment of manufacturing risk. Operational, product, workforce, environmental, cyber and insurance questions need owners before completion, with specialist legal or technical advice where necessary.

6. Contract terms that change the exposure

Commercial terms can change the manufacturer's exposure even when the physical work looks familiar. Escalate contracts containing uncapped or one-sided indemnities, liquidated damages, consequential-loss provisions, extended warranties or guarantees, recall or rectification obligations, customer tooling or stock responsibilities, intellectual-property commitments, safety-critical supply, non-UK law or obligations continuing after the contract ends. Boards should also set value and concentration thresholds that trigger review before signature.

For a new export market, establish the applicable contractual, regulatory and product-liability requirements rather than assuming the UK position travels with the product. Obtain legal and insurance review where indemnities, jurisdiction or customer requirements are unfamiliar.

7. If the growth assumption fails for 30 days

Test what happens if the new contract, machine, acquisition or market delivers no usable output or cash for its first 30 days. Identify committed wages, materials, finance payments, subcontracting costs and customer penalties; then establish which dependencies share the same assumption. This exposes whether growth has created a temporary funding gap, a concentrated programme or a permanent increase in the business's minimum cost base.

8. How InduX examines this

The same sequence runs through every chapter in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Fenwick's fourteen months through it. The change is three decisions taken close together: the new contract, the acquisition, the export order. The exposure includes the indemnity clause signed without full review, an unvalidated inherited tool now in production, and a product liability question that now has to be answered under two legal systems rather than one. The dependency is the acquired tooling and the single new customer contract, sitting on top of each other in the same defective batch. The impact, once the defect surfaced, was a claim broader than ordinary negligence liability, a legacy product tail with no history behind it, and a cross-border question nobody had prepared for. The controls worth checking, in hindsight, were whether the contract had been reviewed by anyone other than sales, whether inherited tooling had been re-validated before use, and whether the acquisition due diligence had asked about historic liability at all. The defensibility question was whether Fenwick could show any of that had been considered before the defect forced it to be.

The response, once understood this way, is rarely a single fix. For a business in Fenwick's position it typically means a legal review of the indemnity terms already signed, a validation programme for every piece of inherited tooling still in production, a proper liability history exercise on the acquired product lines, and a standing rule that any contract above a certain value gets a risk review before signature, not after a claim.

9. What this feeds into

Growth changes more than the picture this chapter has covered. It moves the values a business should be insuring, often faster than anyone realises, which is exactly the ground Chapter Five, Financial Exposure, covers next in full. And the rapid hiring that usually comes with growth creates its own dependencies on people, which Chapter Six, People & Workforce, takes on directly. Neither chapter needs reading before this one makes sense, but both extend it.

10. Five board questions

Which of the growth decisions made in the last twelve months has not yet had a proper risk conversation attached to it?

Do we know, clause by clause, what liability we have contractually accepted under our largest current customer contract?

If we have acquired a business in the last three years, has its historic liability and product history actually been catalogued?

Where has growth put a single customer, machine, site or supplier in a position to materially affect the whole business if it failed?

What is our standing rule for reviewing risk before a significant contract is signed, rather than after it causes a problem?

11. One immediate exercise

Pick the single most significant growth decision your business has made in the last twelve months, a contract, an acquisition, a new market, a major machine, or a period of rapid hiring. Find out whether anyone reviewed its risk implications at the time it was made, rather than at the next renewal. If it involved a contract, find the indemnity clause and read it properly. If it involved an acquisition, check what liability history came with it. If nobody can answer these quickly, that gap is this chapter's exercise, and it is worth closing before the next growth decision adds another one on top of it.

CHAPTER FOUR

Operational Resilience

How long can the business keep delivering when a critical dependency becomes unavailable?

Marlow Manufacturing is a composite scenario. It is not a real business, and its events, timescales and outcomes combine recurring manufacturing patterns to demonstrate how connected dependencies can affect production and delivery.

Marlow is a seventy-person injection moulder, supplying flame-retardant polymer housings to industrial and automotive customers. Two things sat behind its two highest-value product lines, neither of which anyone thought of as fragile until the week both of them failed. The flame-retardant compound used on its largest contract came from a single supplier, the only one on the customer's approved materials list qualified to that specification. And the tool for that same contract only fitted one press in the building, the sole large-tonnage machine Marlow owned capable of running it.

In the space of three weeks, the supplier had a fire that took its own production offline for months, and Marlow's press suffered a hydraulic failure that needed a part with a fourteen-week lead time from overseas. Marlow could not simply switch material suppliers, because the customer's specification did not allow it without a lengthy requalification process. It could not simply run the tool on another press, because no other press in the building had the tonnage. Two production lines that had never gone down together went down together, for reasons that had nothing to do with each other and everything to do with the fact that neither had a second option behind it.

Marlow's machines were well maintained. Its supplier was a good one. Nothing here was a failure of standards. It was a failure to ask, before either event happened, what would actually happen if either of these single points failed on its own.

The two failures were not connected in any technical sense. They were connected only by the fact that Marlow's highest-value contract had been allowed to depend on two separate single points of failure at once, without anyone stepping back to look at the contract as a whole rather than as two unrelated supply arrangements.

Operational Resilience asks: What could stop production or delivery, how long would recovery genuinely take, and could the business withstand that period?

1. Why operational resilience behaves differently in manufacturing

Manufacturing output is often physically bound to particular assets, materials, processes, approvals and places. Losing one machine, supplier, site or system can therefore remove the ability to produce a specified part even when demand, people and commercial intent remain.

Substitution is also slower and harder than it looks from outside the business. A specialist press cannot be hired for the afternoon the way a laptop can be replaced. A qualified material cannot be swapped for a chemically similar one without requalification, and where a customer specification

names an approved supplier, as it did for Marlow, substitution is not a technical decision at all, it is a contractual one that takes months to clear.

Manufacturing resilience has to be planned before disruption. A workaround for a failed press, unavailable material or lost process may still need to satisfy customer specifications, safety requirements, validation and certification; those constraints cannot be waived simply because the planned route is unavailable.

Growth adds to this rather than reducing it. A business running comfortably below capacity has somewhere to absorb a shock. The same business running flat out against demand, exactly the position growth tends to create, usually has nowhere left to put the load when one part of the system stops.

The consequences also travel further than the immediate loss of output. A missed delivery on a just-in-time contract can trigger penalty clauses, damage a relationship a business spent years building, or hand the volume to a competitor who happened to have spare capacity at the right moment. The cost of a stopped line is rarely limited to the days it stayed stopped.

2. Map the dependencies behind delivery

Operational continuity depends on the same nine categories used throughout this handbook:

- **Customer and programme:** which orders, approvals or revenue streams depend on the operation;
- **Site:** where production, storage, utilities, testing or dispatch must occur;
- **Technology:** machinery, tooling, processes and utilities required for output;
- **Digital system:** software, networks, control systems, data and access needed to operate;
- **Specialist person:** competence, authority or knowledge that is thinly spread;
- **Supplier and material:** external capability, components, consumables and approved specifications;
- **Geography:** countries, regions, transport routes or clusters shared by apparently separate sources.

The grouping is for readability, not a separate framework. The critical question is which dependencies support the same output and therefore share a failure path. A fallback person must have the competence, authority, access and recent experience to act. A fallback supplier must be capable, willing, available, approved where necessary and independent of the same upstream or geographic source.

Connected production also creates operational-technology dependency. Remote access, central scheduling, vendor support and shared control systems should be mapped, segmented where appropriate and supported by tested isolation and recovery procedures. A system that cannot be operated or restored without one vendor, credential or individual is an operational dependency even when the machine itself remains intact.

3. The supplier you never named

A single-source supplier is not automatically a problem. Plenty of manufacturers run well on one trusted source for a given material or component. The problem is not having the supplier. It is not knowing, in any structured way, what happens the day that supplier cannot deliver.

This gets harder rather than easier as a business grows and its customers become more demanding. Marlow's flame-retardant compound was not single-sourced by accident or by cost-cutting. It was single-sourced because the customer's own specification approved exactly one supplier, and using an unapproved one would have breached the contract even if a chemically equivalent alternative existed. The constraint sat with the customer, not with Marlow, which is precisely why it is easy to overlook. Nobody chose this dependency. It arrived attached to winning the business.

Knowing which suppliers sit in this position is not a procurement exercise, it is a resilience one. A short list of the handful of suppliers whose failure would actually stop production, checked against whether an alternative exists and how long qualifying it would take, tells a business more about its real exposure than most of what appears on a standard supplier scorecard.

Where a customer specification is the reason for the single source, the conversation worth having is with that customer, not just the supplier. Asking whether a second approved source could be qualified in advance, before it is needed rather than during a crisis, is a request most customers will take seriously once the risk is explained in their own terms rather than yours.

4. The machine that has no backup

Every manufacturer depends on machinery. Far fewer have asked which specific machine, if it broke down tomorrow, would stop a product line with no fallback at all.

Marlow's press was maintained properly and had never given trouble before the failure that stopped it. The issue was never how well it had been looked after. It was that nothing else in the building could do what it did, and the part that failed had a lead time measured in months rather than days. A well-maintained machine with no backup is still a single point of failure. Maintenance reduces the chance of failure. It does not remove the consequence of failure when it happens anyway.

The practical question is not whether a machine might break, because eventually most machines do. It is what the business does in the gap between the breakdown and the fix, whether that gap is measured in days because a spare exists or an alternative machine can be adapted, or in months because neither does. Knowing which answer applies to your own most critical machine, before it fails, is the entire point of this section.

A critical spares list, held for the handful of components with the longest lead times on the handful of machines that matter most, is one of the least expensive pieces of resilience a manufacturer can build. It does nothing to prevent a failure. It does a great deal to shorten the gap that follows one.

5. Recovery time is the real number

Almost every conversation about operational resilience eventually lands on the same figure, and it is rarely the one anyone expected. It is not whether a failure can be fixed. It is how long fixing it actually takes, measured against how long the business can survive without producing.

Marlow's flame-retardant material and its damaged press both had a genuine fix. The supplier would eventually rebuild. The press part would eventually arrive. Neither fact helped in the meantime, because the business was losing production and, with it, customer confidence and cash, for every week the gap lasted. A theoretical recovery that takes four months is not meaningfully different from no recovery at all if the business cannot survive four months without its two highest-value lines running.

This is the number worth establishing honestly for your own most critical dependencies, in weeks, not in the abstract language of "we'd sort something out." Chapter Five picks this up directly from the financial side, because the recovery period is not just an operational question, it is the number an insurance recovery period is meant to match, and very often does not.

Most businesses have never actually timed this. They know roughly how a normal disruption plays out, a late delivery, a short breakdown, because normal disruptions happen often enough to build a feel for them. Nobody has a feel for an event that has never happened before, which is exactly the kind of event a single point of failure produces when it finally fails.

6. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Marlow's two failures through it. The change, in this case, was not a single decision but the accumulation of two contract-driven constraints, an approved single supplier and a single qualified press, both accepted without complaint because both came with winning valuable business. The exposure was production entirely dependent on two assets with no second option behind either. The dependency was the material and the machine, sitting on the same contract, so that either failure alone would have been serious and both together were close to existential. The impact was two production lines down simultaneously, for months, against customers who had their own delivery obligations to meet. The controls worth examining afterwards were whether an alternative material source had ever been explored even informally, and whether any other press in the building could have been adapted, at a cost, to run the tool. The defensibility question was whether Marlow could show its customers, honestly, that this risk had been considered in advance rather than discovered in the moment, and whether the two dependencies had ever appeared together on the same page of any review, rather than being managed by two different people who had never compared notes.

The response that followed was not about avoiding single-source relationships altogether, which is often not realistic. It was about knowing, before the next equivalent contract is signed, exactly what the fallback plan is for the assets that contract depends on, and pricing that risk into the decision rather than finding it afterwards. In practice that meant a second material source put through

qualification even though it was not yet needed, and a costed option to adapt a smaller press as a genuine fallback rather than a theoretical one.

7. What This Feeds Into

The recovery time this chapter has focused on becomes a financial question the moment it is understood operationally. Chapter Five, Financial Exposure, picks up exactly where this chapter leaves off, on whether the business's financial protection is actually sized to the recovery period its own operations require. And the systems point touched on briefly in section two, the operational technology and digital systems increasingly running production directly, gets its full treatment in Chapter Eight, Emerging Risk. Neither chapter needs reading first, but both assume the honest recovery-time work in this one has already been done.

8. Five board questions

Which single machine, if it failed tomorrow, would stop a product line with no realistic backup?

Which suppliers, named or not, would stop production if they failed, and do we actually know who they are?

If our most critical single point of failure happened today, how many weeks before production genuinely resumed?

Where has a customer specification or contract quietly locked us into a single source we would not otherwise have chosen?

Have we ever tested one of these failure scenarios properly, or only assumed we would manage if it happened?

9. One immediate exercise

Identify the one machine, the one supplier and the one site your business could least afford to lose. For each, write down honestly how long recovery would actually take, not how long you would like it to take. Where you cannot answer with any confidence, that is the gap worth closing first, before the next contract adds another single point of failure on top of the ones you already have.

CHAPTER FIVE

Financial Exposure

Insured is not protected

Aldercroft Engineering is a composite scenario. It is not a real business, and its values, events and outcomes are illustrative. It combines recurring manufacturing patterns to show how financial exposure can extend beyond an insurance schedule.

Aldercroft is an eighty-five person structural steel fabricator. A fire in its machine shop destroyed most of the plant on that side of the building. The claim, on paper, looked straightforward. It was not. The sum insured on plant and machinery had been set years earlier and indexed by a small percentage every year since, while replacement costs for the machines actually lost had risen far faster than that indexation ever captured. The insurer's loss adjuster valued the true replacement cost at nearly double the declared sum insured, and applied the condition of average that sits in almost every commercial property policy. The claim was settled at roughly half of what it would have cost to replace the plant like for like.

That was the first gap. The second appeared once the rebuild timeline became clear. Aldercroft's business interruption cover ran for twelve months, a figure nobody had revisited since the policy was first taken out. Clearing the site, agreeing the claim, ordering replacement machinery with lead times measured in months, installing and commissioning it, and rebuilding the customer relationships lost during the gap took the better part of two years. The insurance payments for lost profit stopped at month twelve. The business was still not back to full production.

The third gap was smaller but telling. A batch of components in production at the time of the fire had already shipped to a customer with a dimensional fault nobody had caught before the disruption began. The customer's product liability claim was straightforward, Aldercroft's insurance covered the injury exposure. What it did not cover was the cost of the components themselves, which Aldercroft had to replace at its own expense, because a product liability policy insures the damage a product causes, not the product.

Aldercroft was insured, comprehensively, by any reasonable standard. Premiums had been paid on time for years, and nobody at the business had ever been told the cover was inadequate. It was not protected, in the sense that actually mattered the day the fire happened, and the difference between those two words turned out to be worth several hundred thousand pounds the business did not have.

Financial Exposure asks: Which losses, commitments or interruptions could the business not absorb, and what assumptions support the protection it believes it has?

1. Why "Insured" And "Protected" Are Different Claims

An insurance policy answers a narrow question. Is there a contract that will pay out if a defined event happens, and have all its conditions been met. It says nothing, on its own, about whether the amount it

pays is enough, whether the timescale it covers is long enough, or whether the type of loss that actually happens is even the type the policy was written to respond to.

Chapter One flagged this gap. This chapter is where it gets worked through properly, because it is the single most common way a manufacturer discovers, at the worst possible moment, that having a policy and being financially protected were never quite the same thing.

The gap is rarely visible before a claim, because a policy that has never been tested looks identical whether it is right or badly short. Aldercroft's schedule read the same, renewal after renewal, whether the sum insured was accurate or half of what it should have been. Nothing about the paperwork changed. Only the underlying value did.

2. The financial exposure map

Financial exposure is the part of disruption, liability or change the business would have to fund or absorb. It includes:

- property and machinery replacement values;
- lost gross profit, continuing fixed costs and additional expenditure during interruption;
- debt, lease payments and lender covenants;
- customer penalties, rework, recall and rectification costs;
- uninsured or excluded losses and policy excesses;
- working-capital pressure, stock peaks and payment delays;
- concentration of turnover or profit in one customer or programme;
- contingent interruption through suppliers, utilities, systems or sites;
- the cost of retaining people and capability during recovery.

Insurance may finance part of this exposure. It does not define the exposure and it cannot replace a credible recovery or cash plan.

3. How underinsurance may work

Where the wording contains an average, underinsurance or other valuation provision, a claim may be reduced when the declared value or sum insured is inadequate. The mechanism and outcome depend on the policy wording, valuation basis, facts and any special terms; it should not be assumed to operate identically across policies.

For illustration only, if a value that should have been declared at £2 million is insured at £1 million and the wording applies average proportionately, a £400,000 covered loss could be reduced to £200,000 before any excess. The calculation demonstrates the exposure; it does not predict how a particular claim will be settled.

4. Where values quietly drift

Nobody sets out to be underinsured. It happens through a small number of routes that repeat across almost every manufacturer who discovers the gap the hard way.

Values get set once and then indexed by a modest percentage each year rather than genuinely reassessed, while the actual cost of building materials, machinery and specialist labour moves in steps that indexation rarely keeps pace with. Buildings get insured at market value rather than rebuild cost, and for an industrial property the two can diverge sharply, because rebuild means demolition, debris removal, professional fees and construction to current building and environmental standards, considerably more demanding than whatever applied when the building went up decades earlier. Plant gets valued as it currently sits rather than as it would actually be replaced, installed, commissioned and running, a figure that includes delivery, rigging and electrical work the original purchase price never had to cover, and for older or bespoke machinery the honest answer is often that like-for-like replacement is not possible at all, only an upgrade at a materially higher cost. And stock gets insured at a comfortable average rather than its actual peak, which matters enormously for any manufacturer with seasonal demand, because a loss does not wait for the quiet part of the year.

None of these routes involve anyone behaving carelessly. They are simply what happens to a number nobody has a routine reason to revisit, on a renewal cycle that rewards a quick answer over an accurate one, year after year, until a loss finally forces the question someone should have asked long before.

5. The maximum indemnity period is a recovery assumption

Where business interruption cover forms part of the response, two assumptions need deliberate review. The first is the insured measure of earnings or gross profit, which may differ from the figure used in financial accounts and may treat wages or other costs differently under the wording. The second is the maximum indemnity period: the maximum period for which the policy may pay following an insured interruption, subject to its terms.

Set that period against the full recovery route - debris removal, permissions, reinstatement, ordering, delivery, installation, commissioning, qualification, customer approval and the time needed to recover lost demand. The period should follow a credible recovery plan, not a historic default. Whether cover responds at all, and for how long, remains policy-specific.

6. What product cover may and may not fund

A defective product can create at least three different financial categories: injury or damage caused to other people or property; the cost of repairing or replacing the manufacturer's own defective work; and recall, withdrawal, inspection, rework or customer-consequence costs. Policy treatment varies. Product liability may respond to some third-party consequences subject to its wording, while the defective product itself and recall expense may require different arrangements or remain with the business.

The board should identify the plausible cost categories first, then ask an authorised adviser which are insured, restricted, excluded or retained. A product-policy label is not enough to answer that question.

7. The cost that sits outside the shortfall

The obvious cost of underinsurance is the gap between what a claim should pay and what it actually pays. The real cost is what that gap does to the business afterwards.

A manufacturer receiving sixty per cent of a genuine rebuild cost cannot simply rebuild. It can borrow the difference, assuming it can still borrow at all with a broken order book and a serious loss on the books, at a time when a lender is looking at a business with a damaged site and no clear evidence of when normal trading resumes. It can rebuild smaller, permanently reducing its own capacity and handing volume to competitors who happened not to have the fire. Or it can stop. None of these are dramatic failures in the moment the loss happens. The insurance did not fail loudly. It simply turned out not to be enough, and by the point anyone discovered that, every remaining option already carried a cost that would not have existed if the numbers had been right from the start.

8. Map the dependencies behind the numbers

A financial figure becomes meaningful only when its dependency is visible. Customer or programme concentration affects revenue and margin; a site or technology dependency affects recovery cost and time; a specialist person may determine whether output can restart; suppliers, materials, systems and geography can create contingent interruption outside the premises.

Model the credible 30-day cash effect of losing the most consequential dependency: cash receipts delayed or lost, continuing wages and fixed costs, finance commitments, additional expenditure, customer consequences and the point at which available liquidity becomes constrained. Record which parts are controlled, retained, financed or potentially transferred and which remain uncertain.

9. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Aldercroft's fire through it. The change, in hindsight, was years of rising replacement costs and growing production volume that the insured values never caught up with. The exposure was a sum insured on plant sitting at roughly half its true replacement value, a twelve-month indemnity period against a realistic rebuild of close to two years, and a product policy that was never going to cover the components themselves. The dependency was the machine shop as a single site holding the bulk of Aldercroft's specialist plant, with no second location to fall back on. The impact was a claim settled at a fraction of the true cost, interruption payments that stopped a year before production genuinely recovered, and an unrelated but simultaneous product cost that had to be met from the business's own funds. The controls worth examining were whether plant values, the indemnity period and the gross profit calculation had ever been independently assessed rather than simply indexed forward. The defensibility question was whether Aldercroft could have shown, before the fire, that any of these figures had been properly tested.

The response, once understood this way, was not a single renewal conversation. It was an independent valuation of plant at true replacement cost, a rebuild assessment for the buildings, an indemnity period extended to a figure that matched the operational recovery time established honestly rather than assumed, and a separate look at whether product recall cover was worth arranging given the customers Aldercroft now supplied. None of that was expensive relative to the shortfall it closed. It simply required treating the figures on the schedule as live numbers rather than inherited ones.

10. What this feeds into

Getting these figures right the first time is one thing. Being able to show, after a loss, that they were properly assessed rather than simply carried forward is another, and it is exactly what Chapter Seven, Claims & Defensibility, takes on next. The renewal conversation where these numbers should actually get tested is covered in Chapter Nine, working with advisers and the insurance market.

11. Five board questions

When was our plant and machinery sum insured last independently assessed against true replacement cost, rather than indexed forward from a previous figure?

Is our buildings sum insured based on rebuild cost, including demolition, fees and current building standards, or on a market value that means something different?

Does our business interruption gross profit figure treat wages as variable, and if so, is that actually realistic for this business?

If we lost our most critical site tomorrow, would our indemnity period genuinely cover the time it would take to rebuild, re-equip and win back the customers lost in the meantime?

Do we know, specifically, what our product liability policy would and would not pay for if a defect reached a customer next month?

12. One immediate exercise

Take your current sum insured for plant and machinery and compare it, honestly, against what it would cost today to have equivalent equipment delivered, installed and commissioned. Do the same for your indemnity period against the recovery time you established in Chapter Four. If either number falls short, that gap is this chapter's exercise, and closing it costs a great deal less than discovering it during a claim.

CHAPTER SIX

People & Workforce

Where does the business critically depend on people?

Whitlock Toolcraft is a composite scenario. It is not a real company, and its events and outcomes are presented to demonstrate how different workforce changes can expose the same underlying dependency.

Whitlock is a fifty-person tool and die maker, building and maintaining the moulds its customers use to produce their own parts. One senior toolmaker had set up and maintained the company's two most complex tools for thirty years. Nobody had ever asked him to write down how he did it, because he had always been there to do it himself, and asking felt unnecessary while he was. When he gave notice to retire, entirely properly and with several months' warning, Whitlock discovered that the knowledge behind setting those two tools to the tolerances its biggest customer required existed almost nowhere else, not in a manual, not in a training record, and only partly in a set of handwritten notes that made sense to nobody but him.

At almost the same time, a busy quarter had pushed Whitlock to bring in five agency machinists over eight weeks, faster than its normal hiring and training process could manage. Each one received the standard site safety induction. Fewer received the process-specific sign-off the more demanding jobs actually required, because the priority was getting machines running, not slowing down to complete paperwork that felt like it could wait. A quality escape on a customer contract was later traced to an agency machinist running a process step he had never been formally certified on, one that everyone on shift had assumed someone had already checked. The immediate cost was a batch of components pulled back off the customer's production line before they were fitted, plus the weeks of goodwill it took to convince that customer the fault had been found and genuinely closed. The underlying cause was smaller than the reaction to it. A sign-off checklist existed. Nobody had made following it, for that specific process, mandatory rather than assumed.

Neither event was unusual. A long-serving employee retiring with proper notice and a business flexing its workforce to meet demand are both ordinary parts of running a manufacturer. Together, they showed the same gap from two different angles: how much of Whitlock's actual capability lived in people's heads and hands rather than anywhere the business could see, document or transfer.

People & Workforce asks: Which capabilities, decisions, relationships or authorisations depend on particular people, and what would happen if those people became unavailable or the workforce changed faster than the operating system could safely absorb?

1. The People & Workforce exposure map

People exposure extends beyond whether enough employees are available. It includes:

- critical technical and process knowledge;
- competence, assessment and authorisation;

- supervision and management capacity;
- fatigue, overtime and shift design;
- recruitment, retention and realistic replacement times;
- health, safety and occupational exposure;
- licences, certifications, approvals and delegated authorities;
- customer, supplier and regulator relationships held by named individuals;
- access to machinery, digital systems, passwords and technical records;
- succession and emergency cover for leadership and specialist roles.

The important question is not simply, “Who might leave?” It is, “Which outputs, decisions or relationships would become unavailable with them?”

Manufacturing competence is often inseparable from a particular machine, material, process, customer specification or operating environment. That makes gaps visible quickly: in stopped production, unsafe work, failed tolerances, delayed delivery or an inability to release work with confidence.

Workforce exposure can rise without anyone joining or leaving. Overtime, additional shifts, compressed training, inexperienced supervision and repeated cover for absent colleagues can gradually reduce the margin for error. Headcount may look adequate while the effective capacity to supervise, check and recover from mistakes has weakened.

2. Why people risk behaves differently in manufacturing

Manufacturing competence is often inseparable from a particular machine, material, process, customer specification or operating environment. That makes gaps visible quickly: in stopped production, unsafe work, failed tolerances, delayed delivery or an inability to release work with confidence.

This makes tenure a double-edged asset. A long-serving employee is usually the most capable person in the building at what they do. They are also, often without anyone deciding it should be this way, the only person who can do it, which means their departure, whether planned or sudden, removes a capability rather than simply a headcount. This is not a comment on the individual's replaceability as a person. It is a structural observation about how their knowledge sits inside the business, concentrated rather than distributed, and available only through them rather than encoded somewhere the business itself owns. A manufacturer can employ someone for thirty years, value them entirely appropriately, and still have built no independent way of running what they run once they are no longer there to run it.

Many specialist manufacturing roles can be difficult to recruit and slow to replace. The realistic replacement period should be tested for the particular competence, process and labour market involved rather than inferred from a contractual notice period. For some specialist roles, reaching equivalent independent competence may take a year or longer even if the vacancy itself is filled much sooner.

The business should compare that recovery period with the time the operation can tolerate reduced capability. The issue is not whether a person is ‘replaceable’; it is whether the capability, authority and

judgement attached to the role can be restored before production, quality or customer commitments are materially affected. Set against that, the cost of running a structured handover while the outgoing person is still available, paid time spent training a second person, a few weeks of lower unit output while two people work a process one used to run alone, looks modest. Few manufacturers make that comparison explicitly, because the handover cost is visible and immediate while the cost of not doing it only shows up later, in a vacancy that will not close or a quality problem that traces back to nobody having checked.

3. The knowledge that only lives in one head

Whitlock's toolmaker did nothing wrong. He built genuine expertise over decades and used it well every day he was there. The failure, if there was one, belonged to the business, which never built a system for capturing what he knew while he was available to explain it.

This is the pattern behind most key-person dependency. It rarely looks like negligence at the time. It looks like a business trusting its most experienced people to keep doing what they do well, without a parallel habit of documenting how they do it, training a second person alongside them, or treating their eventual departure, retirement, illness or simply a better offer elsewhere, as a certainty to plan for rather than a risk to hope against.

The practical fix is not complicated, though it does take deliberate effort. It means identifying, honestly, which roles in the business are genuinely irreplaceable on short notice, and building a second person's competence alongside the first while there is still time to do it without pressure. It means treating a senior employee's knowledge as a business asset to be transferred, not a personal talent to be admired from a distance. None of this required Whitlock's toolmaker to do anything differently. It required the business to start the transfer years before his notice period began, rather than during it.

None of this requires a formal knowledge-management system built for its own sake. It can be as simple as pairing the outgoing person with a successor for a fixed number of weeks, recording the process on video while they talk through what they are doing and why, or asking a deliberately naive set of questions, what would go wrong here, how would you know, what would you check first, and writing the answers down where a second person can find them. The point is not to produce a perfect manual. It is to move enough of what one person carries in their hands and their judgement into a form that survives their absence, whether that absence is planned, as Whitlock's was, or sudden, as it is for many businesses that never get the retirement conversation this one did.

4. Rapid recruitment creates its own risk

Growing headcount quickly, whether through agency labour, fixed-term contracts or a fast permanent hiring push, solves a capacity problem and creates a different one in the same move. Chapter Two set out why competence needs to map directly onto the equipment and process actually in use, checked rather than assumed for agency and temporary labour. This is where that gap does the most damage.

A generic safety induction is not the same thing as process-specific competence, and the two get conflated under pressure in a way that feels reasonable in the moment. Everyone on shift assumes

someone has checked that a new operator is genuinely signed off for a specific process, and because everyone assumes it, sometimes nobody actually has. Supervision thins in exactly the period it matters most, because the people best placed to supervise new starters closely are usually also the people most needed to keep production moving during the busy period that caused the hiring in the first place.

None of this argues against agency labour or rapid hiring, both of which are often the right commercial response to a genuine spike in demand. It argues for treating the checking itself as non-negotiable regardless of how much pressure the business is under, because the pressure is precisely when the checking is most likely to be skipped.

There is also a liability question sitting underneath this that is easy to miss under production pressure. Using an agency does not transfer responsibility for a worker's health and safety away from the host business that directs and supervises their work day to day. Both the agency and the host carry duties toward an agency worker, and in practice it is the host who controls the site, the equipment and the process, and who therefore holds most of the practical responsibility for confirming someone is competent to do what they have been asked to do. A signed agency contract stating the worker is qualified is not the same thing as the host having checked it, and if something goes wrong, the paperwork the host actually holds, not the paperwork the agency says exists, is what gets asked for first.

5. Where this meets the long tail

Some workforce exposures surface years after the work occurred. Record retention should reflect the type of information, applicable regulation, contractual need and credible long-tail exposure. Some occupational-health records have specific statutory retention periods, while other training and employment records should follow a documented retention schedule. Historic employers' liability insurance information is prudent to preserve because claims may arise long after the work occurred.

6. If a critical capability were unavailable for 30 days

The first impact may be an inability to set, approve, release or troubleshoot particular work. Output may be reassigned, delayed or run under less experienced supervision. Overtime and management intervention can increase while other work receives less attention. By 30 days, the business may face backlog, missed delivery commitments, subcontracting cost, quality uncertainty and pressure on the remaining workforce.

Map the connected dependencies: the specialist person, the technology they understand, the customer or programme supported, any digital access or records they control, the site or shift where competence is concentrated and supplier relationships held through that person. People dependency rarely exists alone.

7. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Whitlock's two events through it together. The change was a planned retirement on one side and a rapid agency recruitment drive on the other, arriving within the same few months. The exposure was tooling knowledge that existed in one person's head and one set of personal notes, alongside process competence assumed rather than checked for temporary labour. The dependency was the two complex tools behind the company's largest contract, and the specific process step behind the quality escape, both resting on individuals whose actual competence had never been formally captured anywhere the business could see. The impact was a live risk of losing the ability to run its most valuable work the moment the toolmaker's notice period ended, and an actual quality failure traced directly to an uncertified operator running a process nobody had confirmed he could run. The controls worth examining were whether any knowledge transfer had been attempted before the retirement was announced, and whether the agency onboarding process actually distinguished between general site safety and process-specific sign-off. The defensibility question was whether Whitlock could show, for either the toolmaker's replacement or the agency machinist, exactly what competence had been confirmed and when.

The response was not a single fix. It was a structured handover period with the retiring toolmaker, paid deliberately as transfer time rather than treated as a formality, a written and photographed record of the tool-setting process that had never existed before, and a firm rule that agency labour cannot run a named list of critical processes without a specific, checked sign-off, however urgent the shift feels at the time. In the composite scenario, Whitlock introduces a structured handover, creates an accessible record of the critical setting process and prevents agency workers from performing named critical activities until their competence has been assessed and authorised. The purpose is not to eliminate dependence on skilled people. It is to prevent one person's availability - or one unchecked assumption - from determining whether the business can operate safely and reliably.

8. What this feeds into

The evidence this chapter has been building toward, training records, sign-offs, documented competence, becomes the material Chapter Seven asks whether you could actually produce under scrutiny. Neither chapter needs reading first, but the workforce discipline here is what gives Chapter Seven something real to point to when a claim, an audit or an investigation eventually asks the question.

9. Five board questions

Which roles in this business would be genuinely difficult to replace on short notice, and does a second person hold any of that same competence today?

Has anyone who holds critical process knowledge been asked to help transfer it, or only assumed to be available indefinitely?

Could we show, for every agency or temporary worker currently on site, exactly which processes they are trained, assessed and authorised to perform?

Where has supervision thinned during a recent busy period, and did anyone notice at the time or only afterwards?

If our most experienced person left with normal notice tomorrow, what would we actually lose that a handover period could not recover?

10. One immediate exercise

Name the one person in your business whose departure, tomorrow, with no notice at all, would cause the most damage to what you can actually produce. Ask what specifically they know that nobody else does, and whether any of it exists anywhere else, written down, recorded, or held by a second person. If the honest answer is nowhere, that is this chapter's exercise, and it is worth starting while that person is still there to help close it.

CHAPTER SEVEN

Claims & Defensibility

Could you prove the risk was properly managed?

Sherbourne Engineering is a composite scenario. It is not a real business, and its timescales and outcome are included to demonstrate the consequences of evidence that exists but cannot be retrieved quickly.

Sherbourne is a sixty-person precision fabricator running three CNC lines and a small electroplating bay. A fault in a control cabinet on its busiest line started a fire late one evening. The sprinkler system activated, the fire brigade attended, and the line was out of action for eleven weeks while replacement parts were sourced and the electrics recertified. Sherbourne submitted a business interruption and machinery damage claim the next morning, confident it had done nothing wrong and expecting the process to be straightforward.

It was not. The insurer's loss adjuster asked two things almost immediately. First, whether Sherbourne had told the insurer, at its last renewal, that it had moved a portion of its electroplating work to a night shift with a smaller supervisory team, a change nobody at Sherbourne had thought to mention because nobody had connected an operational scheduling decision with an insurance disclosure obligation. Second, whether the annual electrical testing the policy required on exactly the type of control cabinet involved had actually been carried out on schedule.

It had been carried out. Sherbourne's maintenance contractor had tested the cabinet ten months earlier, within the required period. But the certificate existed only as an email attachment in a maintenance manager's inbox, the underlying test record sat in the contractor's own system rather than Sherbourne's, and nobody could immediately say which of the three CNC lines the surviving paper file in the site office actually referred to. What should have been a same-day confirmation became six weeks of emails between Sherbourne, its broker, its maintenance contractor and the loss adjuster, while the business absorbed the cost of the disruption without knowing whether the claim would be paid in full, in part, or at all.

It was, eventually, paid close to in full. The night shift change turned out not to be material enough to have altered the terms offered, and the testing had genuinely been done. Sherbourne had not, in the end, done anything wrong. What it had done was fail to organise the evidence of having done things right, and that failure cost it six weeks of uncertainty on top of eleven weeks of disrupted production, at exactly the point the business could least afford either.

Claims & Defensibility asks: If a decision, control, incident or loss were challenged tomorrow, could the business produce reliable evidence showing what happened and how the risk was managed?

Defensibility does not mean proving that nobody was at fault or guaranteeing that an insurance claim will succeed. It means being able to provide a coherent, reliable and traceable account that an informed external party can test.

1. Who might ask for the evidence?

A manufacturer may need to demonstrate its decisions and controls to a customer investigating defective or delayed work; a regulator or enforcement authority; an auditor or certification body; an injured employee or third party; a lender, investor or acquiring business; an insurer or loss adjuster; a court, tribunal or legal adviser; or its own board trying to understand what failed.

The underlying evidence may be similar, but the question being asked will not be. Defensibility therefore begins with reliable operational records, not a file assembled only for insurance.

2. Why claims behave differently in manufacturing

Most manufacturers experience their insurance policy twice: once at renewal, when it is a document to be negotiated and signed, and once at claim, when it becomes evidence to be tested. The gap between those two moments is where this chapter sits. A policy is not a guarantee of payment. It is a conditional promise, and the conditions are rarely tested until something has already gone wrong and the business is under the most pressure to get a fast answer.

This matters more in manufacturing than in many other sectors because manufacturing losses are physical, technical and often expensive to prove. A fire, a flood, a machine breakdown or a contamination event does not just cause damage, it destroys some of the evidence of what caused it, at the exact moment the insurer most needs to understand what happened. A claim on a manufacturing site is rarely a simple statement of loss. It is a technical investigation, and the business's own records are usually the only independent account of what its equipment, its people and its processes were actually doing before the loss occurred.

The uncomfortable truth is that being right is not the same as being able to show you are right, quickly, to an external party who was not present when the decisions were made and therefore has to understand the incident through the wording, physical evidence and records available. Sherbourne had complied with its policy. What it had not done was build a habit of keeping the proof of that compliance somewhere it could retrieve in hours rather than weeks. That gap, between actual compliance and demonstrable compliance, is what this chapter is about.

3. Five recurring points of friction

1. **How the risk was described.** Whether material information was included when the insurance, contract or approval was arranged.
2. **Coverage, responsibility and causation.** Whether the event falls within the relevant policy, contract or legal responsibility, and what actually caused the loss.
3. **Whether required controls operated.** Whether inspections, maintenance, safeguards, supervision and other applicable conditions were completed.
4. **Proof and valuation of the loss.** Whether the business can substantiate damaged property, lost output, additional costs, contractual consequences and mitigation.
5. **Notification, preservation and cooperation.** Whether the right parties were informed promptly, evidence was preserved and information supplied appropriately.

Before a non-consumer insurance contract is entered into, renewed or varied, the Insurance Act 2015 requires the insured to make a fair presentation of the risk. Broadly, this means disclosing every material circumstance the business knows or ought to know, or providing enough information to put a prudent insurer on notice that further enquiries are needed. The presentation must be reasonably clear and accessible, and material representations must be substantially correct. Whether a fact is material depends on whether it would influence a prudent insurer's judgement.

The effect of warranties and terms reducing a particular risk depends on the wording, the breach, the circumstances and sections 10 and 11 of the Insurance Act 2015. Late notification consequences also depend on the policy and prejudice caused. These are matters for appropriately qualified legal or insurance advice, not assumptions.

4. What defensible evidence looks like

Defensibility is not a single document. It is a habit of generating and keeping the ordinary records a well-run manufacturer produces anyway, in a form that can be found again when someone asks for it under pressure.

The records that matter most tend to share three features. They are dated, so a business can show what was true at a specific point in time rather than what is generally true now. They are contemporaneous, created at the time the thing happened rather than reconstructed afterwards, because a maintenance log written up two months late in response to a claim reads very differently to an adjuster than one completed the same day. And they are retrievable by someone other than the one person who happens to remember where they are, which is the exact gap Chapter Six addressed from the workforce side.

In practice this means maintenance and inspection certificates held in a system the business itself controls, not only in a contractor's inbox. It means training and competence sign-offs kept somewhere searchable, not filed away and forgotten the moment the induction is complete. It means incident and near-miss reports written up on the day, however minor the event seemed at the time, because a pattern of near misses is often exactly what a challenged claim asks a business to show it was managing. None of this needs to be elaborate. A shared folder with a consistent naming convention, checked and tidied quarterly, closes most of the gap that cost Sherbourne six weeks.

5. The long tail: employers' liability

Occupational-disease and other long-tail claims may arise years after the work occurred. The policy relevant to an employers' liability claim may be the one that applied during the alleged exposure, not the policy in force when the claim arrives. Preserve historic employers' liability insurance information and the supporting exposure, health-surveillance, training and incident records required by law or a documented retention schedule.

Since October 2008 there has not been a general legal requirement to keep expired employers' liability certificates for 40 years, although HSE advises employers to retain a complete record as far as possible. Some health records have specific long statutory periods, including 40 years in specified COSHH circumstances. Lost history can be difficult to reconstruct, so long-term preservation should be deliberate, secure and consistent with data-protection obligations.

6. Working with external parties after an incident

The applicable policy, contract and law determine who must be notified, when and in what form. Establish those routes before a loss and use them promptly once the threshold is met. Nominate one internal coordinator to maintain a chronology, preserve documents, track questions and costs, and keep facts separate from assumptions. The purpose is not to build a case against another party; it is to maintain an accurate position while several investigations may be running at once.

7. The first 30 days

Safety, emergency response and any regulatory or legal notification come first. Preserve the scene and evidence where safe and lawful; notify relevant parties under the applicable contract or policy; establish a controlled chronology; record decisions and costs; and separate known facts from assumptions. The objective is not to construct a preferred narrative. It is to protect people, meet duties and preserve reliable evidence while the position is still changing.

Map where evidence depends on a specialist person, contractor, digital system, site or supplier. A maintenance certificate held only in a contractor's inbox or system access controlled by one person is an evidence dependency as well as an operational one.

8. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Sherbourne's fire through it. The change was a night shift restructure introduced without anyone connecting it to the insurance disclosure it should have triggered. The exposure was a genuine compliance record, the electrical test, that existed but could not be produced quickly, alongside an operational change that had never been assessed for materiality. The dependency was a single maintenance contractor holding the only accessible copy of a certificate the business itself should have controlled. The impact was eleven weeks of production loss extended by a further six weeks of claim uncertainty, at real cost to cash flow and management time. The control question was whether Sherbourne had ever audited where its compliance evidence actually lived, rather than assuming it existed somewhere. The defensibility question was whether the business could produce, on request, proof of the thing it already knew to be true. The response was concrete: certificates and test records pulled into a system the business itself owns, a standing instruction to flag any operational change at the next renewal rather than waiting to be asked, and a single named contact for any future claim from day one.

9. What this feeds into

The evidence habits this chapter has set out rest directly on the workforce records Chapter Six described building, and on the growth and change discipline Chapter Three set out for keeping a policy's picture of the business current. Chapter Nine, on working with advisers and the insurance market, picks up the other side of this relationship, what a renewal conversation should actually cover so that fewer of these five reasons arise in the first place.

10. Five board questions

If we had a significant loss tonight, could we produce our current compliance and maintenance evidence within twenty-four hours, or would we be searching for it?

Has every operational change in the last twelve months been checked against what we last told our insurer?

Do we know, without asking around, exactly where our employers' liability history is held, and does it cover the business's entire life?

If a claim were challenged tomorrow, who in this business would be the single point of contact, and would they know where everything lives?

Are we keeping records because we have decided they matter, or because nobody has yet asked us to produce them?

11. One immediate exercise

Pick one significant risk this business insures against, a fire, a flood, a major machine breakdown, and ask what evidence you would need to produce within forty-eight hours of a loss to support a claim without delay. Find out now whether that evidence exists, where it is held, and who could retrieve it. If the honest answer involves a phone call to a contractor or a search through someone's inbox, that gap is this chapter's exercise, and it costs far less to close today than it did for Sherbourne to discover it live.

CHAPTER EIGHT

Emerging Risk

What is changing around the business?

Redmoor Manufacturing is a composite scenario. It is not a real business, and its events and outcomes combine recurring patterns to show how a developing risk can cross operational, digital and financial boundaries.

Redmoor is a ninety-person metal fabricator that had, over several years, automated much of its shop floor. Networked programmable logic controllers ran its cutting and forming lines, a central system scheduled jobs across them, and that scheduling system was connected, for convenience, to the same office network that ran email and the finance system. Nobody had ever formally decided to connect production to the office network this way. It had simply grown that way, one integration at a time, each one sensible on its own.

A phishing email opened by someone in the finance team gave an attacker a foothold in the office network on a Thursday afternoon. By Friday morning, ransomware had spread from the office systems into the scheduling system that controlled the cutting and forming lines, and Redmoor's entire production floor stopped. It took two weeks to rebuild the affected systems, verify the machines were safe to restart, and confirm that customer and employee data taken during the attack, which the attacker threatened to publish unless a ransom was paid, needed to be reported.

The incident began with a convincing email that led to compromise on the office network. What made the operational consequence possible was the connection between that environment and production systems, combined with the absence of a current view of how those dependencies had developed.

Redmoor did have a cyber insurance policy, and it responded well to the parts of the incident it was designed for, incident response, forensic investigation, legal advice on the data breach and the ransom negotiation itself. What it was not designed for, and had never been reviewed to cover, was two weeks of stopped production. The policy capped business interruption connected to any physical or operational disruption at a small fraction of what Redmoor actually lost in unmade parts and missed deliveries. Redmoor's separate property and business interruption policy, the one that would normally have picked up a production stoppage of this kind, excluded losses connected to a cyber event outright, a clause added at a renewal two years earlier that nobody at Redmoor had asked about at the time because production loss and cyber cover felt like two entirely separate conversations.

Two policies, each pointing at the other, left the single largest cost of the incident sitting in the gap between them.

Emerging Risk asks: What is changing around or through the business, how could it reach operations, and what would trigger action?

1. Why emerging risk behaves differently

Emerging risk is defined by uncertainty and movement, not by whether an insurer has named or priced it. The issue may be new technology, a regulatory change, a climate or geopolitical pathway, a novel material, a changing workforce practice or a dependency whose consequence is not yet understood. Insurance is one possible consequence of the analysis, not its starting point.

For manufacturers, the fastest-moving version of this over the past decade has been the convergence of information technology and operational technology, the office network and the machines on the factory floor. These used to be entirely separate worlds, run by different people, on different systems, with no reason to connect. Efficiency, remote monitoring and centralised scheduling have connected them almost everywhere, and the insurance products written to cover each side have not always been rewritten to reflect that the line between them has largely disappeared.

This is not a reason to distrust automation or connectivity, both of which have made most manufacturers more competitive, not less. It is a reason to treat the question of what is covered, and what sits between two policies rather than inside either one, as something that needs actively checking rather than assuming.

2. IT/OT convergence: cyber risk on the production line

Information technology, the office network, email, finance systems, customer records, is what most people mean by cyber risk, and it is what most cyber insurance was originally built to cover. Operational technology, the programmable controllers, sensors and control systems that actually run machinery, is a different environment with different risks. An IT outage costs time and data. An OT compromise can stop physical production, damage equipment, or in the worst cases create a genuine safety incident on the shop floor.

Redmoor's mistake was not buying cyber insurance. It was assuming that a policy written around IT risk automatically extended to the OT consequences of the same incident, without checking. This is an easy assumption to make, because the attack that caused both kinds of loss was a single event, one phishing email, and it does not feel intuitive that the insurance response to it should split cleanly along a technical line nobody outside IT usually thinks about.

Recent joint guidance from the UK's National Cyber Security Centre and the US Cybersecurity and Infrastructure Security Agency has pushed manufacturers specifically toward maintaining a clear, current inventory of their operational technology and the systems connected to it, adopting recognised standards such as IEC 62443 for industrial control system security, and treating the boundary between IT and OT networks as something to be deliberately designed and monitored rather than left to grow organically. The appropriate controls depend on the architecture and risk. NCSC guidance supports maintaining an accurate OT asset inventory, understanding connections and remote access, controlling boundaries, planning safe isolation and recovery, and obtaining specialist support where internal competence is insufficient. No single control guarantees that an incident remains on the office side of the network.

Compliance point: Cyber Essentials baseline controls

Scope and market affected: UK organisations of any size, particularly those handling customer or supplier data digitally or holding contracts that require it. Certification is voluntary, not a legal requirement, but is increasingly specified by insurers, customers and public sector contracts as a baseline expectation.

Current at: 8 September 2026.

Official source: the National Cyber Security Centre's Cyber Essentials scheme, published at nsc.gov.uk, updated for assessments created after 26 April 2026.

Evidence the business should retain: current Cyber Essentials certification or self-assessment records, evidence of multi-factor authentication across cloud services, and records showing critical security updates were applied within the scheme's required timeframes.

Trigger for specialist advice: renewing or arranging cyber insurance, tendering for a contract that specifies a security certification, or after any incident that exposes a gap between actual practice and certified controls.

3. Insurance boundaries require verification

Property, machinery-breakdown and cyber policies may treat cyber-related physical damage and interruption differently. Some wordings exclude, restrict or affirm particular cyber events. The result depends on the exact policies, endorsements, cause and facts; the business should map the boundary with its adviser rather than infer it from the product name. The effect for a policyholder is that a genuinely physical loss, a stopped production line, spoiled work in progress, can fall outside a property policy the moment a cyber event is anywhere in the chain that caused it, even where the property damage itself looks identical to any other operational disruption.

This does not mean cover is unavailable. It means the cover has to be bought deliberately, usually as a specific extension within a cyber policy that responds to business interruption and physical consequences arising from an OT compromise, rather than assumed to exist inside a general property policy or a generic cyber policy bought primarily for data breach costs. The only way to know whether that gap exists in a specific business's own arrangements is to ask the direct question, at renewal, of both policies at once, rather than reviewing each one separately as though the other did not exist.

4. A practical signal-and-trigger discipline

For each developing issue, record:

- **Signal:** what has changed or been observed?
- **Pathway:** how could it reach the business?
- **Dependency:** what would it affect?
- **Pace:** how quickly could the position change?
- **Consequence:** what could happen, including within 30 days?
- **Trigger:** what evidence or threshold would require action?

- **Ownership:** who monitors it and when is it reviewed?

Classify the current response as **Observe**, **Investigate** or **Act**. This is an attention decision, not a probability score. Growth & Change begins with a decision the business is making; Emerging Risk begins with a changing condition that may require the business to decide.

Cyber Essentials version 3.3 has applied to new assessments since 27 April 2026. It is a useful baseline for in-scope IT infrastructure, but certification does not by itself establish that industrial control or operational-technology risk is adequately managed.

5. What a personal-data breach may require

Not every cyber or security incident is a reportable personal-data breach. Where a personal-data breach is likely to result in a risk to people's rights and freedoms, the organisation must notify the ICO without undue delay and, where feasible, within 72 hours of becoming aware. If the risk to individuals is high, affected people may also need to be informed without undue delay. Every personal-data breach should be documented whether or not it is reported.

The practical control is an incident route that distinguishes operational disruption, data protection, safety, contractual and insurance obligations, identifies the people authorised to decide and records the basis for each decision.

6. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Redmoor's incident through it. The change was years of gradual, individually sensible integration between office IT and shop-floor OT, with nobody responsible for reviewing the connection as a whole. The exposure was a single phishing email with the ability, once inside the office network, to reach systems that controlled physical production. The dependency was a scheduling system that sat on both sides of a boundary nobody had deliberately drawn. The impact was two weeks of stopped production, a data breach requiring regulatory assessment, and a gap between two insurance policies that each assumed the other covered the loss. The control question was whether the IT and OT networks had ever been reviewed together, by someone who understood both, rather than separately by whoever happened to own each system. The defensibility question was whether Redmoor could show, to its insurer and to the ICO, exactly what had been assessed and when. The response was a formal network segmentation project, a specific cyber extension bought to cover OT-related business interruption, and a standing instruction to review IT and OT exposure together at every renewal rather than as two separate conversations.

7. What this feeds into

The single point of failure this chapter has described, one scheduling system sitting across an undefended boundary, is the same pattern Chapter Four set out for physical dependencies, applied to a digital one. Chapter Nine, on working with advisers and the insurance market, is where this gets

tested properly, in the specific questions a renewal conversation should ask about where cyber cover ends and property cover begins.

8. Five board questions

Do we know exactly where our office network connects to the systems that actually run our machines, and who reviewed that connection last?

If a cyber incident stopped production tomorrow, would our cyber policy, our property policy, both, or neither actually respond to the lost output?

Who in this business is responsible for noticing when something has changed enough to need a fresh conversation with our broker, rather than waiting for the next renewal date?

Could we show a regulator exactly what we assessed and decided the last time we suspected a data incident, even one we judged did not need reporting?

What is the newest piece of technology, regulation or market change we have adopted in the last year that nobody has yet asked an insurance question about?

9. One immediate exercise

Ask whoever manages your IT systems and whoever manages your production equipment to describe, separately, where the two networks connect. If the two descriptions do not match, or either person struggles to answer precisely, that mismatch is this chapter's exercise, and it is worth resolving before an incident forces the answer out under far worse conditions.

CHAPTER NINE

Working With Advisers and the Insurance Market

Does your adviser understand the business as it operates today?

Hartley Foundry is a composite scenario. It is not a real business, and its events and outcomes combine recurring patterns to show how an external adviser relationship can become a concentrated dependency.

Hartley is a seventy-person iron foundry that had used the same regional broker for twelve years, and the same named contact at that broker for almost as long. Over that time, the relationship had become genuinely useful, the contact understood the business, knew its history, and could explain Hartley's risk to an underwriter in terms that made sense. Nobody at Hartley had ever thought to ask what would happen if that particular person left, because the relationship had always simply worked.

The contact retired in the spring, seven weeks before Hartley's renewal date. The handover to a new account handler at the same brokerage consisted of a shared drive folder and a short introductory call. The new handler had never visited the site, had not been part of any of the previous conversations about Hartley's recent investment in a second melting furnace, and worked from a renewal submission that was, in substance, the same document that had been lightly updated every year for over a decade.

When the submission reached the market, one underwriter queried an unexplained increase in machinery values and asked how the new furnace changed output, protection and recovery. The query, and the days it took to answer properly once the new handler realised how much context was actually missing, pushed the renewal to within a few days of expiry before terms were agreed. Hartley got cover, but on terms noticeably tighter than the previous year, and with none of the negotiating room a business normally has when a renewal is handled with enough lead time to test more than one insurer.

Nothing about this requires bad faith by the adviser or deliberate withholding by Hartley. The business had simply allowed a current understanding of its risk to depend on one person inside an external firm. The scenario ends with that dependency exposed: Hartley must rebuild the risk record it should have owned, determine what the market needs to know and decide how the adviser relationship will be governed in future.

1. Why this relationship behaves differently

A strong renewal starts with an accurate picture of the business, owned by the manufacturer and tested by an adviser who understands how to present it.

An insurance policy is not a product a business buys once and then owns. It is renegotiated every year, usually annually, against a market that moves independently of the business itself. The quality of that renegotiation depends almost entirely on how well an underwriter understands the risk being priced, and an underwriter only understands what the submission in front of them actually says.

This makes the adviser relationship different from most of the other supplier relationships a manufacturer manages. A broker is not simply a distribution channel for a product. A good one is the mechanism by which a business's own risk management, everything the rest of this handbook has described, actually reaches the people who decide what it costs to insure against. A submission that undersells the business's genuine controls costs it money at renewal. A submission that misses a material change, as Hartley's did, costs it credibility exactly when credibility matters most.

The relationship can decay quietly, the same way Hartley's did, precisely because nothing visibly goes wrong until the renewal that exposes the gap. A transactional relationship, refreshed once a year with minimal engagement in between, can look identical to a genuinely well-managed one for years, right up until the year it does not. This is precisely why the failure is so easy to miss from inside the business. A board reviewing its insurance arrangements each year sees a renewed policy, a premium, and a broker who has always come through, and has no natural prompt to ask whether that outcome reflects a genuinely current understanding of the business or twelve years of momentum quietly running on an increasingly out-of-date picture.

2. The manufacturer owns the risk picture

The board owns the accuracy of the business description, the decisions it makes and the exposure it deliberately retains. A broker or intermediary can challenge, advise, translate the risk for insurers and arrange cover. The insurer assesses and prices the proposal. A risk engineer, lawyer, accountant, cyber specialist or other adviser may address questions within their competence. None of those roles removes the manufacturer's ownership of the underlying facts or decisions.

A strong relationship should produce a current description of the operation; stated demands and needs; documented material changes; explanation of market approach and recommendation; clarity on uninsured or retained exposures; an agreed claims-support service; and a record the business can retain independently of one contact.

3. What a good adviser relationship looks like

The difference between a broker relationship that protects a business and one that quietly erodes is rarely about the individual's competence. It is about how much the relationship depends on one person's memory rather than the business's own records, and how much contact happens outside the single annual renewal conversation.

A relationship worth relying on usually involves the adviser knowing about material changes close to when they happen, a new piece of equipment, a new customer contract, a change in shift patterns, rather than trying to reconstruct a year of change in the fortnight before renewal. It usually involves more than one point of contact on both sides, so that neither a retirement at the brokerage nor a departure inside the business leaves the relationship starting from nothing. And it usually involves the business treating its own renewal preparation as something it owns, rather than something it hands over and waits to hear back about.

None of this requires distrusting a good broker. It requires recognising that even a good broker can only work with what they are told, and building a habit of telling them, rather than assuming twelve years of history means nothing needs saying again. In practice this is rarely a large undertaking. Most

manufacturers already hold a monthly or quarterly leadership meeting where operational changes get discussed. Adding one standing question, whether anything decided this quarter is worth flagging to the broker, costs a few minutes and closes most of the gap that otherwise only surfaces at renewal, when there is far less time to close it properly.

4. Preparing for renewal: evidence the business should own

Much of this handbook has been building toward a single moment, the point where a business has to show an underwriter, in a form they can actually use, that its risk is genuinely understood and managed. A renewal prepared well in advance draws on records that should already exist rather than being assembled under deadline pressure.

Current sums insured and rebuild or reinstatement values, checked against real cost rather than carried forward from last year, as Chapter Five set out. An accurate, current description of what the business actually does, including any growth, new sites, new equipment or new customer concentration since the business was last described to the market, the gap Chapter Three addressed directly. Maintenance, inspection and statutory compliance evidence for the equipment and premises being insured, in the form Chapter Two described building. Training and competence records for the workforce, particularly anywhere agency or temporary labour is used, from Chapter Six. A clear claims history with an honest account of what happened and what changed afterwards, since an unexplained claim reads far worse to an underwriter than one with a documented response attached, the defensibility Chapter Seven set out. And, increasingly, a clear answer to how the business's cyber and operational technology exposure is actually covered, given the gap Chapter Eight described between products that each assume the other responds.

None of this needs to be produced from scratch every year if it has been kept current throughout. The businesses that find renewal straightforward are rarely the ones with the least risk. They are the ones that never let this evidence go stale between one renewal and the next. A single renewal file, updated as each piece of evidence changes rather than reconstructed once a year, turns this from an annual scramble into a short conversation about what, if anything, has genuinely moved since the last update.

5. Choosing and checking an adviser

Most manufacturers choose a broker once and rarely revisit the decision, which makes it worth choosing carefully and checking periodically rather than treating the relationship as permanent by default. A firm carrying on regulated insurance distribution must be appropriately authorised, exempt or acting as an appointed representative within its principal's permissions. Check the firm and permissions through the FCA Firm Checker or Financial Services Register. The information and recommendation required will depend on the service and customer classification; ask the adviser to state its market scope, service basis and recommendation rationale in writing.

Compliance point: broker and adviser authorisation

Scope and market affected: UK businesses arranging commercial insurance through a broker or intermediary, under the Financial Conduct Authority's regulation of insurance distribution.

Current at: 8 September 2026.

Official source: the Financial Conduct Authority's Financial Services Register and guidance on how to check a firm or individual is authorised, published at [fca.gov.uk](https://www.fca.gov.uk).

Evidence the business should retain: the adviser's Firm Reference Number, a copy of their status disclosure and terms of business letter, and a dated note of when this was last checked, particularly after a change of named contact or a change of ownership at the brokerage itself.

Trigger for specialist advice: appointing a new broker, being told of a merger or acquisition affecting the current one, or any uncertainty about who is actually placing cover and standing behind it.

Checking the register takes a few minutes and confirms a firm is currently authorised rather than relying on a logo or a name that sounds familiar, which matters more than it used to given how often brokerages have changed hands through consolidation in recent years. A business that has not checked this since it first appointed its adviser, particularly if that appointment happened more than a few years ago, may be checking a firm that no longer exists in the form it was originally chosen.

6. An early renewal timetable

- **120-90 days:** confirm ownership, timetable, objectives and material business changes.
- **90-60 days:** validate values, recovery assumptions, claims information, dependencies and evidence; resolve specialist questions.
- **60-30 days:** approve the market presentation, understand the adviser's market approach and compare terms, exclusions and service - not premium alone.
- **30-14 days:** close outstanding information, agree decisions and record retained gaps before binding.
- **After placement:** verify documents, conditions and actions; update the business-owned file and carry material-change triggers into the year.

7. Understanding the market approach

Identical risk can price very differently from one year to the next, for reasons that have nothing to do with anything the business itself has done. Insurers move in and out of manufacturing sectors, tighten or loosen their appetite for particular exposures, and respond to their own claims experience across an entire industry rather than any single policyholder. Market capacity, insurer appetite, claims experience, inflation and reinsurance conditions can influence price and terms. They do not explain every movement, so the adviser should separate external market factors from changes in the risk, presentation, cover and competition actually obtained.

This is not a reason to assume a price increase reflects poor risk management, and it is equally not a reason to assume a competitive quote means nothing needs improving. It is a reason to ask an adviser directly, at renewal, what is happening in the market generally as well as what is happening to the business specifically, so that a board understands which part of any change in terms it can influence and which part it cannot. A good adviser will usually be willing to explain this distinction plainly if asked directly, because it is in their interest for a client to understand what they are actually being charged

for. A business that never asks tends to draw the wrong conclusion either way, either assuming a rising premium proves it is being penalised for something it has not actually done, or assuming a falling one proves a problem has been solved when the market simply softened underneath it.

8. Market scope, remuneration and conflicts

An adviser describing a fair analysis should consider a sufficiently large number of contracts within the relevant market or sector to support a personal recommendation; other advisers may use a limited panel or a single insurer. Ask what was considered, what was not, and why the recommendation meets the manufacturer's demands and needs.

Commercial customers can request written disclosure of commission received by an insurance intermediary, including associated remuneration covered by the FCA rule. Ask separately about fees, premium-finance income, ownership relationships, facilities or delegated authority that may influence the route to market. A commercial arrangement is not automatically improper; undisclosed or misunderstood incentives are the governance issue.

9. Claims support should be explicit

Before a loss, establish who will notify insurers, preserve evidence, coordinate loss adjusters, track costs and challenge a disputed interpretation. Claims advocacy varies significantly by adviser and policy. The expected service should be recorded rather than inferred from the relationship.

10. How InduX examines this

The same sequence applies here as everywhere else in this handbook.

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

Run Hartley's renewal through it. The change was a retirement at the brokerage, arriving without warning close to Hartley's own renewal date. The exposure was a year's worth of operational change, principally the new furnace, that had never been formally recorded anywhere the business itself controlled. The dependency was a single external contact holding the only real understanding of Hartley's risk profile. The impact was a renewal pushed to the wire, tighter terms, and no real time to test the market against more than one insurer. The control question was whether Hartley had ever kept its own record of what its broker knew, rather than trusting that knowledge to survive a personnel change at another company entirely. The defensibility question was whether the business could have produced, at seven weeks' notice, a complete and current account of its own risk without depending on anyone else's memory. The response was a standing internal renewal file updated through the year rather than assembled each spring, a second named contact agreed with the brokerage, and a diary reminder to flag any material operational change to the adviser within weeks of it happening rather than waiting for the next renewal to mention it.

11. What this feeds into

Every earlier chapter's evidence, values from Chapter Five, operational accuracy from Chapter Three, compliance from Chapter Two, workforce records from Chapter Six, claims defensibility from Chapter Seven, cyber and OT clarity from Chapter Eight, converges here, in the submission a broker actually takes to market. Chapter Ten closes the handbook by setting out what it looks like to run this as a single connected discipline rather than ten separate concerns.

12. Five board questions

1. Could we brief a replacement adviser without relying on our current contact's memory?
2. Do our stated insurance demands and needs reflect the business's actual dependencies, recovery priorities and ability to retain loss?
3. Do we understand the adviser's market scope, recommendation rationale, ownership relationships and complete remuneration?
4. Which important needs remain uninsured, underinsured or deliberately retained, and has the board accepted them consciously?
5. What claims support would we receive after a major incident, and have we tested that expectation before needing it?

13. One immediate exercise

Ask the adviser, without first supplying a fresh business description, to summarise what it currently understands about the manufacturer's operations, recent changes, critical dependencies, recovery priorities and principal retained risks. Compare that account with the current board view. Record every material difference, assign an owner and determine whether it requires immediate notification, renewal action or specialist advice.

CHAPTER TEN

Building a Risk-Resilient Manufacturing Business

A practical operating model for leadership

Oakfield Industrial is a composite scenario. It is not a real company, and its events, timescales and outcomes are illustrative.

Eighteen months ago, Oakfield had a difficult year. A near miss on underinsured machinery, caught only because a broker happened to ask the right question at renewal, arrived within months of a senior estimator giving notice with almost nothing about his role written down anywhere else. Neither event caused a loss. Both left the leadership team badly shaken by how close they had come, and by how little they had actually known about their own exposure until something forced them to look.

What changed after that was not, in the first instance, more insurance. Oakfield's board set a standing ten-minute item at each quarterly meeting, what has changed in the business since we last discussed this, and began keeping a single running file of the evidence, compliance and maintenance records, training and competence sign-offs, supplier and site information, that earlier chapters of this handbook have each described building. They asked their broker to be told about material changes as they happened rather than at renewal. They mapped which of their machines, systems, people and suppliers the business would genuinely struggle to replace on short notice, and started closing the gaps deliberately rather than hoping none of them broke. Oakfield began with simple governance: a named owner, a recurring leadership review and a controlled record of findings and actions. Some improvements were inexpensive; others required investment, specialist input or difficult decisions. The purpose was not to make risk management effortless. It was to direct limited time and money toward the dependencies capable of causing the greatest harm.

None of this made Oakfield's risk disappear. When the supplier later suffered a fire in the composite scenario, Oakfield still faced disruption. However, leadership began from a known dependency, an existing decision route and an alternative already under investigation. Preparedness reduced uncertainty and response time; it did not remove the loss or guarantee continuity. That distinction - between predicting an event and preparing for its consequence - is what this final chapter is about.

1. Resilience is a discipline, not a document

Nine chapters have each looked at a different way manufacturing risk actually behaves, growth outrunning the picture an insurer holds, a single point of failure stopping production, a policy that pays less than expected, a business that cannot function without one person's knowledge, a claim that gets challenged for want of paperwork, a risk changing faster than the market can price it, and a broker relationship nobody had checked was still sound. None of these are really about insurance. Insurance is what a business buys once it has already decided how much of its risk it can absorb itself and how much it needs to transfer. Getting that decision right depends entirely on how well the business understands its own risk in the first place, and that understanding is not something a policy document can provide.

A resilient manufacturer is not one that has eliminated risk. It has built the habit of noticing change, asking what that change exposes and deciding what must happen before an incident, claim or renewal forces the question. That discipline requires time, ownership and, in some cases, investment or specialist support. Its value is not that every loss becomes cheaper or easier. It is that leadership begins from a known picture, named priorities and evidence it can use, rather than starting from zero.

2. The six pillars, one business

This handbook uses six connected pillars:

- **Growth & Change:** Has the business's understanding of risk moved with its decisions, growth and operating model?
- **Operational Resilience:** What could stop production or delivery, how quickly would the effect spread, and how would the business recover?
- **Financial Exposure:** Which losses, commitments or interruptions could the business not absorb, and what assumptions support the protection it believes it has?
- **People & Workforce:** Which capabilities, decisions, relationships or authorisations depend on particular people, and can workforce change be absorbed safely?
- **Claims & Defensibility:** Could the business demonstrate what it knew, decided, controlled and did if externally scrutinised?
- **Emerging Risk:** What is changing around or through the business, how could it reach operations, and what would trigger action?

A single incident rarely respects these boundaries. Oakfield's supplier fire touched operational resilience directly, but it also tested whether the business's financial exposure had been correctly sized, whether its claims evidence was in order, and whether the change had been properly flagged to its broker. The six pillars are a way of asking better questions, not six separate boxes to tick once a year and then forget. This is also why reading this handbook chapter by chapter, as most manufacturers will, can create a slightly misleading impression that each pillar sits in its own compartment. In practice they overlap constantly. A growth decision, Chapter Three's territory, almost always has a workforce consequence, Chapter Six's, and frequently an operational resilience one too, Chapter Four's, all from a single decision to take on a larger contract or open a second site.

3. The nine dependencies underneath them

Use the exact nine dependency lenses throughout the business:

1. Customer
2. Programme
3. Site
4. Technology

5. Digital system
6. Specialist person
7. Supplier
8. Material
9. Geography

These categories are a practical lens, not a claim that every dependency fits perfectly into one box. A contract may create exposure around a customer or programme; a machine is normally examined under technology.

The risk is often in the chain

A manufacturer may appear to have several separate dependencies when they all support the same commercial outcome. One customer may depend on one programme, produced on one machine, operated by one specialist person, using one material from one supplier in one geography. Looking at each dependency separately understates the concentration.

Record the output, customer or programme supported; connected dependencies; available substitute; time until disruption becomes material; credible 30-day consequence; control or contingency; evidence that it works; owner; and next review date. Two suppliers are not genuine diversification if both obtain the same material from the same geographic source. Two machines are not independent if both rely on one control system or specialist setter.

4. Running the diagnostic in practice

The InduX diagnostic follows seven connected stages:

CHANGE → EXPOSURE → DEPENDENCY → IMPACT → CONTROL → DEFENSIBILITY → RESPONSE

- **Change:** What is different?
- **Exposure:** What new or altered risk does that create?
- **Dependency:** What does the business now rely on?
- **Impact:** What could happen, particularly within the first 30 days?
- **Control:** What currently reduces the likelihood or consequence?
- **Defensibility:** Could the business demonstrate that the control exists and operates?
- **Response:** Should the business understand, reduce, retain, finance, transfer, prepare for or escalate the risk?

Use qualitative attention labels rather than a universal score:

- **Priority review:** serious 30-day consequence, concentrated dependency and weak or unproven control.
- **Review recommended:** material exposure with partial protection, unresolved uncertainty or ageing evidence.

- **Lower current concern:** relevant controls and alternatives exist and have been appropriately tested.
- **Not assessed / insufficient information:** leadership cannot yet support a conclusion.

These labels indicate the need for attention. They are not predictions, actuarial ratings or proof that a risk is adequately controlled.

5. Ownership and the minimum viable risk discipline

One senior leader should be accountable for ensuring the process operates. Ownership of individual risks and actions should remain with the people closest to them - operations, finance, engineering, commercial, people or technology. Accountability is centralised; knowledge and action are distributed.

- **Board or owner-director:** approves priorities and retained exposure.
- **Named executive lead:** maintains the process and review cadence.
- **Risk or action owner:** delivers each control or decision.
- **Evidence owner:** keeps critical records current and retrievable.
- **Specialist adviser:** addresses questions outside internal competence.

When something changes: record it, identify affected pillars and dependencies, decide whether immediate review or specialist advice is needed, and assign an owner and date.

Monthly: update priority actions, record new uncertainty or evidence gaps, and escalate any materially worse 30-day consequence.

Quarterly: ask what has changed, review consequential dependencies, examine overdue actions and untested assumptions, and confirm what leadership is deliberately retaining.

Annually: revalidate dependency mapping, financial assumptions and critical controls; reconcile them with contracts, certifications, financing and insurance; and review external support.

After an incident, near miss or major decision: record the learning, update assumptions and controls, and test whether it affects other sites, lines, programmes or dependencies.

Track four board-level indicators: priority actions overdue; critical dependencies without a tested contingency; critical controls without current evidence; and material changes awaiting assessment. These show whether the discipline operates; they do not measure total risk.

6. Twelve-question manufacturing resilience self-audit

Growth & Change

1. Can we identify every material business change made during the last twelve months?
2. Did each one trigger a proportionate review of exposure, dependencies and required response?

Operational Resilience

3. Do we know which operations would stop first if a critical dependency became unavailable?
4. Have the most important alternatives, recovery procedures or workarounds been tested?

Financial Exposure

5. Have we modelled the 30-day cash effect of losing our most consequential dependency?
6. Does leadership understand which losses are reduced, retained, financed or transferred?

People & Workforce

7. Which critical capabilities depend on one person, shift or thinly spread team?
8. Can we demonstrate who is trained, assessed and authorised for critical work?

Claims & Defensibility

9. Could we retrieve evidence that our most important controls operated within 48 hours?
10. Have we tested how the business would preserve evidence, track costs and communicate after a serious incident?

Emerging Risk

11. What three developing risks are we actively monitoring, and what would trigger action?
12. Who owns each one, and when will the evidence next be reviewed?

Any “unsure” response is a finding, not a failed test.

7. The first 90 days

Days 1-30: appoint the executive lead; complete the self-audit; map the most consequential dependency chains; and select no more than three priority findings.

Days 31-60: run a 30-day operational and financial stress test for the leading dependency; confirm owners, actions, evidence and deadlines; and escalate specialist questions.

Days 61-90: test one contingency or incident scenario; verify at least one assumed alternative supplier, process, person or system; integrate the change trigger and quarterly review into existing governance; and record what remains deliberately retained or unresolved.

Where a finding requires regulated insurance advice, legal interpretation, engineering assessment, cyber expertise or another specialist judgement, seek an appropriately authorised or qualified adviser. The manufacturer should still own the decision, the evidence supplied and any exposure deliberately retained.

A manufacturer can begin internally using this self-audit and diagnostic. Risk360 provides a more structured way to work through selected business changes and preparedness indicators. Neither route replaces professional advice where the circumstances require it.

8. Five board questions

1. What has materially changed since this board last reviewed the business’s risk?
2. Which connected chain of dependencies could create the greatest operational or financial impact within 30 days?
3. Which priority control or contingency are we assuming works without having tested it?
4. Which exposure are we deliberately retaining, and has the board consciously accepted the consequence?

5. What three improvements will be completed during the next 90 days, by whom and with what evidence?

9. One immediate exercise

Select one material change from the last three months and run it through all seven stages of the InduX diagnostic. Map connected dependencies rather than stopping at the first one. Identify the credible 30-day consequence, classify the current level of attention and assign one decision or action to a named owner with a completion or review date. The objective is not to close every gap immediately. It is to ensure that no material finding remains ownerless or invisible.

Every successful manufacturing business changes. The question is not whether change can be stopped, or whether every loss can be predicted. It is whether the business notices when its assumptions no longer match reality - and whether leadership acts before an incident makes the difference visible for them.

Start with one change. Follow the dependency. Test the consequence. Decide what happens next. Then repeat.

Glossary

Control - A measure intended to reduce the likelihood or consequence of an exposure.

Defensibility - The ability to provide a coherent, reliable and traceable account of what the business knew, decided, controlled and did.

Dependency - A customer, programme, site, technology, digital system, specialist person, supplier, material or geography on which an important output relies.

Exposure - The way a change or condition could create operational, financial, legal, workforce, customer or other consequence.

Maximum indemnity period - In relevant business interruption wording, the maximum period for which the policy may pay following an insured interruption, subject to its terms.

Operational technology (OT) - Hardware and software used to monitor or control physical equipment and industrial processes.

Preparedness label - A qualitative indication of the need for attention: Priority review, Review recommended, Lower current concern, or Not assessed / insufficient information. It is not an actuarial rating or prediction.

Time to consequence - How quickly the absence or failure of a dependency begins to create material harm.

Time to recover - How long it takes to restore an acceptable level of output or capability.

Time to survive - How long the business can withstand the operational and financial effect before a critical threshold is crossed.

Source and verification notes

This edition uses official primary sources for legal and regulatory orientation. Key references include:

- Health and Safety Executive guidance on risk assessment, PUWER, LOLER, pressure systems, agency workers, health surveillance and employers' liability insurance.
- UK legislation, including the Health and Safety at Work etc. Act 1974, the Management of Health and Safety at Work Regulations 1999 and the Insurance Act 2015.
- Department for Business and Trade and Office for Product Safety and Standards guidance on UKCA and CE conformity routes.
- Environment Agency and GOV.UK guidance on environmental permitting and waste duty of care.
- Information Commissioner's Office guidance on personal-data breach notification, employment records and accountability.
- National Cyber Security Centre guidance on operational technology and Cyber Essentials requirements for IT infrastructure, version 3.3.
- Financial Conduct Authority Handbook and guidance on commercial insurance distribution, market scope, remuneration and checking a firm's status.

Readers should verify the current position at the time of use. Sources may change after the verification date and the correct requirement depends on the manufacturer's circumstances.

About InduX

InduX helps UK manufacturers understand how business change alters risk. Its six-pillar framework, dependency mapping and Risk360 diagnostic are designed to make hidden concentration, weak assumptions and unresolved evidence visible before an incident or renewal forces the question.

<https://indux.co.uk>

Industrial Exchange (InduX) is a trading name of South Suntreso Limited, registered in England and Wales, company number 16122395.

© 2026 South Suntreso Limited. All rights reserved.

First edition published September 2026.