MULTIPLEWORKS

Executive Briefing

No. 02

Built for a world that's leaving

An executive briefing on technology resilience.
Worth half an hour and a coffee, on a quiet morning.

June 2026

CONTENTS

THE IDEA IN BRIEF

The exposure. Your business relies on technology it does not control far more than anyone has measured, and most of that reliance was never chosen. Cloud, critical services, and now AI have become load-bearing while everyone still treats them as conveniences. What happens if one of them disappears? For twenty years, nobody needed to ask.

Why now. That is changing fast. Governments are redrawing the map around digital sovereignty, providers can restrict or reprice access at will, and AI is volatile enough that a model you rely on this week can be gone the next. The stable, open world your business was built for is being taken apart, and the business has not kept up.

The cost of waiting. A dependency you have not planned for fails on its own timetable, usually at the worst moment. Most continuity plans, if they exist, have never been tested, so they are not really plans. The board has always been accountable for surviving the bad day. What has changed is how likely that day has become.

What to do. Start with the few things the business cannot fail to do, whatever happens, and trace each down to what it depends on. Protect those few properly and build them to survive losing any single supplier or system; let the rest break. Then rehearse a couple of bad days on paper, before one arrives for real. If you cannot say what next week looks like when a system you rely on disappears at the end of this one, that uncertainty is the gap, and the rest of this briefing is how to close it.

The assumption beneath everything

Every business runs on a belief it has never examined: the technology will be there tomorrow. For two decades that was true often enough to stop being a belief and become the floor everyone stands on without looking down. The cloud was its high point. A business could hand the hard parts of its technology to a few global providers, retire the infrastructure it used to run, and buy capability as a service that was always on. It gave up control, and with providers that vast and a world that open, that looked like no loss at all.

What the cloud delivered was not the resilience it implied. It was sold on never going down, and it produced concentration. A handful of providers now sit beneath a startling share of the world's commerce, so when one has a bad day the failure does not stay contained; it ripples across every business standing on the same foundation. The arrangement meant to remove single points of failure manufactured a few enormous ones and called them reliability. The pattern is about to repeat: each wave of technology promises to relieve a dependency and ends up deepening it, while the language of resilience hides a growing exposure.

The deeper problem is that the assumption was never a decision anyone took. Businesses outlive the conditions they were designed for, and an organisation built on one decade's assumptions carries them into the next whether or not they still hold. Because nobody chose the assumption of stable, global, always-available technology, nobody has noticed it starting to fail, and the most exposed businesses are often the most confident their technology is sound. There are now two kinds of business: the one whose confidence is earned, having looked hard at what it depends on and what losing it would mean, and the one whose confidence is only the comfort of never having looked. The question worth carrying as you read on is plain and uncomfortable. Which of those two is yours?

The ground is shifting, deliberately

Until recently, instability in technology was commercial and technical. A supplier might fail, a system break, a contract be renegotiated on worse terms. These are familiar risks, and established businesses know how to carry them. The new instability is political and deliberate, and it comes from the people who write the rules everyone else works within. Digital sovereignty has moved from a phrase in a policy paper to government action, and the consequences land on any business that assumed the technology map would stay as it was.

In April 2026 the French government ordered every state ministry to plan its migration off Microsoft Windows and Office, across roughly two and a half million workstations, calling it a strategic necessity to reclaim its digital independence from foreign control. The order reaches well beyond the desktop into collaboration tools, databases, and AI platforms, all to be replaced over the coming years with sovereign or open alternatives. France is unusual in the scale of its commitment, not its direction; other governments are building national AI capability and sovereign computing for the same reason. Separately, a leading provider was recently required to withdraw its most capable models at short notice, for reasons outside any customer's control, and access did not return because customers asked.

What matters for a board is not any single one of these events, each arguable on its own terms. It is the direction they share. The bodies now reducing their dependence on technology controlled elsewhere are governments, which set the conditions every business works inside. When the rule-makers themselves are hedging against open, always-available technology, it has stopped being a safe default for anyone. A business built for seamless global access has found itself built for a world being dismantled around it, without changing a line of its own design.

Artificial intelligence, at the worst possible moment

Onto this shifting ground, businesses are laying their heaviest dependency yet, and faster than they have ever taken one on. Artificial intelligence is being woven into how work gets done, usually without anyone deciding it should become essential, and faster than anyone asks the questions a critical dependency deserves. It arrived dressed as a productivity tool, and was adopted as one. While everyone admired the convenience, it began to carry load, and many businesses now depend on it more than they would admit if asked.

The trouble is that AI is the least suitable candidate imaginable for an unexamined dependency. It is an emerging technology with almost none of the stability of the infrastructure it is replacing. The models change between releases, the suppliers appear and consolidate, the pricing is rewritten as providers hunt for a model that pays, and the regulation is being laid down as everyone walks across it. It is also, as the recent withdrawals show, more exposed to political decision than ordinary software has ever been. A business leaning on it takes a deeper dependency than usual, on a more volatile technology, in a less stable world, which is about the worst combination one could arrange on purpose.

The discipline this calls for is simple to state and rare to find. The depth of a dependency should match the maturity of the thing depended on and the stability of the world around it. A settled technology with interchangeable suppliers and predictable terms can be allowed to become essential, because the cost of any one part failing is low and recovery is understood. A technology that can be changed, repriced, restricted, or withdrawn at short notice should be kept to a dependency the business can survive losing, at least until the ground beneath it settles. Most organisations have this backwards, giving their deepest and least reversible reliance to the technology least able, right now, to bear the weight.

The dependency nobody decided to take on

The dangerous dependencies are rarely the ones a business chooses deliberately, with a paper before the board and a risk assessment attached. They accumulate one convenient decision at a time, until the organisation is leaning hard on something it never decided to lean on. This is how AI enters most businesses: a team adopts a tool because it makes a task faster, the tool proves useful, more work flows through it, the old way falls into disuse, and one day a capability the business depends on rests entirely on a service it does not control and never chose.

The most damaging part is what happens to the fallback. When a manual process is replaced by an automated one, the old way is rarely kept against the day it might be needed again. It is retired. The procedure is forgotten, the people who knew it move on, and the muscle that once did the job by hand wastes away. As long as the replacement works, none of this shows, and it looks like pure progress. The cost appears the day the replacement is gone and the business reaches for the fallback it assumed it still had, and finds an empty space where the capability used to be.

This is why an undesigned dependency is more dangerous than a designed one of the same depth. A designed dependency comes with a decision, and usually a plan for what happens if it goes away. An undesigned one comes with neither: its depth is unknown until failure tests it, and the failure arrives with no rehearsed response. The first task of resilience is to map the invisible dependencies, while the question is still an exercise rather than an emergency. If you cannot say, as you read this, how many your own business now carries, you have already found the gap this briefing exists to close.

When the bill becomes the risk

Continuity is usually imagined as availability, whether the system is up or down. But availability is only half of it. A capability the business can no longer afford is as lost to it as one switched off, and the economics of the technologies now becoming essential make affordability a live continuity risk. This is newer than it looks, because the last generation of enterprise software taught the opposite lesson. A licence, once bought, made further use effectively free, so heavier use only improved the return. A generation of leaders learned to treat technology cost as fixed and solved.

Artificial intelligence, priced as most of it now is by consumption, reverses that. When you pay for each unit of use, cost rises with usage instead of flattening against a licence. The more the business relies on it, the more it costs, and there is no ceiling to stop the curve. The exposure has already produced casualties. Large, sophisticated organisations have burned their entire annual budget for these tools within months. Nothing went wrong; the tools worked, usage climbed, and the cost climbed with it while no one watched the total until it had run away. Some have responded by withdrawing the tools their own people had come to depend on, a continuity failure arriving through the finance function rather than the technology one.

The lesson for the board is that a dependency can be severed by price as cleanly as by a power cut, and the cut can come from your own side of the table as readily as the supplier's. A provider that triples its prices, a usage pattern that outruns the budget, a commercial model that shifts as a vendor reaches for profit, any of these takes a capability away as surely as a technical failure or a regulatory ban. Designing for resilience therefore means designing for affordability under stress: knowing what a critical capability costs at full use rather than pilot volume, finding the ceiling before the spending does, and keeping a fallback the business can still pay for when the economics turn.

Resilience is a board accountability

Resilience has always been the work that goes unrewarded until the day it matters more than anything else in the building. It is invisible while it works, it costs real money to maintain, and it can never be shown off, because its whole purpose is that nothing happens. So it loses almost every budget argument, it is the first thing cut when money is tight, and most organisations do far less of it than they would admit in a boardroom. The incentives run the other way. The executive who builds a flashy new capability is rewarded; the one who makes sure the business could survive losing one is, at best, not blamed when the bad day never comes.

The resilience that does exist is largely theoretical. Continuity plans get written, approved, filed, and never exercised, so nobody finds out the fallback no longer works until they reach for it in earnest. A plan that has not been tested is only a document about resilience, and the gap shows on the one day it is too late to close. The manual process retired when the system came in, the colleagues who knew how to run it and have since left, the workaround that depends on a second service that is also down, these look like nothing on paper and prove decisive in practice.

It has always been the board's to answer for; what has changed is how long the neglect can continue. For years it was treated as back-office detail, and a stable world made that affordable: the bad day stayed hypothetical and the bill for being unready never came due. A fragmenting world and a deepening reliance on volatile technology are ending that grace.

British Airways learned this in 2017. A contractor at one of its data centres switched off a power supply and then reconnected it in an uncontrolled way, sending a surge through the systems and physically damaging them; the failover that should have caught the failure went down too. Every flight was grounded, around seventy-five thousand passengers were stranded, the bill came to roughly eighty million pounds, and the chief executive spent the following days answering for it in public. The cause could hardly have been more trivial, or the consequence more public.

A board that goes on treating continuity as somebody else's technical concern, a matter for the people who look after the servers, will discover at the worst moment, with customers waiting and the operation stalled, that it was its own concern all along.

Start from what you cannot fail to do

Everything so far has worked upward, from the dependencies a business has accumulated to the outcomes they put at risk. There is a second view, and it runs the other way, from the top down. A board tends to find it the more natural of the two, because it begins with the business rather than the technology, and with one question: what must this organisation be able to do, no matter what?

The list is usually short and, once written down, surprisingly stable: serving customers, taking payment and meeting obligations, honouring its commitments and the law, and protecting whatever would be ruinous to lose. These are the outcomes that have to survive, and they barely change with the nature of the threat. A flood, a failed supplier, a withdrawn AI model, a cyber incident, a budget that overruns, all threaten the same handful of outcomes by different routes. The threats are many and unpredictable; the things that must not fail are few and knowable.

I once watched a firm lose its entire SAP environment to its own fire-suppression system. The halon discharged as designed, but the rack happened to sit directly in front of the jet, and the blast meant to protect the room took the system out instead. No risk assessment would ever have listed that, because no one could have imagined it. Keeping SAP running, though, was on the list from the first day, and an outcome you have resolved to protect is covered however absurd the cause turns out to be.

That asymmetry is what makes the top-down view powerful: it turns an impossible question into a manageable one. You cannot defend against every disruption that might come, and the attempt is how resilience budgets get exhausted and abandoned. You can name the few outcomes the business could not survive losing, and ask three plain questions of each: how long could we last without it, what would each day of its loss cost, and what does it depend on to keep working? The first two size the problem. The third is where top-down meets bottom-up, because tracing an outcome down through the people, processes, suppliers, and technology beneath it is how you find the dependencies that actually matter.

This is also where grading begins, because a dependency takes its importance from the outcome it serves rather than from the technology itself. The model that drafts marketing copy and the model inside the system that takes customer payments might be identical, but one sits beneath an outcome you could lose for a week and the other beneath one you could not lose for an hour. Protect them accordingly. A board that starts from the technology

drowns in the list of things that could go wrong; one that starts from the few things it cannot fail to do knows exactly which to care about, and has turned resilience from an unbounded anxiety into a finite, fundable piece of work.

The grading you already do

Every business already knows how to do this; it has only ever done it for storage. Not all data is worth the same to protect, and storage settled that so long ago that no one notices it was ever a question. The information a business needs instantly sits on fast, costly storage, copied continuously and recovered in minutes. The records it must keep but rarely touches sit on slow, cheap, archival storage that might take a day to retrieve and is none the worse for it. Between them runs a graded scale, and beside it a second for recovery: how fast each kind of data must come back, and how much of it the business can stand to lose. This is mature, unglamorous work, so routine that few would file it under resilience at all. It is just how storage is done.

The discipline beneath it is worth lifting out, because it is more useful than the example that taught it. You classify what you hold by how much it matters, then match the protection, and its cost, to the classification. You do not guard everything to the highest standard, which is ruinously expensive and mostly unwarranted, nor leave everything at the lowest, which is reckless with the things that matter. You grade, and you put your protection where the grading says it belongs. Applied to storage this is so familiar as to be invisible; applied to almost anything else it is missing entirely.

The strange thing is that the graded discipline every organisation practises on its data has almost never reached its compute, its services, or its AI. A tool used for something trivial and a model woven into a customer-facing operation are adopted with the same absent level of resilience thinking, as though their failures carried the same consequences. Few have asked which of their AI dependencies are critical enough to deserve a guaranteed fallback, and which are convenient enough that an outage would be an irritation rather than an incident. The question storage answers as a matter of course, how much does this matter and what protection does that justify, has never been put to the technologies now becoming load-bearing.

Extending that discipline is the practical core of resilience, and it dissolves the objection that resilience cannot be afforded. It is affordable because it is graded. You protect the few dependencies whose failure would halt the business, accept thinner protection for the many that would merely inconvenience it, and make the choice on purpose rather than by accident on the worst day. For most organisations the task is less to invent a new discipline than to take the one they already apply to their data and point it, with equal seriousness, at the compute, the services, and above all the AI they now depend on. The discipline is already in the building. It has simply never been pointed at the things that now matter most.

Designing a business that can take the hit

Resilience is building the business so it can lose a piece of technology and carry on. It means designing for the bad day rather than spending the good ones assuming it away. It is not a retreat from technology, which would only trade one exposure for another and give up the advantage good technology brings. A handful of principles separate the organisations that absorb a disruption from the ones it knocks flat, and none requires predicting which disruption will come, only accepting that one will.

The foundation is that the business should keep running, more slowly and less conveniently but running, when any critical system is taken away. That needs a fallback that actually exists, is kept in working order, and is used often enough to work on the day it is needed. Next, own what is genuinely yours, above all the knowledge and data the business runs on, in a form you control rather than locked inside one supplier's platform, so that changing providers is switching rather than rebuilding from nothing. And let no single supplier, or single country, hold a capability the business cannot do without, so that no decision taken in a boardroom or a government department you will never sit in can take the whole thing offline at a stroke.

A final principle is easily lost in the enthusiasm around a powerful new technology: keep the human capability to do the work alive even as the machine takes more of it on, so that losing it costs speed rather than the ability to function, and so someone who understands the work is still there to judge whether it is being done well. None of this is exotic, and none of it is expensive set against the exposure it addresses. It is the work an assumption of permanent stability let everyone skip, and the arriving world will not let anyone skip it much longer.

Rehearse the bad day

A business cannot know which disruption will come, and trying to forecast the exact future is a trap, because the one that arrives is usually the one nobody listed. The discipline that works is rehearsal rather than prediction. You take a few concrete, plausible scenarios, each standing for a class of failure, and work each one through to the end as though it had happened.

The scenarios have to be specific enough to bite. 'Something goes wrong' is too vague to act on. 'Our primary AI provider withdraws its models on a Friday with a day's notice', 'a regulator restricts a service we depend on and gives us thirty days to comply', 'our main cloud region is down for forty-eight hours', 'the renewal price doubles when we cannot absorb it', each is concrete enough to walk the business through and watch what breaks. Run one properly and in an afternoon it shows which outcomes stall, which fallbacks still work or have quietly rotted, and where the gaps sit, all of it in conditions you chose rather than ones chosen by events.

A handful of well-chosen scenarios covers most of the real exposure, which is what makes the approach affordable as well as effective. Because different threats converge on the same few critical outcomes, you do not need a scenario for every disruption, only enough to press each outcome that matters from a couple of directions. Three or four, taken seriously once a year, teach a business more about its resilience than any document ever written and filed.

This is the method that turns every principle in this briefing from an intention into something tested, and it is the cheapest insurance there is. Running a scenario on paper costs a meeting and an uncomfortable afternoon; meeting the same one unrehearsed costs the incident itself, at full price, with the clock running and no quiet room in which to think. The businesses that come through a serious disruption are usually the ones that had already lived through it once, deliberately, while it was still imaginary.

The questions worth asking now

Most executives reading this are not starting from a blank page. The tools are in use, the dependencies are forming, and the assumption of a stable, open technology world is already built into systems designed when it still held. The work is to look clearly at what exists and ask the questions the years of stability allowed everyone to postpone. They are uncomfortable by design, because their value lies in not being able to answer them with confidence.

The first is the plainest. If a technology the business depends on became unavailable at the end of this week, withdrawn by a supplier, restricted by a government, priced beyond reach, or simply broken, what would the following week look like? Would the business slow down, or stop? And is that answer one the leadership has tested and chosen, or drifted into while the world was still holding still? The second follows. Where has the organisation come to rely on technology it does not control without ever deciding to, and has anyone mapped those quiet dependencies before circumstances map them first? The third is the most easily deferred and the least safely ignored. Of the continuity arrangements the business believes it has, how many have been tested in the last year under conditions resembling a real failure, and how many are documents that have never met the day they were written for?

A business that can answer these three well is, almost by definition, resilient, because answering them requires having done the work resilience consists of. A business that cannot is not necessarily in trouble, but it is carrying a risk it has not measured, in a world busily withdrawing the stability that risk was betting on. The comfortable response to a briefing like this is to file it under things that happen to other, less careful organisations. The truer one, for most who have read this far, is to notice they have just been asked three questions about their own business and could not confidently answer any of them.

The advantage of asking now, while the questions are still an exercise to be done with a coffee on a quiet morning, is that the answers can be acted on calmly and cheaply. The first answer is the cheapest of all, because mapping where you are exposed costs nothing but the willingness to look. The alternative is to meet the same questions later, as an emergency, when they cost far more to answer, when the explaining is done to a board or a regulator rather than to a coffee cup, and when there is no quiet morning in which to think.

The figures and examples in this briefing reflect the position at the time of writing. In a landscape moving this fast, the specifics will date; the argument they illustrate will not.

MultipleWorks helps organisations understand where they have become dependent on technology they do not control, and how to design and test for resilience as the ground shifts beneath them. If this briefing reflects something your leadership team is working through, we would welcome the conversation. Get in touch at hello@multipleworks.com.hk.

References

The French government directive. Direction interministérielle du numérique (DINUM), interministerial directive of 8 April 2026 requiring every ministry to prepare plans to migrate from Microsoft Windows and Office and to reduce extra-European digital dependencies by autumn 2026, covering operating systems, collaboration tools, databases, and artificial intelligence platforms. Reported by The Next Web and others, 2026.

The British Airways outage. International Airlines Group, statements on the May 2017 British Airways data-centre failure, including the gross cost estimate of approximately £80 million and the attribution to a power disconnection and uncontrolled restoration. Reported by BBC News, June 2017.

SUGGESTED CITATION

MultipleWorks (2026d). Built for a world that's leaving: An executive briefing on technology resilience. MultipleWorks Limited, Hong Kong. Available at multipleworks.com.hk/briefings.

v1.0 · June 2026