A GUIDE FOR EXECUTIVES · GULF EDITION

BCM without illusions

How to protect your company's cash flow

A practical map of business continuity for organisations in the GCC: from spotting your weak points to running the business on the day the systems go down.

---

risk-place.org · Business continuity and resilience practitioners working to ISO 22301, NCEMA guidance and international good practice

CONTENTS

Check your resilience in 5 minutes — free diagnostic at risk-place.org/quiz.html

Introduction. What this handbook is about and how to use it

This handbook is written for the owner and the top leader of a company — and for the team that will be acting alongside you when something goes wrong.

Most leaders have a P&L, a budget and a sales plan. The financial result is visible. Operational resilience — the ability to keep cash flowing when a key link breaks — is almost always a blind spot. It stays invisible in the reports until the moment it is too late.

Business Continuity Management (BCM) is exactly the systematic view of that blind spot. It is not about backup servers and not about 'the IT people'. It is about how your business keeps earning on the day its usual support disappears.

How to use this handbook

This is a map and a way of thinking, not a folder of ready-made forms. The job of this book is that, once you finish it, you clearly see where your business is exposed

and what to do about it, step by step. The full method — with templates, worked examples and assessment — is taught in the ERGP programme at risk-place.org.

The logic follows a leader's path: first assess where you would break, then work through your critical processes (BIA), then strengthen the two halves of resilience — IT infrastructure and people — prepare for dealing with regulators, and finally know what to do on the very day the systems go down.

PART I

Why cash flow is fragile

The executive's blind spot and the true cost of a disruption

1.1 The owner who cannot see the weak points in cash flow

A company's cash flow is not a line in a report. It is a living, fragile structure assembled from strategy, partner agreements, client relationships, internal processes and IT infrastructure. Like a tower of dominoes, it holds as long as every piece stays in place.

Let one major risk materialise — a software failure, a cyberattack, the loss of a key supplier or a data breach — and the structure starts to collapse. Recovery takes days, weeks, sometimes it never happens. And the blow lands not on the IT department but on your wallet and your reputation.

A dangerous assumption. 'If we have IT staff, the business is protected.' Support staff update software, watch the antivirus and restore files. But the question 'what happens to incoming cash if this fails?' is analysed by no one in most companies. That is the direct responsibility of the owner and the CEO.

1.2 What business continuity is — and why it is not about 'the IT people'

Business continuity is the ability of a company to respond effectively to a disruption, keep delivering its product or service at an acceptable level, and restore broken processes as fast as possible. The key words are 'product', 'level' and 'time', not 'server'.

Backups, antivirus and a spare communication channel are means. Continuity is a management goal: how much money you lose per hour of downtime and how quickly you return to normal. That is why the business, not only IT, must own the topic.

Mini-case: 'we had IT, we had antivirus — the business stopped for a week.' A trading company in the region thought it was protected: it had a system administrator and a licensed antivirus. Ransomware came in through an ordinary employee's email and encrypted the shared drive with orders and inventory.

Backups were being made, but no one had ever tried to restore them — recovery took six working days. What was lost was not files, but a week of shipments and two customers.

1.3 How much a day of downtime costs

Until the cost of downtime is calculated, the topic stays abstract. The math is simple: take revenue for a period, divide by the number of working days, and add what does not show up in revenue directly — contractual penalties for missed obligations, customer churn, overtime spent on recovery, reputational damage.

That number changes the conversation. 'It would be nice to work on continuity' is one thing. 'Every day the key process is down costs us a specific sum, and we know how much' is another.

Downtime has a price — and it is rising. In a 2024 survey, more than half of organisations said their most recent significant outage cost over $100,000, and one in five said it cost more than $1 million (Uptime Institute, Annual Outage Analysis 2024). For a company running on thin margins, a single serious disruption is a question of survival, not convenience.

Mini-case: the number that ended the argument. A services company calculated the cost of a day of downtime for its core process for the first time — it came to about two months of profit. After that, the debate over whether a backup channel and exercises were worth it settled itself: the cost of preparing was less than one day of losses. Calculate the cost of downtime before the outage, not after.

1.4 The four threats that most often bring a business down

In practice, cash flow tears at four typical points. We will look at each in more detail later; for now let us just name them — almost certainly at least one will ring true.

- **A key employee leaves.** The person the process rests on resigns or relocates — and with them go the arrangements, the passwords and 'how it actually works'. In the Gulf, where teams are international and mobile and a departure can mean leaving the country within weeks, this risk is sharper than almost anywhere else.
- **A cyberattack.** Ransomware, a breach, a hacked mailbox. What stops is not 'IT' but sales, shipments and payments. The region's visibility makes it a target: energy, logistics and finance in the GCC are attacked deliberately, not

incidentally.

- **A key supplier or logistics link fails.** One supplier of materials, components or a service drops out — or a strait, port or air corridor closes and the whole chain halts. Recent years have shown how quickly regional shipping disruptions translate into empty warehouses.
- **Software becomes unavailable.** CRM, ERP or a cloud service stops working — for a technical reason or an external one, including sanctions-driven cut-offs of foreign vendors. If your critical system is run from another jurisdiction, its availability is not entirely your decision.

Where to start. Do not try to close all four at once. Take the threat that is most real for you and follow it through the whole book — assess it, price it, strengthen it. One closed vulnerability is worth more than a beautiful plan covering everything at once.

PART II

Step 1. Assessment: where you will break

Points of risk concentration and critical weak spots

2.1 How your cash flow works and where risk concentrates

The first step is to break down which operational and management links make up your cash flow. Where exactly money enters the company, and what has to work for it to keep entering.

Then find the points of risk concentration — places where everything rests on one thing: one server, one supplier, one sales channel, one key person. These are single points of failure. Until you see them they look reliable; on the day of an outage they are exactly what stops the business.

The link stress-test — three questions. For every critical link ask yourself: 1) If this function stops in the next hour, can I restore it? 2) How much will that cost? 3) Will customers and partners find out? If you cannot answer even one of these calmly, you have found a weak spot.

2.2 Critical weak spots: how to see them

A weak spot is not an abstract 'risk in general' but a concrete dependency whose break hits the money. Look at dependencies on people, on systems, on counterparties and on external conditions. Often the most dangerous weak spot is the one everyone knows about but never gets around to fixing.

Here it is important not to drown in lists. At this stage you need not a perfect table but an honest picture: three to five places whose failure would stop the money coming in. Start strengthening with those.

Mini-case: the recipe held by one supplier. A food producer bought a key ingredient from a single overseas supplier for years — and the supplier also held the refined recipe for the blend. When the supplier doubled the price and missed deadlines, it turned out there was no one to switch to quickly: the company could not reproduce the recipe on its own. The weak spot was not 'the supplier' but the

lost control over its own know-how.

PART III

Step 2. BIA to ISO 22301

Business impact analysis in plain words

3.1 What a BIA is — and why you need one

A BIA (Business Impact Analysis) is the heart of the international continuity standard ISO 22301 — and of the Gulf frameworks built on it. In essence it answers one question: which processes are so critical that we could not survive their loss, and how fast must they be brought back?

A BIA moves the conversation from 'what if' to numbers. For each critical process you determine how fast losses grow during downtime and how long an interruption is tolerable. Priorities are born from the BIA: what we restore first, and what can wait.

Two key measures. RTO (Recovery Time Objective) — the target recovery time: how quickly the process must be working again. RPO (Recovery Point Objective) — the acceptable data loss: how much recent data you can afford to lose (for example, 'no more than the last hour'). RTO answers 'how fast', RPO answers 'how much data can be lost'.

3.2 What to look at and how it is done

A BIA is built around critical processes, not departments. For each you look at what it produces, which people, systems and suppliers it depends on, how fast damage grows if it stops, and the maximum tolerable period of disruption (MTPD).

Logic matters more than form. You do not need to analyse everything: 80% of the risk usually lives in 20% of the processes. Your goal is to single out those few and understand their dependencies in depth — down to the specific person, server and counterparty.

The mark of a good BIA. After it you can answer a board question in one sentence: 'If process X goes down tomorrow, we lose this much per day, we bring it back in this many hours, and here is what we need to do it.' If you cannot, the BIA

is not finished.

Mini-case: the process only one person knew. While mapping critical processes it emerged that customer invoicing rested entirely on one employee and their personal file. When they were on leave abroad, invoices did not go out. The BIA did not add new risks — it made visible a dependency the company had not noticed for years. That is the point of the analysis: to surface the critical before it breaks.

PART IV

Step 3. Strengthening: IT resilience

Backups, cloud, concentration risk and the traffic-light model

4.1 Backup and recovery

Half of resilience is technical. And technical continuity begins not with the backup itself but with tested recovery. A backup you have never restored is hope, not protection.

The baseline is the '3-2-1' rule: three copies of the data, on two different media, one of them off-site (in another data centre or the cloud). For truly critical systems you add a standby site you can switch to. But any scheme only has value if it has an RTO and RPO proven in exercises.

- Run test restores regularly, not just backups.
- Keep at least one copy physically and logically separated — so ransomware cannot reach it too.
- Know your RPO for each system: how much data you would actually lose in a failure right now.

4.2 Distributing server capacity and the cloud

Where to keep systems — in your own server room, in the cloud, or a hybrid — is not a matter of fashion but of managing risk. Each option has its own weak spot: your own room depends on one location, its cooling and its power — no small matter in a climate of fifty-degree summers; the cloud depends on the provider and the link; a hybrid is harder to run but more resilient.

The practical principle is geo-redundancy and no single point of failure. Critical services should not depend on one data centre, one link and one provider. In the Gulf add one more question: data-residency rules. Several regulators require certain data to stay in the country — check before you architect, not after. Ask your IT team directly: 'If this data centre or this provider goes down for a day — what exactly stops for us?'

What the business, not only IT, should keep in view. A list of critical systems

and where they are hosted; provider contracts and the guaranteed service levels (SLA); whether there is a standby site and where; who brings a system back on a Friday, and on what phone number. This is management information, not technical detail.

4.3 Concentration and third-party risk

A very modern continuity risk is concentration — how much of your business rides on a single outside party. One cloud provider, one SaaS platform, one critical software vendor. When they fail, you fail — usually without warning and through no fault of your own.

This became painfully clear in July 2024, when a faulty update to a widely used security product took millions of Windows computers offline at once — grounding flights and disrupting hospitals, banks and retailers around the world, Gulf carriers and airports included. Not one company's mistake, yet thousands of businesses stopped together.

The answer is managing dependency. Map which critical systems rely on a single provider or a single vendor's software, and ask what happens if that provider has a bad day. For key positions, plan ahead: an alternative provider or region, the ability to run in a degraded mode, and above all the right to get your data out.

- Keep a register of critical third parties and how dependent you are on each.
- For every critical dependency, hold a 'plan B': what replaces it, and how long a switch would take.
- The right corporate actions: secure data-export rights and keep critical data where you — not a single vendor — control it.

Mini-case: when one vendor's outage stopped the shift. A company ran its whole operation on a single cloud platform with no fallback. When the provider had a multi-hour outage, order-taking, shipping and billing stopped at once — there was simply nowhere else to go. What saved the day was an old habit of exporting key data locally each week. The lesson was expensive: critical data must live somewhere you control, and no single provider should be able to take your entire business offline.

4.4 The traffic-light model: Red, Amber, Green

To keep recovery from turning into chaos, the response is split into three sequential stages — like traffic lights. This is our working model, proven on real

operations: it is simple, easy to explain to a team, and convenient to build a plan around.

Red — first hours: stop the bleeding and notify. Immediate response: activate the team, alert the right people, contain the damage and switch on the communications plan. The goal is to stop the bleeding.

Amber — temporary mode: keep running while you recover. Move to workarounds and backup capacity, hold a critical minimum of output, and work the plan with clear roles and deadlines.

Green — full recovery. Return systems and processes to normal, restore data, review the incident and update the plans for next time.

The strength of the model is that everyone on the team understands which stage the company is in and what is required of them right now. First stop the bleeding, then live temporarily, then return to normal. You cannot skip stages.

What to check on your IT this week. Do you have a fresh backup of critical data — and when did you last actually restore it? How many systems depend on a single data centre or a single provider? Which critical software comes from a single vendor, and what happens if it has a bad day? Three honest answers will already show where you are exposed.

PART V

Step 3. Strengthening: people

Why a team 'stops working' in a crisis — and how to change that

5.1 What stress does to people

The second half of resilience is people. You can have perfect plans and standby servers, but in a real disruption the decisions are made by living people under pressure. And here biology takes over.

Acute stress triggers the 'fight or flight' response: the brain switches from rational analysis to fast, intuitive reactions. In everyday life that saves us; in managing a crisis it raises the chance of a fatal mistake. A person either rushes about doing unnecessary things or freezes.

Why it matters to know this in advance. You cannot 'pull yourself together' by willpower at the moment reason has already switched off. But you can train simple actions in advance so they fire automatically, and distribute roles so that not everything depends on one overloaded person.

5.2 How people behave in a crisis

People react differently, and predictably. Some rush to act without a plan ('fight'), some avoid and withdraw ('flight'), some freeze and wait for orders ('freeze'). None of these is malice — it is physiology, and it must be worked with, not punished.

In an international team — the norm for the Gulf — add one more layer: in a crisis people fall back on their first language and their home culture's habits of hierarchy. A plan that assumes everyone will challenge the boss in fluent English at 3 a.m. is a plan on paper. Keep crisis instructions short, visual and rehearsed.

A separate danger is accumulated burnout. The World Health Organization classes it as an occupational phenomenon with three signs: emotional exhaustion, growing cynicism and reduced effectiveness. A burned-out team enters a crisis with no reserve — and breaks faster.

Early warning signs in a team. Emotional exhaustion, cynical jokes about work

and clients, a drop in quality from normally strong people. This is not 'laziness' but a signal that the reserve is running out. Spotting them is the leader's job before a crisis, not during it.

5.3 How to prepare the team

People's readiness is built on three things: clear roles, practice and simple self-regulation techniques. Roles — so everyone knows their patch and does not wait for one person to decide. Practice — so the right actions become automatic.

- **Tabletop exercises:** play out a server failure or an office lockout 'on paper'. Only a live rehearsal shows what works in the plan and what does not.
- **The 10-second pause:** a short pause before a decision under pressure noticeably lowers the chance of a critical error. A simple technique you can use right in a meeting.
- Do not suppress the team's emotions — channel the energy of stress into concrete tasks from the plan.

Mini-case: an exercise that saved a real day. A company ran a tabletop exercise on a warehouse outage — on paper, who does what in the first hours. It surfaced that half the team did not know where the plan was or whom to call. Two months later a real outage happened, and the same team handled it calmly: they had lived through it once already. Exercises are cheap, and they return the confidence you cannot buy in the moment of a real incident.

PART VI

Regulators and public authorities

Whom to report to in an incident — and how to prepare

6.1 Whom to report to, and about what

In a serious incident the technical side is only half the job. The other half is communication: whom you must and should tell, what, and by when. Silence or muddled messages often do more damage to reputation and to relations with the authorities than the outage itself.

Decide the circle of recipients and their order in advance. Typically: the relevant regulator and supervisory bodies (within the required deadlines), the customers and partners whose obligations are affected, your insurer, the media if needed, and always your own staff, so they do not learn the news from the outside.

The rule of deadlines. Many regulators set hard deadlines for reporting an incident — sometimes hours, not days. You need to know these before the incident: looking them up during the outage is too late. For regulated sectors in the Gulf this is a mandatory requirement, not goodwill.

6.2 How to prepare in advance

Crisis communication is not improvised — it is prepared. The minimum: a communications plan with a list of recipients and deadlines, a designated single spokesperson, and pre-agreed message templates for typical scenarios — in English and Arabic, because in the Gulf your regulator, your staff and your customers may each expect a different language.

- **One voice:** in an incident, one designated person speaks for the company — not whoever happens to be in the chat or on social media.
- **What not to do:** before the facts are clear — do not admit fault, do not name damage figures off the top of your head, do not promise deadlines you do not know.
- Keep current contacts for regulators, key customers and the insurer at hand — including on paper, in case the systems are down.

Mini-case: five voices instead of one. In the thick of an incident several employees wrote about the problem at once — each their own version. Customers and partners received contradictory messages, and the anxiety grew larger than the outage itself. The lesson is simple: a designated single spokesperson matters more than the speed of individual reactions, and a pre-agreed message template saves hours at the crucial moment.

PART VII

What regulators in the Gulf require

Standards and requirements — and why they matter to the business

7.1 The framework you are expected to know

Business continuity in the GCC is not only common sense — for a growing number of organisations it is a formal requirement. Knowing this framework is doubly useful: it protects you from findings and penalties, and gives you a ready-made structure to get organised.

- **ISO 22301** is the international reference for business continuity management systems, and the backbone that regional requirements build on. Certification to it is increasingly requested in government and enterprise tenders across the Gulf.
- **United Arab Emirates:** NCEMA — the National Emergency Crisis and Disasters Management Authority — maintains the national business continuity standard (AE/SCNS/NCEMA 7000) and expects critical entities to run documented continuity programmes aligned with it.
- **Saudi Arabia:** SAMA requires banks and financial institutions to operate a business continuity management framework, with defined governance, BIA, testing and reporting; the Personal Data Protection Law adds duties around data incidents.
- **Financial centres and sector regulators** — the Central Bank of the UAE, DFSA in the DIFC, FSRA in ADGM, the Qatar Central Bank and their peers — publish their own operational-resilience and outsourcing rules for firms under their supervision.

Requirements evolve; before building internal documents, verify the current edition of each framework that applies to you. The point of this chapter is not the reference numbers but the pattern: in the Gulf, continuity is moving from good practice to supervised obligation.

The business view. A recognised standard is not paperwork for its own sake. A working ISO 22301 or NCEMA-aligned programme earns credibility in tenders, with banks, insurers and large customers — you show that you are a reliable partner. In a region that competes on trust and long-term relationships, that is a

commercial asset, not a compliance cost.

PART VIII

Working when everything is down

A practical scenario: no CRM and no ERP tomorrow

8.1 Scenario: tomorrow your CRM and ERP are down

Picture it literally: you arrive in the morning and your main systems — CRM, ERP — do not work. It does not matter why: an outage, an attack, an external cut-off. The only question is: how do you keep earning today, not in a week? The good news is that you can prepare for this, and the preparation costs almost nothing.

The point of a temporary mode is not to reproduce every system function but to hold a critical minimum: take an order, ship it, get paid, keep the customer. Everything else can wait.

- **Communication:** a pre-agreed backup channel — a work messenger or chat, a list of key people's personal phone numbers, group calls. Arrange this before the crisis.
- **Orders and records:** a simple offline form (a spreadsheet or a paper sheet) where orders, shipments and payments are written down while the systems are gone. It is transferred into the system afterwards.
- **Roles:** who takes orders, who confirms shipment, who keeps in touch with customers, and who with IT. One person should not carry it all.
- **Customers:** a short, honest message — 'we are working in a special mode, we are taking orders, delays are possible' — beats silence.
- **Return to normal:** when the systems come back, everything done by hand is carefully reconciled and entered — so no data from the transition is lost.

Mini-case: half a day on paper instead of a week of downtime. A trading company lost access to its CRM and telephony for a day because of a provider outage. Instead of sending people home, the manager moved order-taking to a messenger and a simple spreadsheet, handed out roles and warned key customers. The company worked at about seventy percent of normal — instead of zero. In the evening the data was entered into the system. The difference between 'we prepared' and 'we panicked' was worth a day's revenue.

ROADMAP

90-day resilience roadmap

Resilience is not built in a day, but not over years either. A realistic horizon for first results is three months.

Days 1–30. See

Calculate the cost of a day of downtime. Find three to five critical weak spots and single points of failure. Take the free self-assessment at risk-place.org.

Days 31–60. Strengthen

Test a restore from backup. Close the most dangerous dependencies — in IT and in people. Assign roles for an incident.

Days 61–90. Rehearse

Run a tabletop exercise on the failure of a key system. Draft a communications plan — in both working languages — and a traffic-light response plan. Agree on backup channels.

About the authors

This handbook was prepared by the risk-place team — practitioners in business continuity and operational risk. Our work is grounded in international good practice and the ISO 22301 standard, with hands-on experience implementing continuity programmes at large industrial and service organisations and training leadership teams to act under uncertainty and stress. We speak the language of money, not paperwork.

WHAT NEXT. This book is a map. Practice is what helps you walk the route faster and without mistakes.

The ERGP programme — the Executive Certificate in Enterprise Resilience Governance — works through assessment, BIA, resilient operations, crisis leadership and governance in 94 chapters and six modules, fully available in Arabic

and English, with a verifiable certificate. Details: risk-place.org

And you can start right now: the free diagnostic shows your readiness level in 5 minutes — risk-place.org/quiz.html