WHITEPAPER



CLOUD AND PLATFORM OUTAGES

Coordinate anyway.

A practical playbook for organisations that depend on cloud platforms and digital services. How to maintain communication, coordinate incident response and protect critical operations when essential technology becomes unavailable.

T+00:00 T+00:02 T+00:05 T+00:10 T+00:15

INSIDE

Contents

What to do when the platforms you depend on, and do not run, stop working: reaching people, holding the response together and proving afterwards what you decided.

Foreword 03

Executive summary 04

Introduction 05

01 **Anatomy of a platform outage** 06

02 **Incident scenarios** 12

03 **Lessons and what to change** 18

04 **Practical toolkit** 22

Checklist and glossary 36

Use cases and buyer's guide 41

Sources 43

FOREWORD

The day the tools we coordinate with were the thing that broke.

On Thursday 12 November 2026 a Microsoft 365 incident took Exchange Online away from organisations in dozens of countries at once, and degraded a dozen services that lean on it.[1,2] It was not an attack. Microsoft found no indication of malicious activity, apologised, and traced the failure to a defect of its own.[3,4] The detail is on the next few pages, and it is worth reading closely, because almost nothing about it was exotic.

What makes it worth a document like this is not its size. It is what it removed. Mail stopped. Files and calendars went with it for many. For teams whose response plan starts with an email, the first question was not what to do, but how to tell anyone.

That is an awkward problem, because every plan assumes a way of reaching people, and most plans assume the same way. This whitepaper is about the gap that opens when the assumption fails: how to reach your people when the usual channels are the thing that failed, how to hold a response together across a night and a handover, and how to prove afterwards what you decided and when.

The case study is Microsoft 365 because it is recent and well documented. The problem is not Microsoft's. Any cloud platform, identity service, business application or network you depend on and do not run can do the same thing to you, and the work that protects you is the same in each case.

The Crises Control team
Crises Control

EXECUTIVE SUMMARY

When the coordination layer fails, you need an alternative designed to avoid the same dependencies.

01 · THE OPERATIONAL CHALLENGE

- A cloud productivity platform carries email, chat, files, sign-in and often the security console, so one fault removes several things at once.
- The same platform usually carries the incident process itself: the call tree, the plan, the ticket queue and the bridge invitation.
- Outages of this kind last for hours or days, not minutes, so they cross shift handovers, weekends and time zones.

02 · INCIDENTS AND RISKS TO PREPARE FOR

- A platform or collaboration outage that takes email, chat and files down together.
- An identity or single sign-on failure that locks staff out of everything at once.
- A supplier change gone wrong, which is a more likely cause than an attack.
- A cyber incident where you must deliberately take the usual channels off the air.

03 · KEY RECOMMENDATIONS

- Keep one alerting route that is hosted elsewhere and signs in separately, and test it.
- Decide now who may declare an incident when nobody can email, and how they are reached.
- Write the fallback for each critical service on paper, with the numbers printed, not linked.
- Treat the backlog as part of the incident: restored is not the same as caught up.
- Rehearse a two-day outage with a handover, not a one-hour one.

04 · WHAT EFFECTIVE COORDINATION LOOKS LIKE

- The response team is reached and acknowledges in minutes, on channels that do not share the failure.
- One log of decisions, messages and tasks that survives the outage and is readable afterwards.
- Customers, staff and regulators hear something useful early, on a channel you still control.
- Every decision and timing can be shown to a board, an auditor or a regulator.

05 · HOW CRISES CONTROL SUPPORTS THE RESPONSE

Most organisations can send alerts. Few can manage incidents end-to-end.

Crises Control is an Operational Incident Coordination Platform. It reaches people by push, SMS, voice and email, delivered through carriers that do not depend on your productivity platform, tracks who has acknowledged, and escalates on no response under the rules you configure.

Response plans run as owned tasks in one incident workspace, leaving a timestamped record that can be reviewed and exported. The platform supports your response plan; it does not replace it.

INTRODUCTION

Three questions you cannot answer with email.

Written for IT and technology leaders, business continuity and resilience managers, CISOs, risk and compliance professionals, incident and emergency teams and operations directors. On 12 November 2026 the question was not whether an organisation held a continuity plan. The question was whether it could reach the people named in that plan, agree what to do, and show afterwards what it had decided.

Can we reach everyone?

A route to the whole response team that does not sign in through the platform that has failed, tested against a current contact list rather than a stale one.

Can we run it for two days?

A response that survives a shift handover, a night and a weekend, with deputies named and a handover note that does not live in a mailbox.

Can we prove it?

Decisions, timings and messages recorded as they happen, so the board, the auditor and the regulator get the same account.

PART 01

When the platform goes.

One supplier fault can remove email, chat, files, sign-in and your security console at the same time. Understand the shape of it before it happens.

01 What happened on 12 November 2026

02 What regulators expect of a dependency you do not run

03 One service, mapped before you need it

CASE STUDY · MICROSOFT 365 GLOBAL OUTAGE

What happened on 12 November 2026

> **The failure took away the tools most organisations use to run an incident.**

Microsoft's preliminary post-incident report puts the start at 16:20 UTC on Thursday 12 November 2026, 11:20 Eastern, and the end at 16:10 UTC the next day: just under twenty-four hours. Microsoft tracked the incident under the combined reference EX1550543/MO1550821.[5] The advisory named Exchange Online, Teams, SharePoint Online, Microsoft 365 Copilot and Defender XDR, and said users in any region could be affected.[1] Problems were reported in at least 70 countries, against a user base Microsoft puts at more than 430 million.[4,6] Impact was not uniform: Exchange Online was worst hit, while Teams chats and meetings worked for most of the outage.[5]

The explanation changed through the day. Microsoft first pointed to elevated infrastructure load, then to servers being unexpectedly unavailable rather than any rise in traffic.[1] On 14 November it described an issue introduced during an earlier network capacity upgrade that surfaced after a recent network change.[3] Its preliminary post-incident report of 15 November, made available to tenant administrators, is more specific: a latent defect in an internal Windows process on the Microsoft 365 networking gateway servers, where competing tasks waiting on shared resources exhausted processor and memory until servers stopped answering.[5,7] That report states it is preliminary and that a final one will supersede it. At the time of writing the final report had not appeared.[5]

It was not an attack. Microsoft said there was no indication of malicious or cyber activity, and that this was a technical issue unrelated to cyber or to AI.[1,3] According to the same preliminary account, three ordinary conditions had to meet for the defect to matter: the fix carrying it had been in place since July, an automated failover had moved Teams traffic onto the gateway shared with Exchange Online about fifteen hours earlier, and roughly 14 per cent of gateway capacity was already down from an unrelated problem. Then automated recovery spread the load, and the boundaries meant to contain a fault did not hold.[5,7]

CASE STUDY · MICROSOFT 365 GLOBAL OUTAGE

What Microsoft said, and who it reached

> **None of them had a fault of their own. They shared a supplier.**

Recovery was gradual, and deliberately so. Microsoft said that given the scale of demand it restored services in stages: “While this took longer, it helped ensure a stable restoration.”[6] By its preliminary report, Asia Pacific recovered at about 11:00 UTC on 13 November, Europe at 11:35 and webmail globally at 14:17, and the incident closed at 16:10.[5,7] Microsoft said on 14 November that it had restored the remaining customers and regions, with enterprise email queues still draining.[3,8,9,10] Recovery varied by service, by organisation and by region, and users were warned to expect intermittent access and delayed delivery meanwhile.[4]

Microsoft apologised. It said it was “deeply sorry for the disruption this incident has caused our customers and partners” and that it took the trust placed in its services “extremely seriously”.[4] It later added that it regretted the impact on those who depend on it.[3] The detail easiest to miss is the one that outlasts the outage: Microsoft extended its email delivery timeout to 48 hours to help protect delayed messages while delivery caught up, and when those messages arrived they carried the time they were delivered, not the time they were sent.[4]

The impact followed the dependency, not the sector. In the UK, GP surgeries, pharmacies, police forces and schools were affected; some surgeries asked patients not to message through their websites, and forces asked the public to telephone instead.[4,11] Bloomberg reported NHS staff turning to paper, calls and texts, and reports of rescheduled operations, while clinical work did not appear to be affected; one London trust said its operations went ahead as planned.[12,13] Japan’s finance ministry could not receive email for most of the day and ministries posted paper notices.[14] The Reserve Bank of Australia lost Outlook across departments while critical operations continued; New Zealand paused visa processing; court systems in Singapore and the United States reported delays.[6,14] In Canada, a Toronto hospital network declared a code grey and said patient care was not affected; governments, courts and rail operators across North America reported delays of their own.[15,16] None had a fault of their own. They shared a supplier.

WHAT WENT WITH IT

One outage, several capabilities

What a productivity platform usually carries, and what stops when it does. Check which of these are true for your organisation.

WHAT FAILS	WHAT STOPS WITH IT	THE USUAL FALLBACK	WHY THE FALLBACK OFTEN FAILS
Email	Approvals, supplier escalation, customer correspondence, alerting that runs over email	Phone and personal mobile	Numbers are held in the mail client that has gone
Chat and meetings	The incident bridge, the war room, cross-team coordination	A dial-in conference number	The invitation and the PIN are in a calendar nobody can open
Files and intranet	Response plans, runbooks, contact lists, call trees	A printed copy	The last printed copy is older than the last reorganisation
Single sign-on	Ticketing, HR, monitoring consoles, the status page you publish	Local or break-glass accounts	Break-glass credentials are untested or expired
Security tooling	Alert triage and threat visibility during a major event	Manual review	The team is already committed to the outage
Status and service desk	Telling staff what is happening and collecting what they see	A mass notification route	It signs in through the platform that failed

REGIONAL NOTES

A dependency you do not run is still yours

These are financial-sector regimes and they bind only the entities within their scope, so most readers of this playbook are not caught by them. They are still the clearest published account of what a supervisor expects of an organisation that depends on a service it does not run. Check which, if any, apply to each legal entity in your group.

REGIONAL NOTE · UNITED KINGDOM

Operational resilience rules require in-scope firms to identify important business services, set impact tolerances and test against them, and firms had to be able to remain within tolerance by 31 March 2025.[17,18] The critical third parties regime lets the regulators oversee providers whose disruption could threaten the stability of, or confidence in, the UK financial system. Four cloud and technology providers were designated from 13 July 2026: Amazon Web Services, Google Cloud, Microsoft and Oracle, through named UK or EEA entities.[19,20] Five more follow from 23 November 2026, taking the total to nine from that date.[21]

IN PRACTICE The five added from 23 November cover financial services, market information, and technology and payment services, so the regime now reaches past cloud.[21] Designation moves none of your own responsibility. From 18 March 2027, FCA rules require an incident to be reported only where a firm reasonably believes it risks intolerable consumer harm, safety and soundness, or market stability.[22,23]

REGIONAL NOTE · UNITED STATES

The FFIEC Business Continuity Management booklet is examiner guidance rather than a rule: it shapes what examiners look for at institutions and their service providers.[24] NYDFS Part 500 is binding on covered entities. Section 500.16 requires written incident response and business continuity and disaster recovery plans, a plan for communicating with essential persons during a disruption, and testing at least annually with the staff and management critical to the response.[25]

IN PRACTICE Section 500.16 also requires current copies of those plans to be accessible during a cybersecurity event. Section 500.17's 72-hour clock runs from determining that a cybersecurity incident has occurred, which a supplier outage with no malicious activity will not usually start.

REGIONAL NOTE · EUROPEAN UNION

DORA, Regulation (EU) 2022/2554, has applied since 17 January 2025 to EU financial entities. Article 11 requires an ICT business continuity policy, response and recovery plans, and periodic testing of continuity plans, expressly including critical or important functions contracted out to ICT third-party providers. Article 14 requires crisis communication plans for responsible disclosure of major incidents to clients, counterparts and the public.[26]

IN PRACTICE Article 14 also requires communication policies for internal staff and for external stakeholders, and at least one person tasked with implementing the communication strategy. Crisis communication is an obligation with an owner, not an optional extra.

REGIONAL NOTE · MIDDLE EAST

The UAE Central Bank's Operational Risk Management Regulation, Circular 1/2026, took effect on 14 September 2026 for licensed financial institutions. Article 3.4 requires each critical operation to be mapped to the people, technology, processes, data, facilities and third-party providers that deliver it. Article 11.8 requires recovery objectives that reflect the tolerance for disruption approved by the Board.[27] Saudi Arabia's SAMA framework applies to SAMA member organisations and to all banks operating in the Kingdom.[28]

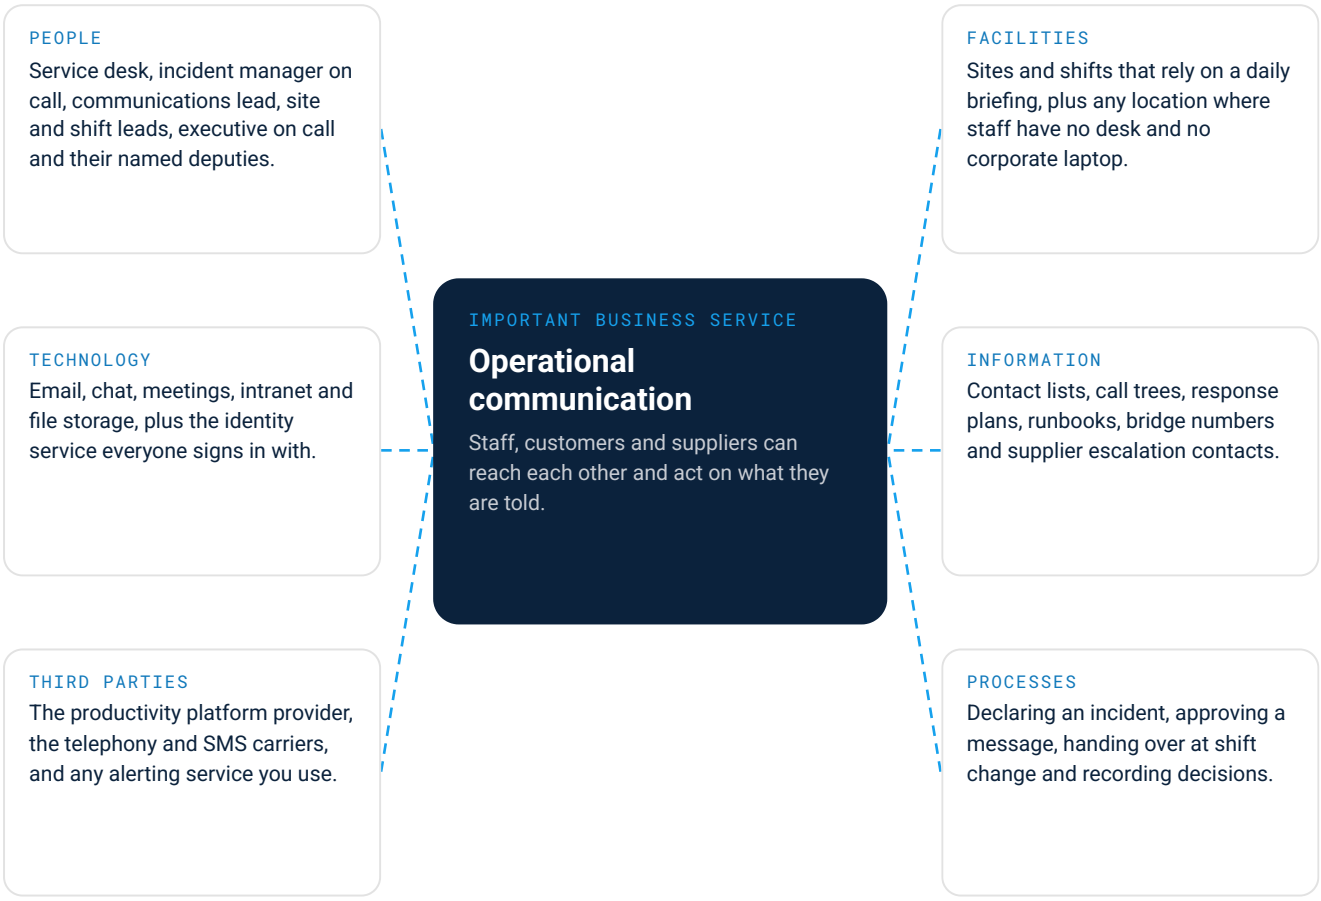
IN PRACTICE Article 3.4's mapping is the dependency map overleaf under another name. Confirm which rulebook applies to each entity, including entities in financial free zones.

A summary for orientation, not legal advice. Requirements depend on sector, size, licence and jurisdiction, and change over time. Confirm the current text with your legal and compliance teams.

DEPENDENCY MAP

One service, mapped before you need it

An illustrative map for day-to-day operational communication, the service almost nobody lists as critical until it fails. Map yours the same way.



TIME TO REACH THE TEAM (ILLUSTRATIVE)	ACKNOWLEDGEMENT TARGET (ILLUSTRATIVE)	FALLBACK BRIEFING (ILLUSTRATIVE)	LIMIT BEFORE ESCALATION (ILLUSTRATIVE)
15 minutes	**80% in 30 minutes**	**Every 2 hours**	**4 hours**

> All values are illustrative. Set your own, approved by whoever owns the service, and record where each number came from.

PART 02

Incident scenarios.

What you do in the first hour decides whether the next two days are coordinated or improvised.

01 Severity, and a first hour on the clock

02 The same afternoon, lane by lane

03 Three severe days and the work that follows

SEVERITY

A severity ladder for platform outages

Agree the triggers in advance so nobody debates the level while staff sit idle. If in doubt, declare higher and step down.

SEV 1

Critical

TRIGGERS

Email, chat and sign-in are unavailable across the organisation, or a critical service has stopped and the fallback is not working. Safety, care or revenue is affected.

WHO IS ALERTED, HOW

Crisis team, executive on call, communications lead, security lead. Board chair informed.

`Push, SMS, voice call · escalate after 5 minutes`

SEV 2

Major

TRIGGERS

One platform service is down, such as email, while others hold. A tested fallback is working but staff are slowed. Several sites or teams affected.

WHO IS ALERTED, HOW

Incident team, service owners, site leads. Communications lead on standby.

`Push and SMS · bridge opened`

SEV 3

Moderate

TRIGGERS

Degraded performance or a single region or team affected. Supplier has acknowledged an incident. No material impact on customers yet.

WHO IS ALERTED, HOW

IT on call and the service owner.

`Push and email`

SEV 4

Watch

TRIGGERS

Supplier advisory, a change window at the provider, or a related incident elsewhere that could spread.

WHO IS ALERTED, HOW

Duty manager logs it and briefs the next shift.

`Email and incident log`

FIRST HOUR

Mail flow stops at 13:30

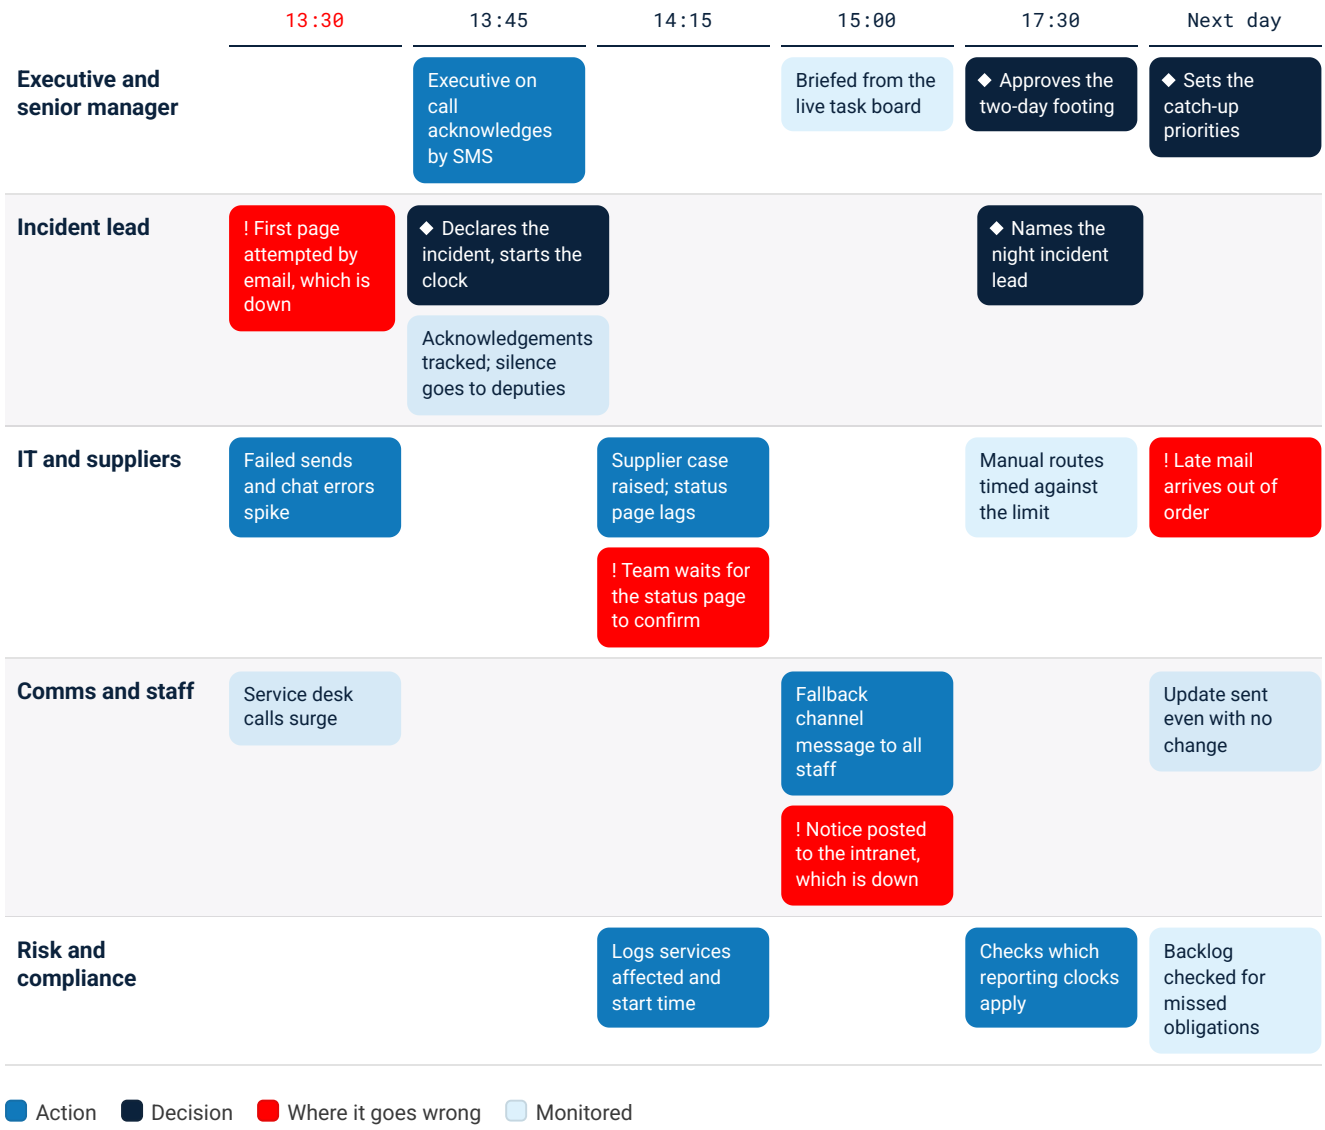
An organisation whose email, chat, files and sign-in all run on one productivity platform. Owners are roles, not names.

CLOCK	ACTION	OWNER
T+00:00	**Reports arrive faster than monitoring** The service desk takes calls about failed sends and a frozen chat client. Nobody can tell yet whether it is local or global.	SERVICE DESK LEAD
T+00:05	**Check the supplier, then stop waiting** Confirm against the provider's status page and an independent outage tracker. Do not wait for the status page to catch up before acting.	IT ON CALL
T+00:10	**Declare, and page out of band** An authorised manager declares the incident. Page the response team by SMS and voice, from a service that does not sign in through the failed platform.	INCIDENT MANAGER
T+00:20	**Open a bridge nobody needs an invitation for** Use the standing conference number printed in the plan. Confirm who is on it, who is missing, and who is covering the missing.	INCIDENT LEAD
T+00:30	**Tell staff what to use instead** One short message on the fallback channel: what is down, what to use, what not to try, and when the next update comes.	COMMUNICATIONS LEAD
T+00:45	**Protect the work that cannot wait** Identify the processes that fail quietly without email: approvals, referrals, bookings, shift confirmations. Put a manual route and an owner on each.	OPERATIONS LEAD
T+01:00	**Set the rhythm, the handover and the record** Fix the update cadence, name the next shift's lead while you still have everyone, and open the decision log.	INCIDENT MANAGER

SWIMLANE

The same afternoon, lane by lane

Who acts, who decides, where it goes wrong and how progress is watched, from the first failed send to the handover.



Illustrative. The red boxes are not recommended steps. Each is a failure seen repeatedly on 12 November: paging over the channel that has failed, waiting for a supplier to confirm what you can already see, posting to a staff channel that shares the outage, and treating late mail as if it had arrived on time. If your plan contains any of them, that is the finding.

SCENARIOS

Three severe days to rehearse

HOSPITAL OR HEALTH SYSTEM

Email and chat fail across every site

THE DISRUPTION

Appointment confirmations and internal referrals stop. Clinical systems keep running.

WHO DECIDES

The on-call director declares and authorises manual working. Each site lead owns their own fallback.

WHERE IT GOES WRONG

Staff fall back to paper with no agreed template, so the work cannot be reconciled.

MONITOR AND TEST

Reconcile paper work daily. Rehearse a full day on paper and time the reconciliation.

PROFESSIONAL SERVICES OR FINANCE

A two-day outage across time zones

THE DISRUPTION

Approvals, client correspondence and deadline-bound filings stall, and offices recover at different times.

WHO DECIDES

A single incident lead holds one view of status for all regions, with a named deputy per time zone.

WHERE IT GOES WRONG

One region declares the all-clear while another is still down, so coordination stops too early.

MONITOR AND TEST

One status board for every region. Rehearse a handover between two time zones.

LOGISTICS, RETAIL OR MANUFACTURING

Shift and site communication goes dark

THE DISRUPTION

Depots, stores and plants lose the daily brief, shift changes and exception handling.

WHO DECIDES

The operations director authorises local autonomy within set limits.

WHERE IT GOES WRONG

Frontline staff have no corporate laptop, so a fallback that assumes a desk reaches nobody.

MONITOR AND TEST

Acknowledgement by site, not by headcount. Rehearse a handover by SMS and voice.

OUTSTANDING ACTIONS

Service restored. The work is not finished.

An illustrative board the morning after recovery. Every open action has one owner, a time and a way to check it is really done.

ACTION	OWNER	DUE	STATUS	HOW IT IS CHECKED
Confirm mail flow and chat are stable for all sites	IT on call	09:00	IN PROGRESS	Send and receive tested from each site
Tell staff the fallback is stood down	Communications lead	09:30	DONE	Message sent on both the fallback and the restored channel
Work the email backlog by priority, not by date	Operations lead	12:00	IN PROGRESS	Oldest unactioned item is under four hours
Reconcile everything handled on paper	Site leads	17:00	OPEN	Paper record matched to the system of record
Identify deadlines and approvals missed in the outage	Risk and compliance	12:00	OVERDUE	Named list with an owner for each
Check security alerts raised while tooling was degraded	Security lead	14:00	IN PROGRESS	Queue reviewed back to the start of the incident
Brief the executive committee	Executive on call	16:00	OPEN	Briefing note filed with the incident record
Preserve the decision log and message record	Incident lead	11:00	IN PROGRESS	Export taken and stored outside the platform
Written explanation and fix from the supplier	Supplier lead	Day 5	BLOCKED	Post-incident report requested and chased

> Review the board at a fixed time until every row is done or handed to the post-incident review. The backlog rows are the ones most often dropped once the service looks healthy.

PART 03

Lessons, and what to change.

What 12 November showed, and the work that follows from it.

01 Ten things to take from the outage

02 Restored, and caught up

03 This week, this month, this quarter

LESSONS

What organisations should learn from this outage

Crises Control's recommendations, not findings issued by Microsoft; the facts they rest on carry their sources. Ten things worth acting on, whoever your provider is.

01 **Incident response has to continue when the platform is gone**

Any plan whose first step is to email the response team has a single point of failure in its first step. Keep one alerting route that does not sign in through your main platform, and prove it with that platform declared unavailable.

02 **Technical restoration is not operational recovery**

Your supplier closes its incident when the service is healthy. You close yours when nothing has been missed. Recovery ran at different speeds by service, organisation and region, so a multinational could hold two contradictory truths at once.[4] One lead holds one view of status, and the all-clear is stated per region and per service.

03 **Alternative channels must be tested, and must not share the dependency**

A fallback that signs in through the identity service that failed, or routes over the platform that failed, is not a fallback. Test the channel, the contact list and the authority to use it.

04 **Know what your critical services actually rest on**

It reached surgeries, schools, police forces, courts, a finance ministry and a central bank, none of which had a fault of their own.[11,14] Map each critical service to its platform, identity service, network and carriers. Your security console belongs on that map: Defender XDR was named in the same advisory.[1] Alert triage degrades exactly when you need it.

05 **Staff need instructions, not reassurance**

Say what is down, what to use instead, what not to try, and when the next update comes. Silence gets filled with personal email and messaging apps, which adds a records problem to an outage.

06 **Someone must own the incident, and someone must be allowed to decide**

Most delay in an outage is waiting for a decision. Agree decision rights in advance, name a deputy for every accountable role, and turn the plan into owned actions with deadlines.

07 **Plan the handover and the out-of-hours response**

It ran just under twenty-four hours, across a night and a shift change, and the backlog ran on from there.[5] A ninety-minute rehearsal that ends with the service coming back does not test the part most likely to fail.

08 **Reconcile the backlog before you call it closed**

Delayed mail arrived stamped with the time it was delivered, not the time it was sent, so the queue does not read in order.[4] Work it by priority, and check for approvals and deadlines missed while nothing looked wrong.

09 **Record decisions and actions as they happen**

The account a board, an auditor or a regulator asks for cannot be reconstructed from memory. Keep the log somewhere the outage cannot reach, and export it before the incident closes.

10 **Test continuity plans against the failures you will actually get**

Not a fire and not a flood. A supplier change that goes wrong, an identity provider that will not answer, or a platform that is healthy for some of your people and not others. Testing had missed this defect because the test machines had fewer processor cores than production.[5,7]

Microsoft 365 is the case study, not the only scenario. The same ten apply to:

Cloud hosting and infrastructure | Collaboration and productivity platforms | Identity and single sign-on | Business applications | Network and connectivity | Third-party providers

AFTER THE ALL-CLEAR

Restored, and caught up

The supplier's incident ends when the service is healthy. Yours ends when nothing has been missed.

	RESTORED	CAUGHT UP
What it means	The platform is healthy and the status page is green.	Every delayed item has been worked and nothing is outstanding.
Who declares it	The supplier, on its own timetable.	Your incident lead, against a checked list.
Email	Sending and receiving works again.	Late and out-of-order mail has been read and actioned by priority.
Approvals and deadlines	The approval route works.	Approvals and deadlines missed during the outage are identified and owned.
Manual work	Staff can stop using paper.	Everything done on paper is reconciled to the system of record.
Evidence	The supplier may publish a report.	Your decision log, messages and timings are exported and stored.

YOUR PLAN

This week, this month, this quarter

Owners are suggestions. Adjust them to your structure.

WHEN	ACTIONS	OWNER
This week	Check that the response team can be reached without the main platform. Print the bridge number and the supplier escalation contacts into the plan.	Head of resilience
This week	List the processes that stop quietly without email: approvals, referrals, bookings, shift confirmations. Name an owner and a manual route for each.	Operations lead
This month	Confirm whether your alerting, ticketing and status page sign in through the platform they are meant to cover. Separate any that do.	CIO or CTO
This month	Agree who may declare an incident when nobody can email, name their deputies, and write down how each is reached.	Head of resilience
This quarter	Run the tabletop in this playbook across a shift handover, with the usual channels unavailable for the whole exercise.	Head of resilience
This quarter	Agree what good looks like for the backlog: who works it, in what order, and who checks for missed obligations.	Chief risk officer, compliance

PRACTICAL TOOLKIT

Eleven tools for your next incident

The first seven are the core set that appears in Crises Control's incident guidance. The last four were built for this subject, for organisations whose incident depends on a platform they do not run. Copy them, adapt them and rehearse with them.

01	**Readiness assessment** Score eight areas from 0 to 3 and find your weakest link.	23
02	**Roles and responsibilities matrix** Who is responsible, accountable, consulted and informed for each activity.	24
03	**Incident decision log** Record every material decision, who made it and why, as it happens.	25
04	**Escalation workflow** Four levels, their triggers, who is alerted and who decides.	26
05	**Communication templates** Messages for the response team, staff, site leads, a supplier and customers.	27
06	**Cloud outage exercise: tabletop and war game** One scenario, two ways to run it: a scored 60-minute war game or a three-hour tabletop.	28
07	**Post-incident review template** Eight prompts that turn an incident into lessons and owned actions.	31
08	**Technology dependency worksheet** Map each critical service to the platform, identity and network it rests on.	32
09	**24-hour outage response planner** What to do in the first fifteen minutes, the first day, and after restoration.	33
10	**Response performance scorecard** Seven measures of how the response went, against targets you set yourself.	34
11	**Improvement register** Findings, corrective actions, owners, due dates and evidence of completion.	35

> Every tool here works on paper. In Crises Control the same roles, escalation rules, templates and logs can be set up in the platform, so they run during the incident and leave a record afterwards.

TOOLKIT 01 · READINESS ASSESSMENT

How ready is your response?

Score each area from 0 (not in place) to 3 (fully in place) for your cloud-dependent organisation, then add up the total.

AREA	THE QUESTION	SCORE 0 TO 3
01 · ACTIVATION **Alert**	Can you reach the right people on more than one channel within minutes, day or night? GOOD LOOKS LIKE Response team, executives and key suppliers reached by push, SMS and voice when email and chat are down.	0 1 2 3
02 · ACTIVATION **Acknowledge**	Do you know who has seen the alert, and does silence trigger escalation by rule? GOOD LOOKS LIKE Acknowledgement tracked against each critical service; no response escalates in minutes.	0 1 2 3
03 · COORDINATION **Escalate**	Are escalation paths set by role and time, not by whoever remembers who to call? GOOD LOOKS LIKE Escalation tied to time lost against a stated limit, not to individual memory.	0 1 2 3
04 · COORDINATION **Coordinate**	Is there one shared workspace and log for the incident, not scattered calls, chats and inboxes? GOOD LOOKS LIKE One incident log shared by IT, operations, communications and compliance.	0 1 2 3
05 · EXECUTION **Execute**	Are your procedures turned into owned tasks with deadlines that someone can track? GOOD LOOKS LIKE Fallback and recovery steps for each critical service run as owned tasks.	0 1 2 3
06 · EXECUTION **See**	Can leaders see tasks, people, locations and blockers live, without having to ask? GOOD LOOKS LIKE A live view of services, tasks and sites that the executive can read without a briefing.	0 1 2 3
07 · ASSURANCE **Prove**	Could you produce a timestamped timeline for a regulator, insurer or inquiry within a day? GOOD LOOKS LIKE A timeline that supports board, auditor and regulator questions, produced on request.	0 1 2 3
08 · ASSURANCE **Rehearse**	Do you exercise the whole response, not just the call tree, at least once a year? GOOD LOOKS LIKE Severe but plausible scenarios rehearsed end to end, including the loss of usual channels.	0 1 2 3

0 to 8
Alerting only
You can send messages. The response still depends on people and luck.

9 to 16
Coordinating in parts
Some steps work. Gaps appear at night, at scale or when systems fail.

17 to 24
Coordinated and provable
You can run the response and show exactly what happened.

TOOLKIT 02 · ROLES AND RESPONSIBILITIES MATRIX

Who decides, who acts, who is told

Most delay in an outage is waiting for someone to decide. Agree decision rights in advance and name a deputy for every accountable role.

ACTIVITY	EXECUTIVE ON CALL	INCIDENT LEAD	IT AND SUPPLIERS	COMMS	RISK AND COMPLIANCE	SITE AND SERVICE OWNERS
Declare the incident and start the clock	I	A	C	I	I	I
Page the response team out of band	I	A	R	C	·	I
Raise and chase the supplier case	·	A	R	·	I	·
Authorise manual or paper working	A	C	C	I	C	R
Approve the first message to staff	I	A	C	R	C	I
Approve the first message to customers	A	C	I	R	C	I
Move the response to a two-day footing	A	R	C	C	C	C
Name the next shift's incident lead	C	A	I	I	·	I
Declare the all-clear, by region	A	R	C	C	C	C
Own the backlog and the reconciliation	I	C	C	I	C	A
Post-incident review and evidence	I	R	C	C	A	C

R Responsible: does the work A Accountable: owns the decision C Consulted I Informed

> Illustrative. One A per row, and every A has a named deputy who can act without waiting. The site and service owners column is where a platform outage differs from a systems outage: much of the response is local.

TOOLKIT 03 · INCIDENT DECISION LOG

Write it down as you go

Every material decision, in one place, with what was known at the time in the rationale. During a platform outage this log is often the only record that survives, so keep it somewhere the outage cannot reach.

TIME	DECISION	OPTIONS CONSIDERED	DECIDED BY	RATIONALE	EVIDENCE OR REFERENCE
14:05	*Move all sites to manual shift handover*	*Wait for mail to return; use personal email*	*Incident lead, with operations lead*	*No supplier estimate; two sites already improvising; manual template exists and was tested in March*	*Bridge note 2; task 7*

Record decisions not to act as well as decisions to act. The scribe keeps the log; the decision owner confirms each entry.

TOOLKIT 04 · ESCALATION WORKFLOW

Escalate on time lost, not on mood

Severity says how bad it is. Escalation says who must decide next, and it runs on how long the organisation has been without its usual channels.

	TRIGGER	WHO IS ALERTED	WHO DECIDES
LEVEL 1 **Degraded service**	Email or chat is slow or failing for some users. A workaround is holding.	IT on call, service owner.	Incident lead, with the service owner.
ESCALATE IF IT IS NOT RESOLVED IN 30 MINUTES, OR A SECOND SERVICE FAILS			
LEVEL 2 **Channels down**	Staff cannot rely on email or chat. The organisation is working around the gap.	Incident team, communications lead, site leads.	Incident lead declares and owns the response.
ESCALATE IF THERE IS NO SUPPLIER ESTIMATE WITHIN AN HOUR			
LEVEL 3 **Running blind**	No reliable estimate for recovery, or a critical process has stopped. Manual working is in place.	Executive on call, risk and compliance, supplier executive contact.	Executive on call, on options set out by the incident lead.
ESCALATE IF RECOVERY WILL CROSS A SHIFT HANDOVER OR A WEEKEND			
LEVEL 4 **Multi-day footing**	The outage will run into a second day, or an obligation has been missed.	Executive committee, board chair, regulators through the compliance lead.	Executive committee: manual working, staffing, customer position and reporting.

Illustrative thresholds. Set yours against the limit you recorded on your dependency map, so the same rule works for a service that can be down for an hour and one that can be down for a day.

TOOLKIT 05 · COMMUNICATION TEMPLATES

Prepare messages in advance

Use these templates as starting points during an outage. Replace the placeholders with confirmed information, keep messages concise and check the final message length before sending. Longer messages may be split into multiple SMS segments, depending on the platform and encoding used. Agree sign-off rules now, and keep these where the outage cannot reach them.

SMS · RESPONSE TEAM

> Incident declared [time]. Email and chat are down. Bridge [number] PIN [pin]. Do not use email. Reply 1 to confirm, 2 if unavailable.

No reply in 5 minutes: call the deputy.

SMS · ALL STAFF

> Email and chat are unavailable. Use [fallback channel]. Do not resend; messages are queued. Next update [time].

Keep it to one screen on a phone.

VOICE · SITE AND SHIFT LEADS

> Read the agreed script: what is down, what to use instead, what to do about anything that cannot wait, and when the next brief comes.

Record who answered and who did not.

EMAIL · CRITICAL SUPPLIER

> Incident affecting [service] since [time]. Within 15 minutes please confirm scope, your reference, a named engineer, an estimate and any workaround. Copy the contract owner.

Send from a route outside the affected platform.

PUSH · CUSTOMERS OR SERVICE USERS

> Some of our services are slower than usual because of a supplier issue. [What to do instead]. We will update here by [time].

Name the supplier only once confirmed.

EMAIL · STAFF, AFTER RECOVERY

> Services are restored. Messages sent during the outage may arrive late or out of order. Check [what] before assuming it was actioned. Report anything missed to [owner] by [time].

This is the message most often forgotten.

TOOLKIT 06 · CLOUD OUTAGE EXERCISE: TABLETOP AND WAR GAME

Cloud Outage Exercise

One scenario, two ways to run it. Pick the format by what you need out of the room.

60-minute war game

DURATION	60 minutes: six eight-minute injects, then a twelve-minute debrief. Brief the room before the timed hour starts
IN THE ROOM	Five to eight, from IT, operations, communications and senior management
FACILITATOR	Anyone outside the response team, plus an observer holding the scoring sheet
WHAT YOU GET	A score out of 40 you can compare with your last rehearsal, and the gaps behind it

Three-hour tabletop

DURATION	Three hours: six injects, including an extended final scenario, a break and a thirty-minute debrief
IN THE ROOM	Executive on call, incident, IT, operations, communications, risk and compliance, and two site or service owners
FACILITATOR	Head of resilience, with an independent observer
WHAT YOU GET	Written findings and gaps for your resilience evidence, from a discussion that crosses a handover and a second day

THE SCENARIO, FOR BOTH FORMATS

Your productivity platform fails without warning. Email, chat and shared files stop across every site. The supplier's status page shows nothing yet, and there is no restoration estimate and no stated cause.

- Keep email, chat and the intranet out of the room throughout. That is the point of it.
- Nobody looks up the plan. If they cannot recall it, that is the result, and the facilitator does not help.
- Run it with the people who would really be called, not stand-ins who are easier to book.

RUNNING IT, WHICHEVER FORMAT YOU PICK

What the observer writes down

Not what people said, but what was decided, by whom, and when. Note every decision that was discussed and never made: those are the findings.

The debrief

Three questions. What would we have done differently with more time? What did we assume was available that was not? Which of these gaps already has an owner?

Afterwards

Score it or write it up while it is fresh, then put anything that needs fixing into the improvement register, Toolkit 11, with an owner and a date. An exercise with no register entry changes nothing.

THE RUNNING ORDER

60-minute war game

Before	Brief the room: the scenario, the rules, who is playing which role. This sits outside the timed hour
00:00	Injects 1 to 6, eight minutes each, back to back. Read it, let the room work, stop on the decision
00:48	Debrief and scoring, twelve minutes. Score it before anyone leaves the room

Three-hour tabletop

00:00	Brief: the scenario, the rules, roles, and what the observer will record
00:15	Injects 1 to 3, twenty minutes each. Let the discussion run; stop it on the decision
01:15	Break, ten minutes. The incident does not pause, so the room comes back to a handover
01:25	Injects 4 and 5, twenty minutes each
02:05	Inject 6 as two sessions of twelve minutes: the all-clear, then the day after
02:30	Debrief, thirty minutes. Write the findings up before anyone leaves

TOOLKIT 06 · CLOUD OUTAGE EXERCISE: TABLETOP AND WAR GAME

The injects

Six injects, in order, shared by both formats. The war game runs to the exercise clock in the first column; the tabletop runs to the scenario clock beneath it and takes the session times set out on the previous page. What a good response looks like is the “5 looks like” column on the scoring sheet overleaf: read it before you facilitate, not during.

EXERCISE AND SCENARIO TIME	THE INJECT	THE FACILITATOR ASKS	THEY MUST DECIDE
EXERCISE 00 to 08 min SCENARIO 13:30	**1. Initial disruption** The failure is reported. Nobody can tell yet whether it is local or global.	How do we know this is not just us? Who is told first, and on what channel?	Whether to declare, who declares, and on whose authority.
EXERCISE 08 to 16 min SCENARIO 14:15	**2. Communication breakdown** Staff have started groups on personal messaging apps. A manager has emailed a customer from a personal account.	Which channel is the approved one, and who says so? What do we tell people about records?	The approved staff channel, the first message, and who signs it off.
EXERCISE 16 to 24 min SCENARIO 14:45	**3. Operational impact** Approvals, documents and the intranet are unreachable. Two processes have quietly stopped and nobody has noticed.	Which processes fail silently without email? Who is checking, and how often?	Which work continues manually, who authorises it, and how it will be reconciled.
EXERCISE 24 to 32 min SCENARIO 15:30	**4. Supplier uncertainty** Still no restoration estimate. Customers are asking, and a journalist has called the switchboard.	What can we say that is true? What must we not say, and why?	The customer line, the external line, who approves them, and the update cadence.
EXERCISE 32 to 40 min SCENARIO 17:30	**5. Overnight handover** The outage will run overnight. Half the room is due to finish, and the lead has been on since the first report.	Who takes it from here, and what does the incoming lead need in their hands?	The night lead, the handover method, and the trigger for waking an executive.
EXERCISE 40 to 48 min SCENARIO Next day	**6. Partial restoration, and the backlog** Mail is back for some offices, not others. Delayed messages arrive out of order and an approval deadline passed in the night. **Tabletop:** run this as two sessions, the all-clear and then the day after.	Who gives the all-clear, and for whom? What is still outstanding?	Whether to stand down, who works the backlog and in what order, and who checks for missed obligations.

TOOLKIT 06 · CLOUD OUTAGE EXERCISE: TABLETOP AND WAR GAME

Scoring it, and what to do after

The score belongs to the war game and is only useful against your own previous score. The measures beside it belong to the tabletop, where the output is written findings rather than a number. The guidance underneath applies to both.

AREA	WHAT YOU ARE SCORING	0 LOOKS LIKE	5 LOOKS LIKE	SCORE
Incident activation	How quickly it was declared, and whether anyone was clear who could declare it.	Twenty minutes of debate about whether this counted.	Declared on evidence inside ten minutes, by a named role.	☐☐☐☐☐☐
Communication	Whether staff received one clear instruction on a channel that was actually working.	People learned from each other on personal apps.	One approved message, one channel, inside twenty minutes.	☐☐☐☐☐☐
Decision-making	Whether decisions were made, by whom, and on what information.	Everything waited for an executive nobody could reach.	Decisions taken at the right level, deputies used.	☐☐☐☐☐☐
Task ownership	Whether actions left the room with owners and deadlines.	Actions were discussed, not assigned.	Every action has an owner, a deadline and an escalation.	☐☐☐☐☐☐
Leadership visibility	Whether leaders could see status without interrupting the response.	Repeated requests for a briefing.	Leaders read status from one shared view.	☐☐☐☐☐☐
Shift handover	Whether the response survived the change of team.	The incoming team started again from scratch.	A named lead, and a handover held outside the platform.	☐☐☐☐☐☐
Recovery management	Whether the all-clear and the backlog were handled deliberately.	Stood down when the status page turned green.	All-clear by region and service, backlog owned.	☐☐☐☐☐☐
Evidence and accountability	Whether you could show afterwards what was decided and when.	Reconstructed from memory days later.	A timestamped log, exported before closing.	☐☐☐☐☐☐
			TOTAL OUT OF 40	

0 to 15

The response ran on improvisation. Work the readiness assessment on the first toolkit page before running this again.

16 to 29

The parts work and the joins do not. It is usually the handover, the backlog, or who was allowed to decide.

30 to 40

Coordinated and provable. Run it again with a harder inject: an identity failure, or a supplier that never gives an estimate.

WHAT THE TABLETOP MEASURES INSTEAD OF SCORING

01
Time to reach the full team

02
Time to the first message to all staff

03
Whether the handover held

04
Decisions logged, each with an owner

> The total is an internal exercise score for comparing one rehearsal with the next. It is not an industry benchmark, not an assessment against any standard, and not a certification of anything.

TOOLKIT 07 · POST-INCIDENT REVIEW TEMPLATE

Measured against the time you lost

Hold it soon, ideally within two weeks, while memories and logs are fresh. Blame-free, evidence-led, and honest about the backlog.

INCIDENT	DATE	SEVERITY	REVIEW LEAD

01 The timeline

When did staff first notice, and when was it declared? How long between the two, and what did the gap cost?

02 Reaching people

How long to reach the full response team out of band? Who could not be reached, and why?

03 Decisions and decision rights

Which decisions were made, by whom, on what information? Did any wait for someone with authority?

04 Working without the platform

Which processes stopped quietly? Which fallbacks worked, and which had never been tested?

05 The handover

Did the response hold across the shift change and overnight? What was lost in the handover?

06 Staff, customers and service users

What were they told, when and on what channel? Who did the message not reach?

07 The backlog

What arrived late or out of order? What was missed, and how did we find it?

08 Actions and evidence

What changes, who owns it and by when? What does this mean for our plans and our next test?

TOOLKIT 08 · TECHNOLOGY DEPENDENCY WORKSHEET

What each critical service actually rests on

One row per service, filled in with the people who run it rather than from the asset register. The column that catches organisations out is the third: a platform can be healthy while the identity service in front of it is not.

CRITICAL BUSINESS SERVICE	CLOUD PLATFORM OR TECHNOLOGY	IDENTITY, NETWORK AND TELECOMS IT RESTS ON	IMPACT IF UNAVAILABLE	ALTERNATIVE WORKING ARRANGEMENT	SERVICE OWNER
Customer payments	*Cloud payments platform*	*Identity provider sign-in; site broadband; SMS for one-time codes*	*No card payments at 40 sites; cash only; queues within the hour*	*Offline card terminals, printed fallback script, cash float released*	*Head of retail operations*

Keep a dated copy outside the platforms named in column two. A dependency map stored in the file service it describes is no map at all.

TOOLKIT 09 · 24-HOUR OUTAGE RESPONSE PLANNER

The first 24 hours, and the day after

A planning skeleton for an outage with no restoration estimate. The windows are illustrative planning intervals, not recovery targets and not a promise about when anything comes back. Set yours against your own tolerances and shift pattern.

WINDOW	PRIORITIES	DECISIONS TO MAKE	WHO OWNS IT
First 15 min	Confirm the failure is not local. Check the provider's status page and an independent tracker, then stop waiting for either.	Is this an incident, and who declares it?	Service desk lead, IT on call
First 30 min	Page the response team out of band. Open the standing bridge. Confirm who is on it and who is covering the missing.	What severity, and who is the incident lead?	Incident manager
First hour	One message to all staff: what is down, what to use, what not to try, when the next update comes. Open the decision log.	What is the approved staff channel, and who signs messages off?	Communications lead
Hours 2 to 4	Identify the processes that stop quietly without email. Put a manual route and a named owner on each. Raise the supplier case and chase it.	What continues manually, and who authorises it?	Operations lead, supplier lead
Hours 4 to 8	Brief customers and service users on a channel you still control. Check which reporting clocks apply to you.	Do we move to a multi-day footing, and who decides?	Executive on call, risk and compliance
Hours 8 to 12	Name the next shift's lead while everyone is still available. Write the handover somewhere the outage cannot reach.	Who leads overnight, and what wakes an executive?	Incident manager
Hours 12 to 24	Hold the update cadence even when there is nothing new. Start listing what has been deferred, missed or done on paper.	What have we already missed, and who is told?	Incident lead, risk and compliance
After restoration	Give the all-clear explicitly, by region and by service. Work the backlog by priority. Reconcile manual work. Export the record before closing.	Are we restored, or are we caught up?	Incident lead, service owners

> The windows are illustrative planning intervals for your own rehearsal, not recovery targets and not a prediction of when any supplier will restore service. On 12 November 2026 the provider restored in stages over roughly a day, deliberately, to keep the restoration stable.

TOOLKIT 10 · RESPONSE PERFORMANCE SCORECARD

What the response actually cost you

Seven measures you can take straight from your own incident record. There are no targets printed here on purpose: set them yourself, against your tolerances and your own previous incidents.

INDICATOR	WHAT IT TELLS YOU	YOUR TARGET	THIS INCIDENT	GAP AND ACTION
Time to declare	The gap between the first credible report and a declared incident. It is usually the largest delay you can actually remove.			
Time to reach the response team	From declaration to the last named responder reached on a channel that was working.			
Acknowledgement rate	The share of those alerted who confirmed, and by when. A delivery receipt is not an acknowledgement.			
Time to assign critical actions	From declaration to every critical action having one named owner and a deadline.			
Overdue actions	How many actions passed their deadline, and how long the longest one ran before anyone noticed.			
Time to complete the handover	From the start of the shift change to the incoming lead holding current status without asking.			
Time to reconcile the backlog	From restoration to the point where delayed work has been worked and missed obligations are listed with owners.			

No benchmarks are printed here. Published averages come from different sectors, different severities and different definitions, and a target you did not set yourself is not a target. Use your own tolerances and your own last incident.

TOOLKIT 11 · IMPROVEMENT REGISTER

Findings, and who closed them

The review on the previous page is worth nothing without this one. A block per finding, carried into your risk and audit process until the evidence column is filled in.

WORKED EXAMPLE

FINDING
Response team paged by email

WHAT HAPPENED
First page went out on the failed platform; 40 minutes lost before anyone noticed

OPERATIONAL IMPACT
Declaration to full team: 52 minutes

ROOT CAUSE OR CONTRIBUTING FACTOR, WHERE ESTABLISHED
Call tree step one was a distribution list on the affected platform

CORRECTIVE ACTION
Move first page to the out-of-band route; remove the distribution list from step one

OWNER
Head of resilience

DUE DATE
31 Jan

STATUS
In progress

EVIDENCE OF COMPLETION
Test page sent with platform declared unavailable; timings recorded

FINDING 01

FINDING

WHAT HAPPENED

OPERATIONAL IMPACT

ROOT CAUSE OR CONTRIBUTING FACTOR, WHERE ESTABLISHED

CORRECTIVE ACTION

OWNER

DUE DATE

STATUS

EVIDENCE OF COMPLETION

FINDING 02

FINDING

WHAT HAPPENED

OPERATIONAL IMPACT

ROOT CAUSE OR CONTRIBUTING FACTOR, WHERE ESTABLISHED

CORRECTIVE ACTION

OWNER

DUE DATE

STATUS

EVIDENCE OF COMPLETION

> Keep this with the post-incident review, not in a separate file. A finding without an owner, a date and evidence is a note, and notes do not survive the next reorganisation.

CHECKLIST

Could you run an incident without your platform?

Tick what is true today. Anything blank is your next action.

- [ ] **An alerting route independent of your main platform**
  Hosted outside your estate, with its own sign-in and a current contact list.
- [ ] **Contact details that survive the outage**
  Mobile numbers for the response team, held somewhere that is not the mail client.
- [ ] **A standing bridge number and PIN**
  Fixed, printed in the plan, and usable without a calendar invitation.
- [ ] **Offline copies of the plan and the call tree**
  Dated, owned and refreshed on a schedule.
- [ ] **Declaration authority agreed for a comms failure**
  Named people and deputies, with how each is reached written down.
- [ ] **A staff channel that does not share the dependency**
  Agreed in advance and tested with a real message.
- [ ] **A backlog procedure**
  Who works it, in what order, and who checks for missed approvals and deadlines.
- [ ] **A two-day rehearsal in the last twelve months**
  Including a shift handover, with the usual channels unavailable.
- [ ] **Break-glass access to the consoles you need**
  Emergency accounts for an identity failure. Untested break-glass is the same as none.

A platform outage and an identity failure are not the same incident: one leaves telephony and sign-in working, the other locks you out of a platform that is still running. Test each route the way it would fail. Printed call trees, standing PINs and break-glass accounts hold personal data or open a way in. Store and review them accordingly.

GLOSSARY

Words worth knowing

Productivity platform

The bundled service that carries email, chat, meetings, files and often sign-in for an organisation.

Out-of-band

A route that does not rely on the systems affected by the incident. What qualifies depends on the failure: a platform outage, an identity failure, a cyber incident and a telecoms outage each rule out different routes.

Single sign-on

One identity service used to sign in to many systems. Convenient, and a single point of failure.

Break-glass account

An emergency account that works when the normal sign-in route does not. Only useful if tested.

Code grey

A hospital alert used in parts of Canada for a loss or failure of infrastructure, including IT and communications.

Backlog

Work that queued during an outage and arrives afterwards, often late and out of order.

Critical third party

UK term: a provider designated by HM Treasury whose disruption could threaten the stability of, or confidence in, the UK financial system. Four cloud and technology providers were designated from 13 July 2026, and five further suppliers covering financial services, market information, and technology and payment services from 23 November 2026, taking the total to nine from that date.

Severe but plausible scenario

A demanding but realistic disruption used to test whether an organisation can keep going.

Shift handover

The point at which responsibility passes to a new team. In a long outage it is the most common place for information to be lost.

All-clear

The explicit statement that normal working has resumed. Best given per region, and separately from the supplier's status page.

WHY COORDINATION, NOT JUST NOTIFICATION

Most organisations can send alerts. Few can manage incidents end-to-end.

When email, chat and files fail together, sending the alert is the easy part. Getting through the next two days depends on who acknowledged, which fallbacks started, what customers and regulators were told, and whether you can prove it all afterwards.

WHERE INCIDENTS GO WRONG

01	**Fragmented information**	Different teams work from different information.
02	**Delayed decision making**	Leadership lacks visibility, so decisions wait.
03	**Uncoordinated response**	Several teams work independently on the same problem.
04	**Lack of visibility**	Critical actions become difficult to track.
05	**Lack of accountability**	Actions and decisions are poorly documented.

THE WHOLE INCIDENT, IN ONE PLATFORM



Detection is the one stage Crises Control does not perform. Incidents are raised by your people, through SOS, the app or the web, or by a monitoring system where an integration has been configured. The platform coordinates everything from the alert onwards.

> Crises Control is an **Operational Incident Coordination Platform**. It helps organisations maintain operational control during incidents by combining alerting, coordination, execution, visibility and accountability in one platform.

HOW THE MODULES WORK TOGETHER

An Operational Incident Coordination Platform when the platform is what failed

Here is the outage in this playbook, one module at a time. Crises Control runs on Microsoft Azure, which is a different service from Microsoft 365; it supports your response plan rather than replacing it, and it does not detect an outage on its own.

01 START **Ping**

The response team, executives and site leads are alerted by push, SMS and voice, delivered through carriers that do not depend on your productivity platform; no response escalates to deputies under the rules you configure.

02 COORDINATE **Incident Manager**

An authorised user declares the incident and activates the platform outage response plan, with roles and a single communication log in one workspace.

03 EXECUTE **Task Manager**

Fallback channels, manual working, the supplier chase and the staff message become owned tasks with deadlines, tracked to completion.

04 OPERATE **Control Centre**

Leaders see who has acknowledged, which sites are affected and what is outstanding in one live view, without asking for a briefing.

05 PROVE **Audit & Reporting**

Messages, acknowledgements and task activity are timestamped automatically, where supported by the platform. Decisions are recorded by the incident team as they are made. The resulting record can be reviewed and exported.

06 ASSIST **CRAiG**

CRAiG (beta) retrieves live incident information, response plans, tasks and reports in plain language, within each user's own permissions. It does not decide or act.

COORDINATION
Everyone works from the same incident, roles and log.

VISIBILITY
Leaders see progress without chasing for updates.

EXECUTION
Every action has an owner, a deadline and an escalation.

ACCOUNTABILITY
One named owner for every decision and task.

EVIDENCE
A reviewable record, built as the response happens.

WHAT IS STANDARD, WHAT IS CONFIGURED

Clear about what the platform does.

Some capabilities work as standard. Others depend on how your team sets up plans, groups and rules, or on connecting your own systems. Your demonstration and implementation plan confirm what applies to you.

STANDARD

Included in how each module works.

- SOS press-and-hold alert from the mobile app, raising an incident and sharing location where available
- Ping alerts by push, SMS, voice, email and web alerts, with one-tap acknowledgement
- Incident Manager workspace, roles and a single communication log
- Task Manager tasks with owners and deadlines when a plan is activated
- Control Centre live view of status, acknowledgements, tasks and locations
- Audit & Reporting: messages, acknowledgements and task activity timestamped automatically, reviewable and exportable as PDF or CSV
- CRAiG (beta): answers questions and summarises your incident data within each user's permissions. Available in the mobile app and the web portal

CONFIGURED

Set up by your team to match your plans.

- Response plans, task lists and message templates per incident type
- Groups, sites and audiences to target
- Escalation rules and timings when people do not respond
- Cascading from one channel to the next
- Response options people can choose when they acknowledge

INTEGRATIONS

Connected to your other systems, where configured.

- Directory sync with Microsoft Entra ID, SCIM or scheduled file import
- Single sign-on with Entra ID, Google Workspace or Okta
- The Microsoft Teams app, enabled by your administrator
- Triggers from monitoring or other systems, set up as an integration project

Crises Control does not detect incidents on its own; the note on the lifecycle page explains how they are raised. On dependencies: the platform is hosted on Microsoft Azure, which is a separate service from Microsoft 365; SMS and voice are delivered through independent carriers; push reaches the Crises Control mobile app. Microsoft Teams and email are available as alerting channels, so neither should be relied on as the fallback during a Microsoft 365 outage. Where you have configured single sign-on through an identity provider, that sign-in route depends on it, so keep a route that does not. No platform can be promised through every possible hosting, authentication, connectivity or telecommunications failure, and nothing here replaces your continuity strategy, your technical recovery processes or the people who lead your response.

USE CASES

How organisations use Crises Control in cloud and platform outages

Cybersecurity and IT incidents

USE CASE

TRIGGER Cyber breach, cloud outage, network disruption, unavailable applications or email outage

Send targeted alerts, provide staff instructions, launch a structured incident where required, assign response actions and share service updates.

Ping | Incident Manager | Task Manager | Control Centre | Audit & Reporting

Business continuity and disaster recovery

USE CASE

TRIGGER Planned disaster recovery test or real-world business disruption

Activate response plans, notify teams across locations, distribute instructions, coordinate recovery actions and retain exercise or incident records.

Ping | Incident Manager | Task Manager | Control Centre | Audit & Reporting

Disaster recovery and system downtime

USE CASE

TRIGGER A hospital information system or other critical technology service is disrupted

Communicate recovery or downtime procedures, distribute tasks across departments and coordinate the response. Automated plan triggering depends on configured integrations and workflows.

Ping | Incident Manager | Task Manager | integrations where configured | Audit & Reporting

IT and operational disruption notifications

USE CASE

TRIGGER An IT service or business application becomes unavailable

Notify relevant IT teams and employees, provide updates and coordinate any required follow-up actions.

Ping | Incident Manager | Task Manager | Audit & Reporting

Operational communications resilience

USE CASE

TRIGGER Normal communication channels are unavailable, degraded or insufficient

Use independent communication channels to reach internal teams and relevant external contacts, track delivery and acknowledgements, and retain an auditable record.

Ping | Incident Manager | Audit & Reporting | integrations where configured

Out-of-hours incident response

USE CASE

TRIGGER An overnight or weekend emergency requires on-call managers or contractors

Alert the on-call response group, track acknowledgement, escalate non-response and coordinate actions until the incident is resolved.

Ping | Incident Manager | Task Manager | Control Centre

Multi-location and international incident coordination

USE CASE

TRIGGER An incident requires response from teams in different countries or operating locations

Mobilise geographically dispersed teams, keep relevant stakeholders informed and maintain a shared record of incident activity.

Ping | Incident Manager | Task Manager | Control Centre | Audit & Reporting

Testing, exercises and audit evidence

USE CASE

TRIGGER Notification tests, evacuation drills, disaster recovery exercises or scenario simulations

Run planned tests, track delivery and acknowledgements, record participation and retain evidence for internal review or audit.

Ping | Incident Manager | Audit & Reporting

These show how the platform can be applied to different operational scenarios. Capabilities marked 'where configured' or 'where available' depend on each organisation's set-up.

BUYER'S GUIDE

Eight questions to ask any incident platform supplier

Many tools can send an alert. These questions separate a notification tool from a platform that coordinates the whole response.

01 **Does it stop at the message?**
Ask what the platform does after the alert is sent: tasks, log, escalation, recovery.

02 **Can it prove who acknowledged?**
Delivery is not acknowledgement. Ask how silence is spotted and escalated, and what you must configure.

03 **Does it run your procedures?**
Ask whether your plans become owned tasks with deadlines, or stay as documents.

04 **Is there one shared picture?**
Ask what leaders see live, and whether it needs someone to update a spreadsheet.

05 **Does it work when your systems do not?**
Ask how teams coordinate if email, chat or your network is down.

06 **Can frontline staff raise the alarm?**
Ask how a colleague on site escalates, with their location where available, in one action.

07 **What evidence does it leave behind?**
Ask for a sample timeline of the kind a regulator, insurer or inquiry would want.

08 **Can you rehearse in it?**
Ask whether you can run exercises on the same platform you would use for real.

SOURCES

Sources and notes

Sources accessed 16 November 2026; every link last checked 18 November 2026. The technical account and the incident times here follow Microsoft's Preliminary Post Incident Report of 15 November 2026, made available to tenant administrators in the Microsoft 365 admin center under EX1550543/MO1550821. That report states it is preliminary and will be superseded by a final one. Microsoft's quoted words are cited to the outlets that carried them. Regulatory summaries are for orientation, not legal advice. Last reviewed: November 2026.

01 Lawrence Abrams, Massive Microsoft 365 outage disrupts Outlook, Teams, Copilot worldwide, BleepingComputer, 12 November 2026, www.bleepingcomputer.com/news/microsoft/massive-microsoft-365-outage-disrupts-outlook-teams-copilot-worldwide/

02 SecurityWeek, Microsoft 365 outage disrupts Outlook, Exchange Online and Teams, 12 November 2026, www.securityweek.com/microsoft-365-outage-disrupts-outlook-exchange-online-and-teams/

03 Microsoft issues apology following global service outage, SABC News, 14 November 2026, sabcnews.com/sabcnews/microsoft-issues-apology-following-global-service-outage/

04 Tom Gerken, Laura Cress and Cachella Smith, Why is Microsoft's outage lasting so long and will emails be lost?, BBC News, 14 November 2026, www.bbc.com/news/articles/cwvg238zrnm5o

05 Microsoft, Preliminary Post Incident Report for Microsoft 365, incident EX1550543/MO1550821, report dated 15 November 2026, Microsoft 365 admin center. Start 16:20 UTC 12 November, end 16:10 UTC 13 November; a final report supersedes it

06 Laura Cress, Microsoft outage 'stark reminder' of how much we rely on tech, BBC News, 14 November 2026, www.bbc.com/news/articles/cxzd9171m51qo

07 Cyber Press, Microsoft reveals Windows bug that crashed over 80% of servers during Microsoft 365 outage, reporting the preliminary Post Incident Report, 16 November 2026, cyberpress.org/microsoft-reveals-windows-bug/

08 Microsoft 365 Status, update confirming restoration of the remaining customers and regions, 14 November 2026, x.com/MSFT365Status

09 Microsoft confirms all services are restored after global outage, Sky News, 14 November 2026, news.sky.com/story/microsoft-confirms-all-services-are-restored-after-global-outage-13612219

10 Washington Post, Microsoft says services fully restored after hourslong global outage, 14 November 2026, www.washingtonpost.com/business/2026/11/14/microsoft-outage-outlook-email-365-recover/8e33ace4-e00b-11f1-912c-4fbb80d73f7f_story.html

11 The Guardian, Microsoft outage hits hospitals, GPs, schools and police across UK, 13 November 2026, www.theguardian.com/technology/2026/nov/13/microsoft-outage-hits-nhs-hospitals-gps-schools-and-police-across-uk

12 Ashleigh Furlong, NHS Staff Turn to Paper Notes After Sweeping Microsoft Outage, Bloomberg, 13 November 2026, www.bloomberg.com/news/articles/2026-11-13/nhs-forced-to-reschedule-procedures-after-microsoft-outage

13 King's College Hospital NHS Foundation Trust, Microsoft outage, 13 November 2026, kch.nhs.uk/news/microsoft-outage/

14 Microsoft outage hits Outlook, Teams and 365, BusinessTech via Bloomberg, 13 November 2026, businesstech.co.za/news/technology/884139/microsoft-outage-hits-outlook-teams-and-365/

15 CP24, Microsoft outage: Toronto hospital network declares Code Grey, 12 November 2026, www.cp24.com/local/toronto/2026/11/12/code-grey-declared-at-toronto-hospital-network-amid-microsoft-outage/

16 Yahoo News Canada, Governments, courts, trains and hospitals hit by Microsoft outage in North America, 12 November 2026, ca.news.yahoo.com/governments-courts-trains-hospitals-hit-144549869.html

SOURCES

Sources and notes, continued

Regulatory and standards references. Article-level addresses are given throughout wherever one could be verified; where a publisher blocks automated checking, the entry gives the publication, title, date and domain instead. Two sites here refuse automated checks and were opened in a browser: the CBUAE Rulebook and the FFIEC IT Handbook. Regulatory summaries are for orientation, not legal advice.

17 UK FCA, PS21/3: Building operational resilience, March 2021, www.fca.org.uk/publication/policy/ps21-3-operational-resilience.pdf

18 UK PRA, SS1/21 Operational resilience: Impact tolerances for important business services, as updated, www.bankofengland.co.uk/prudential-regulation/publication/2021/march/operational-resilience-impact-tolerances-for-important-business-services-ss

19 The Critical Third Parties (Designation) Regulations 2026, SI 2026/777, made 8 July 2026, in force 13 July 2026, legislation.gov.uk/uksi/2026/777/made

20 HM Treasury, UK financial system strengthened with new safeguards for major technology providers, 10 July 2026, gov.uk/government/news/uk-financial-system-strengthened-with-new-safeguards-for-major-technology-providers

21 The Critical Third Parties (Designation) (Amendment) Regulations 2026, SI 2026/1203, made 16 November 2026, in force 23 November 2026, legislation.gov.uk/uksi/2026/1203/made

22 UK FCA, PS26/2: Operational Incident and Third Party Reporting, March 2026. In force 18 March 2027; notification thresholds at SUP 15.18.6R, www.fca.org.uk/publication/policy/ps26-2.pdf

23 UK FCA, FG26/3: Operational Incident Reporting, March 2026, www.fca.org.uk/publication/finalised-guidance/fg26-3.pdf

24 FFIEC, IT Examination Handbook, Business Continuity Management booklet, November 2019. Guidance to assist examiners, not a rule, ithandbook.ffiec.gov/it-booklets/business-continuity-management/

25 New York DFS, 23 NYCRR Part 500, sections 500.1, 500.16 and 500.17, as amended 1 November 2023, www.dfs.ny.gov/cybersecurity/23-NYCRR-Part-500

26 Regulation (EU) 2022/2554 (DORA), Articles 11, 14 and 64, OJ L 333, 27 December 2022, eur-lex.europa.eu/eli/reg/2022/2554/oj/eng

27 Central Bank of the UAE, Operational Risk Management Regulation, Circular 1/2026, Articles 3 and 11, effective 14 September 2026, rulebook.centralbank.ae/en/rulebook/operational-risk-management-regulation

28 SAMA, Business Continuity Management Framework, Circular 381000058504, 27 February 2017, rulebook.sama.gov.sa/en/business-continuity-management-framework

BOOK YOUR PLATFORM OUTAGE DEMONSTRATION

Run the incident when email and chat are gone.

Most organisations can send alerts. Few can manage incidents end-to-end. In a live, 30-minute demonstration built around a cloud and collaboration platform outage, you will see how our incident management software, an Operational Incident Coordination Platform, takes one incident from first alert to final report. Bring your own scenario.

WHAT YOU WILL SEE

01	REPORT	Start where you would start: a failure reported by your service desk, or raised from the app.
02	PING	Watch the team reached on push, SMS and voice, and see the acknowledgements come back in real time.
03	INCIDENT MANAGER	Declare, activate a response plan, and see roles and the communication log open in one workspace.
04	TASK MANAGER	Turn that plan into owned actions with deadlines, then see what happens when one of them stalls.
05	CONTROL CENTRE	Take the executive's seat: who responded, what is outstanding, which sites are affected.
06	AUDIT & REPORTING	Export the record at the end of the session and read what it captured, and what it did not.
07	CRAIG	Ask CRAiG (beta) a question about the incident you have just run.

Book your free demo ▶

Explore Crises Control for IT incident response

+44 (0) 20 8584 1385 · contactus@crises-control.com

SCAN TO BOOK

© 2026 TRANSPUTEC LTD. ALL RIGHTS RESERVED · LAST REVIEWED: NOVEMBER 2026