BEYOND blue

OPERATIONAL RESILIENCE POLICY

Scenario Testing Lessons Learnt

It has been nearly two years since the UK financial regulators introduced a comprehensive package of operational resilience regulations in March 2021. Amongst the requirements placed on financial institutions was the need to set impact tolerances against its important business services (IBS), and then test the ability of the institution to remain within those tolerances through a process of scenario testing using severe but plausible scenarios.

Having been at the heart of scenario testing for some time now, I thought it might be useful to share some insights, some challenges, and some ideas for the future.

SO WHAT DO THE REGULATIONS *ACTUALLY* SAY?

The FCA regulations[1] require that...

SYSC 15A.5.3

A *firm* must carry out scenario testing, to assess its ability to remain within its *impact tolerance* for each of its *important business services* in the event of a severe but plausible disruption of its operations.

SYSC 15A.5.4

In carrying out the scenario testing, a *firm* must identify an appropriate range of adverse circumstances of varying nature, severity and duration relevant to its business and risk profile and consider the risks to the delivery of the *firm*'s *important business services* in those circumstances.

The first challenge is that the methodology for undertaking such testing was not specified, nor was the actual definition of the terms "severe" or "plausible" when applied to selecting scenarios.

So just how severe should these scenarios be, after all it is always possible to construct a scenario which breaks any organisation.

[1] FCA Handbook

The PRA supervisory statement[2] gives us a few further clues over regulatory expectations

6.1...

> "Impact tolerances assume a disruption has occurred, and so testing the ability to remain within impact tolerances should not focus on preventing incidents from occurring. The PRA expects firms to focus on recovery and response arrangements."

6.2...

> "Firms should identify the severe but plausible scenarios they use for testing. When setting scenarios, firms could consider previous incidents or near misses within the organisation, across the financial sector, and in other sectors and jurisdictions."

Elsewhere in the PRA supervisory statement firms are instructed to accept that one or more preventative controls has failed, leading to a disruptive event, and the focus to be on testing recovery and response arrangements. This is not to devalue the work which needs to happen to understand (and assess) preventative controls, but rather to reinforce the attention paid to testing recovery and response based on the assumption that a disruptive event has occurred. Put another way, we pay less attention to likelihood of occurrence, and more to the impact if that scenario does occur.

The second part gives us a strong indication that plausibility can, in part, be linked to the occurrence of similar incidents or near misses within the financial sector or beyond; and that we should be open to such events occurring within our organisation irrespective of our protective controls.

SYSC 15A.5.6

SO WHICH SCENARIOS SHOULD I CONSIDER?

The FCA regulations also give us guidance on the range of scenarios which we are expected to consider, namely...

In carrying out the scenario testing, a firm should, among other things, consider the following scenarios:

Corruption, deletion or manipulation of data critical to the delivery of its important business services;

Unavailability of facilities or key people;

Unavailability of third party services, which are critical to the delivery of its important business services;

Disruption to other market participants, where applicable

Loss or reduced provision of technology underpinning the delivery of important business services.

[2] PRA - SS1/21 Operational resilience: Impact tolerances for important business services

This hints at a standard set of scenarios, and indeed the Operational Resilience Collaboration Group (ORCG) compiled such a set which included some 24 impact scenarios. Each impact scenario effectively describes a form of disruption which may have multiple possible triggers. For example, the destruction of a building might happen because of a fire, flood, subsidence or explosive event – but the consequence is similar – namely the physical destruction of the building and disruption of all business operations which require access to assets within the building.

Building on this, we can suggest a taxonomy which covers most, if not all, scenarios which might impact operational resilience.

EVENT		IMPACT	TRIGGER – EXAMPLES
Property	Denial	Access to a building is denied due to events at that building or in vicinity	Evacuation, Government action, Snowstorm, Storms, Transport disruption
	Destruction	A building is physically damaged or destroyed	Flood, Fire, Subsidence, Earthquake, Radiological hazard, Explosive Event
Technology	Infrastructure Destruction	Irreversible damage to, or destruction of infrastructure	Fire, Flood, Power Surge, Theft, Sabotage, EMP
	Infrastructure Failure	Failure of an infrastructure component or system	Hardware Fault, Network Outage, Power Outage, Operator Error
	Storage Corruption	Large scale corruption of data across mass storage	Disk Array Failure, Operator Error, Malicious Action
	Application Failure	Failure of an application	Memory Fault, Disk Error, Logic Error, Race Condition, Failed Upgrade, Operator Error, User Error, Malicious Action
	Application Corruption	Corruption of application code or configuration data	
Information	Denial	Destruction of data whether accidental or deliberate	
	Corruption	Corruption or manipulation of data	
	Breach	Theft or accidental mishandling of data	User Error, Malicious Action
People	Temporary Absence	Absence of key personnel for a period of time	Sickness, Pandemic, Strike, Transport Disruption
	Permanent Absence	Absence of key personnel on an ongoing basis	Fatality, Major Injury, Departure
Third Party	Critical Infrastructure failure	Disruption or failure of critical national infrastructure	Any of above trigger events
	Critical third party failure	Disruption or failure of critical third party	Any of above trigger events
	Critical third party compromise	Compromise of the systems of a critical third party handling bank information or providing bank services	Malicious action

ARE THERE THINGS WHICH I *SHOULDN'T* CONSIDER?

How severe should my scenarios actually be? There seem to be some examples of scenarios which would break, not just the financial institution, but perhaps the economy as a whole. Are these ruled out of scope or should they still be considered? The regulations don't give us an easy answer, but they do give us some indications that certain scenarios may indeed be so severe that they may drive us out of tolerance and that regulators may both expect and potentially accept that fact. They also suggest that we should focus on a single causal event, but do encourage us to be more demanding in our scenario testing as our services become more resilient over time – a process of continual challenge and maturity.

Typical issues which may be judged "too severe" and perhaps "implausible" might include:

National critical infrastructure disruption with wide area effect – for example a major failure of the national grid, of our core telecommunications infrastructure, or large scale disruption to transport infrastructure

A major coronal matter ejection/solar storm leading to wide area disruption of electrical grid, satellite and terrestrial communications

An outbreak of major hostilities including attack on the territory of the UK itself

But I suspect regulators will argue that unfortunately a major cyber attack is both severe and plausible, even if that attack might only succeed because multiple cyber security controls have failed. There seem all too many successful ransomware attacks today to exclude that possibility.

AND HOW SHOULD I DO THIS?

So I have a potential set of scenarios, with plausibility based primarily on the occurrence of similar events inside the organisation (or near misses), within the sector or beyond. So how can I best test these scenarios?

It is at this point, that we need to have a means of exploring the consequence of the scenario, perhaps varying by the duration of the disruptive event. For example, if a particular business critical application has failed, what is the downstream impact of that event on an IBS.

Our mapping work under the Operational Resilience regulations can give us a view of the single points of failure within the organisation, namely those assets which if unavailable or damaged might severely impact IBSs – and perhaps cause a period of disruption or loss of service integrity which might exceed one or more impact tolerances. Of course, our mapping may also identify points of failure which impact multiple IBSs, for example disruption to bank IT infrastructure.

The next step is to generate scenarios in which those points of failure are disrupted, and model the downstream cascade of consequences which flows from that. Our chaos engineer has pulled the plug on a key system, disconnected a major network cable, or closed a key building. Perhaps we will be brave enough to do exactly that in future, but for the moment our scenario testing remains just an exercise.

Those downstream consequences begin to build as we extend the period of the outage, from hours, to days, perhaps even longer. We need to understand how those consequences grow and potentially multiple over different time frames (which also helps inform just where we breach the impact tolerance), by suggesting your systems are down for 6 hours, 24 hours, 72 hours, maybe longer. In addition to technical recovery time, we also need to account for the additional time required to process and clear backlogs, and we can truly say we are back in business.

ARE WE ABLE TO RESPOND TO A SCENARIO?

The scenario impacts identified in our taxonomy can also be mapped to the response and recovery measures which we might adopt which aim to mitigate the impact of that scenario. We can create a mapping with the typical response measures which most financial institutions have in place, for example...

EVENT		RESPONSE MEASURE
Property	Denial	Business continuity plan, alternate premises and disaster recovery, working from home plans
	Destruction	
Technology	Infrastructure Destruction	Redundant communications and IT systems, disaster recovery, emergency sourcing
	Infrastructure Failure	Redundant communications and IT systems, disaster recovery
	Storage Corruption	Redundant storage, backup/recovery processes, disaster recovery
	Application Failure	Redundant IT architectures, disaster recovery
	Application Corruption	Backup/recovery processes, disaster recovery
Information	Denial	Redundant storage/IT architectures, disaster recovery
	Corruption	Backup/recovery processes, data integrity checks
	Breach	Access, credential and crypto key management
People	Temporary Absence	Key person plans, cross-skilling and augmentation arrangements
	Permanent Absence	
Third Party	Critical Infrastructure failure	Redundant infrastructure provision, emergency power and water provisioning
	Critical third party failure	Alternate supplier arrangements, exit strategies, step in rights
	Critical third party compromise	Alternate supplier arrangements, exit strategies etc

This identifies that hundreds of scenario variations, in fact can boil down to a more manageable number of response and recovery measures that need to be tested. Next comes the question of how robust our response and recovery measures are, just how long would it take to recover services – to invoke the business continuity plan, to activate the disaster recovery site, to restore the application. Each of the events has its own set of recovery actions – from the near instant reconfiguration of active-active applications replicated across data centres, to the more pedestrian restoration of many servers post a ransomware attack, or the move to alternate premises as a business continuity plan is activated.

We need to estimate and validate these timescales as best we can, and then ask ourselves are we within impact tolerance? There may be more than one answer to that when the typically shorter FCA tolerances around customer harm are compared with the longer PRA tolerances around financial and market stability. Often we will find ourselves out of tolerance in the most complex scenarios (data corruption, cyber attack, third party failure) – sometimes due to additional time required to analyse and understand the nature of the incident before recovery can begin, – but we may still be able to explore how to reduce those recovery times even if not to within a demanding 24 hour tolerance.

WE ARE BEYOND TOLERANCE, *OR ARE WE?*

So we are beyond impact tolerances, what can we do next? To mitigate customer harm, explore customer treatment strategies such as emergency access to cash and other ways of identifying and supporting vulnerable customers and how effective these strategies might be.

We now have a view of the invocation times for response and recovery strategies, and the time taken to activate customer treatment strategies. All of this gives us comfort (or perhaps not) in our ability to deal with the events we have just assumed.

This also hints that the answer to the question of whether a scenario drives us out of tolerance is more nuanced and more complex than it may first seem. For example...

- At what point does the impact tolerance clock start ticking? Cyber scenarios drive extended forensics timescales that in some cases may alone drive us beyond impact tolerance. Can we assume that we immediately rebuild systems and complete forensics in parallel?
- At what point does the impact tolerance clock stop ticking? Do we exclude time taken to test and establish confidence in the security of our rebuilt systems post a cyber incident?
- How do we handle cases in which a partial restoration of service can be achieved, perhaps focussed on vulnerable customers to reduce harm?
- What if the scenario results in intermittent service rather than a complete disruption, and how do we model that, and does that drive us out of tolerance?
- How do we deal with confidentiality/data breach scenarios if they don't disrupt service to customers but can cause harm through exposure of sensitive personal or market data?

In short, this is not yet an exact science, and all of us are learning along the way including, I suspect, the regulators.

BUT WHAT ABOUT THE TRIGGERS...

The Operational Resilience mindset, is very different to the Operational Risk mindset, which organisations have been embedding for years. Many people find it hard to accept that the event might have actually occurred and that the protective (and detective) controls they have put in place have actually failed. They may also find frustrating that the scenario tests focus on response and recovery measures, and in doing so, offer little insight into where protective (or detective) controls might need to be improved to mitigate the causes of that scenario.

The question is, can we do anything to help inform these debates, without detracting from the focus of scenario testing on response and recovery? Perhaps.

Once we have understood which “Impacts” will drive us out of tolerance, we can start to ask ourselves which causal events might create that impact, and then focus on what we could do to reduce the likelihood (and impact) of that event occurring. For example, if we know that a certain building is critical to an IBS because of the IT hosted in that building, and we know that our business continuity plans may be inadequate, then we can start to explore the “vulnerability” further.

The Operational Resilience focus is on considering actions to accelerate response and recovery - for example a disaster recovery capability, or a replication of key systems across multiple sites, or perhaps a contingency measure or workaround we can invoke quickly. But we can also focus the attention of the Operational Risk community on the actions they might take to improve protective and detective controls, for example improved fire precautions or even a relocation to a site which is beyond the flood plane.

WHAT DOES THE FUTURE HOLD?

Scenario testing will continue, and of course we will also over time look to improve the fidelity of the environments we undertake those tests within. Moving from paper based assessments, through simulations and exercises, to the ultimate goal of live testing where we feel confident enough to “pull the plug”. We will continue to debate how much is enough, and quite where the boundary of severe but plausible may lie. We will argue over which assumptions we should make over various recovery time components, and find ourselves often unable to robustly quantify some of those elements. We will have clashes of culture between classic risk based assessments using likelihood as a factor, and the operational resilience culture of assuming an event occurs and exploring the consequences.

But in doing so we hope to have rebalanced the scales. The focus has shifted from preventing events from happening, to assessing the response and recovery measures we need to have in place to help deal with more severe events. Ultimately, that is exactly what the regulators sought to achieve, even if we will end up accepting that many scenarios will remain beyond acceptable impact tolerances for some time yet.

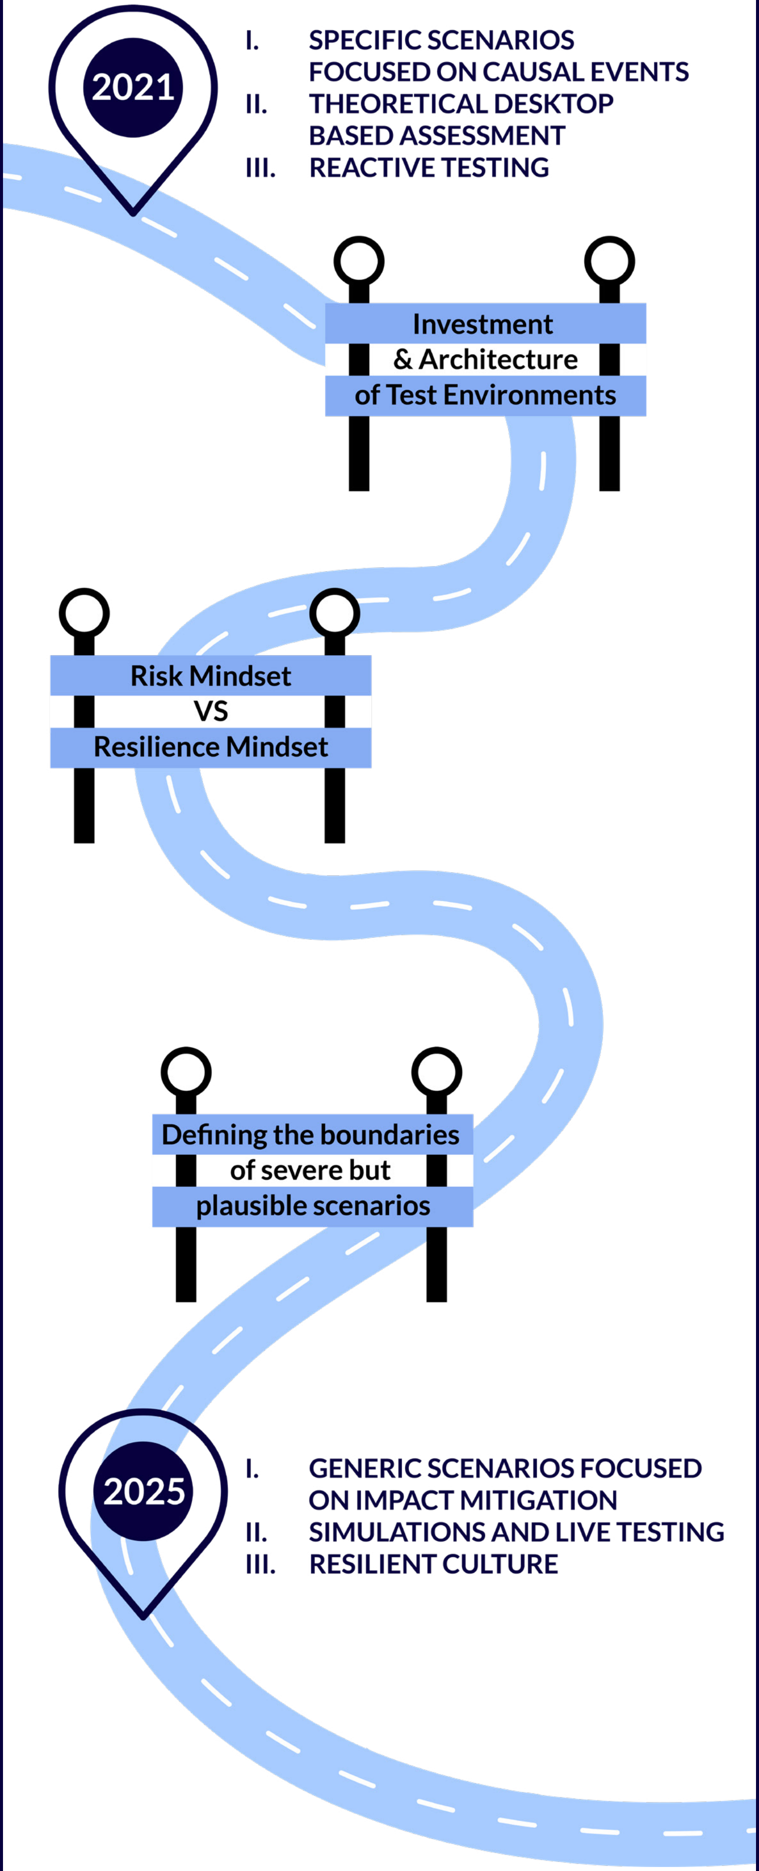
2021
Investment & Architecture of Test Environments
Risk Mindset VS Resilience Mindset
Defining the boundaries of severe but plausible scenarios
2025

2021

I. SPECIFIC SCENARIOS FOCUSED ON CAUSAL EVENTS
II. THEORETICAL DESKTOP BASED ASSESSMENT
III. REACTIVE TESTING

2025

I. GENERIC SCENARIOS FOCUSED ON IMPACT MITIGATION
II. SIMULATIONS AND LIVE TESTING
III. RESILIENT CULTURE