**Why Vendor Risk and Business Continuity Can't Live in Separate Silos**

*Michael Carpenter and Steve Fochler*

A bank gets a call early, before most employees are at their desks. A payment processor is experiencing minor service degradation. Transaction times are running a little longer than usual, but it's nothing critical, and the vendor says they're working on it.

By mid-morning, the institution's digital banking platform is rejecting payments because a short timeout window can't tell the difference between a slow transaction and a failed one. The call center says the wait time climb is well past normal. The fraud detection system, tuned for typical timing, starts flagging delayed payments and blocking good customers. The core processor's nightly reconciliation generates a wave of exceptions that someone now must research by hand.

A minor vendor degradation just became an institution-wide operational event. So how does this happen, and what can banks do to close the gap?

**1. The Cascade Is the Real Risk**

A third-party incident doesn't stay contained to the vendor. It moves through whatever depends on it. Downstream systems waiting on that vendor's output can stop working, and customer channels start to show stress before the vendor even declares an incident.

It's the type of risk due diligence doesn't pick up. Due diligence will tell you how reliable the vendor is, but it doesn't look at how a payment processor slowdown could impact other functions. That's the role of business continuity and understanding that impact influences how you manage your vendor relationships.

**Key takeaway**: A vendor incident's blast radius rarely stays limited to the function the vendor provides. Map the dependencies before an incident forces the question.

**2. The Accountability Gap**

A bank's ability to fix the vendor's problem is limited, but the consequences land entirely on the institution. The vendor diagnoses, prioritizes, and resolves the issue on its own timeline, while customer impact, regulatory exposure, and financial loss accrue to the institution in the meantime.

That gap isn't a failure of vendor management. TPRM reduces the likelihood of an incident and limits its impact, but incidents are still going to happen. A mature program plans for that reality ahead of time instead of hoping nothing breaks.

**Key takeaway**: TPRM can't eliminate the accountability gap, only shrink it. From incident response plans to addressing timelines for fixes in the vendor contract, banks need to use all the tools at their disposal to prepare for potential problems and limit the fallout.

**3. Map Vendors to Functions, Not Just to Risk Tiers**

No vendor operates in isolation. Each one connects to systems, data flows, and other third parties that keep a bank running. A vendor that several other systems rely on needs tighter uptime commitments than a standalone service, regardless of where it sits in a risk tiering matrix. A vendor sharing infrastructure with other critical providers needs insurance limits that account for the the downstream impact of its failure.

This is where TPRM and business continuity must talk to each other. One program identifies which functions are critical and how badly an outage would hurt; the other identifies which vendors sit behind those functions. Neither has the full picture without the other.

**Key takeaway**: A vendor's risk tier and a vendor's operational importance aren't the same measurement. A bank needs both before it knows where to focus.

**4. Detect Before the Vendor Tells You**

There's a difference between monitoring and reporting. Reporting is a dashboard showing uptime. Monitoring goes further, tracking whether a critical workflow completes within time tolerances. If the first sign of a problem is the vendor's formal incident notice, a bank's visibility depends entirely on the vendor's own assessment, timeline, and framing of the scope.

Real detection comes from watching your bank's internal signals: response times drifting, error rates moving outside normal ranges, reconciliations breaking in ways that don't match historical patterns, customer complaints clustering around a process. None of that requires waiting on the vendor to say anything.

**Key takeaway**: Monitoring should be built around a bank's critical processes, not just vendor reports. This gives the bank the best opportunity to proactively detect issues and implement incident response plans quickly.

**5. Recovery Isn't "We're Back Up"**

There are three separate jobs during a vendor incident.

Eradication is the vendor fixing the actual problem. That's their responsibility, not the bank's.

Containment is a shared responsibility. The vendor works to stop the problem from spreading on its end, while the bank acts on its own side without waiting to know the root cause: turning

on backup processes, rerouting volume elsewhere, and cutting off vendor access if it's a security issue.

While the vendor may handle its own verification, the bank shouldn't assume or blindly trust it. When the vendor says "we're back up," that's a claim, not proof. The bank has to test it, reconcile the numbers, and compare results from before and after the incident. Only then can anyone call it resolved.

Recovery also needs to be measured against something concrete — the service levels in the contract, not a vague sense that things seem fine.

Then there's the liability cap: the dollar amount the vendor owes if an outage happens. Most contracts just use whatever number the vendor proposed. Where possible, that number should align with the full financial impact of an incident. Estimating that impact means looking at outage duration realistically: the vendor's own history of outages, the outage history of similar third parties offering the same service, and the vendor's maximum tolerable downtime (MTD) as a stand-in for what an unplanned event would likely look like. Most institutions haven't done this math, and without some coaching or training on the approach, trying to match liability caps to real outage impact can end up feeling more theoretical than achievable.

**Key takeaway**: A vendor's word that a problem is resolved isn't confirmation that it's resolved. Recovery gets tested and reconciled independently, against what was contracted, and liability limits should align with the full financial impact of a vendor incident where possible.

**Closing the Loop**

After an incident is resolved, your bank needs to understand how its incident response plan worked and identify opportunities to do better next time.

One straightforward version of an After Action Report is naming three things that worked and three that didn't. Assign a person, an action, and a deadline for each finding. This ensures that what gets surfaced results in real change, not just an observation.

In the end, vendor incidents aren't really about the vendor. They're a test of whether your bank has done the work — mapping which functions actually depend on that vendor, building containment steps that don't wait on the vendor's timeline, monitoring internal signals instead of relying on the vendor's own incident notice, and verifying "resolved" instead of accepting it.

Make sure your TPRM program and business continuity function work together hand in hand. It's the best way to protect your bank from the impact of a vendor incident.