Joshua Rosenberg: Thrive in any environment - strengthening resilience through risk management

Remarks by Mr Joshua Rosenberg, Executive Vice President and Chief Risk Officer of the Federal Reserve Bank of New York, at the GARP Global Risk Forum, Federal Reserve Bank o New York, New York City, 7 November 2019.

* * *

As prepared for delivery

Good morning and thank you for the opportunity to speak to you today. Let me begin by saying that the views I express are my own and do not necessarily represent those of the Federal Reserve Bank of New York or the Federal Reserve System.

The ability to thrive in any environment, in both normal and chaotic conditions, is an organizational imperative. We can strengthen organizational resilience—"the ability to prepare and plan for, absorb, recover from, and more successfully adapt to adverse events"—through risk management.[1] That's because the tools of risk management help us make plans and take action today in the face of an uncertain tomorrow. In fact, we can deliver even more value to our organizations by defining resilience as a goal of risk management.

Today, I will talk about some of the challenges to achieving resilience and how risk management can be part of the solution. While there are many aspects of resilience, I'm going to choose two challenges and their connection to risk management: organizational silos and the complexity of control.

Organizational Silos

The first challenge is organizational silos.[2] Threats do not respect organizational boundaries. Some, like cyber threats, actually exploit them. Hurricanes don't care about your organizational chart, and cyber attackers would actually like a copy of it to use it against you. Since threats aren't siloed, our defenses can't be siloed either.

One of the insights of enterprise risk management is that good outcomes rely on integration across organizational boundaries and types of risk. That's why it's not just risk management, its *enterprise* risk management.[3] We can use approaches from enterprise risk management to create bridges across silos.

To start, enterprise risk management focuses on an organization's objectives and the critical processes that support them. Critical processes span organizational units in a connected chain of activities, starting with inputs delivered by suppliers, to value-added transformations within the organization, to the final delivery to customers. So, to strengthen resilience, we identify the risks to process outcomes and then design processes and controls to address those risks.

Since silos weaken resilience, you may want to look in your organization at some of the hotspots where enterprise-wide coordination is essential. I'd start with incident response. Are you ready for complex threats that require teamwork across organizational units? Another challenging area is external dependency management. Do you have a coordinated and coherent approach to understand and manage vendor exposures across the full spectrum of risks, including information security, data privacy, compliance, business continuity, and credit?

Risk silos are as problematic as business-line silos.[4] While the specialization of risk disciplines has benefits, common negative side effects are insufficient communication and coordination, followed by information gaps and ineffective decision making.

Take the case of a vendor that provides critical services and is on the brink of bankruptcy. Are credit risk analysts, operational risk and business continuity professionals, and the business affected working together to share information, insights, and potential responses? Or, are silos preventing information from getting to the right people at the right time to prepare and act?

Risk silos can result in siloed decisions, and looking at risks one by one is incoherent. We could drive operational risk to zero by having no operations (I guess we can all go home early), but then strategic risk rises to infinity. Since most meaningful decisions involve balancing risks, to get to the right decision, we need a risk picture that covers all relevant risk types and engages all relevant risk experts.

Similarly, real world vulnerabilities do not have to conform themselves to a particular risk type. Think about weaknesses in identity and access management. This is a driver for risks ranging from cyber, to compliance, to operations, and more. This is also a type of risk that can't be mitigated solely through technical controls put in place by information security experts; that is, a siloed risk response is insufficient. So, stronger collaboration and coordination across risk disciplines is one solution to the negative effects of risk silos.

Another specific pain point created by risk silos is the flourishing, and perhaps overgrowth, of risk assessments. Most organizations have many risk reviews: often a different review for each type of risk. Some focus on assets, others focus on organizational units, and yet others look at specific parts of processes. The result can be a set of assessments with potentially inconsistent coverage and disconnected recommendations.

So, an appealing alternative to counter that trend and bridge siloed assessments is to integrate risk assessments. A unified approach can bring together risk and business professionals to understand and manage the risks to critical processes. The idea is to create a common picture of gaps and then design coordinated action plans that can target improvements in resilience.

Let's focus on business continuity to see how an integrated risk approach can strengthen resilience.[5] Business continuity programs have traditionally been designed to restore critical physical or technology infrastructures that can be disrupted by threats like extreme weather or loss of power. So, a business continuity plan might involve arranging for a backup data center if the primary data center goes down due to a hurricane. Within an organization, different business areas may develop their own plans that are focused on their specific business needs and priorities.

An integrated business continuity plan is one that is designed to protect critical organizational processes and their essential assets from the full range of relevant threats. An integrated approach begins with organizational objectives. This bridges a gap between what good looks like to an individual business and what it looks like to the organization as a whole. Among other things, an integrated approach analyzes and plans for the impacts of disruptions on suppliers and customers from an operational as well as reputational perspective.

Integration broadens the focus from the availability of infrastructure (e.g., the servers are up and the facilities are accessible) to successful organizational outcomes (e.g., the organization is able to deliver products and services). In the current threat environment, achieving resilience is as much about protecting the integrity and confidentiality of information as it is about maintaining availability. An integrated plan delivers coordinated and coherent responses to incidents, including ones of significant scale, complexity, and surprise that can overwhelm normal control systems.[6]

The Complexity of Control

The second resiliency challenge is the complexity of control. To improve resilience, we seek to control outcomes, and risk management has a lot to say about control. A simple story about

control goes like this: when we are setting up a business process, we look for sources of problems and start adding controls. Mainly, we'll create controls to prevent problems.

But, recognizing that there will be times when prevention doesn't work, we also put in place controls to detect problems and then correct them. Once the process is up and running, if problems come up that we didn't expect, we add additional controls to prevent those problems from happening again.

I wanted to go through some examples of how this doesn't always work well in practice. From that, we'll see that a systems perspective can explain why our attempts at control sometimes fail and what we can do about it.[7]

A key insight from systems thinking is that controls are part of a system. Controls interact with other controls and with the production process. One implication of connectedness is that changes to any part of the system can spillover and affect other parts. Controls that work well separately may not work well together.

That's one reason why controls created in isolation and added incrementally can create unexpected results. It's a weakness of control systems that have grown through an ad-hoc approach of "see a problem, add a control."

Let's walk through a scenario where controls don't behave as expected or intended.

On Sunday, you run a business resumption exercise to test the ability of your primary data center to fail over to your backup data center. The test appears to be a success, and you have some useful lessons learned to make the failover process run better next time. But, when you come to work on Monday morning, something is wrong. It turns out the test corrupted customer account data. Your call center is overwhelmed with questions and complaints, because your billing system sent a blast email with overdue payment notices to all of your customers.

In a simple world, controls solve problems, they don't create them. In the real world, the control and production systems are linked, so the failed execution of a control can, in and of itself, disrupt operations. In other words, controls can cause harm.[8]

On Tuesday morning, a hurricane damages your main office and floods the backup generator that powers the emergency communications system. Because the emergency communication system is offline, employees do not know they should report to the backup work location, so operations are out for the day.

In a simple world, controls are protection, they don't need protection. In the real world, both the production and control system can be damaged and degraded. So, the control system itself must be designed to operate under stress, and its performance must be monitored.

It's been a tough week so far, but you've made it to Wednesday. At noon, your information security officer tells you that a security scan has flagged a possible cyber breach. You cut off all network connections, send non-essential employees home, and wait until the end of the day until diagnostics are complete. The news at the end of the day is good: it turns out that the alert was a false positive. But, the organization lost half a day of productive activity.

In a simple world, controls detect problems immediately, diagnose with perfect accuracy, and result in corrective actions that are instantly and fully effective. In the real world, detection takes time, diagnosis is not always accurate, solutions take time to implement, and they don't always work. That said, we must make decisions—including ones about whether and how to execute controls—based on the information available at the time.

Control design and execution take place in a world in which there is uncertainty about whether

there is a problem, what the problem is, what to do about it, when to do something, and what the effects of what we do will be. We can make better decisions when we understand and manage that uncertainty.

There is also uncertainty about threats. We will never be able to envision the characteristics of all future causes of disruption. Even when we plan effectively, we will be surprised again. And, like any other system, the control system will exhibit performance variability. For these reasons, we can't rely solely on prevention. Resilience requires the ability to withstand and recover when threats make it through preventive barriers, so a strong control system must incorporate preventive, detective, and corrective controls.

Perhaps because we are overoptimistic about prevention, detective and corrective controls sometimes get second-class billing. For many threats, time is not on our side; the impact grows as the detection time increases. That's the case with "dwell time" for cyber breaches, which is the amount of time attackers are on your network before you detect them.[9] It's important to decide how long you are willing to wait to find out that you have a problem, and then use that as a requirement to design detective controls.

Underinvesting in corrective controls is problematic too. For example, how would you feel after a ransomware attack, if you find out that your backed-up data is intact but it will take months and millions of dollars to restore? Or, if your first corrective action is the wrong solution and makes the problem worse? It takes time and resources to develop the tools, procedures, and skills to accurately diagnose causes, contain the immediate threat, and implement a long-term solution. Will you be satisfied with the performance of your corrective controls when you need them?

There's also a bias towards automated controls that is not always well considered. Both manual and automated controls have a place in a control system. While machines are better than people at repetitive tasks with clear decision criteria, people are better than machines (for now) at tasks that are ambiguous and require judgment and flexibility. When you automate controls, are you preserving the knowledge and skills of staff who need to understand the system well enough to solve problems when automation fails?

So, we've seen how controls are part of a dynamic system in an environment of uncertainty. We can't know whether our controls are sufficient unless we look at the behavior of the control system as a whole. And, we can't achieve our desired outcomes unless we design for them.[10]

Conclusions

Today, I've highlighted two challenges to resilience: organizational silos and the complexity of control. And, I've proposed a set of solutions based on risk management tools: a focus on organizational objectives, end-to-end management of critical processes, integrated risk management, and a systems approach to control. These all share the common characteristic of being integrative.

Resilient components do not necessarily add up to a resilient whole. So, organizational resilience is an enterprise goal that is achieved through coherent and coordinated enterprise solutions.

A resilient organization is able to make better plans, decisions, and actions that deliver desired outcomes over a range of conditions. To strengthen resilience, we seek to tame the problematic present and constructively move forward into the hazy future. The tools of risk management can help get us there.

---

1 National Research Council, 2012, Disaster Resilience: A National Imperative, Washington, DC: The National Academies Press.

2 Gillian Tett, 2005, The Silo Effect: The Peril of Expertise and the Promise of Breaking Down Barriers, New York,

NY: Simon and Schuster.

[3] Committee of Sponsoring Organizations of the Treadway Commission, 2017, Enterprise Risk Management: Integrating with Strategy and Performance.

[4] Philippa Girling, 2013, "Operational Risk and Convergence" in Operational Risk Management: A Complete Guide to a Successful Operational Risk Framework. Hoboken, NJ: Wiley. Ramy Farha, Allen Meyer, Evan Sekeris, and Elena Belov, 2018, Non-Financial Risk, Convergence, and Integration, Oliver Wyman. Steve Culp and Chris Thompson, 2016, The Convergence of Operational Risk and Cyber Security. Accenture. Nader Mehravari, 2013, Resilience Management through use of CERT-RMM and Associated Success Stories, 2013 IEEE International Conference on Technologies for Homeland Security.

[5] Richard Caralli, Julia Allen, David White, Lisa Young, Nader Mehravari, and Pamela Curtis, 2016, CERT Resilience Management Model Version 1.2: Service Continuity, CERT Program: Carnegie Mellon. Rico Brandenburg, Tom Ivell, Evan Sekeris, Matthew Gruber, and Paul Lewis, 2019, Striving for Operational Resilience: The Questions Boards and Senior Management Should Ask, Oliver Wyman.

[6] David P. Moynihan, 2009, The Network Governance of Crisis Response: Case Studies of Incident Command Systems. Journal of Public Administration Research and Theory, 19, pp. 895–915.

[7] Peter Senge, 2006, The Fifth Discipline: The Art and Practice of the Learning Organization, New York: Doubleday. Nancy Leveson, Nicolas Dulac, David Zipkin, Joel Cutcher-Gershenfeld, John Carroll, and Betty Barrett, 2013, "Engineering Resilience into Safety Critical Systems" in Resilience Engineering: Concepts and Precepts, Eds. Erik Hollnagel, David Woods, and Nancy Leveson. Dorchester, UK: Dorset Press. Resilience Engineering: Sidney Dekker, Erik Hollnagel, David Woods, and Richard Cook, 2008, New Directions for Measuring and Maintaining Safety in Complex Systems, Lund University School of Aviation.

[8] Jonathan B. Wiener, 1998. Managing the Iatrogenic Risks of Risk Management, Risk, Health, Safety, and the Environment, 39, pp. 9-82.

[9] M-Trends 2019, FireEye Mandiant Services Special Report.

[10] Nancy Leveson and John Thomas, 2018, STPA Handbook. Matthew Leitch, 2011, Intelligent Internal Control and Risk Management: Designing High Performance Risk Control Systems, Burlington, VT: Gower Publishing Company. Deloitte, 2009, Mining Safety: A Business Imperative.