Operational resilience for financial markets participants

Explore the benefits of investing in recovery before you need it

IBM

Contents

Executive summary

The financial services industry, its regulators, the media and public have become obsessed with eliminating failure or outages. Any organization that experiences such an event is publically held to account with their reputation in the marketplace at risk of being tarnished. There is an expectation, certainly in the minds of commentators and the media, that any failure should be recoverable via an automatic failover to an alternative system or facility. In actuality, this type of event does occur at frequent intervals but most people are unaware because recovery is seamless and very effective.

However, there are occasions when either this is not possible or for business reasons it is not the most appropriate path to follow, or both. The marketplace is complex with many interconnected system dependencies. Therefore in the event of a major failure, a number of factors have to be considered when determining the appropriate recovery actions, both technical and commercial, often with regulatory consequences if wrong decisions are made. For example, in the world of financial exchanges, there is a primary obligation to offer fair, orderly and transparent markets at all times.

In meeting this obligation there are many elements to consider. These include member firms' connectivity to the secondary facility and the regulatory body, as well as synchronization of the order book across the production and backup environments. One must also consider the technical and commercial issues, such as the ability to ensure effective and timely communication to the marketplace combined with member firms' readiness and compliance with reopening the market to conform with the regulatory requirements. It becomes a very complex decision to make while under extreme time pressure from the media, customer and regulators to reopen the market and restore the norm.

Therefore the reduction of the risk of any incident occurring is one that should be actively pursued through appropriate investments, despite the law of ever-diminishing returns and the expense incurred. The investment case should be based on an acceptable business impact assessment approved by the responsible business leaders along with a view of the likelihood of any particular failure event occurring. For example, a power failure that has previously occurred and caused disruption may happen again unexpectedly during or after recovery. Is it appropriate to have a "belt and suspenders" approach to high availability? The answer, in many cases, is yes.

Given the ever-increasing complexity of the challenge, we would argue that you should expect issues to actually increase. Therefore, an organizations' ability to prepare and recover from an incident needs senior executive focus. This will help ensure the appropriate level of planning, scenario development and testing, ownership, governance, practiced procedures, and internal and external strategies (such as for communications) are documented. There must also be an appropriate level of business impact analysis conducted and approved by the business leadership. This is an area where if your organization fails, it will have no defense. These areas should be invested in and your customers, the media and regulators will all expect you to have invested and tested. If you have not, it will be a difficult position to defend after an adverse event.

The financial services sector (FSS) can look at other industries for lessons that can be learned in its desire for resiliency and recovery capabilities, best practices and approaches. The airline adage, "don't tell me about the cost of safety, tell me about the

cost of failure," is true, and can be applied to the FSS. In addition, because of the fiduciary responsibilities of financial markets, participants must take into consideration the potential for reputational damage, financial cost and other negative implications that may occur around any given incident. The analogy may be extended further to the failure to recover successfully being similar to a damaged plane successfully landing. The publicity turns positive with both the skills of the pilot to identify the critical issue and manage the risk by following proper procedures, and the availability of critical information and the visualization of the data that facilitates the recovery are applauded. Airlines invest significantly in testing and training pilots in a wide variety of potentially dangerous scenarios so that during the recovery process, decision-making becomes second nature. This is facilitated through detailed documentation of the recovery processes that should be followed in any given incident to assist individuals and ensure consistency while under intense pressure. It is an increasing area of focus by regulators who are mandating that FSS organizations provide evidence of the implementation of proactive supervisory control mechanisms (tools, processes, data and governance) as part of senior management's responsibilities to help smoothly steer an organization in its daily activities.

In summary, you should invest in prevention, expect failure and invest in recovery. The keys to recovery are robust planning, governance, procedures, monitoring, event correlation, communication and regular realistic testing. This should be combined with regular reviews and assessments of readiness.

If you fail to prepare, then prepare to fail

The challenge of resiliency and security has always and will continue to be a significant issue at the forefront of people's minds who are interested in running high-profile operations. The challenge over the last few years, however, has been regularly becoming more difficult. A primary cause is the growth of technology and its complexity combined with the need to maintain a real-time or always-on organization to enable participants to survive and compete in the modern marketplace. The ultra-low latency and colocation challenges have compounded the complexity and technical challenges at the same time as exchanges are experiencing an ever-increasing threat of a security breach or cyber attack. All of this can lead to a very real life threatening event for an organization.

These challenges are now being addressed with specific regulation, as in the case of the United States Securities and Exchange Commission (SEC) with its regulation entitled *Systems Compliance and Integrity* (SCI) that was passed at the end of 2014 and will need to be complied with no later than the last quarter of 2015. The regulation, while focused on the US and the self-regulatory organizations (SROs), will ultimately lead to wider industry adoption as a best practice. The SCI regulation will be seen, at least in the US, as the benchmark for preparedness. All firms should review these regulations and see if your company would be compliant with the appropriate sections relevant to your organization, regardless of geographic location.

Globally, this regulatory focus can be evidenced by the International Organization of Securities Commissions' (IOSCO) recent release of two industry consultation reports aimed at further enhancing the ability of financial markets and intermediaries to manage risks, withstand catastrophic events and swiftly resume services in the event of disruption.

The consultation report *Mechanisms for Trading Venues to Effectively Manage Electronic Trading Risks and Plans for Business Continuity* from IOSCO provides a comprehensive overview of the steps trading venues take to manage the risks associated with electronic trading and the ways they plan for and manage disruptions through business continuity plans. As technology continues to evolve, exchanges will need to continuously adapt to these changes. The report provides recommendations to help regulators ensure that trading venues are able to effectively manage a broad range of evolving risks. It also proposes sound practices that should be considered when developing and implementing risk mitigation mechanisms and business continuity plans aimed at safeguarding the integrity, resiliency and reliability of an exchange's critical systems.

IOSCO's second consultation report, *Market Intermediary Business Continuity and Recovery Planning*, proposes standards and sound practices that regulators could consider as part of their oversight of all business continuity and recovery planning by market intermediaries. The supporting release material points out that "these sound practices may also prove useful to intermediaries who are developing and implementing business continuity plans."

The rationale for the reports has been articulated because of several recent disruptive events and emerging threats in major international financial markets. These events highlighted the need to examine and identify the key measures and arrangements in place at trading venues and market intermediaries to restore their critical functions should a disruption occur. Both reports contain lists of recommendations and sound practices.

It is clear that there is now significant focus on this domain from regulators and the media and the public. It is more critical than ever before that there is a significant proactive focus from organizations to ensure that they remediate the chances of an incident as much as possible including a reliable capability to recover in a timely and efficient manner.

Why is operational resilience required?

This may seem to be an obvious question, but it is worth considering why resilience is important. By any reasonable measure, resilience is not optional. The challenge is how to make it work in your organization.

1. Regulation requires it and penalties can result from failure to comply. Every industry has its own regulation and governance, ranging from generic health and safety measures through to industry-specific capabilities. Increasingly, governance is reported and publicized and any failures can attract significant fines or penalties.

2. Customer perception is adversely affected by poor service. Whether caused by the publication of governance failures or through tangible service outages, a poor IT service can cause substantial reputational damage.

3. Incidents can cause errors in processes and mistakes or losses in data. If your data is lost or corrupted, this can have a noticeable impact on the people and processes which use your services. If transactions or processes are not completed, duplicated or left in an uncertain state, rework and investigation can take time and resources that can compound the reputational damage.

The alternatives to preparation can be very public, expensive and potentially create systemic risk

There have been an increasing number of outages or issues with major exchanges and their member firms or participants that are causing concern to the wider industry and specifically their regulators. Trust and reliability are key elements of any successful financial services organization. The primary concern is the increasing number of outages and the negative publicity this generates.

The industry has been going through significant change over the last 10 years, both from a business and technology prospective. The drive for low-latency trading across all asset classes has added to the complexity the industry was already dealing with. The pace of change and the rapid deployment of new technologies do not come without risks. As the industry negotiates this period of technology change, participants are increasingly required to respond and adapt to identified challenges. However, this could potentially impact their ability to protect against new and unidentified threats, while systemic risks are growing and are more difficult to anticipate. As the following examples illustrate, there have been a number of issues that show the issues the industry faces are real and pertinent.

NYSE

An issue at the New York Stock Exchange (NYSE) in July 2015 required the market to close at 11:32 AM for over three hours, although it eventually reopened for the close. While not significantly disruptive to market execution—NYSE today only processes less than 20 percent of the market and the symbols can be traded on many other US venues—it was clearly a major issue. The incident was not helped by the fact it coincided

with similar technical issues at BATS Global Markets a few days before and United Airlines on the same day causing initial concern over a cyber attack. US President Barack Obama and the SEC were briefed and the story ran all day on the major news outlets. A number of commentators from the industry expressed disappointment that the NYSE could not failover automatically. Interestingly, the decision was made not to switch over to their back-up facility. This was articulated as the most effective and efficient approach. The disaster recovery (DR) center was positioned as the right option in the event of a major disruption such as a terrorist attack or natural disaster. The challenge not to switch was described as one of connectivity and prioritizing being ready for the close. The initial statement on the issues appeared to be the rollout of new software that caused the problem; NYSE needed to close the market in order to resolve the problem.

SGX

In November 2014, the Singapore Exchange (SGX) suffered a power outage that led to the exchange's production center losing all power. However the exchange was able to recover after several hours and operate a fair, orderly and transparent market. The outage was a major concern to the exchanges community and its regulator, the Monetary Authority of Singapore (MAS). After the incident, a MAS review resulted in the SGX agreeing to spend SGD 20 million on remediation efforts and is restricted from increasing charges imposed until these efforts have been completed.

NASDAQ's Facebook IPO

While all eyes were on the market on 19 May 2012, the day of Facebook's IPO, the exchange that handled the listing, NASDAQ OMX, faced major challenges. NASDAQ knew that this would be the largest IPO and individual participation would be unprecedented. Its systems to handle this type of trading had been tested in preparation for this event. However, a major unexpected glitch caused havoc for hours. Following the Facebook IPO, Tabb Group's report, *IPO Survey: Market Barometer*, found that the impact of Facebook's IPO on investor confidence was almost as great as the market-wide Flash Crash on 6 May 2010. The exchange was also fined significantly by the SEC for the issues faced at the time of the Facebook IPO.

Euronext

Euronext prompted an angry response from clearing brokers when some of its servers remained running after a technical glitch shut down three client portals. The incident allowed some customers to continue trading, while others were cut off from the market. The decision raises questions about how an exchange should react to technical issues when they are required to operate a fair, orderly and transparent market throughout the recovery process.

Bloomberg

Bloomberg's trading terminal experienced a prolonged outage in June 2015, resulting in the postponement of a treasury auction in the UK. Service was fully restored as of 4 PM London time. "We experienced a combination of hardware and software failures in the network, which caused an excessive volume of network traffic," the company said in a statement to financial TV network CNBC. The statement continued, "This led to customer disconnections as a result of the machines being overwhelmed. We discovered the root cause quickly, isolated the faulty hardware and restarted the software. We are reviewing our multiple redundant systems, which failed to prevent this disruption."

The pattern cannot be ignored

What is becoming clear is that although each of the incidents above is related to a single organization, they did or could have had an impact at a systemic level which is the real concern. The regulations and consultation paper (SEC SCI and IOSCO) referred to above include specific activities to address these concerns such as industry-wide testing and active participation in the generation of a business continuity plan (BCP) with the involvement of all internal and external stakeholders.

What is operational resilience and how does it relate to operational risk?

Operational risk is defined in the Bank for International Settlements' (BIS) Basel II document as "the risk of a change in value caused by the fact that actual losses, incurred for inadequate or failed internal processes, people and systems, or from external events (including legal risk), differ from the expected losses." The BIS Basel II categorizes Operational Risks into the following high-level categories:

Internal Fraud: Misappropriation of assets, tax evasion, intentional mismarking of positions and bribery

External Fraud: Theft of information, hacking damage, third-party theft and forgery

Employment Practices and Workplace Safety: Discrimination, workers compensation, employee health and safety

Clients, Products, and Business Practice: Market manipulation, antitrust, improper trade, product defects, fiduciary breaches and account churning

Damage to Physical Assets: Natural disasters, terrorism and vandalism

Business Disruption and Systems Failures: Utility disruptions, software failures and hardware failures

Execution, Delivery, and Process Management: Data entry errors, accounting errors, failed mandatory reporting and negligent loss of client assets

Failure to meet Basel II requirements has a direct impact on capital allocation, regardless of any additional costs or fines or both associated with non-compliance. Operational resilience focuses on the business disruption, systems failures, and execution, delivery and process management aspects.

Managing operational resilience through the identifications, assessment, monitor and control, and mitigation phases is an insurance policy against an extremely broad range of risks and is not generally a direct revenue-generating activity. However, these activities are often regarded as lower-priority items, especially in economic down cycles and when an organization is going through significant change, such as a growth strategy.

Planning for things that may go wrong requires a conscious choice. The natural activities of a business and its IT organization do not, in themselves, create resilient systems and processes. Too many conflicting pressures affect organizations, making it difficult to strike a balance between the time, cost and opportunity challenges.

Being resilient requires more than simply the ability to recover from failures. Of course, this helps because failure is more likely now than in previous times, so recovery excellence is important. Resilience also requires the ability to prevent failures and take action to avoid them. To use an analogy, you need to be good at fire prevention rather than just fire fighting. You need to be good at putting out fires when they break out, but you need to be even better at spotting where fires might occur and removing those likelihoods.

Resilience relates to your business context because the degree of acceptable risk will vary among different organizations. Resilience should be proportionate to the impact on your organization. This means the effort and cost of preventative measures should relate to the likelihood and impact of events that may occur. Chemical plants, banks and government organizations will all have different views on the need to avoid incidents, as well as the amount of effort and resources they are willing to invest to achieve this.

To be resilient, your organization needs a combination of both preparation and demonstration with accountable governance. Preparation puts in place proactive measures to avoid incidents or speed recovery from them or both. Demonstration comes from the events that prove the success of your preparation. Planning is essential to establishing your services on a secure

footing and, as always, it is less expensive to solve resilience issues in advance at the design stage. When a recovery capability is deployed, it needs to be checked and tested to ensure it is effective. This checking may be a planned test or it may be observing your capability in action as it limits or avoids service impact. These considerations form the cornerstone of a robust "three lines of defense" resilience model.

Why is operational resilience difficult to achieve?

The resilience of organizations is threatened by constant change and increasing complexity. Some of the major factors that can affect resilience are listed in Table 1.

Dimension	Description
Mergers and acquisitions	Private sector organizations are subject to discontinuous change when they acquire or dispose of companies or divisions. Public sector organizations can be reorganized and restructured in line with changing political control. In either case, the result is a set of fragmented IT and business capabilities, with varying standards and processes.
Skills	There is often a lack of breadth and depth in skills to address the needs of the business to run and change.
Outsourcing	In general, cost savings has been the primary driver for outsourcing, which can place quality and service delivery at risk.
Culture	Is the company culture appropriate, and is the leadership setting and acting with the necessary tone?
Complexity	Complexity is everywhere and continues to get worse.
Projects	Most organizations seek to drive change in line with their strategic goals through projects. While change control is often a feature of major projects, resilience aspects such as functionality or quality can be sacrificed in return for time to market advantages.
Incremental change	A large amount of change occurs on a small scale with incremental additions or alterations. These changes can often attract less oversight and control than major projects.
Human error	Mistakes happen, whether through lack of training, poor control or just carelessness. These are not always noticed at the time.
False positives	Poorly designed controls can generate warnings and errors that are not correct—the "cry wolf" syndrome—which works to discredit the controls.
Inadequate testing	Testing is expensive and time consuming, leading to a tendency to do as little as possible, or to create tests that are easy to pass.
Process failures	Lack of mature processes can cause risk to increase, through poor change control, inadequate configuration management and lack of root-cause elimination..
Lack of focus	By not receiving proper attention, systems can become degraded or overlooked.
Operating model	Core operations functions that reside in a single location or site (for example often resulting from migration to "captive" organizations in a simple lift-and-shift approach) can become a single operations point of failure.
Poor global infrastructure	Global epidemic outbreaks (for example swine flu or Ebola) often prohibit travel to destinations necessary to provide operational support. In these cases, an absence of global technical resources can introduce operational resiliency risk.

Table 1. Factors that affect business resiliency.

Every organization has to deal with challenges like these. Some may be periodic, such as mergers and acquisitions or major projects, while others are present continuously and require repeated attention. Very few organizations can afford to prepare fully for every eventuality, making it vital to have conscious, repeated assessments and actions to deal with the most pressing possibilities.

Suggested approach and method

Accepting the need for resilience, and the difficulties inherent in achieving it, your organization can take practical and consistent steps to prepare and develop this capability. The resilience framework from IBM outlines the approaches you can take. It shows the different areas of investigation for resilience capability and the different activities to be undertaken for these areas. The framework, shown in Figure 1, contains three complementary methods that can be used:

- **Resilience assessment:** Question-based checks to understand how resilient you are.
- **Resilience proving:** Principles and a scorecard to demonstrate your resilience.
- **Resilience maturity:** A matrix to check how well you manage resilience.

Your choice of method depends on what drives you to act.

- Major incident raises concerns over resilience? Choose a resilience assessment.
- A drive for process improvement or best practice or both? Review your resilience maturity.
- Regulatory or audit concerns driving a focus on governance and assurance? Undertake resilience proving.

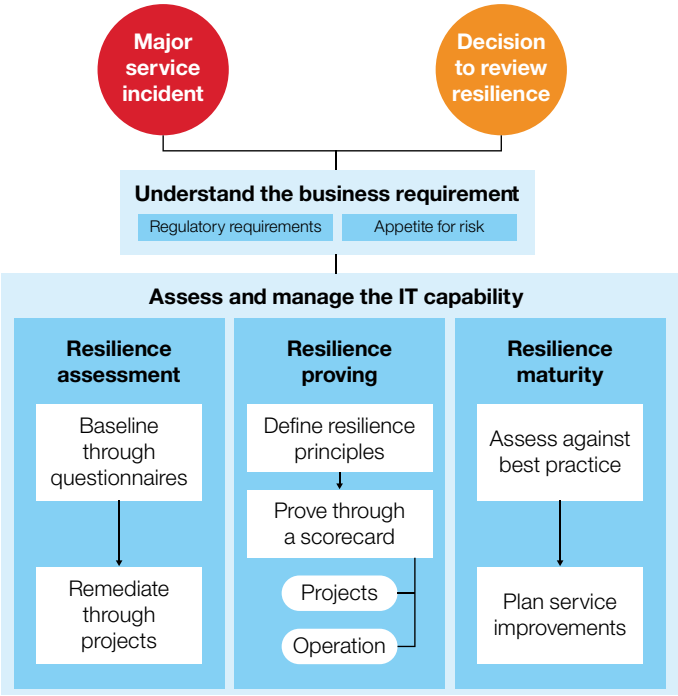

Figure 1. The resilience framework from IBM.

The framework dimensions and activities

When you assess your resilience, where should you look? Each of the above methods requires a way to drill down into the organization, its capabilities and performance. The resilience framework dimensions break down the scope of a typical IT organization into themed topics and subtopics. This helps organize the data systematically, while making it easier to identify the people in the organization who can answer the questions and give their input. Every organization is different so you will need to amend this typical breakdown to map onto the different support functions across your specific IT structure.

For each of these framework dimensions, what should you consider? Breaking down the dimensions into preparation and demonstration activities can help you consider the topic across the full lifecycle of activity, from design to BAU operation and continual service improvement. This approach blends two different strands of resilience thinking.

Preparation

- Setting out risk-appropriate measures to deliver the required resilience.
- Creating designs and templates that enable resilience.
- Making pre-production estimates of performance and capacity requirements.

Demonstration

- Carrying out pre-production and business as usual testing to prove capability.
- Reporting on the technology and service management processes that manage incidents that occur.
- Undertaking governance and assurance of overall resilience capability.

Resilience assessment: How resilient are you?

Using the framework dimensions and activities, draw up a set of questions to elicit the required information. Using a mixture of closed and open questions, you should gather responses, documentation, reports, pictures and any relevant materials that illustrate the resilience position.

- Observations and concerns should be identified through workshops, interviews and document review.
- Key findings will be derived from the most significant concerns.
- An assessment of the likelihood and potential impact of IT resilience risks should be undertaken to help determine the priority of recommendations.
- A set of prioritized recommendations can then be derived from findings and risks.

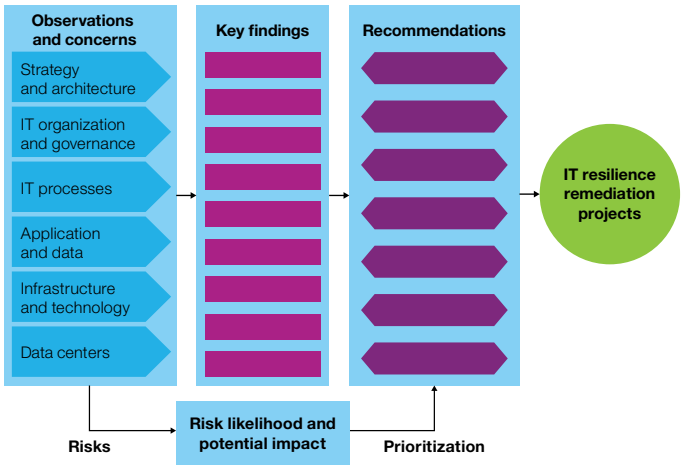

Figure 2. The resilience assessment process.

Resilience proving: How do you know you are resilient?

What is resilience proving? Simply stated, it is a set of activities that allows you to demonstrate a level of resilience in the IT systems that are being or have been deployed within your organization. World-class organizations conduct resilience proving as part of the upstream design / build / test activities prior to a solution being deployed. They also continue resilience proving on a regular basis as part of standard operational management in the downstream world after system deployment.

In order to conduct resilience proving, you need to have a suitable, objective scoring mechanism. This enables you to measure your solutions against a set of resilience principles that are valued by your organization. These principles should be based on generic good practices that have been customized to meet your needs. Examples of principles are listed in Table 2.

Project stage	Principle
Analysis	All applications and infrastructure must be classified in the appropriate business resilience service tier.
Design	Resilience must be designed from the business system or application downwards. An analysis of the potential scenarios under which the system could fail must be performed as part of requirements analysis and solution design. All solutions will only be designed using standard resilience patterns. Applications on shared infrastructure must have sufficient isolation to ensure appropriate availability.
Build	Applications must not contain hard-coded references to infrastructure components.
Test	Applications and infrastructure will only enter production when it has been proven that all technical (non-functional) requirements have been addressed.

Table 2. A detailed explanation of resilience proving steps.

Within the project lifecycle, scoring must be able to demonstrate that the resilience principles are being upheld by the solution design and have been achieved prior to production acceptance. Outside of the project lifecycle, the downstream environment can use evidence of resilience or its absence as generated by IT service management, disaster recovery and operational management processes. A scoring approach can be developed that shows how resilience of the solution is being managed and maintained.

Resilience maturity assessment: Can you manage your resilience?

If you compare your position against generic resilience maturity statements, you can make a judgment on areas of good practice or areas for improvement. The assessment compares current practice against a set of markers for each maturity level, enabling you to identify where you score in each area.

1. Compare your deployed resilience capability against the matrix to establish the current position
2. Use the matrix to identify a target maturity level and work out the actions needed to close any gap between the current and target levels.

Conclusion

With the growing demand for 24x7, always-on services, the importance of resilience is greater than ever within our industry. Any failures can bring unwelcome attention from the news media and draw unwanted scrutiny from industry regulators. At the same time, when the power of social media means the visibility, impact and cost of unavailability is higher than ever before, there truly is nowhere to hide.

IBM can help organizations determine their resilience capability by carrying out a structured assessment, undertaking resilience proving and scoring their maturity of processes to help manage resilience. IBM has proven that by adopting a conscious approach to preparation and demonstration, organizations can achieve higher degrees of resilience. Equally important, IBM also helps organizations demonstrate this capability to business and regulatory stakeholders.

A change of focus is required. Resilience can often be classed among the non-functional requirements, which, experience shows tend to be given less focus than functional requirements. Yet a proper resilience approach is your best assurance of functional performance. Without it, the real business goals of your organization will likely be difficult or impossible to achieve.

Maturity level	Unfocused	Aware	Capable	Mature	World class
Resilience design: Identify	Little or no risk assessment done. No risk appetite defined. No linkage from IT systems to business use.	Basic risk assessment done. Basic risk appetite defined. No linkage from IT systems to business use.	Reasonable risk assessment done. Reasonable risk appetite defined. High-level linkage for specific IT systems.	Full risk assessment done. Basic risk appetite defined. High-level linkage for all critical IT systems.	Full risk assessment done. Full risk appetite defined. Detailed linkage from IT systems to business use.
Resilience design: Prevention	Little or no standard design patterns or templates. Little or no understanding of likely business volumes or performance characteristics.	Limited use of standard design patterns or templates. Basic performance testing for application and system changes.	Standard design patterns or templates exist but not mandatory. Reasonable performance testing. High-level business volumes obtained.	Standard design patterns or templates exist with working waiver system. Reasonable performance testing. High-level business volumes obtained.	Standard design patterns or templates rigorously used. Comprehensive performance testing. Full business volumes obtained.
Resilience validation	No pre-production resilience testing carried out. No pre-production performance and capacity profiling done. No regular BAU recovery tests carried out.	Limited pre-production resilience testing carried out. Basic pre-production performance and capacity profiling done. Irregular or partial BAU recovery tests for selected systems carried out.	Reasonable pre-production resilience testing carried out. Reasonable pre-production performance and capacity profiling done. Regular BAU recovery tests for selected systems carried out.	Full pre-production resilience testing carried out. Reasonable pre-production performance and capacity profiling done. Regular BAU recovery tests for all critical systems carried out.	Full pre-production resilience testing carried out. Full pre-production performance and capacity profiling done. Comprehensive BAU recovery tests carried out.
Resilience action	Little or no automated recovery, or HA provision. No analysis of what incidents show about resilience capability. No root-cause analysis to identify and remove problems.	Basic automated recovery, or HA provision. Limited analysis of what incidents show about resilience capability. No root-cause analysis to identify and remove problems.	Reasonable automated recovery, or HA provision. Reasonable analysis of what incidents show about resilience capability. Basic root-cause analysis to identify and remove problems.	Automated recovery, or HA provision, for all critical systems. Full analysis of what incidents show about resilience capability. Reasonable root-cause analysis to identify and remove problems.	Automated recovery, or HA provision, for all relevant systems. Full analysis of what incidents show about resilience capability. Full root-cause analysis to identify and remove problems.
Resilience assurance	No resilience reporting. No service improvement program based on lessons learned.	Basic resilience reporting. Basic service improvement program based on lessons learned.	Reasonable resilience reporting. Reasonable service improvement program based on lessons learned.	Full resilience reporting. Reasonable service improvement program based on lessons learned.	Full resilience reporting. Full service improvement program based on lessons learned.

Table 3. A description of how resilience improves as the process matures.

For more information

To learn more about operational resilience for financial market participants, please contact your IBM representative or IBM Business Partner, or visit the following website: **ibm.com**/industries/financialmarkets

About the authors

Philip Enness is the Global Lead for Markets Infrastructure (Exchanges, ECNs, ATS, CDS, Regulators and Central Banks) at IBM, based out of Singapore. He has extensive financial markets knowledge gained in a 25-year career in the industry. He has worked for IBM in EMEA, North America, Asia Pacific and China across a range of roles, combined with 17 years in investment banking, primarily as a trader covering foreign exchange, fixed income and proprietary trading. He is a regular speaker at industry conferences on a wide range of subjects ranging from industry and technology trends. Philip is also a regular contributor to IBM's strategic thinking on its Financial Market strategy.

Martin Jowett is a Distinguished Engineer and Executive IT Architect. He is the GBS CTO for Banking & Financial Markets Europe, CTO for Performance & Availability Engineering, a Member of IBM Academy of Technology, a Fellow of the British Computer Society and is based in the UK.

Andrew Graham is an Executive Architect with the IBM Global Financial Markets team based in London, advising clients on business technology, architecture, technical strategy and innovation as it applies to the wider financial markets industry. With 17 years' experience in consulting, sales and delivery he is certified as a Chartered Engineer, and has worked with sell side, buy side and market infrastructure institutions. He has recently been awarded his Level 3 CISI Investment Operations Certificate.

Prashant Jobanputra is an Executive Partner within the IBM GBS FSS consulting team, where he leads the finance, risk and fraud consulting business. He has worked extensively in the consulting and financial markets industry, including global risk and control positions based in the UK and Asia.

© Copyright IBM Corporation 2015

IBM Sales and Distribution
Route 100
Somers, NY 10589

Produced in the United States of America
September 2015

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at **ibm.com**/legal/copytrade.shtml

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF wNON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

The client is responsible for ensuring compliance with laws and regulations applicable to it. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the client is in compliance with any law or regulation.

Please Recycle

FDW12346-USEN-00