Committee on Payments and Market Infrastructures

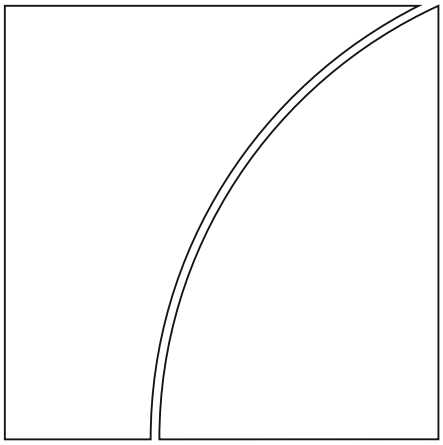

Cyber resilience in financial market infrastructures

November 2014

BANK FOR INTERNATIONAL SETTLEMENTS

This publication is available on the BIS website (www.bis.org).

ISBN 92-9131-988-6 (print)

ISBN 92-9131-989-3 (online)

Contents

1. Introduction

Cyber attacks*[1] against the financial system are becoming more frequent, more sophisticated and more widespread. Given the critical role that financial market infrastructures (FMIs) play in promoting the stability of the financial system, the Committee on Payments and Market Infrastructures (CPMI)[2] has sought to understand the current cyber risks faced by FMIs and their level of readiness to effectively deal with worst case scenarios.

The Committee established a working group (WG; see Annex 2) to analyse the relevance of cyber security* issues for FMIs and their overseers within the context of the *Principles for Financial Market Infrastructures* (PFMIs).*

WG members carried out a stocktaking exercise in the form of interviews with FMIs, their participants, providers and other relevant stakeholders in their respective jurisdictions. The primary objective of these interviews was to better understand FMIs' abilities and perspectives in the field of cyber resilience.* From the interviews, the group learned that: (a) cyber resilience is steadily becoming a top priority for FMIs; (b) FMIs are addressing the risks to broader financial stability posed by cyber threats* to their own systems, notwithstanding the challenges of doing so; (c) FMIs consider a two-hour recovery time objective (2h-RTO*) in the context of an extreme cyber event to be challenging (and may take several years to achieve), but some believe there are a number of feasible solutions that can be explored or may already be in place to minimise recovery times; and (d) FMIs support the regulatory community in providing the impetus for communication and coordinated action in the pursuit of effective solutions.

In addition, although senior management at FMIs increasingly considers cyber resilience a top priority, industry leaders also generally believe that current efforts to move the industry towards the achievement of faster target recovery objectives need to be stepped up given the growing threat to the financial sector. The aim of this document is to describe some of the evolving practices and concepts that FMIs are applying and discussing in their approach to cyber resilience, and also to lay the foundation for the work necessary to strengthen financial stability by enhancing cyber resilience in the FMI industry as a whole.

The report is guided by the assumption that FMIs must be able to recover operations quickly and settle activities by end-of-day even in extreme scenarios. Given the state of FMI cyber resilience and the industry perspectives shared by the FMIs interviewed, authorities concluded that coordinated action and possibly guidance in addition to PFMI Principle 17 may be justified.

[1] Terms marked with an asterisk are defined in the glossary (Annex 1).

[2] Prior to 1 September 2014, known as the Committee on Payment and Settlement Systems (CPSS). Please note that references to reports published before that date cite the Committee's old name.

2. Basic assumptions

The following three assumptions form the basis of this report: (i) the PFMIs constitute the starting point of the analysis; (ii) settlement finality is not affected; and (iii) the 2h-RTO and end-of-day settlement requirements embedded in Principle 17 of the PFMIs are expected to be met.

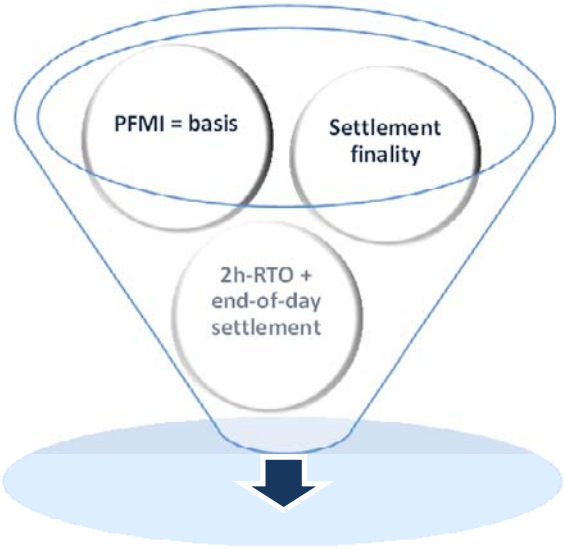

2.1 The PFMIs as the starting point

The PFMIs are international standards for FMIs, developed by the CPMI and the Technical Committee of the International Organization of Securities Commissions (IOSCO), applicable to SIPSs,* SSSs,* CSDs,* CCPs* and TRs.* The overall objective of the PFMIs is to ensure that FMIs promote stability and efficiency in the financial system. The PFMIs are in the process of being implemented in many jurisdictions. Principle 3 requires a sound risk management framework for comprehensively managing legal, credit, liquidity, operational and other risks. Operational risk and governance are specifically addressed in Principles 17 and 2,* respectively. Cyber risk falls within this domain. Any issues or recommendations connected to cyber resilience in FMIs are expected to be handled within the context of the PFMIs.

In line with the PFMIs, the minimum standards for cyber risk management are not expected to vary by type of FMI. However, the specific type of FMI and/or the corresponding impact of a cyber attack on the financial system may influence the specific approach or tools needed to meet those minimum standards, in terms of:

1. whether the FMI contains information that is not available anywhere else and/or whether it registers basic ownership rights (eg in the case of CSDs and certain TRs) that – in the event of loss – would create major ownership issues;
2. the magnitude of the disruption to the financial system that an FMI could generate, as determined by its position in the transaction chain* and/or the size and number of participants it serves; and/or
3. whether an FMI enjoys a (near) monopoly and therefore positions itself as the only option for the fulfilment of its services.

2.2 Settlement finality

The finality of settlement is a legally defined moment, ie the irrevocable and unconditional transfer of an asset or financial instrument, or the discharge of an obligation by the FMI or its participants, in accordance with the terms of the underlying contract.

The finality of settlement is important for the stability of the financial system. Credit, liquidity and legal risks are allocated among the parties to payment and securities transactions based on the principle of finality. The liquidity condition of financial institutions and their customers depends on the certainty of the assumption that transactions that are considered final will remain final. While erroneous data could result from an extreme cyber event, assurance of the finality of those transactions is necessary to maintain financial stability.

The WG concluded that if there were a case where a recipient had no legal right of acquisition based on an underlying claim and it was decided that the transfer order might have to be reversed by an entry of opposite magnitude and effect to offset such invalid or unauthorised transaction, the original settlement by the system would remain protected.

2.3 The 2h-RTO objective is relevant for cyber resilience

As FMIs play a critical role in domestic and international financial stability, the PFMIs explicitly require that an FMI have a business continuity plan that addresses events which threaten to significantly disrupt operations, including events that could cause a wide-scale or major disruption. Although some parts of the PFMI text are written in terms of recovery from physical threats, the need for rapid resumption of key services in response to cyber attacks is equally important. Therefore, Principle 17 on operational risk is also intended to encompass cyber security (see excerpts below).

PFMI Principle 17:

"*An FMI should identify the plausible sources of operational risk, both internal and external, and mitigate their impact through the use of appropriate systems, policies, procedures, and controls. Systems should be designed to ensure a high degree of security and operational reliability and should have adequate, scalable capacity. Business continuity management should aim for timely recovery of operations and fulfilment of the FMI's obligations, including in the event of a wide-scale or major disruption.*"

PFMI Principle 17, Key Consideration 6:

The business continuity plan "*should be designed to ensure that critical information technology (IT) systems can resume operations within two hours following disruptive events*". Moreover, the plan *"should be designed to enable the FMI to complete settlement by the end of the day of the disruption, even in the case of extreme circumstances*".

PFMI Principle 17, Explanatory Note 3.17.13:

"A business continuity plan should have clearly stated objectives and should include policies and procedures that allow for the rapid recovery and timely resumption of critical operations following a disruption to a service, including in the event of a wide-scale or major disruption."

Although FMIs have identified challenges to achieving a 2h-RTO in an extreme cyber scenario, senior managers understand and support this objective. Section 4 of this report discusses in more detail the concepts and practices that some FMIs believe may assist in shortening recovery times.

It is worth noting that a 2h-RTO could involve trade-offs with other aspects of cyber security and resumption. For example, in some cases, ensuring a 2h-RTO may mean that forensic analysis of the attack, needed to preserve the integrity of the evidence collected and to ensure that it can be used effectively in a legal case, cannot be completed as easily or comprehensively as in the case of a long closure of systems. While forensic analysis may be postponed, creating the conditions to perform it post-event is a responsibility that cannot be dismissed.

3. Why are cyber risks special?

The PFMIs identify operational risk (and the cyber risk belonging to it) as one of the core risks confronting FMIs. Because cyber risk is a relatively new, highly complex and rapidly evolving phenomenon, it can be very difficult to manage. Cyber attacks may take the form of persistent malicious action by third parties intent on creating systemic harm or disruption, with concomitant financial losses. It may be extremely hard to determine the extent of an event, how to remedy it and how to recover. The very unpredictability of cyber risk dictates the urgency of having a proper approach in place to manage it.

Over the past several years, cyber threats have emerged as a growing systemic risk to FMIs. There are a number of reasons for this: (i) the role of technology in the provision of financial services has deepened; (ii) the degree of interdependency and interconnectedness between operators in financial markets is very high and growing; and (iii) both attackers and their motivations have become more diverse, bringing fresh threats from unexpected sources. Attackers now include "hacktivists", who seek merely to disrupt activity; cyber criminals motivated by financial gain; terrorists aiming to cause political and financial instability; and nation state-related actors attempting to interfere with or gain access to sensitive information, or to cause systemic instability. The biggest challenge in making FMIs cyber-resilient is managing their complexities and interdependencies by proactively addressing failures, adopting effective resilience techniques, and resolving problems through cooperation.

Attackers are also using increasingly sophisticated methods. For instance, in recent years a new class of intrusion, known as an advanced persistent threat (APT),* has emerged and continues to evolve. At the same time, entry points through which an FMI could be attacked are multiplying and include counterparties, vendor products, employee workstations and rogue employees. Social engineering* often serves as a means to deliver malware* into IT systems (eg spear-phishing*).

Given their sophistication, range of motivations and pervasive scope, cyber attacks can present unique challenges to FMIs' operational risk management frameworks. In some cases, the risk management and business continuity protocols used in the event of physical attacks are ineffective or could actually exacerbate a cyber attack. For example, automated backup systems that may help preserve sensitive data in the event of a physical attack on a head office could be as vulnerable to a cyber attack as the primary system(s), and might in some instances help the malware propagate faster.

4. Integrated approach to cyber resilience

An integrated approach to cyber resilience is a means to ensure the survivability of FMI operations, even if services have to be conducted in a degraded state. Survivable operations are designed to absorb the shock of an attack without systems breaking down totally. In order to cope with the idiosyncrasies of cyber attacks and enable services to resume, FMIs typically follow an integrated approach based on the adoption of a cyber resilience framework developed internally or adapted from a more generic framework – examples being: the NIST framework,* published in February 2014; the World Economic Forum's cyber resilience approach,* published in January 2014; and the 2013 MITRE framework.* External

consultants or auditors are sometimes hired to advise on the implementation of a cyber resilience framework.

Although the frameworks differ in terms of their exact setup and categorisation, an integrated approach to cyber resilience typically covers three broad dimensions:

1. *Scope*: Generally, FMIs' cyber resilience frameworks address a number of scenarios that may result from a cyber attack, including a confidentiality breach, an availability breach and an integrity breach.

2. *Cyber governance*:* The framework covers not just an FMI's IT infrastructure, but also people, processes and communication.

3. *Range of measures*: It is essential for an FMI to apply a wide variety of controls to effectively (i) prevent a cyber attack from occurring, (ii) detect an attempted or successful attack, and (iii) resume services at pre-agreed levels after an attack.

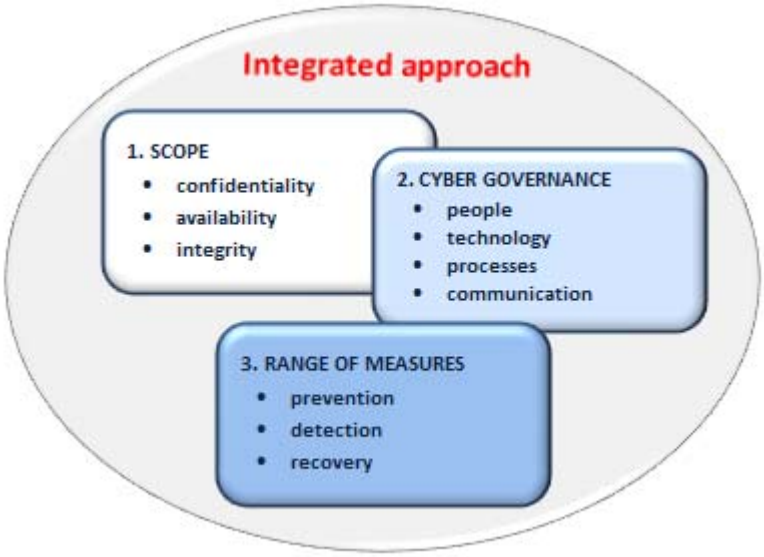

4.1 Scope

Generally, FMIs' cyber resilience frameworks aim to address three broad scenarios:

1. A confidentiality breach, which involves confidential information being stolen.

2. An availability breach, where the services provided by an FMI are inaccessible or unusable upon demand by authorised entities (eg because the channels of communication between an FMI and its participants and other organisations are unavailable) but where the systems per se are still intact.

3. An integrity breach, which is the corruption of an FMI's data or systems affecting the accuracy or completeness of the information and processing methods (and which could also impact the availability of services).

The focus of the majority of cyber attacks continues to be on compromising confidentiality (eg stealing sensitive data) and degrading system availability (eg DDoS* attacks). However, more recently, the risk of attacks impacting the integrity of either the software or the data (or both) of an FMI

has been receiving increasing attention. Three generic risk scenarios are briefly illustrated below in increasing order of severity of impact on an FMI's operations and on the financial system.

The main purpose of this illustration is to highlight the diversity of challenges to the resilience of FMIs in different scenarios. Measures intended to protect against a physical incident (such as local system duplication and remote data centres) may not be equally effective against cyber attacks (see also non-similar facility in Section 4.3.3). This could especially be the case with regard to integrity.

SCENARIO 1	SCENARIO 2	SCENARIO 3
Confidentiality breach	**Availability breach**	**Integrity breach**
• Confidential data stolen through cyber attack. • Ability to provide services not necessarily impaired. • Attack may serve as the initial phase of a more sophisticated scenario. • Can be difficult to recognise and mitigate in a timely manner. • Damage to FMI's reputation.	• Unavailability of services through eg denial-of-service attack. • Impact on eg communication between FMI and participants, support to participants, FMI's updates on availability of services, communication with suppliers (market feeds) and information exchange with counterparties. • Effect of disruption on participants and financial market worsens the longer the downtime.	• FMI's core data or systems are corrupted in a cyber attack. • Integrity of FMI's information or systems no longer trusted. • Backup systems possibly corrupted as well. • Initially, systems may seem to process normally. • Decision should be taken whether to stop service in order to restore systems to a trusted state. • Time to detect and analyse problem could be considerable. • Time needed to restart service delivery on the basis of a clean situation could be substantial. • Impact may be systemic since participants' positions within FMI could be blocked and no longer trusted. • Could trigger a loss of confidence in the financial markets, eg due to disputes or confusion about ownership rights and financial positions. • Possible knock-on effects on other FMIs, participants and their customers and markets, including liquidity and credit effects.

4.2 Cyber governance

FMIs that effectively manage cyber resilience do so in part because they have implemented a comprehensive governance framework. They acknowledge that cyber resilience is not just about information and communication technologies. Rather, it has a broader impact and relevance for these organisations. Four general areas are covered in a governance framework: people, technology, processes and communication.

4.2.1 People

An integrated approach means that an FMI's entire staff – from operational to senior management as well as board level – are involved in the two key components of cyber resilience: security and business resumption. Effective cyber resilience requires that cyber risks be comprehensively addressed within the FMI's risk management framework.

The lead taken by an FMI's senior management is an important factor in cultivating a strong level of awareness of and commitment to cyber resilience throughout the organisation. Generally, FMIs are aware that cyber risk is not only an operational issue but rather an enterprise-wide risk that threatens their viability as going concerns. Cyber resilience is accordingly an enterprise-wide issue, and internal auditors may play a significant role in confirming the efficacy of cyber risk initiatives and policies, and in ensuring that an FMI attaches appropriate importance to cyber resilience.

A consistent feature of FMIs with a sound information security policy in place is a governance structure which ensures that information security is considered in all aspects of the business and a culture that recognises its importance. An integrated approach means that cyber resilience is treated as part of the core business, and is not something tacked onto existing tasks. In such organisations, staff at all levels, including top management, undergo targeted training on a regular basis to increase overall awareness and enhance preparedness to deal with a range of cyber threat scenarios. Organisation-wide awareness is crucial. With a high level of awareness across all employees, it is more likely that a potential victim will report a suspected instance of cyber infiltration and that the appropriate incident management process will be triggered in time to mitigate damage.

4.2.2 Technology

Attackers generally identify vulnerabilities to bypass intrusion detection tools. They exploit the gaps or seams that exist between subsystems and business processes in complex environments. They also benefit from situations where legacy software lacks sufficient security support. Currently, cyber attackers exploit vulnerabilities in an individual institution's IT infrastructure to unleash damage across the financial industry, as extensively reported in the press. As a precaution, most FMIs give prominence to IT in their cyber governance and endeavour to make their systems more resilient, implementing layered cyber resilience measures to counter threats. Examples of IT governance measures, such as access controls, are given in Section 4.3.

4.2.3 Processes

In an integrated approach, the implications of cyber resilience from an operational risk perspective should be properly assessed as part of the decision-making process at board level (covering eg new services, products, IT investments and an FMI's organisational structure). Some FMIs have introduced clear cyber resilience-related processes, including identification of responsibilities and accountabilities. One such process is risk acceptance. As part of operational risk management, it includes input and analysis on cyber resilience and business continuity from relevant staff at all levels, including business units, internal audit, the chief information security officer and the board.

4.2.4 Communication

Given how interconnected FMIs are with their participants, other FMIs, service providers and third-party vendors, effective channels of communication between them are essential. However, information-sharing can be hampered by the difficulty of maintaining trusted relationships with a broad range of security teams, which are often based overseas. Sufficient trust between an FMI's security teams and its counterparties is crucial for them to be comfortable with sharing sensitive information.

Most FMIs appear to be striving to ensure the continuity of their information and communication channels both in normal and stressed circumstances. In the event of a cyber attack, timely communication with stakeholders, including relevant authorities, is critical to resuming operations and preventing the attack from spreading.

4.3 Range of measures

The third dimension of an integrated cyber resilience approach is the range of security measures to be taken, mainly in the area of IT. There is no silver bullet that effectively protects against all cyber attacks. Rather, FMIs are diversifying their investments across different categories (prevention, detection and recovery) of cyber security measures and tools. The breakdown of security measures into prevention, detection and recovery components is not strict, since measures are mutually reinforcing and serve several purposes at once.

For instance, it is common for solutions to combine prevention and detection functions (eg anti-virus software both detects and isolates malicious code). Likewise, steps to enhance recovery can include process segmentation (which can strengthen prevention) combined with frequent checkpoints, validation and reconciliation (which can strengthen detection).

In the past, cyber resilience was focused on prevention, and measures often involved retro-fitting IT security solutions to existing systems. An increasing amount of attention and indeed investment are now being devoted to improving monitoring, detection and recovery capabilities.

These investment decisions are multifaceted. The cyber threat landscape is constantly evolving, and the industry must contend with the complexity and high costs of eventual solutions as well as the increasing likelihood and severity of extreme cyber events. Although committed and aware of the importance of information security, FMI management may, in some cases, face challenges in making the investments that may be necessary.

Some FMIs have hinted, however, that near-term steps can be taken to progress towards making 2h-RTO and end-of-day settlement a reality during an extreme cyber event. The measures necessary are likely to require investments in a combination of prevention, detection and recovery techniques. These three elements, in the context of 2h-RTO, are mutually reinforcing and must be considered jointly. The following sections detail examples of the key practices, concepts and strategies that many FMIs are adopting to build up cyber resilience, reduce service downtime and ensure end-of-day settlement.

4.3.1 Prevention

FMIs recognise that many of the following prevention measures are basic elements of a strong information security programme. They also view these elements as key steps towards a quick resumption of operations. Nevertheless, no FMI today can confidently assume that cyber attacks can be prevented. The security hypothesis is that FMI systems are compromised and that FMIs need to develop intruder detection capabilities (see Section 4.3.2).

Identification

FMIs are improving their understanding of the business context, of the resources that support critical functions and of the related cyber risks, thereby focusing and prioritising their efforts, consistent with their risk management strategy and business needs.

Awareness

Approaches to raising cyber awareness on all levels within an organisation (see Section 4.2.1) include staff training and threat analysis – for example, in the context of combating social engineering – and provide a basis for building and developing an effective security framework.

Defence in depth

Practices commonly referred to as "defence in depth" encompass network security management. This is the process of layering systems and system components and building firewalls, so that if one component is compromised, it does not give the attacker access to another. Internet-facing applications, such as desktop e-mail, are considered to be at greatest risk and are therefore segregated from core system components.

Prevention of malicious activity

Malicious activity may be prevented through the use of anti-virus solutions as well as analysis of web services and infrastructure to identify vulnerabilities that attackers may exploit to inject malicious code. Such analysis includes monitoring and inspection of suspicious web-based e-mail and traffic for

malicious code, DDoS attacks and any attempts by hackers to capture user details. Once suspicious traffic is identified, it can be blocked and action undertaken to nullify the source of the threat.

*Reducing attack surfaces**

An important part of prevention is curtailing the points through which an attacker can gain access to an FMI's network. Practices include limiting the number of internet gateways, whitelisting software and isolating critical parts of the network.

Application development

Prevention can also be instilled during the software development life cycle for applications developed in-house. The development process should include enforcement and testing of secure software coding standards to limit the number of vulnerabilities introduced into production systems.

Testing and application management

Security audit and penetration testing involve the use of advanced analysis and simulated attacks to ensure compliance with security standards and to identify vulnerabilities in existing security arrangements. Penetration testing is regularly carried out by FMIs, both internally and in collaboration with external consultants.

Application management includes application whitelisting and security-related patching. Application whitelisting ensures that only approved applications are installed on servers and workstations, thereby minimising the risk of an attacker installing malicious applications within the operating environment. Patching addresses vulnerabilities in a system by applying updates to applications and operating systems in a timely manner as they become available. The discontinuation of support for a particular application leaves it more vulnerable to attack.

Access control

Access control measures are essential to prevent unauthorised access to data and/or systems. Administrative privileges are limited to those staff members who really need to have them; senior staff are alerted whenever a user seeks privileged access. By minimising the number of users with administrative privileges and following the principle of "least privilege",* this approach aims to limit both the risk of an attack from within an organisation and the number of entry points for an external attacker. As a detection measure, access logs and alerts are used to monitor privileged access and identify unusual activity.

Infrastructure control and development

The design of the IT infrastructure can have significant implications for security management. As controls on the infrastructure are tightened up, preventive and proactive measures are taken. *Virtual machines (VMs)** or *virtual desktop image (VDI)** are technical measures that enable staff to access the desktop which is hosted in a centralised server. As desktop security and data protection are centralised, security patches are easily deployed and granting or denying access to a specific user is more straightforwardly resolved.

VMs can be used to create non-persistence on networks, servers and workstations. A VM can be reset to a known "golden state" to effectively remove any malicious software installed by an attacker. This makes it difficult for an attacker to establish a persistent presence on an FMI's network, and to conduct reconnaissance or propagate through a network as the environment is constantly changing. This process can also aid in recovery.

The deployment of cryptographic defences, such as the encryption of sensitive data (from basic HTTPS protocols to more advanced VPN services), is considered good practice, as it protects the data

from unauthorised modification or access when transmitted through an untrusted environment (eg the internet).

4.3.2 Detection

Prevention is important, but prevention alone is not sufficient. Many organisations make the policy assumption that their defences have been breached and an attacker has already infiltrated their systems. The ability of an FMI to quickly detect and evaluate the scope of an attack is an important aspect of its ability to reduce its recovery window. Most FMIs implement a number of controls to detect anomalous activities in a timely manner while aggregating and correlating data from multiple sources and sensors. In addition, they maintain and test detection processes and procedures to verify the effectiveness of their protective measures.

Monitoring

When data are corrupted, detection, reconciliation and recovery processes tend to differ depending on whether the corruption occurs before or after the data are received and processed by the FMI. Technical monitoring by security or network operations is not enough. Measures need to be put in place to enable stakeholders to detect abnormal transactions (eg anomalies in terms of value, counterparty, time of day) across customers by way of an extra layer of monitoring complementing technical defences. In addition, technicians need to understand an FMI's business day operations, business-critical systems and "normal" system behaviour. Event correlation across monitoring systems helps detect processing anomalies. The implementation of clearly defined event escalation and decision-making processes streamlines an FMI's response capability and reduces the length of time an adversary can maintain a foothold.

Technically based and business-based monitoring solutions intercommunicate using a common taxonomy so that issues are clearly understood and action can be quickly initiated. Having system behaviour monitored by key customers increases the likelihood of abnormal conditions being swiftly reported. Coordination with relevant stakeholders also helps prevent cyber attacks from spreading.

Use of checkpoints

The time required for diagnostics and triage could be reduced through greater application and process segmentation and more intensive use of checkpoints and validation techniques, possibly involving FMI participants. Building such mutually reinforcing designs into systems improves protection and detection but also an FMI's recovery speed.

Other practices

Emerging practices include heuristic* monitoring to detect anomalies, such as abnormal usage of an application or abnormal transaction behaviour on the part of a customer or business partner. FMIs are moving away from flat networks and monolithic applications and introducing segmentation to reduce the "lateral" movement of attackers within IT systems: an example of the latter is the "kill chain"* process developed a few years ago by a security company to track an intruder's movements, which erects barriers at each link in the chain to block attempts to syphon data out of the network. This is now being more widely applied by different organisations.

4.3.3 Recovery

Cyber resilience is the capacity to ensure the survivability of operations, even if services have to be delivered in a degraded state, eg be limited to priority transactions. Survivable operations are designed to absorb the shock of an attack without systems totally breaking down. "Recovery" therefore encompasses both the resumption of activities at a level which is considered "good enough for a certain period of time" and full recovery, ie an eventual return to full service. The following sections set out examples of techniques that are used by FMIs to resume operations in the event of a successful cyber

attack. It may be noted that FMIs do not necessarily apply all the mentioned techniques; in practice, certain combinations of measures are deployed. In addition, practices in this area are constantly evolving.

Recovering operations to a "normal" or "good enough" processing state

Robustness to integrity attacks is important, as an inability to quickly resume operations in a stable state may cause systemic risk and could potentially be transmitted to the wider financial system. Even if recovery as such is quickly achieved, that does not necessarily imply cyber resilience. An FMI that manages to resume operations within two hours may simply be recovering to the vulnerable state which had permitted the attack to succeed in the first place.

With layered technology, it may in some instances be possible to partially resume services – that is, to recover some functionality while still remediating other compromised system components. In the event that intraday recovery of critical components is not possible, many FMIs could extend operating hours beyond the normal end-of-day, on a case by case basis, taking into account linked systems and interdependencies.

FMIs could maintain sufficient operating capacity in their core systems to be able to process a day's volume within a relatively short time frame. Accordingly, unless a cyber incident is so severe that it persists beyond a day, the combination of extended hours and redundant capacity may permit the end-of-day target for completion of processing to be met as long as participants made themselves available to process transactions. If an incident does extend beyond a day, fallback manual processing alternatives might have to be invoked where these are feasible.

Rolling back to the uncorrupted "golden point"

An important step towards recovery is to identify and define when the breach took place, and if the IT environment, data and/or applications have been corrupted in any way. This knowledge is critical in determining an uncorrupted "golden point"* (also referred to as a "golden copy"), from which affected IT environment components, data and/or applications can be restored to the state they were in prior to the attacker's presence.

Effective processes may involve capturing transaction and position data in near real time and storing them outside the system, which could help overcome obstacles to identifying the golden point. Frequent reconciliation against such records could assist in identifying corrupted or fraudulent transactions, either to detect a cyber intrusion or to help establish the golden point. To that end, FMIs could have arrangements in place with an independent third party, or with participants or their customers, to reconcile transactions from the point at which trust is re-established.

FMIs could in this way securely maintain a copy of all raw messages (as distinct from data), periodic independent reconciliations of participants' positions, and periodic snapshots of data instances. A technology that supports this is the "Write Once, Read Many times (WORM)"* device, which can be used as a reference to establish transactions lost between the backup time intervals. And once those transactions are established, FMIs need to have the capability to reprocess them in a timely manner.

Failing-backward or failing-forward

"Failing-backward"* recovery restores the operational environment to the last "trusted" state, ie the state known to be uncompromised. "Failing-forward"* recovery selects a less than fully trusted point from which to continue operations while seeking to restore trust in those IT environment components that remain corrupted. The latter concept is attractive to the extent that it limits service downtime and in that some theoretical work has been done on methodologies to restore "trust on-the-fly".*

Failing-forward recovery, however, is not yet widely recognised as viable. Where it does acquire viability is if failing-backward could not be accomplished in a reasonable time frame. Consequently, the

applicability of failing-forward recovery centres on whether an FMI can afford the downtime required to restore the IT environment supporting its critical operations to a trusted state or whether the FMI should continue to provide critical services in a less than trusted IT environment.

Non-similar facility

The "non-similar facility" (NSF)* seeks to replicate the core functionality of an FMI but using technology different from that used by the primary facility. Accordingly, in the event of a cyber intrusion that has managed to compromise the core systems of the primary facility, an FMI could recommence operations using the non-similar facility, assuming the NSF itself is not compromised. There are examples where different systems covering different functionalities (eg wholesale and retail payment systems) have been adapted to function as each other's backup system in the event of an integrity incident. Generally, an NSF does not provide the full functionality of the primary system, but does allow the most important transactions to be pursued.

An NSF could create a backup of an FMI's data to facilitate resumption of operations after data corruption, with services running independently of the FMI's primary system (and hence remaining uncorrupted). This may require an independent communication channel. One possibility could be for an FMI's participants (or other holders of its data) to send their data directly and separately to two different facilities (eg the primary facility and the NSF) so that each could serve as a base for failing-backward. NSFs could also provide a failing-forward option, eg a cold site fires up and starts processing from a defined point onwards.

However, NSFs may, depending on their design, entail increased operational complexity and cost, which may undermine the resilience of the mechanism and reduce its agility. Also, they may give rise to new forms of attack. NSFs are not a silver bullet, but may, subject to the specific demands of an FMI, serve to (at least partially) improve the survivability of an FMI's operations.

5. Sector-wide considerations

Given the extensive interlinkages and interdependencies in the financial system, adequate cyber security practices at the level of the FMI do not necessarily ensure cyber resilience in the markets it serves. In particular, the markets' resilience is dependent not only on that of the FMI, but also on that of interconnected FMIs, of critical service providers and of the participants. A cooperative or coordinated approach to cyber security is therefore important, in particular to share information on threat intelligence and forensics.

In most jurisdictions, there is an active intra- and inter-industry dialogue on cyber resilience issues. Domestically, these dialogues typically involve FMIs, other financial institutions and critical service providers in both financial and non-financial industries, and in both the public and private sectors. In some jurisdictions, authorities have established trusted network and information-sharing arrangements to promote a best practice resilience approach.

Achieving market-wide timely recovery of operations also imposes challenges on testing. Traditional isolated testing implicitly assumes that all other players operate as usual. By removing that hypothesis, integrated exercises, involving other FMIs and participants, take complexity to a level very few FMIs have experience of.

Nevertheless, the full recovery or survivability of one entity may very well depend on another entity's ability to survive or even rely on data kept in their participants' databases. In such cases, coordination is warranted to achieve the ultimate goal of end-of-day settlement. Furthermore, recovery might require extended hours to enable settlement by end-of-day. The ability to extend hours varies for each FMI, but the common objective is to meet the deadlines imposed by other infrastructures or

counterparties. Coordination and communication would be helpful in exploring when or how these upper limits might be made more flexible in order to enable settlement on the day of disruption.

It has to be noted that many FMIs operate in multiple jurisdictions and across multiple time zones, which may make communication and coordination difficult. As a result, regulators, supervisors and overseers are expected to cooperate, both domestically and internationally, as appropriate, in promoting the safety and efficiency of FMIs (PFMI Responsibility E). In this respect, however, several challenges at the organisational, legal and confidentiality level need to be addressed.

For example, FMIs with global operations have expressed concern over their inability to share classified information within their own organisations. Law enforcement or intelligence agencies may provide classified information to FMI management with security clearances in the country of the authorities. That information, however, cannot be shared with anyone who does not have a security clearance from the country of the authorities that provided the information. FMIs are encouraging authorities to take an active role in helping FMIs to address these challenges through coordination efforts between the public and private sectors.

Annex 1: Glossary

2h-RTO	Recovery time objective of two hours.
advanced persistent threat (APT)	A category of cyber threat directed at a specific target, which is typically business-related or political. Attackers display a high degree of stealth over a prolonged phase of the targeted operation. The attack objectives typically extend beyond immediate financial gain, and compromised systems remain in service even after key systems have been breached and initial goals reached.
attack surface	Refers to the vulnerabilities in the broad categories (software, hardware, network and human) which allow an attacker to enter a system and potentially cause damage to the FMI. A smaller attack surface means the FMI is less exploitable and an attack less likely. However, reducing attack surfaces does not necessarily reduce the damage an attack can inflict.
BCBS	Basel Committee on Banking Supervision.
CCP	Central counterparty.
CSD	Central securities depository.
cyber attack	An attempt to infiltrate that might result in a circumstance or event having an actual adverse effect on cyber security.
cyber governance	A structured approach covering the four broad categories of people, technology, processes and communication to enhance an FMI's cyber resilience.
cyber resilience	The ability to anticipate, absorb, adapt to and/or rapidly recover from disruption caused by a cyber attack.
cyber security	A broad concept for which there is no consensus definition. In this report, the term refers to strategies, policies and standards encompassing the full range of threat reduction, vulnerability reduction, deterrence, international engagement, incident response, resilience, and recovery activities and policies regarding the security of an FMI's operations.
cyber threat	A circumstance or event with the potential to intentionally or unintentionally exploit one or more vulnerabilities in an FMI's systems resulting in a loss of confidentiality, integrity or availability.
DDoS – distributed denial-of-service	In a DDoS attack, an adversary directs a flood of illegitimate service requests to overwhelm the targeted computer or network in an attempt to make the resource unavailable to intended users, thereby seizing control of multiple systems by infecting them with malware.
failing-backward	A recovery procedure that restores the operational environment to the last "trusted" state known to be uncompromised.
failing-forward	A recovery procedure that selects a less than fully trusted ("good enough") point from which to continue operations while seeking to restore trust in the IT environment components that remain corrupted.
GCE	G10 Group of Computer Experts.
golden point	Also referred to as a "golden copy", the point from which affected IT environment components, data or applications can be restored to the state they were in prior to the attacker's presence.
heuristic	The term generally pertains to an experimental (often trial and error-based) method of problem solving, especially when an algorithmic approach is impractical. Applied to cyber security, it refers to the ability to adapt frameworks of rules to actively monitor log-based information sources that can identify malicious code and/or abnormal behavioural patterns in systems and services.

kill chain	Borrowed from military jargon. In the cyber security context, the term refers to the structured sequence of an intrusion (eg define target, build or acquire tools, research employees, test for detection, deploy, launch initial intrusion, etc), the identification of which informs actionable security intelligence by contextualising security events around the attacker and/or the attack.
least privilege	Refers to the principle of limiting access to the minimal level that will allow normal functioning. Applied to employees, the principle of least privilege translates into giving people the lowest level of user rights that they can have to perform their duties and functions.
malware	Short for malicious software, is any software used to disrupt the normal operation of an information system that adversely impacts its confidentiality, availability or integrity.
MITRE framework	MITRE Corporation, *Cyber resiliency assessment: enabling architectural improvement*, Deborah Bodeau and Richard Graubart, May 2013.
NIST framework	National Institute of Standards and Technology, *Framework for improving critical infrastructure cybersecurity*, February 2014.
NSF – non-similar facility	Replicates the core functionality of an FMI but using technology different from that used by the primary facility.
PFMIs	*Principles for Financial Market Infrastructures*.
PFMI Principle 2	Governance: *"An FMI should have governance arrangements that are clear and transparent, promote the safety and efficiency of the FMI, and support the stability of the broader financial system, other relevant public interest considerations, and the objectives of relevant stakeholders."*
SIPS	Systemically important payment system.
SSS	Securities settlement system.
social engineering	In the context of cyber attacks, refers to the methods used to gain the trust of a victim and manipulate them into providing confidential information, downloading attachments, sending money or engaging in similar acts.
spear-phishing	In contrast with conventional phishing, where numerous e-mails are sent out randomly, spear-phishing targets e-mails at selected groups of people with something in common. As with conventional phishing, potential victims are often asked in the e-mails to click on a link that will direct them to a genuine-looking website which asks for confidential and demographic information. Alternatively, clicking on a link will download malicious codes or malware to the victim's computer.
TR	Trade repository.
transaction chain	In this report, the term refers to the sequence of events and linkages covering the entire life cycle of a financial transaction, involving multiple entities, from origination to final settlement.
trust on-the-fly	An approach to system recovery which aims to restore trust in all, or at least a minimum set of, IT environment components while the system remains in operation.
VDI – virtual desktop image	The image of an individual's user interface in a virtual environment. The VDI is hosted on a centralised remote server rather than locally. It is also separated from the physical machine to provide an isolated operating system for users.
VM – virtual machine	A software-based implementation of a machine (eg a computer) that executes programs like a physical machine.
World Economic Forum's cyber resilience approach	World Economic Forum, *Risk and responsibility in a hyperconnected world*, January 2014.
WORM – Write Once Read Many times	A data storage device in which information, once written, cannot be modified.
WPSI	GCE Working Party on Security Issues.

Annex 2: Members of the working group

Chairman (Netherlands Bank)	Coen Voormeulen
Reserve Bank of Australia	Ashwin Clarke Mark Manning
National Bank of Belgium	Nikolaï Boeckx Jorke Kamstra Yves Vandenbosch
Bank of Canada	Christian Belisle
ECB	Frans Rijkschroeff
Bank of France	Claudine Hurman Clement Martin
Deutsche Bundesbank	Sylvia Tyroler
BaFin (German Financial Supervisory Authority)	Christoph Ruckert (representing IOSCO)
Reserve Bank of India	Kashiap Balakrishnan
Bank of Italy	Luigi Sciusco
Bank of Japan	Mitsu Adachi (representing the BCBS)
Bank of Korea	Heejun Yoo
Bank of Mexico	Victor Manuel De La Luz Puebla
Netherlands Bank	Ewoud van Bentem (WPSI Chair) Raymond Kleijmeer Bram van der Meulen (WPSI member)
Bank of Russia	Nikolay Geronin[1] Dmitry Krutov Andrey Kurilo
Sveriges Riksbank	Pär Karlsson
Swiss National Bank	Frédéric Bos Marco Cecchini
Central Bank of the Republic of Turkey	Cemil Ulu Özgür Şanlı
Bank of England	Roz Horton Freddie Hult Ed Kelsey
UK Financial Services Authority	Farrukh Nazir (representing the BCBS)
Board of Governors of the Federal Reserve System	Ken Buckley (GCE Chair) Jeffrey Marquardt Stuart Sperry
Federal Reserve Bank of New York	Lawrence Sweet
IOSCO Secretariat	Rohini Tendulkar
Bank for International Settlements	Emanuel Di Stefano Bezerra Freire Klaus Löber Can Bülent Okay

The group's work has also benefited significantly from the contributions of Peter Kah (Deutsche Bundesbank, WPSI member), Paul Biles (Bank of England), Carlos Conesa, Pan Ng and Ayn du Bazane (CPMI Secretariat).

[1] The working group members duly respect the memory of the late Nikolay Geronin.