SEPTEMBER 2021

CYBER THREATS AND DATA RECOVERY CHALLENGES FOR FMIS

This paper has been produced by firms participating in the industry working group ("IWG") sponsored by the CPMI-IOSCO Working Group on Cyber Resilience, including representatives from DTCC, Euroclear, LCH, the Federal Reserve Bank of New York, TMX Group, and the Reserve Bank of Australia.

Although the CPMI-IOSCO Working Group on Cyber Resilience has sponsored this IWG, this paper represents solely the views of members of the IWG on industry themes and does not represent the views of the CPMI, IOSCO, any of their member organisations, or the firms who participated.

While the IWG does make recommendations with respect to enhancing available tools for data recovery, reconciliation, and replay, the IWG is not suggesting that individual FMIs implement all tools identified. Rather, it suggests that firms should evaluate or adapt solutions that work best for their operating/technology model and potentially implement items on an appropriate timeline.

DTCC
ADVANCING FINANCIAL MARKETS. TOGE

RESERVE BANK
OF AUSTRALIA

euroclear

LCH AN LSEG BUSINESS

TMX THE FUTURE IS YOURS TO SEE.®

26751-WK0922021

EXECUTIVE SUMMARY

Technology drives how markets operate and innovate, and for the global marketplace, reliance on data for service offerings is critical. In fact, the operation of Financial Market Infrastructures ("FMIs") is based on the use and trust of data, and to perform effectively, FMIs must strive to keep their transaction and position data intact. However, there is no standard approach to identifying the types of data that need to be protected, nor the manner in which that data should be protected.[1] In 2019, CPMI-IOSCO sought closer engagement with the industry on this important topic by creating an independent Industry Working Group ("IWG") to evaluate how this issue impacts FMIs. This discussion paper is the result of the IWG's analysis.

An underlying challenge for FMI members of the IWG[2] is that, due to the interconnectedness of the global financial industry and the large number of entities in the ecosystem, the introduction of a data issue, which circumvents or passes through existing detective and preventative controls, could result in contagion effect. This contagion has the potential for widespread service disruptions and unexpected consequences. When facing a cyber-attack, traditional data replication strategies designed for physical or non-cyber disruptions have the potential to spread corrupted data to backup databases, including those within data bunkers and backup data centres. To tackle this challenge, the IWG sought to identify available tools to address data recovery and validation issues, draw out key lessons and principles for use of those tools, and identify areas that would most benefit from further industry collaboration. These objectives are particularly important as the cyber threat vector is quite different from other environmental disruptions. Additionally, elimination of a cyber threat may require a more complex and coordinated response to contain and eradicate the threat.

The IWG identified five key themes that influenced deeper analysis:

- While the two-hour recovery time objective ("RTO") documented in regulatory guidance remains a target objective, when dealing with a data integrity issue, there is a trade-off between speed of recovery and accuracy of recovery. The actions available depend upon the FMI's individual legal and operational environment.
- Recovery capabilities of existing systems were typically designed with physical and non-cyber outages in mind and may not be as effective in maintaining data integrity in the face of a cyber-attack.
- Interconnections between firms increase the potential impact of a data integrity compromise across the industry.
- Recovery from a data integrity breach requires a high degree of trust in the available backup data copies as well as coordination within the settlement ecosystem.
- When considering the recovery objective, the definition of critical services can vary across FMIs and across scenarios.

[1] Organizations like Sheltered Harbor are working with the industry to help formally define frameworks for financial data protection via defined file formats and protective controls.

[2] The members of IWG represented Real Time Gross Settlement ("RTGS") operators, Central Counterparties and Central Securities Depositories.

The analysis conducted by the IWG demonstrates that effective solutions and tools vary from firm to firm.[3] At the individual FMI level, there is no single tool that provides solutions for recovery against all data integrity[4] scenarios, and as firms typically use a number of tools as part of their resilience design, no single best practice offering could be drawn from the analysis. Additionally, the effectiveness of each tool varies according to the scenario's cause, the criticality of the event and the technologies in use, requiring FMIs to adopt a toolkit approach.

The analysis concluded that the footprints of many business systems have grown over time and provide related services with internal integration points that create dependencies. It is possible that some interdependencies could enhance protection from certain types of hardware and software failures, but they are more likely to propagate data corruption and complicate data recovery processes. This complexity contributes to the need for flexibility in the setting of recovery objectives. For some systems, substantial modernisation of solution architecture, including software and hardware, may be required to further enhance recovery capabilities.

The analysis also confirmed that some tools may not be feasible for all FMIs. Notably, the use of a "non-similar system"[5] as a resumption tool was found not to be realistic due to the high cost of development and maintenance, as well as the high degree of operational risk when such a tool is used.

In forming its conclusions, the IWG was able to identify potential tools for use in major data corruption scenarios. Specifically, these tools would allow for the storage and recovery of critical data,[6] outside of an FMI's on-premises systems, and were seen as an area of opportunity. Additional recommendations based on the analysis suggest that there is tremendous benefit in wide-scale industry coordination as, due to the interconnectedness of the marketplace, the most severe scenarios will likely impact many entities, making recovery in isolation impossible. Relatedly, the need for industry exercises to test the feasibility of procedures across all data events is required. Clear guidelines for minimising contagion as well as providing support for resilience against concentration risk can also help increase overall industry resilience.

[3] The IWG analysis used 'heat maps' to identify existing and potential data protection and validation tools as well as assess their effectiveness against categories of scenarios. Tools were assessed on a range of factors, keeping the CPMI-IOSCO Principles for Financial Market Infrastructure ("PFMI") (see https://www.bis.org/cpmi/publ/d101a.pdf) and associated cyber guidance top of mind.

[4] While this paper focuses on data scenarios caused by cyber events, the tools identified within this paper could also be used in data scenarios caused by non-cyber events; e.g. system bugs.

[5] A non-similar system is generically described in Section 6.3.1 of the CPMI-IOSCO Guidance on cyber resilience for financial market infrastructures, which states FMIs should consider the "possibility to resume critical operations in a system that is technically different from the primary system or in a system that performs those operations and completes settlement in a non-standardised way."

[6] For purposes of this paper, critical data is defined to be data that is necessary to provide services to clients at expected levels.

PURPOSE AND OBJECTIVES

In 2019, CPMI-IOSCO sponsored the establishment of three "IWGs" to further investigate key cyber challenges for FMIs. The charter of these industry-led groups facilitates closer engagement between the industry and CPMI-IOSCO to advance the cyber resilience of FMIs.

This paper is the conclusion of the findings of the Data Protection and Validation IWG, which had the following objectives:

1. Categorise different types of data that need to be protected and the potential impacts of a data event;
2. Document current practices and challenges with respect to data recovery options for a range of high-level data scenarios;
3. Identify leading and emerging practices with respect to data protection[7] and validation methodologies, as well as potential areas of strength and weakness; and
4. Identify areas of focus for future industry collaboration, e.g., the promotion of reconciliation and replay capabilities within the industry.

[7] This objective was intended to cover data protection to the extent that these techniques could make an FMI's data recovery steps in a cyber event easier. The working group did not evaluate tools used to ensure data could not be altered during business processing.

APPROACH

To address the first objective, four types of data that could be the subject of a cyber-attack were identified, as described in the table below.

DATA TYPE	DEFINITION
Configuration Data	Information that is required to operate technology including system settings, indexes and user configurations.
Application Data	Source code, processing jobs and scripts that, when compiled, create a system that is used for business processing.
Business Transactional Data	Any transactional information that is accessed, used or modified as part of a business process. With respect to FMI processing, transactional business data is data that initiates a movement of cash or securities between accounts.
Business Reference Data	All data/information that is not transactional in nature and that is required for the FMI to conduct its business. Examples of this type of data include, but are not limited to, counterparty and security identification, pricing and calendar information needed to either record or settle a transaction.

Table 1 – Data Types

While it may be commonly understood that transactional business data is the foundation of how an FMI operates, an impact to any of the four categories data types in Table 1 would be disruptive and affect an FMI's ability to provide services. The data types were also evaluated separately because the impacts to each type could potentially be managed in different ways. For example, since the data types change at different intervals, the recovery options might need to be more dynamic for data that changes more frequently. Configuration Data is the most static data type with changes occurring infrequently, while Business Transactional Data is the least static data type, with updates occurring potentially by the millisecond.

Each data type was then overlaid onto a framework that leveraged an impact-based scenario structure to better understand an FMI's intraday, end-of-day, next day and multiday resiliency capabilities following a cyber event. At the highest level, two types of impact scenarios were evaluated: Impactful ("IMP") scenarios and Extreme but Plausible ("XP") scenarios. IMP scenarios assume that the event has impacted data at a single data centre and therefore is localised. As a result, IMP scenarios can typically be resolved in the same business day through a failover to a secondary data centre. Examples of IMP scenarios are those that have minor impacts to a production service, where the impact to data is limited, or incidents where the response is well understood.[8] Alternatively, XP scenarios assume that operations are impacted in such a way that recovery could occur either by the end of the business day or fall into a subsequent business day or contagion that has manifested itself in multiple, if not all, data centre regions. Examples of XP scenarios include ransomware affecting masses of business data at all data centres, and severe corruption of data whereby an organisation does not trust the output of its business processes.

[8] Use of the term Impactful is not intended to introduce another classification within the industry lexicon. Rather it reflects the need for spectrum of scenarios and differentiates from Extreme but Plausible.

These scenarios were applied for both data availability[9] and data integrity[10] events.[11] Subsequent to this categorisation, a set of recovery, reconciliation and replay tools were identified for each data type and documented in a heatmap. The heatmap assesses the feasibility of the tools against each scenario and includes an indication of whether the tools are readily available to FMIs today.

Details of each of these steps, including the definitions used, are provided in Appendix A. The heatmaps are included in Appendix B.

ANALYSIS AND RESULTS

Four heatmaps were created in the analysis; one for each data type. An analysis of the heatmap results reveals both how effective the identified recovery, replay and reconciliation tools are with respect to a disruptive event, as well as whether the tools are used in practicality by the IWG members.

A horizontal review of the heatmap highlights that the potential effectiveness of tools is generally agnostic to the type of data event. In other words, if a tool is not effective in an integrity event, it will be similarly ineffective in an availability event. It should be noted, however, that XP scenarios typically render certain tools ineffective (e.g. failover to an out-of-region data centre) based on the assumption that data is either corrupted or unavailable at all sites.

A vertical review of the heatmap shows the potential coverage that a specific tool provides under the specific scenario.

For reference, and as an example, an excerpt of the Business Transactional Data heatmap (Tools only used in the Recovery stage) is shown below. In this case, tools that are ineffective in XP Integrity scenarios are typically also ineffective in XP Availability scenarios (horizontal review), while tools identified for the IMP scenarios (vertical review) are typically effective.

[9] Data availability incidents impact a firm's access data within its databases. Examples of such events are malware or ransomware that make the data usable.

[10] Data integrity incidents impact a firm's data in such a way that the outcome of processing such data is unexpected and cannot be trusted. Examples of these incidents include the addition of an extra digit to each cash or securities transfer.

[11] The IWG notes that the tools identified could be used for non-cyber events as well, however, these types of disruptions were not the focus of the evaluation and are not covered in detail in this paper.

TOOL FEASIBILITY RATING	TOOL AVAILABILITY RATING
Feasible	Common
Potentially Feasible	Non-Common
Not Feasible	Rare

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Recovery	TBD1 - Fix data/surgery, rather than recover	Achieve correction of data (rather than recovery) in such a way that isolates the issue and corrects it without further impact to operation.					
	TBD2 - Input correcting or reversing transactions	Use existing mechanisms to send a transaction into the system which acts in a manner such that the original transaction is no longer valid. (Note: volumes may impact the effectiveness of this tool).					
	TBD3 - Failover to production instance at the backup site	Recover data from another production instance located at a backup site and resume operations.					
	TBD4 - Restore data from back-up site to primary site	Recover from a (near) real time copy of business transactional data that was sent to a separate data centre.					
	TBD5 - Recover data from internal backup copy (data stored in data centre environment)	Recover data from an internal backup copy stored in the data centre environment at any site. These can be current or historical versions of business data.					
	TBD6 - Recover from asynchronous backup to separate immutable database (i.e. data bunker, cloud)	Conduct "one time write" backup stored in highly secured and dissociated environment, outside FMI infrastructure (could be cloud based).					
	TBD7 - Recover data from participant or network operator records	Leverage ecosystem contribution to restore services and capabilities.					
	TBD8 - Recover data from buddy FMIs or non-similar systems	Recover business data from a trusted third party or FMI's non-similar system.					

Table 2 – Transactional Business Data Heat Map Excerpt

From a summary perspective, below are the results of the tool feasibility for XP scenarios. When compared to the diagram on the availability of the identified tools at FMIs participating in the IWG, while the majority of tools identified are seen to be Potentially Feasible, in reality, these tools have not been implemented at IWG firms for several reasons, including differing architectural design at the FMIs, the complexity of the solution to implement and test, and the speed in which corruption replicates between data centres or offsite data copies.

TOOL FEASIBILITY SUMMARY (XP SCENARIOS)

Tool Feasibility Summary	
Feasible	
Potentially Feasible	
Not Feasible	
N/A	

	Recovery	Reconciliation	Replay
Application Data			
Business Data			
Reference Data			
Configuration Data			

Table 3 – Tool Feasibility Summary

TOOL AVAILABILITY SUMMARY (ALL SCENARIOS)

Tool Feasibility Summary	
Common	
Non-Common	
Rare	
N/A	

	Recovery	Reconciliation	Replay
Application Data			
Business Data			
Reference Data			
Configuration Data			

Table 4 – Tool Availability Summary

Insights

A close, holistic review of the heat maps revealed the following:

- FMIs have a wide range of tools (as described further in Appendix A) in their tool kit to recover from a data event.
- If a tool resolves an integrity scenario, it will typically work for availability scenarios as well.
- IMP scenarios are generally well covered, and in some instances, the tools are effectively part of business-as-usual processes.
- Although tools exist to address XP scenarios, more attention should be given to the identification of additional areas of opportunity, as the current tool set may not successfully manage this type of incident in an expeditious manner, particularly for transactional business data.
- Configuration data is not a focus for the industry at this time because the tool set is the smallest and the data is the most static.
- While the Cyber Guidance suggests it as an option, the use of non-similar systems for data recovery are not practical.

Additional insights with respect to the heat maps include:

1. **The implementation and availability of the tools depends on a wide range of factors -** The feasibility of adopting certain tools not only depends on the varying technology stacks used by FMIs, but also the rules and regulations of each of the jurisdictions to which they would be applied. This means that, while a tool may seem possible to implement on the surface, many other considerations must be evaluated before it can be put into practice. As an example, the storage of backup copies of data with external providers whose services are offered outside the region may not be feasible due to regulatory requirements to house data within certain geographical regions, potentially resulting in reduced resiliency due to localisation. Additionally, the operational risk associated with implementing such a strategy, particularly with respect to data privacy and third-party risks, may make the tool unpalatable.

2. **The trade-off between speed of recovery and accuracy of recovery, which must be evaluated real time in a cyber event, is important –** While a two-hour RTO remains a target objective for the scenarios analysed in this paper, post-event it is more important to focus on an orderly, safe and controlled recovery, including those scenarios which impact the larger ecosystem, rather than targeting the two-hour recovery objective. FMIs weigh these two risks differently depending on their individual legal and operational environments. Whereas some FMIs can consider resumption of settlement activity with a degree of data loss, or while reconciliation of earlier transactions is ongoing, other FMIs are not able to resume activity until all data is restored and reconciled. Additionally, the action of identifying the point in time at which the FMI is comfortable that its data is accurate could, in and of itself, take considerable time to identify under some conditions. Placing a disproportionate weight on meeting a RTO may also increase other risks, such as resuming operations before the contagion has been fully isolated and eradicated. In most scenarios, where extensive reconciliation and data repair might be necessary to identify the status of all transactions, a two-hour RTO as set out in the Cyber Guidance will be significantly challenging, even with considerable investment in recovery processes and technology. Specifically, in IMP scenarios, the objective of achieving recovery and final settlement on the same day is generally considered to be a more realistic and useful objective.

3. **Tools that are effective in IMP scenarios may not be as effective in XP scenarios –** The heatmaps identify a wide variety of tools that could be leveraged for recovery, reconciliation and replay of data in a cyber event. However, not all tools are equally effective, and some would not be feasible in certain scenarios. By definition, XP scenarios are ones in which data is impacted across all data centres. Thus, in these types of scenarios, all on-line data copies stored within an FMI's environment would have been affected by the event. Data replication patterns exacerbate this issue, causing the corrupted data to spread. Examples of tools that may not be as effective in XP scenarios include: the isolation and correction of corrupt data, inputting correcting or reversing transactions,[12] and client reconciliation mechanisms.

4. **The use of non-similar or buddy systems, as suggested in the Cyber Guidance, is not practical –** The heatmap reveals that, although the use of a non-similar system or a buddy FMI could be effective, the use of such a tool to resume business is not practical. Significant costs and operational complexities would be created by such a tool due to several factors. These factors include, but are not limited to, (i) the difficulties associated with drafting contractual relationships

[12] Some identified tools would not be available to all FMIs regardless of scenario type, especially where such tools would put settlement finality into question.

and data sharing agreements with external parties that safeguard compliance with existing regulatory requirements while meeting the needs of both parties, and (ii) the build and maintenance cost associated with making sure that the differing systems process in the exact same way. These multiple systems would also have different vulnerabilities that would need to be addressed, which could increase operational risk. In addition, the switch between systems would have to be transparent to the client, which may reduce some of the resiliency benefits. Given these considerations, a non-similar system could be costly and introduce more risk to the ecosystem, rather than promote service continuity.

5. **Industry focus should be given to those tools that provide the largest coverage -** Since data types have differing lifecycles, the storage and recovery options identified in the heat map can differ between them. Firms should prioritise the use of tools that are applicable to several data types, including golden and immutable storage copies. Use of other entities in the financial ecosystem to recover data would be difficult when there is a large data impact[13] as FMIs typically have a large client population who play varying roles within the industry. In a pervasive data event, the coordination among clients and other parties required to reconstruct a business day is complex and may not be scalable. Additionally, not all clients have the same capabilities to retransmit data, particularly in instances where the data may have been previously processed. To complete the business day, large sets of transactional and reference data would need to be processed in an expeditious manner. Potential areas of investigation could be to develop a business rules-based mechanism for clients to be able to resubmit transactions,[14] or to build a methodology whereby previously provided client output could be read into systems to update positions behind the scenes.

6. **Interconnections between firms increase the potential impact of data integrity scenarios across the industry -** A large number of external system connections have developed over time between firms in the settlement ecosystem. Third parties are also often a key source of data entering an FMI's ecosystem. These connections generally require fast and reliable data transmission under the assumption that the information source is accurate. When a data integrity breach occurs, these connections have the potential to quickly propagate data corruption to multiple sites, magnifying the scope of the corruption. This could also spread the corruption to potential secondary sources of accurate data, complicating and lengthening the reconciliation process. The same is true where the data repair is inaccurate. In addition, the settlement ecosystem is using, in some cases, the same key suppliers to perform critical activities. As a consequence, this may be a source of a concentration risk, as impact on these key suppliers will likely be cascaded to more than one FMI. Partnerships among FMIs and relevant external parties will be beneficial to identify and deploy response and recovery plans for this type of scenario.

7. **Recovery from a data integrity breach requires a high degree of trust and co-ordination within the settlement ecosystem -** Given the interconnectedness of systems and potential for corruption propagation, data recovery and reconciliation processes may require the exchange of sensitive information between a large group of diverse firms. In many circumstances, participants in the ecosystem are competitors and may face institutional constraints, including potential legal issues,

[13] In the worst-case scenario, members of the IWG agreed that the restoration of service is dependent upon using a prior snapshot/backup copy of data for recovery and then using the ecosystem to try to replicate the business day. This type of recovery is not ideal given the differing capabilities of those involved and the larger implications on the global marketplace.

[14] A business rules-based mechanism would provide FMIs and clients the ability to submit isolated transactions as determined through the use of different business driven queries. For example, the transactions could be a set of identified instructions submitted against a specific client or transmitted within a certain time period.

to collaborating with their peers. Trust and co-ordination has become an issue across the entire event lifecycle: before the event, there is a reliance on clients and third parties to maintain secure systems and data, particularly where third sites or data bunkers are used; during the event, trust and co-ordination are required for information sharing to identify and resolve the issue; and for data recovery, they are needed to source reliable data for repair and reconciliation. A further complication can arise for an FMI when considering the appropriate timing of reinstating a participant that has been suspended due to a compromise.

8. **When considering the objective for recovery of critical services, the definition of critical services can vary across firms and across scenarios -** FMIs prioritise recovery of their "critical services" in accordance with the Principals for Financial Market Infrastructures ("PFMIs").[15] There is no common definition of what a critical service is, and the interpretation of critical services can vary between FMIs and also across scenarios for a single FMI, depending on factors such as the time of day and day of week. Rather than trying to define changeable critical services, the IWG saw more potential in taking a risk-based approach to determining criticality based on prevailing circumstances.

[15] See https://www.bis.org/cpmi/publ/d101a.pdf.

Areas of Opportunity

With respect to the third objective, the IWG was unable to make a single recommendation for practices that would best suit the needs of all FMIs. However, some technologies do offer more promise than others with respect to enhanced cyber resiliency. A few of the options have been reviewed and are noted below, with a high-level summary of strengths and weaknesses.

PRACTICE	STRENGTHS	WEAKNESSES
Use of third-party storage sites or data bunkers	Allows for an off-premises data copy that should not be impacted by the event, given its distance from the data centre.	Introduces operational risk with respect to: Need to connect to the site in a way that does not allow for corruption to replicate to the site ("air-gapping"). In the case of third-party operation of the site, obtaining sufficient comfort of the third-party firm's cyber posture as well as executing and monitoring service-level agreements can be challenging.
Introduction of new technologies (distributed ledger, etc.)	Allows FMIs to rethink how their processes and technology work and start from a greenfield/clean slate.	Not proven in grand scale to enhance resiliency posture. Implementation and adoption take time.
Introduction of reconciliation tools	Provides a mechanism for clients to understand what is currently on the books and records of the FMI. Could be run on a schedule or an ad-hoc basis.	Reconciliation of data needs to be carried out at multiple levels, e.g. original transactions received, post transaction processing, and intraday/end of day positions. Adoption of automated tools by smaller clients is not realistic.

RECOMMENDATIONS AND OTHER OBSERVATIONS

The results of the analysis and associated heatmaps reiterate that recovery from a cyber event is complex, highlighting the practical difficulties of a two-hour cyber recovery for every scenario. A cyber event typically presents itself as a system impact first, which necessitates significant research and forensics to understand the severity of the incident. This analysis can take much longer than two hours, depending on the facts and circumstances surrounding the event, which stresses the importance of strong preventative measures on the part of institutions. The goal of all FMIs should be to recover in such a way that is, first and foremost, safe and conducted as swiftly as possible to avoid introducing instability into the marketplace. While the decisions to invoke recovery tools are not easy ones, given the evolving cyber landscape, the associated threats, and the factors that may be unknown until such an event occurs, it is critical that market participants focus on solutions that provide increased confidence within established risk tolerances.

Accordingly, the IWG makes the following recommendations based on its analysis:

Recommended actions for FMIs

I. Focus on the tool set that is the most harmonised with the FMI's objectives – Each FMI should identify tools that are attainable from a design perspective and focus on the implementation of those that provide the most coverage.[16] In some instances, the tools identified could be viewed as normal operating procedure, either by regulation or by current business practice, but attention should be given to those that provide the largest coverage with minimal impact to the current service offerings. FMIs could leverage the heatmap approach used by the IWG to assess their current capabilities and identify potential gaps.

II. Define logical restore points[17] – FMIs should work with their participants and the larger community to identify restore points that make sense for their business. Individual businesses may have the opportunity to establish, in agreement with their participants, distinct points in time where the community can return in the event of a cyber incident. This would allow the FMI to assess whether it should focus on a rapid resumption of operations or work to identify the last known good point of data. The restore point would then become the logical point in time for the business to restart, allowing participants to move straight into the reconciliation portion of business resumption.

III. Understand legacy technology – Given the continual evolution of technology, FMIs should regularly do a comprehensive evaluation of applications to understand any critical interdependencies and identify opportunities for enhanced resiliency. For example, this analysis may identify the need to redesign applications to be more modular or to rely on an independent

[16] Coverage can be defined by entities in different ways. This could either be across their full suite of services or across their critical business functions. The approach may be influenced by several factors, including architectural design and technical implementation approach principles, or business drivers.

[17] IWG members note that although restore points should be agreed upon in advance, the facts and circumstances surrounding a cyber event may dictate a need to reassess pre-determined restore points during the event. For example, in a ransomware event, whereby data is either locked or destroyed, it may be most beneficial for the ecosystem if the FMI were to use the last copy of data available to it, rather than returning to an agreed restore point earlier in the business day. This action would help ensure finality of transactions processed prior to the cyber event.

set of data, rather than a common data set, in order to prevent contagion. Such changes could help speed up the recovery process and minimise the impact to the business.

Recommended areas for collaboration

IV. An industry partnership should be established to create design principles for housing critical data sets in data bunkers and third sites – Some FMIs currently use third-party sites or off-premise data bunkers to serve as a recovery tool in data impact scenarios. However, as shown in the Appendix B, this is not common practice. One factor for the limited usage may be that there is no established set of "best practice" principles identifying how these solutions could and should be designed. Accordingly, there is an opportunity for the private and public sectors to work together to define a set of principles for the creation of independent sources of data held off network, either in segregated immutable data bunkers or third sites.[18] Once foundational principles are developed, it is understood that they may be used on the part of the regulatory community and adopted into existing guidance.

V. Guidelines for minimising contagion are needed – Complex interconnected ecosystems between firms increase the potential impact of data integrity scenarios across the industry, and FMIs must be confident that the resumption of critical operations is safe and does not risk infecting, or becoming re-infected by, external endpoints. It appears there is currently no consistent approach among FMIs for determining scenarios where it may be appropriate to disconnect or reconnect the FMI from an external endpoint to prevent contagion to or from that endpoint.[19] Given the interconnectedness and global nature of firms, regulatory dialogue on this issue would be beneficial as well, and help establish consistency across jurisdictions, where appropriate.

VI. Standard expectations for assessing parties in the ecosystem should be leveraged, where possible – The use of a common standard to evaluate entities that contribute to the ecosystem, either as a provider or as a client, should be encouraged as a preventive measure. One such common standard for assessment is the Financial Services Sector Cybersecurity Profile.[20]

[18] The IWG encourages principles that do not require FMIs to use data bunkering or third sites as a part of their operational resilience or operational risk solutions; rather these principles should recognise that FMIs need flexibility to develop and implement strategies and tools that best allow them to adequately manage data integrity issues. In addition, the IWG would encourage language to make clear that the design principles would be applicable only in the event that an FMI decides to use this option as part of their risk and resilience strategy and does not suggest this tool is the only solution for FMIs to address data integrity resilience. For example, an FMI could consider storing a copy of the four data types identified in this paper not only within their own on-premises backups, but also with other entities where business/technical connectivity may exist, and the terms of the relationship are mutually agreed. Once these relationships are established and implemented, a copy of data would be stored with that entity, which could be used in of the event that the FMI's on-premises backups are not accessible.

[19] A best practice may be for FMIs to include provisions in their operating rules or governance documents to formalise the ability to disconnect a member/client/third party, if there is a reasonable basis for determining that a disruptive event has occurred and limit the impact to the ecosystem. Conversely, reconnecting participants to the ecosystem is a much harder decision. An FMI needs to be comfortable that the contagion previously evident at the firm is eradicated, and they have the ability to connect in a manner that does not put the larger ecosystem at risk again. The ownership of such decisions has not yet been agreed by the industry and should not be performed alone. It is critical that the public sector develop guidelines on what should be reviewed and attested to before reconnecting a participant to the ecosystem, to provide regulatory certainty underpinning such actions. One potential solution could be for the regulatory community to establish harmonised practices, or rely on a certification process, that could be conducted by an approved list of third parties, who would certify that any issues have been addressed and that it is safe to restore connectivity to the FMI.

[20] https://fsscc.org/Financial-Sector-Cybersecurity-Profile

VII. Conducting industry-wide cyber exercises is critical yet should be coordinated by an independent central coordinating-party – FMIs should undertake cyber exercises with other entities in their ecosystem to rehearse coordination of recovery processes, including the exchange of large and diverse datasets where appropriate.[21] These exercises should be facilitated by independent parties, like industry associations, to bolster wide-scale participation and achieve a result that is independent of any one entity.

Other recommendations

VIII. A common, yet flexible, definition of service criticality is needed with an acknowledgement with respect to prioritisation of resumption[22] – Varying regulatory guidance on operational resilience and broader FMI standards refer to the identification and planning for either critical or important business services. However, the definition of critical or important services can vary across firms and may differ for resilience planning versus recovery/wind-down planning. In addition, the priority order in which to resume individual services can change based on a number of factors. As FMIs offer a range of services, it may be necessary to prioritise and stage the resumption of individual offerings in cyber events - the priority of which may change depending on the time of day, the criticality of operation during the event and the level of risk. As an example, there is currently no consistent industry approach for determining priorities for the staged resumption of critical operations based on the time of day or other time-based factors (day of the week, quarter end, etc.). FMIs should include these factors in their response plans as appropriate. Regulatory guidance should also allow for flexibility around how criticality is identified in practice.

[21] Cyber Guidance: 7.3.1 Coordination. An FMI should, to the extent practicable and possible, promote, design, organise and manage exercises designed to test its response, resumption and recovery plans and processes. Such exercises should include FMI participants, critical service providers and linked FMIs. Where appropriate, FMIs should participate in exercises organised by relevant authorities and in industry-wide tests. Achieving market-wide timely recovery of operations calls for an added dimension to testing exercises. Traditional isolated testing implicitly assumes that all other players operate as usual. Removing that hypothesis helps an FMI to identify plausible complexities, dependencies and weaknesses that may have been overlooked in recovery plans. Accordingly, testing should include scenarios that cover breaches affecting multiple portions of the FMI's ecosystem.

[22] A standard definition of critical data would also be helpful with respect to services.

APPENDIX A – SCENARIO FRAMEWORK

DATA

For purpose of the analysis, four different types of data were assessed:

1. **Configuration Data –** information that is required to operate technology including system settings, indexes, and user configurations.
2. **Application Data –** source code, processing jobs and scripts that when compiled creates a system that is used for business processing.
3. **Business Transactional Data –** any transactional information that is accessed, used, or modified as part of a business process. With respect to FMI processing, transactional business data is data that initiates a movement of cash or securities between accounts.
4. **Business Reference Data –** All counterparty, security identification, and pricing information that is needed to either record or settle a transaction.[1]

SCENARIOS

The Data Protection and Validation Industry Working Group set boundaries for evaluating the tools to use once a corruption/loss has occurred. A three-pronged classification scenario model was used to better understand a firm's intraday, end-of-day, next day and multiday resiliency capabilities following a cyber event. The three categories are defined below.

Impactful Scenarios (IMP):

IMP scenarios are those instances that impact operations under which recovery could occur intraday.[2] Localized in scope with limited impact, such as an impact to a single datacenter region, in these scenarios, a failover to an out-of-region recovery center would be performed. Key characteristics of IMP scenarios include:

- Minor production service, application, system, location or data has been impacted.
- The incident is simple to understand, and the response is well understood.
- Incident may be complex to understand, but the response is well understood.

[1] It is important to note that the data included within the Reference Data category varies by the FMI type. For example, securities pricing data was relevant to Central Securities Depositories but not Central Banks. For this reason, the tools identified were not applicable to all FMIs, given that some of the data in these instances came from external third parties.

[2] Recovery can change based on a variety of factors, including the time of day at which a disruption occurs. Accordingly, recovery time capabilities for this scenario depends, in part, on the exact time that a disruption takes place.

Extreme but Plausible Scenarios (XP):

XP scenarios are those instances that impact operations under which recovery could occur either by the end of the business day or fall into a subsequent business day. XP scenarios assume that the physical hardware and datacenter facilities have not been destroyed and are operable. Typically, these scenarios have the following characteristics:

- On-premise datacenters are impacted across region.
- One of the datacenter regions may be required to be isolated for forensic purposes ("isolated region").
- The resumption activities would need to occur outside of the isolated region.

Extreme but Implausible Scenarios

Extreme but Implausible scenarios assume the event had either a multi-region data centre impact, which requires a physical rebuild of both data centres, or would destroy all available copies of a firm's data. In the first case, the available hardware is destroyed within the data centres at each region and all facilities are inoperable. In the second case, no backup data copy would be available to reload technology or restart operations. These scenarios would have a multi-day impact due to their complexity and cost for mitigation.

For purposes of this exercise, Extreme but Implausible scenarios were not analysed, as they were considered catastrophic and arguably no number of data centres or data backups would be sufficient.

The scenario framework was then overlaid against two different data impacts, which are a subset of the traditional CIA model:

1. **Integrity –** Ensuring that information is accurate. In an integrity data scenario, data is impacted in such a way that it cannot be trusted.
2. **Availability –** Ensuring that data is available to the resources/systems that need them. Availability scenarios include the deletion of databases or ransomware, whereby access to requisite data is compromised.

TOOLS

Once the base cases were identified, a series of tools that could be leveraged for recovery, reconciliation, and replay were documented. Tools vary across the data types for many reasons, including how often the data changes and other potential data sources. Some of these tools have even been documented in regulatory guidance. A complete list of tools evaluated for each data type is shown in Appendix B.

ASSESS FEASIBILITY

The tools were then evaluated against the scenario purely by assessing how feasible the tool would be in its intended use case (recovery, reconciliation, or replay). This rating considers, for each individual FMI, six factors of implementation of the tool in order to determine how realistic it would be to use. These factors are:

1. **Complexity –** how difficult would it be to leverage when needed.
2. **Cost –** how expensive the tool would be to either maintain or implement.[3]
3. **Comfort –** how confident were FMIs in the use of the tool and whether procedures were in place for implementation (i.e. level of maturity of the solution).
4. **Time to implement –** how long would the tool take to be executed once a determination was made to do so.
5. **Expected data loss -** How much data could be lost or not when the tool is leveraged.
6. **Risk of use -** would use of the tool generate undue operational, legal, financial or reputational risk.

[3] Tools that are currently required by a regulatory or other policy standard were considered as business as usual, and therefore cost was not deemed to be incremental for purposes of this standard.

APPENDIX B – HEAT MAPS

The results of the analysis were used to create 'heat maps' in the form of a matrix of tools and scenarios, with color-coded feasibility ratings at the intersection of each tool and scenario. The legend for each heat map is shown in Figure A.

TOOL FEASIBILITY RATING	CURRENT USAGE RATING (TOOL USE)
Feasible	Common
Potentially Feasible	Non-Common
Not Feasible	Rare

Figure A – Heat map Legend

CONFIGURATION DATA

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Recovery	CD1 - Fix data/surgery, rather than recover	Correction of data (rather than recovery) in such a way that isolates the issue and corrects it without further impact to operation.					
	CD 2 - Recover data from internal backup copy (data stored in data center environment)	Recover data from an internal backup copy stored in the data center environment at any site. Speed of deployment will depend on whether automatic or manual deployment methods are used.					
	CD 3 - Recover using pre-configured repair copy	Alternative device is that already config-ured to support business processing and can be made available when required.					
	CD 4 - Recover from asynchronous backup to separate immutable database (i.e. data bunker, cloud)	"One time write" backup stored in highly secured and disassociated environment, outside of the FMIs infrastructure (could be cloud based).					
	CD 5 - Failover to backup site	Leverage a backup site to recover impacted data and resume operations.					

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Reconcilia-tion	CD 6 - Reconcile data against backup/test system	Use a local configuration management system (i.e. tracking versioning and mon-itoring) versus the baseline. The ability to restore automatically or on demand to last stable state would also be required.					
	CD 7: Reconcile with data from secondary site	Use a multi-site configuration management system (i.e. tracking versioning) versus baseline. The ability to restore automatically or on demand to last stable state would also be required.					
	CD 8: Reconcile data from trusted external source (i.e. e.g. data bunker, cloud, golden copy)	Recover a copy of data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e. are “trusted”) prior to entry to the store.					

APPLICATION DATA

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Recovery	AD1 - Failover to back-up system	Recover using a backup system at an alternative site.					
	AD2 - Failover to back-up site	Leverage a back-up site to recover impacted data and resume operations.					
	AD3 – Recover data from internal source code/version control repository (data stored in data center environment)	Recover data from an internal source code/version control repository stored in the data centre environment at any site. These can be current or historical versions of source code. Speed of deployment will depend on whether automatic or manual deployment methods are used.					
	AD4 - Recover data from backup/ test system	Recover data from a non-production backup/test system. Constrained by version in backup/test system, which may not be the same version as in production (may be historical, current or future version).					
	AD5 - Recover data from trusted external source code/version control repository (e.g. data bunker; cloud)	Recover a copy of data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e. are "trusted") prior to entry to the store.					
	AD6 - Fix data/surgery, rather than recover	Correction of data (rather than recovery, e.g., by developing new code) in such a way that isolates the issue and corrects it without further impact to operations.					
	AD7 - Shut-down and restart parts (or all) of the system	Isolate areas of corruption and remove the ability for those processes to operate. (Could be risky if interdependencies are not known).					
	AD8 - Leverage Buddy FMIs or non-similar systems	Leverage data from a trusted third-party or non-similar system.					

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Reconcilia-tion	AD9 – Reconcile data from internal source code/version control repository (data stored in data center environment)	Reconcile data from an internal source code/version control repository stored in the data centre environment at any site. These can be current or historical versions of source code.					
	AD10 - Reconcile data from backup/test system	Reconcile data from a non-production backup/test system. Constrained by version in backup/test system, which may not be the same version as in production (may be historical, current or future version).					
	AD11 - Reconcile data from trusted external source code/version control repository (data stored in data center environment)	Reconcile data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e. are "trusted") prior to entry to the store.					
Replay	Not applicable - Replay concept generally not applicable to application data owing to its static nature						

TRANSACTIONAL BUSINESS DATA

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Recovery	TBD1 - Fix data/surgery, rather than recover	Correction of data (rather than recovery) in such a way that isolates the issue and corrects it without further impact to operations.					
	TBD2 - Input correcting or reversing transactions	Using existing mechanisms to send into the system a transaction which act is a manner such that the original transaction is no longer valid. (Note: volumes may impact the effectiveness of this tool).					
	TBD3 - Failover to production instance at the back-up site	Recover data from another production instance located at a back-up site and resume operations.					
	TBD4 - Restore data from back-up site to primary site	Recovering from a (near) real time copy of business transactional data that was sent to a separate data center.					
	TBD5 - Recover data from internal backup copy (data stored in data center environment)	Recover data from an internal backup copy stored in the data centre environment at any site. These can be current or historical versions of business data.					
	TBD6 - Recover from asynchronous backup to separate immutable database (i.e. data bunker, cloud)	"One time write" backup stored in highly secured and disassociated environment, outside of the FMIs infrastructure (could be cloud based).					
	TBD7 - Recover data from participant or network operator records	Ecosystem contribution to restore the data.					
	TBD8 - Recover data from buddy FMIs or non-similar systems	Recover business data from a trusted third party or FMI's non-similar system.					
Reconciliation	TBD9 - Reconcile data from trusted external source (e.g. data bunker, cloud, golden copy)	Recover a copy of data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e. are "trusted") prior to entry to the store.					
	TBD10 - Reconcile data from network operator records.	Ecosystem contribution to reconcile the data.					

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Reconcilia-tion	TBD11 - Request reconciliation from participant/issuer verification/ obtain participant records.	Each participant should be required store key records to support the recovery & reconciliation efforts.					
	TBD12 - Reconciliation from event logging.	Event logging supports the real time storage of business events/information, event by event (e.g. transaction finality, number of records processed for a certain client, time stamp for end of a certain process). Note: This tool is complementary to other reconciliation tools and is not a complete reconciliation tool on its own.					
Replay	TBD13 - Participant resends transactions.	Participant is responsible for the resubmission of data.					
	TBD14 - Operator replays from internal source.	Leverage an internal source that provides the capability to resubmit transactions from an existing location on the client's behalf.					
	TBD15 - Use "Buddy" FMI – with access to shared data.	Replay missing information from a third party who has been storing information on the firm's behalf.					
	TBD16 - Non-similar system	Replay from FMI's non-similar system.					

BUSINESS REFERENCE DATA

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Replay	RD1 - Fix data/surgery, rather than recover	Correction of data in such a way that isolates the issue and corrects it without further impact to operations.					
	RD2 – Restore data from back-up site to primary site	Recovering from a (near) real time copy of business transactional data that was sent to a separate data center.					
	RD3 - Recover data from internal back-up copy (data stored in data center environment)	Recover data from an internal back-up copy stored in the data centre environment at any site. These can be current or historical versions of source code. Speed of deployment will depend on whether automatic or manual deployment methods are used.					
	RD4 - Recover from asynchronous backup to separate immutable database (i.e. data bunker, cloud)	"One time write" backup stored in highly secured and disassociated environment, outside of the FMIs infrastructure (could be cloud based).					
	RD5 - Recover data from trusted external source (i.e. another correspondent, buddy bank, supervisor, publicly available information, etc.)	Recover a copy of data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e. "trusted") prior to entry to the store.					
	RD6 - Failover to backup site	Leverage a backup site to recover impacted data and resume operations (Assumes corrupted data has not been copied).					
	RD7 - Leverage Buddy FMIs or non-similar systems	Leverage application data from a trusted third-party or non-similar system.					
Reconciliation	RD8 - Reconcile to external trusted source (i.e. publicly available information or trusted party)	Publicly available information or information obtained from a trusted party are used as a starting point for reconciliation of corrupted reference data.					
	RD9 - Reconcile data against backup or test system	Reconcile data from a non-production backup/test system.					

STAGE	TOOL	DESCRIPTION	INTEGRITY		AVAILABILITY		TOOL AVAILABILITY
			IMP	XP	IMP	XP	
Reconcilia-tion	RD10 - Reconcile data from trusted external source (e.g. data bunker, cloud, golden copy).	Recover a copy of data stored at a trusted external/offsite location and within a segregated network. Data are assumed to have been validated (i.e., are "trusted") prior to entry to the store.t					
Replay	RD11 - Operator replays from internal source.	Recover using a backup system at an alternative site.					
	RD12 - Request replay from external source who is the information provider.	Leverage an internal source that provides the capability to resubmit transactions from an existing location on the client's behalf.					