Annex

Good Practices for Addressing Vulnerabilities Related to Operational Resilience

A. Information and Communication Technology (ICT) Risk Management

Along with increased dependence on ICT systems and new technologies as driven by the digitalisation trend, the potential impact of an ICT failure causing operational disruption has increased. To better fortify themselves against this risk, most AIs engaged have worked to incorporate "resilience by design" principles within ICT risk management, focused on the following three key aspects:

1. Eliminating single points of failure

- *Alternative technological pathways or systems* – Major AIs eliminate single points of failure by implementing alternative technological pathways or systems, particularly those that would affect critical operation availability if disrupted. Various tools (e.g. content delivery networks, application programming interfaces (APIs), and integration layers) are also utilised to address potential sources of disruption, enabling redundancy, scalability, and fault isolation to ensure continuous access. Where alternative technological pathways are not feasible, such as when systems supporting multiple critical operations are interdependent, major AIs mitigate the risk of dependency on a single system (e.g. a core banking system). They do so by establishing not only a disaster recovery system but also deploying an alternative system that replicates essential functions, including but not limited to cash withdrawals, account enquiries, and transaction verifications. These measures help add an additional layer of protection for system availability.

- *Auto-switching capabilities* – To enable alternative technological pathways or systems to operate coherently, major AIs integrate auto-switching capabilities where appropriate, allowing for seamless transition to alternate pathways or systems during disruption. These capabilities are designed to detect transient failures (e.g. network glitches, unexpected hardware malfunctions and software failures), and to trigger controlled switching without user intervention. Some AIs also embed associated controls within the design to enhance reliability, including but not limited to robust failure detection mechanisms for real-time monitoring, controlled auto-retry logic (e.g. duplication checks when re-attempting transactions), and timeout mechanisms, along with exceptions monitoring to allow for manual intervention should auto-switching fail unexpectedly.

- *Built-in bypass options* – As an additional safeguard (e.g. in case auto-switching capabilities fail), major AIs also incorporate bypass options with

appropriate risk mitigating measures to ensure critical operations can continue during prolonged system disruptions. These options often include a "break-the-glass" functionality within the system supporting critical operations, enabling manual intervention (e.g. the download of the latest payment transactions for manual approval and processing to support payment operations). Respective risk mitigating measures are pre-agreed to address the associated risks, such as a predefined additional layer of approval matrix to ensure proper authorisation, and reconciliation procedures to maintain accuracy following a bypass. Given that bypass options inherently involve manual intervention which can be resource intensive, these AIs also factor this into their resource allocation planning.

2. Enhancing the resilience and recovery of key ICT supporting assets

- *Active-active configurations at data centres* – To ensure the availability of data centres (which are typically essential to numerous critical operations), many AIs either adopt active-active data centre architectures across geographically dispersed data centres or maintain active-active systems and infrastructure components within data centres using physically segregated or virtualised environments. This enables both systems and infrastructure components to operate simultaneously and share workloads. As part of this, AIs conduct granular assessments of system dependencies and data centre architecture to evaluate the criticality of workloads, such that they can be prioritised for transition to the active-active configuration in a risk-based way.

- *Virtualisation of infrastructure* – To minimise the disruption to critical operations during hardware failure, maintenance or replacement, AIs have virtualised critical supporting assets (e.g. firewalls, load balancers, and database clusters), allowing rapid provisioning, workload mobility and hardware abstraction. Where virtualisation is not feasible, multiple layers of redundancy are implemented. These reduce lead times for recovery, support automated scaling and enable near-instantaneous failover during outages.

- *Modernised and modular design* – A growing number of AIs have modernised their IT infrastructure and systems by adopting virtualised and modular architectures to enable greater agility, faster deployment cycles, and improved fault isolation. This approach has also been applied to systems that represent critical dependencies (e.g. core banking system). A robust system design review is implemented as part of this process to determine which specific functions should be modularised or segregated, ensuring that these measures would not create new single points of failure, security vulnerabilities and/or other dependencies that might in turn trigger cascading outages.

- *Asset and obsolescence management* – Many AIs have established processes and adopted various approaches to manage systems and reduce reliance on

legacy software and hardware. Examples include prohibiting the use of end-of-support systems supporting critical operations, implementing multi-year phased system upgrade programmes, and proactively replacing obsolete hardware underpinning critical operations with modern alternatives well ahead of end-of-life.

- *Diverse technology recovery options* – Major AIs have incorporated diverse technology recovery options into their recovery strategies to enhance readiness to handle various types of disruptions. They have not only established application-based recovery and full-site recovery capabilities (covering all applications within a data centre to another), but have also expanded to a broader range of options, including platform-based and critical operation-based recovery options covering various combinations of applications or systems. These options enable targeted recovery tailored to the scale of the disruption, offering greater flexibility and minimising the impact on unaffected technological components that support other critical operations.

- *Tailored recovery testing approach* – Major AIs have implemented tailored recovery testing approaches to allow more precise and effective validation of actual recovery capabilities under a given scenario. They have moved beyond traditional table-top exercises and disaster recovery testing to a more diverse suite of system recovery tests, including role swap exercises (i.e. temporary switch of roles between the primary and back-up data centres to validate the back-up site's ability to operate as primary and to enable safe fallback) per system or for the full-site conducted at production level, in order to more robustly assess the effectiveness of the technology recovery plan in a live environment. Some have also considered the viability of performing these tests with a surprise element or timeline to assess their recovery capability.

3. <u>Updating ICT risk management frameworks using a "resilience-first" mindset</u>

- *Technology support provided by offices outside Hong Kong* – Some AIs who rely upon or work with their offices outside Hong Kong (e.g. parent banks, subsidiaries, head offices, or other regional offices within the same banking group) have revisited their dependencies to assess whether these offices can meet the AIs' resilience requirements. This is due to shared technology controls or environment, where a failure in one component could disrupt another. These AIs performed resilience-based assessments to determine whether specific systems and infrastructure components require segregation from these offices, or require uplift or adjustment to ensure resilience requirements are met.

- *Change management* – On top of existing change management controls (e.g.

the software development lifecycle and emergency change management), some AIs have adopted tailored change strategies or controls to minimise disruptions arising from changes. These include adopting a "minimum viable product" approach, characterised by smaller and more frequent deployments that inherently reduce the risks associated with each release, enabling more reliable rollbacks and supporting the continuous modernisation of systems to prevent legacy issues. For larger-scale changes, these AIs not only establish a clear project lifecycle for management but also embed resilience as a testing element. In case the changes are driven or managed by the AIs' offices outside Hong Kong, local change committees are established to oversee changes, and have the authority to participate in decision-making processes.

- *Capacity planning* – Major AIs have evolved capacity planning to cover not only underlying systems but also all dependencies supporting each critical operation, reducing bottlenecks and enabling scalability as needed. As such, they are capable of performing capacity planning at a granular level covering all dependent system components, including applications and services (e.g. microservices, batch jobs), integration layers (e.g. APIs, middleware), infrastructure resources (e.g. CPU, memory, storage), data and database workloads/connections. They also perform in-depth analysis to account for non-linear demand across the end-to-end critical operation journey due to the presence of alternative or multiple pathways. For example, the single-entry point for authentication is supported by multiple APIs that route the incoming transactions to subsequent servers, which are also deployed across multiple instances. In addition to performing regular reviews, AIs cater for potentially varied customer behaviour or patterns during different holidays or events (e.g. long holidays, racing days, hot season for travelling, selling of popular concert tickets).

- *Monitoring* – Major AIs have redesigned or introduced new monitoring dashboards and mechanisms to not only cover exceptions monitoring but also proactively flag alerts when early warning signals indicate potential resilience issues, thereby embedding end-to-end visibility. They encompass customer experience indicators, such as payment transactions piling up or customer authentication requiring additional retries or higher than usual response time. Backend signals include specific system slowdown or abnormal functioning, unstable connections with other systems, and situations where capacity has yet to be breached but exceeds usual trends. Such visibility enables auto-scaling, manual intervention or adjustments to other ICT risk management controls as required.

B. Cyber Security Risk Management

As reliance on technology has grown and potentially expanded the surface area for intrusions, the risks posed by cyberattacks to AIs' operational resilience have naturally trended upwards. This landscape is further complicated by geopolitical

uncertainties and threat actors taking advantage of technology advancements. While no major cybersecurity incidents impacting AIs have been observed in Hong Kong so far, many AIs have nonetheless prioritised strengthening cyber resilience, and enhancing risk management capabilities across the full cyber risk management lifecycle.

4. Developing multi-year, threat-based cyber resilience strategies and programmes

- In line with the HKMA's supervisory requirements and expectations[1], many AIs have undertaken multi-year strategic initiatives to holistically enhance and mature cybersecurity controls (e.g. fine-grained access controls and micro-network segmentation). Furthermore, some AIs have established dedicated task forces to proactively monitor and address emerging cyber threats, such as cybersecurity threats posed by quantum computing.

5. Continuously adopting advanced cybersecurity tooling and solutions

- Some AIs are proactively adopting advanced cybersecurity tooling and solutions from the market to bolster their cyber defence, detection and response capabilities. For instance, some AIs have adopted artificial intelligence-powered cybersecurity tools for automated threat detection (including deepfake attacks) and faster response to security events, and cybersecurity solutions that are native to IT environments constructed in public cloud.

6. Reviewing and uplifting Secure Tertiary Data Backup (STDB) arrangements

- To strengthen data resilience against sophisticated cyberattacks (e.g. ransomware attacks), major AIs have taken on board the HKMA's supervisory feedback and implemented a STDB. These AIs continue to explore ways to enhance the efficacy and effectiveness of their STDB solutions in light of evolving cyberattack risks. Examples include setting up dedicated environments and utilities for recovery, reducing expected data loss through more frequent back-ups, and integrating STDB data restoration processes into cyber incident response drills for end-to-end testing.

7. Collaborating closely with key third parties and industry peers

- Apart from strengthening their own cyber resilience, some AIs have placed emphasis on collaboration across their supply chain and the broader financial ecosystem. Practices observed include involving key third parties (e.g. cloud service providers, Security Operation Centres) in cyber drills, and

[1] These include SPM TM-C-1 on "Supervisory Approach on Cyber Risk Management", Cyber Resilience Assessment Framework (C-RAF) 2.0.

participating in industry-wide simulation exercises to fortify preparedness against systemic cyber incidents. In parallel, some AIs also proactively share cyber threat intelligence with industry peers, thereby strengthening the collective cyber resilience of the industry.

C. Third-party Dependency Management

With AIs increasingly engaging third-party service providers (TPSPs) to support their critical operations, it has become imperative that potential vulnerabilities at such TPSPs are also effectively risk-managed. To obtain more reassurance in this regard, AIs have worked to shape their third-party risk management frameworks (TPRMF) around operational resilience considerations, covering five key aspects as set out below:

8. Strengthening governance over TPSPs

- The Board and senior management of major AIs maintain robust oversight of TPSPs (including but not limited to intragroup TPSPs) involved in critical operations to prevent disruptions at TPSPs from impacting delivery. To support this oversight, some AIs have defined clear roles and responsibilities by designating a dedicated workstream or owner as responsible for advancing the ongoing maturity of resilience capabilities within TPSP management. Some AIs have established structured programmes that systematically evaluate their current approaches to managing TPSPs, particularly with regard to the delivery of critical operations. This enables the proper identification of enhancement measures, which are tracked and reported to the Board and senior management as necessary.

9. Integrating operational resilience within TPRMF

- Some AIs have integrated operational resilience considerations (e.g. the achievability of tolerances for disruption and the potential impact on critical operations delivery) into their TPRMF to assess and manage both identified and potential risks before entering into and throughout the lifecycle of a TPSP arrangement. In addition, some AIs have proactively identified and managed sources of concentration risk. This is facilitated by mapping the dependencies on TPSPs for the delivery of critical operations and maintaining up-to-date TPSP registers with some AIs also including the information of key n^{th} parties. Some AIs have applied defined metrics and set thresholds (e.g. more than a specific percentage) to flag situations where a single TPSP accounts for a material proportion of spending or delivery capacity for a critical operation. Where concentration is noted, some have applied mitigation strategies, including to diversify critical operations delivery by engaging multiple vendors.

10. Enshrining operational resilience within contractual remedies

- Major AIs have established legally binding service level agreements (SLAs) with TPSPs. These agreements are set and reassessed regularly to ensure that the TPSP arrangements satisfy the AI's expected level of operational resilience. When discrepancies are identified, some AIs have proactively renegotiated SLAs to enhance alignment with resilience requirements. Examples of revisions include clearly defining the recovery time objectives (RTOs) for services provided by TPSPs to ensure consistency with the AI's own tolerance for disruption, as well as establishing defined turnaround times and protocols for incident response and escalation.

11. Enhancing or establishing tailored third-party risk monitoring/detection controls

- Major AIs have implemented monitoring regimes commensurate with the institution's risk appetite and the criticality of TPSP arrangements under its TPRMF, including regular assessments of performance-related metrics for vendor products and services. Some AIs have integrated evaluations of the TPSPs' resilience capabilities into this monitoring. For example, during regular assessment, they evaluate not only the adequacy of the TPSP's BCP and disaster recovery plans (DRPs), but also assess their actual resilience capability (e.g. downtime). Where feasible, AIs conduct joint testing of BCP and DRPs with TPSPs to validate readiness, identify gaps, and ensure alignment with their defined tolerances for disruption. When early warning signals are detected, the AIs concerned initiate immediate follow-up actions with TPSPs (e.g. issue formal warning notices to TPSPs to prompt timely remediation) and monitor the remediation process. The resilience standards of the TPSPs (e.g. incident record, drill results, remediation) are also considered in renewal decisions.

12. Enhancing readiness for managing terminations

- Many AIs have developed exit plans for planned termination and exit strategies for unplanned termination of TPSP arrangements. Some AIs have further strengthened their readiness for transition by prioritising interoperability across TPSPs. For example, when engaging an alternative TPSP to support a critical operation, the AIs concerned have established standardised, interoperable, and transferable arrangements with the primary and alternative TPSPs respectively. Such arrangements include the set-up of interoperable logical assets (e.g. data, applications, APIs) and physical assets (e.g. hardware), as well as operational processes. This approach enhances AIs' ability to switch between TPSPs with minimal disruption.

D. Business Continuity Planning (BCP) and Testing and Incident Management

With the operational resilience regime setting quantifiable benchmarks (i.e. tolerance for disruption) for the timeliness of recovering critical operations, many AIs have prioritised work to further narrow their RTOs such that they have more buffer and confidence to stay within their tolerances for disruption in the event of more extreme disruptions. To achieve this, AIs have primarily enhanced their BCP and testing arrangements, as well as incident management programmes, to reflect resilience considerations.

13. Reflecting resilience considerations within BCP and incident management

- *Evaluating recovery strategies from time to time* – Major AIs have continued to evaluate their recovery strategies to ensure that they remain "fit-for-purpose" and can meet the established tolerances for disruption under a range of severe but plausible scenarios. Consequently, many AIs have supplemented additional recovery strategies and refined current arrangements. For example, some have introduced contingent staffing on top of work-split and remote working approaches, to enable more timely responses and to support potential manual workarounds.

- *Refreshing the BCP and establishing dedicated plans for critical operations* – On top of dedicated recovery plans which are team- or system-centric, major AIs either refresh the BCP or establish recovery and contingency plans for critical operations to support an end-to-end recovery that offers a critical operation-centric view. These plans typically include critical operation-focused information, such as cross-references to existing related processes (e.g. incident management and crisis communications), enabling all necessary actions to be accessible from a single source and facilitating seamless coordination across teams.

- *Improving business resumption* – Major AIs have improved coordination among business and support functions to support smoother critical operations resumption. The AIs concerned have not only identified dependencies within critical operations but also consolidated dependencies across critical operations to gain a clearer understanding of interdependencies. As a result, AIs are able to coordinate effectively to address dependencies (e.g. pre-coordinating staff resources that support multiple critical operations or markets) and establish appropriate prioritisation (e.g. recovery sequencing) across various critical operations where cross-dependencies exist.

- *Managing incidents with operational resilience as a priority* – AIs have taken steps to make clear that maintaining operational resilience is a key objective of incident management. For instance, some have updated the predefined criteria for the classification of an incident's severity based on the impact to

critical operations. Some have additionally introduced stringent requirements for incident response and recovery, including time-based indicators for each phase within the incident response and recovery cycle covering: (1) incident detection and response, (2) impact assessment, (3) escalation, (4) preliminary root cause analysis, and (5) implementation of rectification measures. These time-based indicators are set to values significantly shorter than the tolerance for disruption, enabling prompt incident response and recovery before any operational resilience expectations were to be breached.