



HONG KONG MONETARY AUTHORITY  
香港金融管理局



Practice Guide on  
**Cloud Adoption**

**January 2026**



## Table of Contents

<b>1. Governance and oversight</b> .....	2
<b>2. Risk assessment and due diligence</b> .....	5
<b>3. Cloud service provider contractual provisions and related controls</b> .....	9
<b>4. Cloud resilience and exit strategy</b> .....	13
<b>5. Cloud security and data protection</b> .....	17
<b>6. Incident management</b> .....	22
<b>7. Ongoing monitoring</b> .....	24
<b>8. Workforce strategy and empowerment</b> .....	26



# Practice Guide on Cloud Adoption

## **Background**

The Hong Kong Monetary Authority (HKMA) adopts an interactive and iterative approach to promoting “responsible innovation” in the banking sector, and has all along been supportive of Authorized Institutions (AIs) progressively incorporating cloud technology within their operations in a prudent manner. To facilitate this, the HKMA has provided various guidance<sup>1</sup> and conducted over 70 supervisory engagements with AIs.

The HKMA has observed remarkable progress in AIs’ adoption of cloud. Specifically, cloud-related projects now account for about 80% of all reportable technology outsourcing initiatives, and roughly one-third to one-half of these involve critical banking systems. The models adopted by AIs have also become more complex, expanding beyond private and public clouds to cover hybrid and multi-cloud environments too.

Against this backdrop, the HKMA considers it timely to expand and deepen the guidance provided on cloud computing through the publication of this Practice Guide on Cloud Adoption (Practice Guide). The contents therein have been developed with regard to insights drawn from the HKMA’s ongoing supervisory engagement with AIs, as well as practices observed internationally.

## **Application**

This Practice Guide comprises both:

- (i) high-level principles from relevant Supervisory Policy Manual (SPM) modules; and
- (ii) good practices that the HKMA has observed through supervisory work, and which AIs may follow to help achieve compliance with the principles.

This “dual-layered” design aims to keep our supervisory expectations up-to-date and relevant for a variety of cloud adoption scenarios and applications, while also providing guidance that is actionable and of implementation reference to AIs.

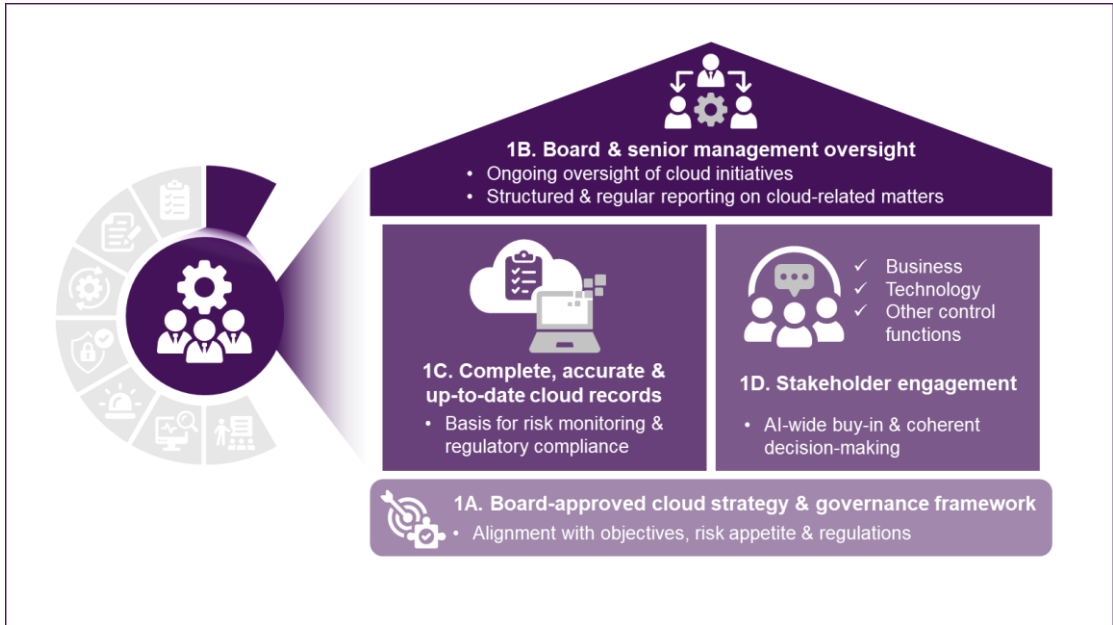
The HKMA expects AIs adopting cloud technology to apply the high-level principles for managing cloud-related risks in a manner that is proportionate and commensurate with the risk profile and criticality of their cloud arrangements, and to adopt the good practices where appropriate for their circumstances. The Practice Guide should be read in conjunction with the HKMA’s other relevant guidance, including but not limited to Supervisory Policy Manual (SPM) modules OR-2 on “Operational Resilience”, SA-2 on “Outsourcing”, and TM-G-1 on “General Principles for Technology Risk Management”. Any questions concerning this Practice Guide can be sent to [cloud@hkma.iclnet.hk](mailto:cloud@hkma.iclnet.hk).

---

<sup>1</sup> Including the 2022 circular on “Guidance on Cloud Computing”, which this Practice Guide supersedes.



# 1. Governance and oversight



## **Principle 1A: Cloud strategy and governance framework**

AIs should establish and maintain a comprehensive cloud strategy, formally approved by the Board, that is consistent with the AI’s strategic objectives, risk appetite, operational resilience targets, and regulatory obligations. If an AI is part of a banking group that has an overarching cloud strategy, the AI should adapt the strategy as appropriate to reflect and cater to local business and regulatory priorities. The AI should also maintain effective communication channels with its group that enable necessary changes or enhancements to the strategies to be made on an ongoing basis and in a timely manner.

A cloud governance framework should be established to translate the strategy into practice by defining oversight arrangements, roles and responsibilities, and risk management approaches proportionate to the materiality of the AI’s cloud adoption. The framework should be integrated with the AI’s enterprise-wide risk and resilience frameworks, adaptable to the distinct responsibilities of different cloud deployment models, and periodically reviewed by the Board and senior management to ensure ongoing effectiveness. Governance processes should be robust, scalable, and subject to continuous improvement.

**Good practices**

- Reviewing the cloud strategy on a defined cycle (e.g. annually and after a material change) to ensure alignment with evolving business priorities, risk exposures, and regulatory requirements.
- Maintaining documented linkages between the cloud strategy and other enterprise strategies (e.g. digital transformation, outsourcing, cybersecurity), ensuring dependencies are visible and traceable.



- Developing a cloud adoption policy under the cloud governance framework that specifies the approval process and implementation responsibilities, sets model-specific control expectations (e.g. access management, data protection, monitoring), and outlines compensating measures where responsibilities rest with cloud service providers (CSPs), which include both external providers and intra-group entities providing the same service.
- Utilising post-migration reviews (e.g. lessons learned from workload cut-over processes), incident analyses (e.g. outages or security events), and cross-project or portfolio reviews, to refine the governance framework and strengthen future adoption.
- Benchmarking cloud governance arrangements against recognised industry standards and regulatory expectations to ensure gaps requiring enhancement are identified.

### **Principle 1B: Board and senior management oversight**

Building on the approved cloud strategy and governance framework, the Board and senior management should exercise effective ongoing oversight of cloud initiatives, including to ensure that cloud-related risks are actively managed following the defined roles and responsibilities. Senior management should implement the cloud strategy by allocating adequate resources, establishing operational controls, and ensuring regular reporting to support effective oversight by the Board. Senior management should also promote open communication and workforce engagement through visible leadership, encouraging discussion of cloud-related matters, and modelling behaviours that foster alignment and resilience.

#### **Good practices**

- Affirming that ultimate responsibility for cloud-related risks rests with the Board and senior management, for example by documenting such responsibility in Board or committee charters, terms of reference or governance manuals, referencing it in key risk management policies, and reinforcing it through senior management performance objectives or key performance indicators.
- Establishing governance structures that enable effective oversight of cloud initiatives, such as setting up appropriate committees, working groups, or designated functions with clearly defined mandates to coordinate cloud-related activities across business, risk and technology domains.
- Embedding sufficient technical expertise within governance and risk management functions to support informed decision-making, ensuring that cloud-related risks are clearly understood, assessed, and managed on an ongoing basis.
- Establishing regular, structured reporting on cloud adoption, supported by analysis tools (e.g. risk assessments, concentration dashboards) that track key indicators and analyses such as adoption progress, availability, concentration exposures, incident trends, risk indicators, and compliance assessments, with



clear escalation of significant issues or breaches to the Board and senior management.

- Promoting transparency and awareness on cloud adoption by ensuring visibility of key cloud-related developments in senior management communications, reinforcing tone-from-the-top messages on governance and risk culture through internal channels (e.g. town halls, newsletters), and ensuring consistent communication of priorities and expectations across the institution.

### **Principle 1C: Cloud record management**

Effective oversight and risk management of cloud initiatives rely on maintaining complete, accurate, and up-to-date records of an AI's cloud resources and arrangements. AIs should ensure that such records provide a sound basis for monitoring risks, managing incidents, supporting planning and decision making, and demonstrating compliance with applicable regulatory requirements.

#### **Good practices**

- Maintaining an up-to-date inventory of cloud assets, including applications, data repositories, infrastructure components, and outsourced services, with the scope and level of detail proportionate to the cloud service model and the AI's degree of visibility or management responsibility, and with a clear indication of each asset's criticality and ownership. The inventory is updated promptly to reflect changes or new deployments, and retains information on decommissioned or terminated arrangements for an appropriate period to support risk monitoring, incident response, and audit or compliance reviews.
- Maintaining a registry of outsourced and cloud-supported functions, covering both critical and non-critical activities. The registry captures information such as the description and purpose of each function, the type and sensitivity of data involved, the identity of the CSP, the physical location of processing or storage, the legal jurisdiction of the CSP, the governing law of the arrangement, and, to the extent contractually feasible, details of subcontracting arrangements.
- Strengthening the registry by documenting application and system interdependencies, maintaining reliable configuration records, and regularly validating internal inventories against information provided by CSPs to ensure accuracy and completeness.

### **Principle 1D: Stakeholder engagement**

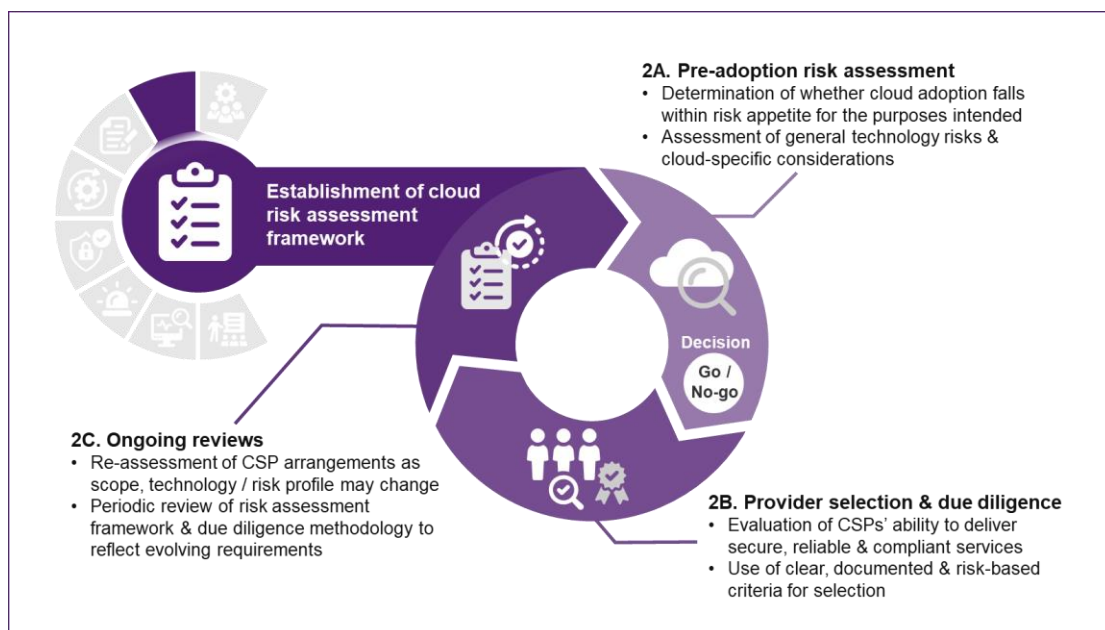
To reinforce the effectiveness of strategy, oversight, and governance structures, cloud governance should be conducted in a transparent and inclusive manner that engages all relevant stakeholders within an AI. The arrangement should help ensure that the perspectives of business, technology, operations, risk, legal, and compliance functions are taken into account, fostering AI-wide buy-in and promoting coherent decision making across functions.



## Good practices

- Establishing a stakeholder engagement plan that specifies which relevant teams will be involved at each stage of the cloud adoption lifecycle, and the processes and communication channels used to coordinate their input and feedback.
- Holding regular multi-disciplinary forums where proposed cloud initiatives, risk assessments, and lessons learned from incidents are shared across functions, with clear documentation of input, challenges raised, and actions agreed.
- Using structured tools (e.g. stakeholder matrices, consultation logs) to evidence that stakeholder views have been considered in forming key cloud adoption decisions, and to ensure that feedback is communicated back to relevant parties.
- Requiring that material cloud decisions (e.g. adoption of new platforms, migration of critical workloads, exit planning) include documented sign-off or endorsement from all relevant functions, proportionate to their roles and responsibilities.
- Embedding periodic stakeholder surveys or feedback sessions to test whether governance processes are viewed as inclusive, balanced, and effective, and using the results to refine processes.

## 2. Risk assessment and due diligence





## **Principle 2A: Pre-adoption risk assessment**

AIs interested in adopting cloud technology should first affirm that it is risk appropriate for them to adopt cloud for the purposes intended. To this end, AIs should conduct a pre-adoption risk assessment that is commensurate with the nature, scale, complexity and criticality of the intended cloud adoption. These assessments should cover both: (i) general technology risks, and also (ii) cloud-specific considerations.

As part of this process, particular attention should be paid to the following factors: data location, jurisdictional and regulatory requirements, reliance on external or internal service providers and their supply chain dependencies, and multi-tenancy arrangements. AIs should also explicitly consider concentration risk arising from reliance on a limited number of CSPs. In Software as a Service (SaaS) models, the reduced direct control over data and operations should also be considered.

The assessment should enable the AI to make an informed decision about whether it is in a position to adopt cloud technology, as well as the risks that may need to be mitigated before implementation.

### **Good practices**

- Using structured assessment tools (e.g. decision trees, whitelists) to identify appropriate cloud services and deployment options based on risk, resilience, and the feasibility of future transition or exit.
- Assessing the AI's internal capability to manage and oversee potential cloud arrangements, ensuring that adequate resources, expertise, and internal governance processes are in place before proceeding with the adoption.
- Developing and mandating the use of a standardised risk assessment template. The template is designed to be comprehensive and cover critical items including: (i) operational dependencies and resilience (e.g. reliance on CSP platforms, availability, subcontractor chains, portability, interoperability, exit arrangements); (ii) security safeguards (e.g. tenant segregation, encryption, access controls, monitoring, incident response); (iii) legal and regulatory requirements (e.g. data residency, cross-boundary data transfers); (iv) contractual protections (e.g. audit rights, liability caps, exit support); and (v) reputational implications (e.g. handling of sensitive customer data, reliance on provider attestations). Completed templates are reviewed and endorsed through the AI's governance structure before implementation.
- Where an AI already has some degree of cloud adoption, reviewing existing service dependencies and resilience by mapping reliance on specific CSP services, subcontractors, and internal interdependencies, and evaluating the current level of concentration and potential vendor lock-in. Scenario exercises (e.g. simulating CSP or subcontractor outages or disruptions) can be used to assess how additional adoption using the same CSP could increase dependence, and whether the resulting risks and impacts would be acceptable



to the AI. Findings can be used to inform adoption decisions and highlight areas requiring additional preparedness before implementation.

- Evaluating jurisdictional and regulatory implications by reviewing where data and services will be located or processed, and assessing alignment with applicable data protection expectations and other relevant regulatory requirements before any adoption decision.
- When assessing SaaS-specific risks prior to adoption, consider factors including reduced control over hosted applications, limitations on direct audit access, and reliance on CSP reporting, and ensure that the proposed arrangements include clear procedures for data extraction and portability.
- Requiring multi-disciplinary governance review of pre-adoption risk assessments, involving business, technology, operations, risk, legal, and compliance functions, with documented approval of any conditions or mitigating measures before progressing to CSP selection or contracting.

## **Principle 2B: Provider selection and due diligence**

Where pre-adoption risk assessments support the use of cloud technology, AIs should proceed to shortlist potential vendors through a robust evaluation and due diligence process that helps the AI determine whether a given vendor would be able to deliver the required services in a secure and reliable manner, and in compliance with applicable laws and regulatory requirements.

In conducting such due diligence, AIs should adopt clear, documented, and risk-based criteria. The extent and depth of the due diligence should be proportionate to the criticality and materiality of the outsourced functions. The evaluation and selection process should also be systematic, and make reference to recognised industry frameworks and independent assurance where appropriate.

To enable better outcomes, AIs should apply a consistent, well-defined methodology during the process that may draw on internationally or nationally recognised certifications, audit reports, and other independent third-party assessments as part of their due diligence to validate the CSP's control environment and compliance capabilities.

### **Good practices**

- Maintaining a documented due diligence framework that comprises consistent evaluation criteria, scoring or weighting methods, and residual risk thresholds.
- Prioritising an “evidence-based” approach to due diligence, which may involve:
  - Verifying the scope and relevance of certifications and audit reports in relation to the intended cloud service model, such as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and SaaS.



- Assessing subcontracting arrangements, especially where cross-boundary or multi-tiered structures may complicate risk management and regulatory compliance.
- Reviewing the legal and regulatory environment of the CSP's key operations and data storage locations to identify jurisdictional or sovereignty risks.
- Assessing CSP resilience and performance using objective metrics such as historical uptime, incident reports, and root cause analysis, supplemented where possible with references from other CSP clients, and confirming that service availability levels support the AI's operational resilience standards.
- Verifying that the CSP's infrastructure and operations provide security and data protection safeguards equivalent to the AI's internal standards, supported by certifications, independent testing, and other third-party assurance over both physical and logical controls.
- Defining risk-based selection criteria that cover key domains such as (i) financial soundness and long-term viability; (ii) organisational resources and technical expertise; (iii) operational resilience and service quality; (iv) information security and data protection; (v) incident response capabilities and track record; (vi) business continuity and disaster recovery arrangements; (vii) subcontracting structures; (viii) regulatory compliance history; (ix) data residency and cross-boundary transfer practices; and (x) strategic alignment with the AI's long-term objectives.
- Forming an evaluation team that includes subject matter experts from relevant areas (e.g. business, technology, operations, risk, legal, and compliance) to ensure their input is incorporated into the selection and decision-making process.

## **Principle 2C: Ongoing reviews of risk assessment frameworks and CSP arrangements**

As cloud arrangements may change in scope, technology, or risk profile over their lifecycle, AIs should not treat risk assessments and due diligence as one-off exercises. Periodic reviews should be undertaken to assess whether initial assumptions remain valid, whether a CSP's security, compliance, or financial condition has materially changed, and whether engagement of the CSP continues to align with the AI's risk appetite and regulatory obligations. These periodic reviews should be conducted at intervals commensurate with the criticality of the cloud arrangements, and whenever significant incidents, regulatory findings, or material changes arise.

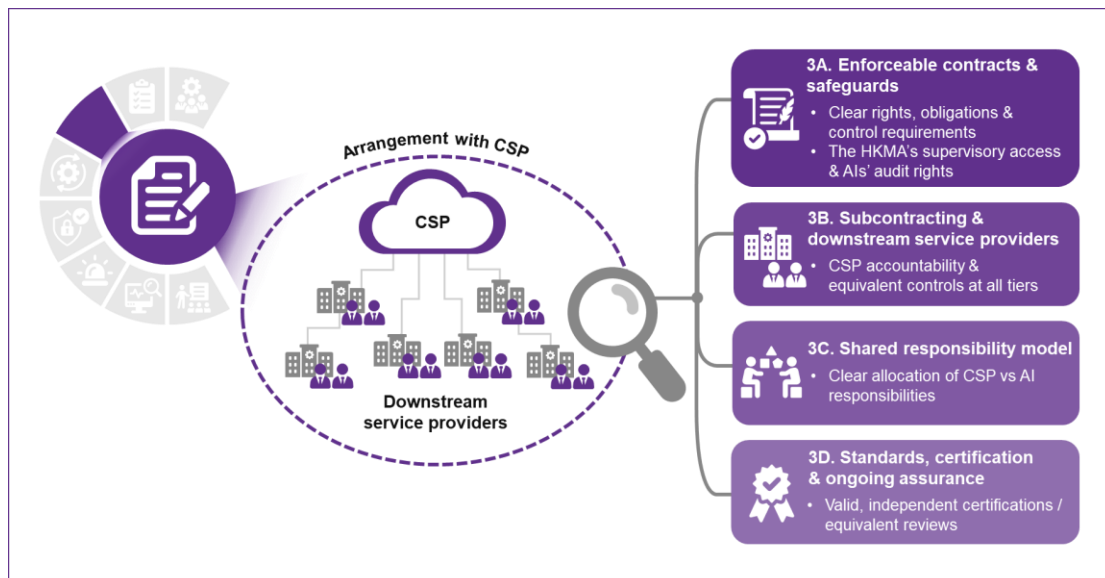
In tandem, the AI's risk assessment framework and due diligence methodology should also be subject to periodic review and timely updates to take into account, as appropriate, changes in business models, technology, regulatory expectations, and the threat landscape, as well as lessons learned from incidents, scenario testing, industry practices, and supervisory feedback.



### Good practices

- Scheduling recurring risk assessments (e.g. annually or more often for critical services) of CSP arrangements, covering certification validity, compliance standing, financial soundness, resilience capabilities, and overall risk posture.
- Conducting ad-hoc reassessments following material events such as CSP service disruptions, regulatory or audit findings, or sector-wide emerging risk alerts.
- Re-performing due diligence reviews if deficiencies are identified, or when entering into or renewing arrangements with a previously assessed CSP, taking into account the nature and materiality of the change.

## 3. Cloud service provider contractual provisions and related controls



### **Principle 3A: CSP contractual arrangements and safeguards**

AIs should establish clear and enforceable contractual arrangements with CSPs. Amongst other rights and obligations, the contractual arrangement should at minimum: (i) provide an AI with the audit and access rights needed for it to effectively oversee and monitor the resilience of the CSP's operations and services; (ii) ensure the HKMA's supervisory access to information stored in the cloud and ability to review relevant risk management controls for the purpose of on-site examinations; and (iii) cover appropriate safeguards around potential subcontracting arrangements (please refer to Principle 3B for more details). AIs should regularly exercise their audit rights to verify that CSPs are in compliance with contractual obligations, risk management standards, and regulatory requirements.



AIs should subject contractual arrangements to regular review to verify that they remain fit-for-purpose, taking into account relevant factors, such as the CSP's performance, regulatory requirements and the latest operating landscape.

### Good practices

- Including, as appropriate, the following elements in contractual arrangements with CSPs:
  - Clearly defined scope of services, roles, responsibilities, deliverables, and follow-up mechanisms to monitor CSP obligations.
  - Service level agreements (SLAs) that:
    - Specify measurable quantitative and qualitative indicators (e.g. Recovery Time Objective (RTO), Recovery Point Objective (RPO), availability targets); and
    - Establish clear criteria for monitoring, reporting, remedial actions, and enforceable consequences for breaches.
  - Documented rights and obligations, including:
    - Obligations for the CSP to address deficiencies identified from audits or assessments, provide evidence of remediation, and cooperate with subsequent verification;
    - Incident notification timelines, reporting criteria, and access to relevant logs and forensic data until resolution;
    - Clearly defined termination rights and conditions; and
    - Exit support, including timely data return or secure erasure upon termination.
  - Expectations for the CSP to establish and regularly test business continuity and disaster recovery arrangements, ensure adequate data retention and forensic accessibility, and to maintain appropriate levels of insurance coverage.
  - Provisions covering security, confidentiality, compliance with applicable laws and regulations, restrictions on data use, and measures for protecting sensitive and personal data where applicable.
  - Clauses defining governing law and jurisdiction, financial obligations, and data location or residency commitments to ensure legal enforceability and regulatory conformity.
  - A minimum notice period (e.g. 12 months) for discontinuation of services supporting critical or important operations.
- Reviewing CSP standard contractual terms during procurement and renewal to identify gaps and negotiate mandatory provisions that address the AI's and regulatory requirements, particularly for critical workloads.
- Establishing expectations for CSPs to cooperate proactively during audit exercises, including providing timely and comprehensive responses to audit enquiries, facilitating access to relevant records, designating competent



subject matter experts to assist in evidence review and clarification, and clarifying or supplementing third-party assurance reports where coverage or detail is insufficient.

### **Principle 3B: Subcontracting management**

As mentioned in Principle 3A, contractual provisions should not just be limited to the direct relationship between the AI and the CSP but should also extend to subcontracting and all downstream service providers, where appropriate. Such provisions should ensure that the CSP remains contractually liable for the performance and compliance of its subcontractors, and that accountability, transparency, and enforceability are maintained across the entire supply chain. AIs should retain rights of notification, approval, and objection over material subcontracting arrangements. In addition, AIs should seek to secure contractually enforceable commitments that equivalent resilience, security, data protection, and regulatory obligations will be applied consistently across every tier of third-party and subcontracting arrangements, to ensure that operational continuity and control adequacy are maintained at all levels. AIs should also put in place effective arrangements to oversee and monitor CSPs' engagement of subcontractors.

#### **Good practices**

- Specifying in contracts which functions may or may not be subcontracted, requiring prior written notification of intended subcontracting arrangements or material changes, obtaining explicit approval where relevant, and preserving the AI's rights to decline or require the CSP to discontinue subcontracting arrangements where these are introduced without authorisation or are judged to create undue risk.
- Ensuring and, where possible, periodically reviewing subcontractors' contractual terms, including SLAs where applicable.
- Establishing structured assurance mechanisms to verify that CSPs effectively oversee their subcontractors, including regular review of how oversight activities, issue escalation, and remediation processes are conducted.

### **Principle 3C: Shared responsibility model**

As part of or leveraging contractual and subcontracting provisions, AIs should agree, maintain and ensure the continued effectiveness of a shared responsibility model with CSPs. This model should clearly distinguish between the CSP's responsibility for the "security of the cloud" and the AI's responsibility for the "security in the cloud". It should set out how responsibilities are divided under different cloud service models (e.g. IaaS, PaaS, SaaS), allocate control ownership in a way that avoids gaps or overlaps, and reflect evidence of the CSP's capability to meet its assigned obligations. The model should also be reviewed and updated as services, technologies, or risk profiles evolve, ensuring ongoing alignment with the AI's risk appetite, security and resilience objectives.



### Good practices

- Documenting responsibilities across IaaS, PaaS, and SaaS by distinguishing the CSP's accountability for "security of the cloud" (i.e. infrastructure, network, platforms, and services operated by CSPs) from the AI's accountability for "security in the cloud" (i.e. workloads, configurations, data, and access managed by AIs), while ensuring the mapping remains current as services and risks evolve.
- For areas where responsibility is shared, specifying how controls (e.g. monitoring, recovery, cryptographic key management) are shared in terms of execution and validation.
- Embedding the agreed roles and responsibilities from the shared responsibility model into the AI's internal policies to guide staff in fulfilling their responsibilities, and ensuring that such records are reviewed and updated to reflect service changes, risk developments or technology evolution.
- Assessing and periodically verifying the CSP's capability to meet its assigned responsibilities by reviewing independent assurance reports, certifications, audit results, or other validated performance indicators.
- Applying a risk-based approach to identify and evaluate gaps or residual risks in the shared responsibility model, ensuring that the scope, depth, and frequency of the review are commensurate with the criticality of the CSP arrangement, and verifying that such exposures stay within the AI's approved risk tolerance.

### **Principle 3D: Standards and certification assurance**

In addition to contractual, subcontracting, and shared responsibility measures, AIs should require CSPs to demonstrate compliance with recognised international and national standards that are proportionate to the nature and criticality of outsourced workloads. Certification or attestation should, where available, be valid, relevant, and independently issued by accredited bodies. Where such certifications or attestations are unavailable, insufficient in scope, or outdated, AIs should consider alternative assurance through reliable third-party assessments or equivalent measures to ensure CSP controls continue to meet regulatory expectations and AIs' own requirements.

### Good practices

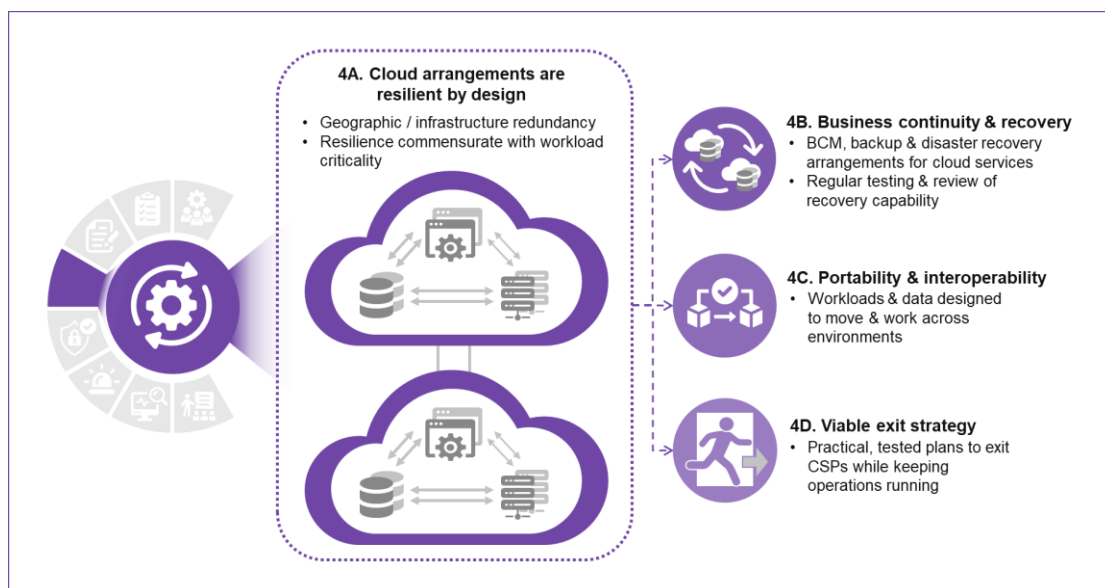
- Benchmarking CSP practices against appropriate industry frameworks to validate the adequacy and consistency of controls across different service and deployment models.
- Reviewing the authenticity, independence, scope, coverage, and validity of the CSP's certifications or attestations to confirm they are current, issued by accredited bodies, and relevant to the services consumed, including verification that they cover the AI's applicable regions, workloads, and data



types, and address key control domains central to cloud security and resilience.

- Mapping certified controls and assurance evidence apply to the AI's operational environment and internal risk assessments to confirm their applicability, relevance, and alignment with the AI's control framework.
- Conducting complementary assessments, where necessary, to validate that critical cloud controls (e.g. workload configurations, identity and access management (IAM), incident response readiness, data protection measures) operate effectively in practice and remain consistent with the AI's risk appetite, and recording any material gaps for follow-up with the CSP.
- Making use of independent assurance reports and other governance materials, and, where on-site examinations are impractical, relying on these sources while recognising their scope and detail limitations, and integrating key findings into the AI's internal assurance and monitoring processes.
- Where available assurance reports or certifications provide insufficient coverage, conducting risk-based audits of CSP arrangements, either by the AI's internal auditors, by appointed independent third parties, or through participation in pooled audit initiatives, to obtain adequate assurance, thereby avoiding sole reliance on CSP self-declarations.
- Establishing processes to monitor the certification lifecycle, requiring CSPs to provide timely notification and evidence where certifications lapse, are withdrawn, or where material deficiencies are identified in controls, and ensuring necessary remediation actions are agreed and tracked.

## 4. Cloud resilience and exit strategy





## **Principle 4A: Cloud resilience**

AIs should ensure that their cloud arrangements are resilient by design and can continue to support their critical operations under severe but plausible scenarios. Resilience arrangements should be commensurate with the criticality of the workloads deployed, incorporating redundancy in geography and infrastructure to minimise single points of failure and enable timely recovery or failover.

To address concentration and geopolitical risks, AIs should assess and mitigate dependencies on individual CSPs or specific geographic locations. More stringent safeguards should be applied where workloads are critical, to ensure that reliance on any single CSP or location does not compromise the AI's ability to maintain critical operations.

### **Good practices**

- Designing critical workloads to run across multiple, geographically dispersed, independently powered data centres with resilient interconnections and seamless failover, so that outages or large-scale disruptions can be absorbed without interrupting critical operations.
- Integrating on-premises and public cloud environments to diversify resilience options, provide fallback capacity, and enhance flexibility.
- Applying resilience measures using a risk-based approach, adopting stronger measures such as deployment across multiple CSP regions or availability zones for critical operations, and regularly reviewing these arrangements to ensure alignment with business objectives and regulatory expectations.
- Making effective use of CSP-native resilience features (e.g. automatic recovery and dynamic scaling of workloads) while ensuring they are correctly configured, tested, and subject to ongoing monitoring.
- Reducing reliance on any single CSP by developing hybrid or multi-cloud strategies, with priority given to critical workloads, supported by consistent governance structures, unified security policies, and staff capable of operating across providers.

## **Principle 4B: Business continuity planning**

Building on the principle of resilience by design, AIs should ensure that a comprehensive business continuity management framework is in place. The framework should be aligned with the AI's strategic objectives, risk appetite, and regulatory obligations, and include continuity plans for cloud services, supporting backup and recovery arrangements, and measures to assess and enhance CSPs' resilience. These arrangements should mitigate the risk of downtime, data loss, and other disruptions affecting cloud platforms that underpin critical operations.

The disaster recovery arrangements should be subject to regular oversight, supported by periodic testing of recovery plans and assessments of CSPs' capabilities.



### Good practices

- Establishing the AI's own business continuity and fallback arrangements for cloud workloads, proportionate to their criticality and risk tolerance. These arrangements should define and align RTOs and RPOs with the AI's approved tolerance levels, and incorporate feasible alternative service options or recovery strategies to meet these objectives in the event of CSP disruptions.
- Integrating CSP continuity capabilities into AI-wide planning by reviewing the CSP's documented arrangements (e.g. business continuity and disaster recovery summaries, testing plans, and related evidence) and aligning testing schedules with the AI's recovery framework.
- Assessing the effectiveness of CSP disaster recovery capabilities by examining practical recovery evidence and test results, rather than relying solely on certifications, to verify that these capabilities remain reliable under severe but plausible scenarios, such as large-scale cyberattacks, prolonged system or infrastructure failures, or data centre unavailability.
- Conducting, where practicable, joint continuity drills with CSPs to validate recovery arrangements and identify areas for improvement.
- Reviewing the business continuity management framework for cloud services on a regular basis, and following major changes in workloads, CSP arrangements, or material risk factors, to ensure continued alignment.

## **Principle 4C: Portability and interoperability**

Beyond immediate continuity arrangements, AIs should also plan for long-term flexibility by designing cloud workloads to be portable (i.e. the capability to migrate applications and data between environments with minimal modification) and interoperable with other platforms (i.e. the capability of different systems or CSPs to exchange and use information reliably) proportionate to their criticality. This approach enables workloads and data to be migrated securely and efficiently across environments when required, providing the flexibility to redeploy critical functions in response to strategic changes, major disruptions, or other severe but plausible scenarios that may affect the current service environment. These aspects should be considered during CSP selection and maintained throughout ongoing arrangements.

### Good practices

- Designing workloads to be portable by adopting open standards for data formats, application programming interfaces (APIs) and software interfaces, enabling seamless migration across different environments (e.g. cloud or on-premises) when needed.
- Operationalising portability and interoperability through containerisation and orchestration technologies, standardised integration layers, API gateways, and embedded cross-environment compatibility testing to ensure that workloads can be deployed and recovered smoothly across different environments and platforms.



- Ensuring that CSPs provide technically feasible and verified mechanisms for backup, migration, and data recovery, and validating that configurations support secure, efficient movement of workloads between environments.
- Integrating the maintenance and validation of interoperability tools into system lifecycle activities such as patching, upgrades, and system retirement, to ensure that interfaces, data exchange mechanisms, and migration pathways remain compatible, functional, and secure as environments evolve.
- Adopting CSP-neutral designs (e.g. distributed architectures, interoperable components) to reduce CSP lock-in and preserve the ability to operate across multiple environments in the long term.

#### **Principle 4D: Exit strategy**

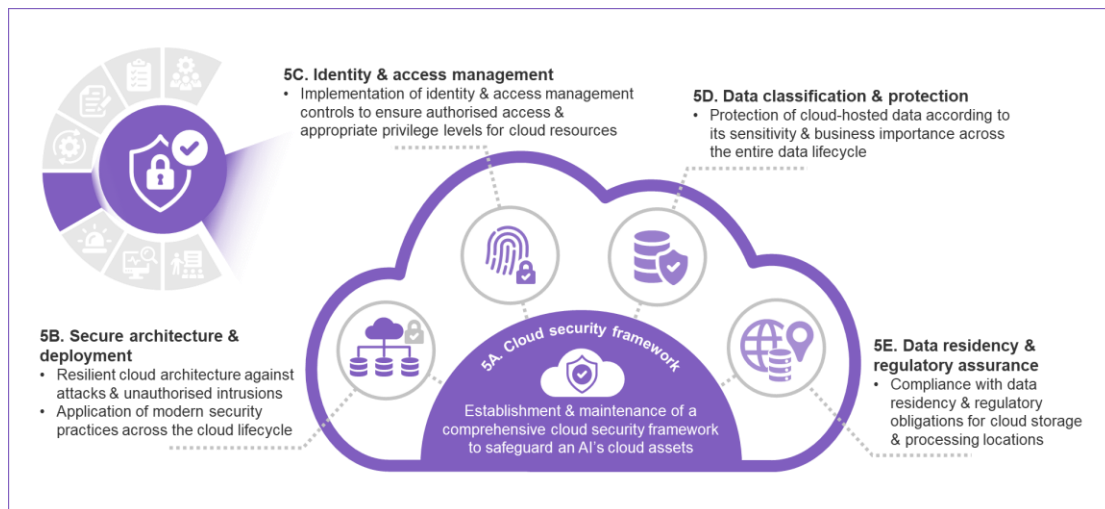
Complementing portability and interoperability measures, AIs should establish viable exit strategies to ensure that their cloud-dependent operations can continue to run even if a given cloud arrangement is terminated. Exit plans should set out objectives, risk indicators, and detailed procedures for contract termination and migration arrangements, together with clearly defined triggers for initiating an exit. The exit plans should be subject to periodic review and testing to ensure their feasibility and alignment with evolving business requirements.

#### **Good practices**

- Conducting business impact and technical assessments to estimate the timelines, costs, and skillsets required for transition, and using the results to identify substitute CSPs, in-house alternatives, and migration tools or standards that enable an orderly exit.
- Designing exit strategies in advance for all cloud services, with particular emphasis on those supporting critical operations, setting out clear objectives, success criteria, roles, responsibilities, indicative costs, and transition timelines to ensure that services provided by CSPs can be terminated and transferred without disruption.
- Taking into account structural limitations, such as intra-group dependencies or cross-boundary arrangements that may constrain exit options, and identifying feasible alternatives to safeguard resilience, such as transitional service agreements, hybrid or multi-cloud configurations, or on-premises fallback environments when direct migration is restricted.
- Testing the feasibility of exit plans on a regular basis through methods such as desktop reviews, walkthroughs, or staff readiness checks. Where proportionate, undertaking stressed exit testing to reconcile backup data across cloud environments and reveal hidden operational risks.



## 5. Cloud security and data protection



### **Principle 5A: Cloud security framework**

AIs should establish and maintain a comprehensive cloud security framework that enables the effective safeguarding of an AI's cloud assets. The framework should clearly set out security measures and assurance requirements, and provide overarching direction for areas such as secure architecture and deployment, IAM, data governance and protection, cryptographic key management, together with other relevant controls. It should be applied consistently across all cloud deployment and service models and scaled appropriately to system and data criticality. The framework should also be subject to ongoing oversight and assurance processes to ensure that it keeps pace with emerging threats, cloud-related technological developments, and shifts in the AI's risk profile.

#### **Good practices**

- Establishing clear roles, responsibilities, and accountability for cloud security management, aligned with the AI's broader governance and risk framework, supported by formal role descriptions, responsibility assignment matrices, and periodic assurance reviews to confirm effective segregation of duties.
- Structuring the framework to address the distinct risk characteristics of different cloud deployment and service models (e.g. public, private, hybrid, multi-cloud, IaaS, PaaS, SaaS), and setting clear expectations for security controls to be applied in each context, such as baseline configuration standards and data segregation requirements.
- Designing and maintaining the framework to ensure that cloud security controls remain consistent and effective across all relevant cloud environments, including, where applicable, hybrid, multi-cloud, and other interconnected configurations, with periodic reviews and timely updates to reflect changing risks and regulatory obligations.



## **Principle 5B: Secure architecture and deployment**

AIs should design and deploy their cloud architecture to be resilient to attacks and unauthorised intrusions. To support this, modern security practices should be embedded across the lifecycle of cloud infrastructure and applications, and include amongst others, effective practices for network design and segmentation, configuration and patch management, securing software and infrastructure, and backup and recovery.

### **Good practices**

#### **Network design and segmentation**

- Designing cloud networks with clear segmentation between development, testing, and production environments, and enforcing strict tenant isolation to minimise the risk of cross-environment compromise or lateral attack pathways.

#### **Application and API security**

- Incorporating established secure development disciplines such as systematic threat analysis and adherence to widely recognised application security principles across the entire lifecycle of cloud applications.
- Protecting APIs and micro-services through least-privilege access, strong authentication, active monitoring, timely decommissioning of unused interfaces, and safeguards for service discovery and mesh functions.

#### **Workload and container security**

- Applying hardened configurations to virtual machines and containers, limiting unnecessary components, and continuously monitoring for vulnerabilities or misconfigurations.
- Securing container environments by restricting administrative access to orchestrators, controlling container registries, and allowing deployment only from trusted image sources.
- Adopting immutable infrastructure approaches to minimise configuration drift and enable consistent redeployment from version-controlled base images.

#### **Automated deployment and configuration management**

- Leveraging Infrastructure as Code (IaC) and automated Continuous Integration and Continuous Delivery (CI/CD) pipelines to standardise deployments, reduce manual interventions, enforce segregation of duties, and preserve infrastructure integrity.
- Deploying Policy as Code tools to automatically validate and enforce compliance with predefined security configurations across cloud components.

#### **Cloud security and policy enforcement**

- Strengthening visibility, threat protection, and policy enforcement across cloud workloads through native security capabilities such as Cloud Access



Security Broker (CASB), Cloud-Native Application Protection Platform (CNAPP), Cloud Security Posture Management (CSPM), and Cloud Workload Protection Platform (CWPP).

#### Patch management

- Implementing a structured patch management process for all cloud resources, clearly defining patching responsibilities between the AI and its CSPs, and ensuring that patches are thoroughly assessed, tested, and deployed in a controlled, phased manner, and are synchronised across production and disaster-recovery environments to prevent inconsistent states or widespread disruption.

#### Backup and recovery

- Enforcing technical backup and recovery controls, including timely restoration of critical resources, replication to alternate or recovery sites, and use of segregated or offline systems (e.g. across CSPs, on-premises, or independent storage facilities), to minimise exposure to CSP-level failures.
- Conducting regular and scenario-based backup and recovery testing to verify backup integrity, completeness and recoverability, supported by post-restoration security validation to ensure recovered environments are hardened before service resumption.

### **Principle 5C: Identity and access management**

Alongside secure architecture and deployment, AIs should also put in place effective identity and access management arrangements to ensure that only authorised users and system components can access cloud resources at appropriate privilege levels. Effective arrangements should clearly set out a framework for managing identities, entitlements, and privileges, enforce segregation of duties, promote strong authentication requirements and zero-trust principles, and also apply to both internal and CSP environments.

#### **Good practices**

- Defining and documenting cloud IAM roles and responsibilities, including configuration rights, with clear designation of accountable business and technical owners (e.g. application and data owners) for each role.
- Maintaining comprehensive IAM policies and procedures for all cloud users and components, with regular reviews and updates, and requiring CSPs, where contracts permit, to support alignment with the AI's IAM provisions, and implementing mitigating controls for any identified gaps.
- Enforcing strong authentication for privileged and sensitive cloud access, including multi-factor authentication (MFA) for management consoles and, where relevant, CSP administrative planes or hypervisor access, protected through hardened endpoints, secure protocols, and end-to-end encryption.
- Implementing zero-trust principles across cloud environments by enforcing least-privilege access, deny-by-default settings, micro-segmentation, and



continuous validation, supported by disciplined access lifecycle management that ensures timely granting, recertification, and revocation of rights at a frequency proportionate to system and data criticality.

- Regularly rotating and securing credentials, ensuring the immediate deactivation of unused or expired credentials, and applying secure processes for credential generation, storage and distribution.
- Implementing continuous monitoring of user access activities across CSP and internal environments, supported by tamper-resistant audit trails and automated alerting mechanisms to detect unauthorised or anomalous behaviour.
- Applying real-time control over privileged activities through granular entitlements, maker-checker functions, and privileged access management (PAM) tools, to prevent unauthorised or high-risk changes.
- Adopting federated IAM solutions, where applicable, to unify identities across cloud and other environments (e.g. on-premises environments), and enhancing authentication security through contextual checks such as device, location, and user-behaviour indicators.
- Designing IAM architecture for scalability and resilience (e.g. automated account lifecycle management, fault-tolerant and high-availability identity services), supporting secure operations across different cloud deployment models, as applicable.

## **Principle 5D: Data classification and protection**

AIs remain ultimately responsible for the security, protection, and integrity of their cloud-hosted data, regardless of the type of cloud arrangements (e.g. outsourced or internally managed) in the process. Accordingly, AIs should systematically classify data based on their sensitivity and business importance, and implement commensurate data protection measures, including layered security controls, as appropriate, across the entire data lifecycle. Areas that should be considered include encryption, cryptographic key management, data loss prevention and secure methods for data disposal.

### **Good practices**

- Classifying data by sensitivity and criticality, with particular attention to confidential business data and customer data, and applying appropriate protection measures such as de-identification, pseudonymisation, or anonymisation where appropriate.
- Encrypting data at rest, in transit, in use, and in backup with strong, up-to-date cryptographic algorithms and key lengths appropriate to its classification, and regularly reviewing cryptographic standards to address emerging vulnerabilities.
- Establishing secure, encrypted channels for migrating and operating servers, applications, and data in cloud environments.



- Maintaining comprehensive cryptographic key management policies and procedures, covering justification, secure generation, unique assignment, rotation, backup, deletion, and regular review throughout the key lifecycle.
- Exercising independent control over encryption keys, where practicable, through approaches such as Bring Your Own Encryption (BYOE), Bring Your Own Key (BYOK), or Hardware Security Modules (HSMs). Avoiding key reuse across platforms, and ensuring that the encryption key management practices are centralised, scalable, and integrated with IAM processes.
- Extending data loss prevention and rights management controls to cloud-hosted data, and applying network segmentation and endpoint protection for user devices accessing cloud services.
- Regularly evaluating CSP multi-tenancy and logical separation controls through technical testing and validation of isolation mechanisms at compute, storage, and network layers, to ensure that tenant-specific protections effectively prevent cross-tenant data exposure.
- Ensuring a secure disposal at the end of a data lifecycle by working with CSPs to implement robust data erasure and proper disposal or reuse of hardware and virtual resources, supported by written confirmations and periodic audits, and obtaining final assurance at service exit that all sensitive data have been irreversibly removed with residual risks minimised.

### **Principle 5E: Data residency and regulatory assurance**

AIs should ensure that data residency and regulatory obligations are addressed when selecting cloud storage and processing locations, in particular for customer data. This includes maintaining visibility of data locations within CSP infrastructures, assessing cross-boundary transfer risks, and managing subcontracting arrangements to ensure compliance with applicable laws and regulatory requirements.

#### **Good practices**

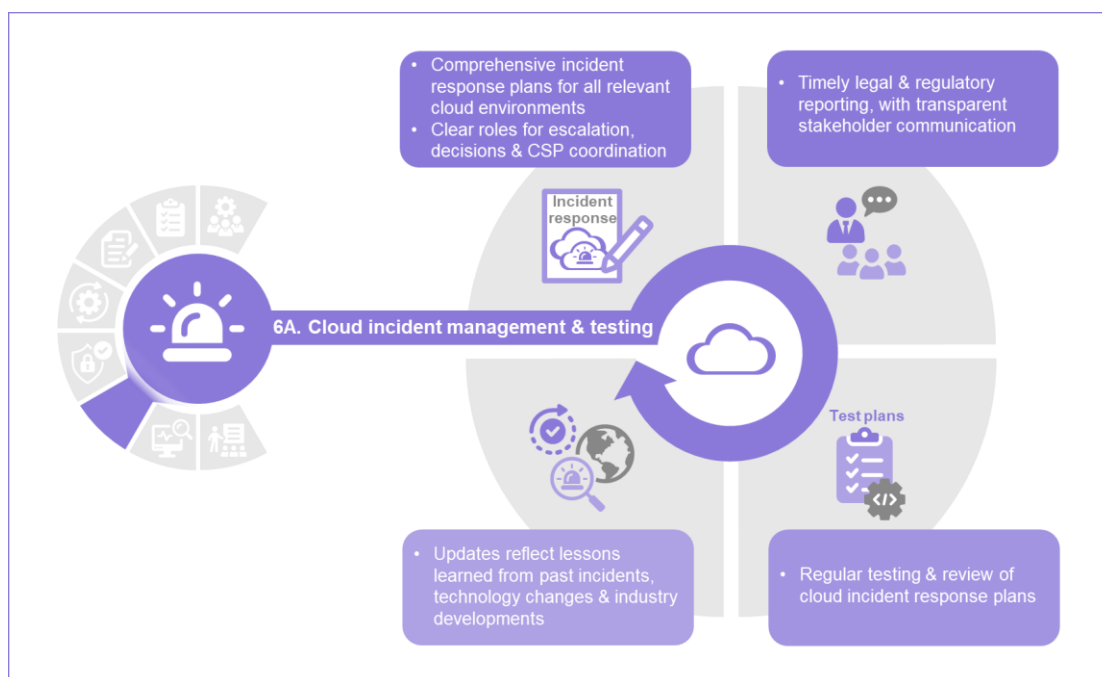
- Conducting periodic jurisdictional assurance reviews of in-use cloud services to confirm that data storage and processing arrangements remain compliant with evolving data-residency laws and regulatory expectations. This may include monitoring legislative changes in relevant jurisdictions, reassessing cross-boundary data risks, and engaging CSPs to validate that storage or processing locations continue to align with approved regions.
- Establishing technical and procedural controls to maintain visibility of actual data locations, such as leveraging CSP-provided location APIs, security dashboards, or data-mapping automation, and validating reported jurisdictions through periodic assurance reviews or third-party audits to evidence compliance with residency restrictions.
- Assessing data-related risks arising from subcontracting chains, particularly confidentiality and compliance risks where subcontractors operate in different jurisdictions, including evaluation of the CSP's due-diligence



process and transparency regarding subcontractor data-handling or further-transfer arrangements.

- Obtaining legal or specialist advice to assess cross-boundary data transfer arrangements, ensuring that mechanisms such as contractual clauses, transfer impact assessments, or approved certification schemes appropriately address applicable data protection and regulatory requirements.

## 6. Incident management



### **Principle 6A: Cloud incident management and testing**

AIs should ensure that their incident response plans adequately incorporate cloud-related risks and dependencies across private, public, hybrid and multi-cloud environments. The plans should clearly set out the actions an AI should take in the event of a disruption involving a CSP, such as outages in cloud infrastructure, security breaches in CSP-managed services, or loss of access to critical cloud-hosted data. The plan should align with and form part of the AI's broader incident and crisis management framework.

Amongst others, the plan should set out who bears responsibility for incident escalation, decision-making, communication, as well as coordination with relevant internal and external parties, including CSPs, to support timely investigation, resolution, and regulatory reporting. AIs should also ensure that all legal, regulatory, and contractual reporting obligations relevant to cloud incidents are fulfilled promptly and accurately, with appropriate support and cooperation from



CSPs, and that communication with affected stakeholders is transparent and consistent with applicable requirements.

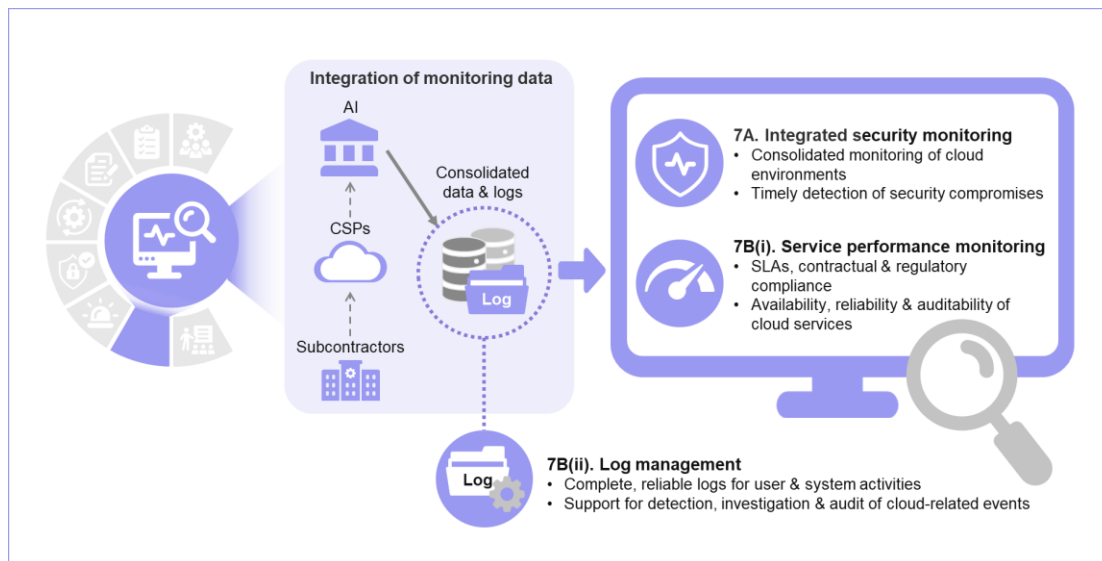
AIs should ensure the continued effectiveness of their incident response plans through regular reviews, testing, and timely updates to reflect lessons learned from incidents, system and technology changes, and relevant industry developments.

### Good practices

- Aligning the incident response plan with the AI's broader crisis management framework through actionable mechanisms (e.g. joint escalation matrices, coordinated response structure with clearly defined CSP contact points, and consolidated dashboards for executive oversight).
- Defining objective escalation criteria (e.g. data-breach severity levels, outage duration, regulatory-reporting triggers) and naming decision owners at each stage to ensure timely and consistent responses.
- Mapping out notification pathways and timeframes for regulatory authorities, law enforcement bodies, customers, and other stakeholders using predefined templates to maintain accuracy and timeliness in communication.
- Maintaining cloud-focused playbooks (e.g. CSP contact information, fallback arrangements, coordination steps for hybrid, multi-cloud, and other interconnected environments with third-party systems) to support rapid coordination and response during cloud-related disruptions.
- Regularly assessing the performance of communication and escalation channels with CSPs to ensure they remain reliable, secure, and able to support timely coordination during incident management and recovery.
- Establishing structured internal reporting workflows and templates that enable effective consolidation and validation of incident information across all relevant teams within the AI, ensuring that any regulatory or stakeholder notifications are accurate and consistent.
- Maintaining up-to-date operational procedures and supporting arrangements (e.g. trained staff, validated contact lists, tested communication channels) to ensure the effective execution of the notification process and readiness to engage with regulatory authorities, law enforcement bodies, and other stakeholders when needed.
- Regularly testing the incident response plan through scenario-based exercises across a range of plausible situations (e.g. security breaches, outages, misconfigurations) and, where applicable, involving CSPs in joint drills to assess coordination in cloud-specific disruptions.
- Using findings from tests and exercises, lessons learned from incidents, and insights from industry-wide events to update the plan, refine communication protocols, and strengthen staff training.



## 7. Ongoing monitoring



### **Principle 7A: Integrated security monitoring**

AIs should conduct ongoing monitoring of their CSP arrangements to ensure that any security compromises are detected in a timely manner. Accordingly, AIs should implement integrated security monitoring across their cloud environments, consolidating data from multiple sources to support timely alerts, centralised visibility, and proactive detection of security threats and other cloud-related anomalies. Monitoring should, where feasible, cover key security and operational layers and extend to CSPs and their subcontractors. Security operations centre (SOC) strategies should make use of monitoring data made available by CSPs and, where possible, subcontractors, in combination with the AI's own data, to enable comprehensive analysis and timely escalation and response to detected events. Monitoring capabilities should be regularly evaluated and optimised to ensure alerts remain reliable, actionable, and effectively integrated into the AI's incident response processes.

#### **Good practices**

- Consolidating and correlating monitoring data from diverse sources (e.g. activity and access logs, network flow records, API audit logs, endpoint alerts, CSP-provided alerts) into a central analytics platform or Security Information and Event Management (SIEM) solution to enable near real-time cross-domain visibility.
- Defining monitoring coverage across all key security layers, such as network (firewall and flow data), workload (host agent telemetry), application (authentication and API activities), and storage (data access and encryption events), with clearly assigned ownership for each layer.



- Integrating CSP monitoring and alerting data with the AI's internal logs using secure APIs or cross-account log ingestion pipelines to achieve end-to-end visibility of cloud environments.
- Operationalising SOC workflows to include cloud-specific triage playbooks, automated alert enrichment, and predefined escalation criteria that link directly to the AI's incident response procedures.
- Periodically calibrating monitoring rules, detection logic, and alert thresholds in collaboration with cloud engineers and threat-intelligence teams to reduce false positives or negatives and align with evolving attack patterns or service configurations.
- Using monitoring outputs to enhance operational procedures, refine controls, and strengthen preventive measures, ensuring that lessons learned from monitoring insights are systematically integrated into incident management, risk mitigation, and continuous improvement processes.

### **Principle 7B: Service performance and log management**

Beyond security, ongoing monitoring should extend to service performance and log management too. For service performance, AIs are expected to monitor compliance with SLAs, contractual obligations, and regulatory requirements, with particular attention paid to critical or important outsourced functions. This includes monitoring availability, reliability, and auditability, supported, where appropriate, by independent validation tools, and maintaining awareness of CSP service updates and lifecycle events to enable timely risk assessment and mitigation.

For log management, AIs should establish cloud-specific approaches that ensure the completeness, integrity, and availability of log data across relevant systems and services. Logs should provide reliable visibility into user and system activities, including privileged or administrative operations where applicable, to support timely detection, investigation, and audit of cloud-related events in accordance with regulatory and institutional requirements.

#### **Good practices**

- Leveraging structured dashboards, automated performance analytics, and periodic service review meetings to assess CSP performance trends, identify early signs of degradation, and escalate recurring SLA breaches or reliability issues for remediation.
- Using independent tools and validation methods (e.g. uptime monitoring, network performance measurement, independent log collection and analytics platforms) to cross-check CSP dashboards and strengthen assurance of service availability, reliability, and auditability.
- Maintaining oversight of CSP operational updates and lifecycle events (e.g. maintenance schedules, major service modifications, disruptions, feature deprecations, decommissioning notices) to support proactive risk assessment



and mitigation, including pre-change impact reviews and post-event reporting.

- Establishing comprehensive log management practices tailored to cloud environments, covering infrastructure, services, and privileged operations, with clear ownership for log generation, collection, retention, and review across internal teams and CSP interfaces.
- Applying encryption, integrity checks, and access controls to ensure logs are tamper-resistant and securely retained, with retention periods defined in proportion to system criticality, and implementing automated alerting for log integrity failures or unauthorised access attempts.
- Aligning log management practices with applicable regulatory requirements, industry standards, and the AI's overall risk management framework.

## 8. Workforce strategy and empowerment



### **Principle 8A: Workforce strategy and resource management**

AIs should maintain adequate in-house competencies and resources at all levels, from the Board, senior management, and staff across the three lines of defence, to carry out effective oversight and operations of cloud computing arrangements. To support this, the Board and senior management should have sufficient technical and managerial knowledge to understand the risks and implications, while operational teams should have the expertise to implement and sustain cloud strategies. Workforce planning should also be responsive and adaptive to the AI's evolving cloud journey. Engagement of external expertise should be strategic and well-controlled, complementing internal capabilities rather than substituting an AI's accountability in managing critical cloud functions.



### Good practices

- Defining and maintaining documented competency frameworks and role-based training pathways that translate the AI's oversight expectations into concrete skill and knowledge requirements for the Board, senior management, and staff within the three lines of defence.
- Conducting regular skills-gap analyses across all relevant functions, including business, technology, risk, legal, and compliance, to identify current and emerging competency needs, with emphasis on technical, operational, and risk-management aspects of cloud adoption.
- Developing and periodically reviewing workforce plans informed by these analyses, ensuring staffing levels and profiles reflect cloud maturity, operational demands, and the AI's strategy and risk appetite, and embedding succession or knowledge retention measures where appropriate.
- Adjusting workforce priorities as cloud adoption advances, including reliance on specialist expertise during initial migrations and greater focus on resilience, optimisation, and risk management in later stages.
- Engaging external expertise strategically where internal skills are not immediately available, under defined scopes that complement internal teams, with competency checks and structured knowledge transfer to maintain capability after contract completion.
- Avoiding sustained reliance on external personnel for core cloud functions, retaining key responsibilities, such as security oversight, monitoring of critical workloads and CSP evaluation, within the AI or, where appropriate, the wider banking group, through periodic review of internal versus external resource balance.

### **Principle 8B: Workforce training and competency development**

To operationalise and sustain the workforce strategy, structured and role-specific training programmes should be established to ensure that the Board, senior management, and staff across the three lines of defence maintain up-to-date knowledge and skills for effective cloud oversight and resilient operations. To achieve this, training should be provided on a regular basis and should address core dimensions of cloud management. These programmes should be continually updated to reflect the latest developments in cloud technology, regulatory expectations and operational needs.

### Good practices

- Designing training programmes with clear objectives, tailored curricula, and measurable outcomes to ensure that all roles involved receive targeted development covering governance, security, compliance, and operations.
- Providing training on a regular basis and ahead of significant cloud adoption or migration events to ensure that staff are equipped for new responsibilities.



- Developing appropriate cloud proficiency, including expertise across multiple CSP platforms where relevant, with training that covers both common functions and CSP-specific features, to ensure effective management and oversight of cloud and outsourcing arrangements, and maintaining a skills inventory to track multi-platform competency levels across teams.
- Regularly reviewing training needs with inputs from management, staff, and, where feasible, CSPs, and ensuring that the review reflects evolving technologies, regulatory requirements, and the AI's business priorities.
- Monitoring relevant professional certifications held by staff and ensuring that they remain aligned with the AI's cloud strategy and operational needs, with renewal tracking and encouragement of advanced certifications aligned to emerging technologies (e.g. containerisation, cloud security automation).
- Using training evaluations, participant feedback, and lessons learned from incidents to continuously enhance training programme content, delivery methods, and frequency.
- Ensuring that external personnel, including CSPs and subcontractors, are subject to training arrangements that build awareness of security, incident response, and applicable regulatory requirements, and verifying that such arrangements are adequate, for example, by reviewing training evidence or including such expectations in contractual terms.