ASIA INTERNET COALITION

25 July 2022

Pengarah
Unit Pakar Risiko dan Penyeliaan IT
Bank Negara Malaysia
Jalan Dato' Onn
50480 Kuala Lumpur

Subject: Asia Internet Coalition (AIC) Submission on Bank Negara Malaysia (BNM) - Exposure Draft on Cloud Technology Risk Assessment Guideline (CTRAG)

The Asia Internet Coalition ("AIC or We") and its members express our sincere gratitude to the Bank Negara Malaysia (BNM) for the opportunity to submit comments and recommendations on **Exposure Draft on Cloud Technology Risk Assessment Guideline (CTRAG)**.

AIC is an industry association comprised of leading internet and technology companies in the Asia Pacific region with a mission to promote the understanding and resolution of Internet and ICT policy issues in the region. In the past, AIC has worked extensively on key policies in Malaysia such as the Review of the Personal Data Protection Act 2010 (2020), the Anti-Fake News Act (2018), and Bank Negara Malaysia's Merchant Acquiring Services (2020) as well as Risk Management in Technology (2018), and has submitted recommendations and best practices to ensure that the industry voice is reflected in the regulatory approach. Our member companies would like to assure BNM that they will continue to actively contribute to the digital economy goals of Malaysia and support the rapid adoption of technology by the financial service institutions (FSIs).

We understand that the proposed guideline complements the Risk Management in Technology (RMiT) policy document to strengthen financial institutions' cloud risk management capabilities. We also recognize the importance of cloud computing services that can help FSIs to reinvent and optimize their relationship with technology, quicken go-to-market access, automate and strengthen security, improve customer experience, and lower costs, compared to traditional IT models. We hope to bring that innovation and security empowerment capability to Malaysia's FSIs.

As such, please find **appended to this letter detailed comments and recommendations**, which we would like BNM to consider when reviewing the Exposure Draft. We are grateful to BNM for upholding a transparent, multi-stakeholder approach in developing the **Draft on Cloud**

Technology Risk Assessment Guideline. We would also greatly appreciate the opportunity to discuss our feedback at BNM's convenience.

Should you have any questions or need clarification on any of the recommendations, please do not hesitate to contact me directly at Secretariat@aicasia.org or +65 8739 1490. Thank you for your time and consideration.

Sincerely,

Paine

Jeff Paine
Managing Director
Asia Internet Coalition (AIC)

Detailed Comments and Recommendations

---

1. Part A, Section 5(c)(ii) – Contract Management

ii) understand the scope of local customer protection legislation and regulatory requirements as well as to ensure that the financial institution's customers receive adequate protection and recourse in the event of a data breach by the cloud service provider; and

Comments

A financial institution's end customers do not have a contractual relationship with the financial institution's cloud service providers (CSPs or any of its other outsourcing providers). Recourse for any data breach should therefore be between the financial institution and cloud service providers.

Cloud computing services operate within a shared responsibility model where for example the cloud service provider is responsible for the security "of the cloud," and a financial institution is responsible for the security "in the cloud." By design, cloud service providers do not have visibility or control of the data stored by a financial institution, therefore they would not be able to distinguish the data uploaded to its services or determine to which user it is attributable.

Therefore, in the unlikely event of a data breach, any recourse should remain with the financial institution as they would be best placed to investigate the data breach, communicate with the end customer, and determine appropriate recourse. It would not be possible for cloud service providers to provide recourse to end customers directly.

Recommendation: We suggest the following amendment –

understand the scope of local customer protection legislation and regulatory requirements to ensure that the financial institution~~'s customers~~ receives adequate protection and recourse, for the benefit of its customers, in the event of a data breach by the cloud service provider

2. Part B, Section 1(c)(i) – Cloud architecture

use immutable infrastructure[3] for deployment to reduce the risk of failure when new deployment of applications enter production by creating a new environment with the latest version of the software. The on-going monitoring of the cloud environment should include automating the detection of changes to immutable infrastructure to combat evolving cyber-attacks;

> *[3] Immutable infrastructure is an infrastructure paradigm where servers are never modified after deployment. The servers are replaced rather than changed.*

Comments

We support this provision and, in fact, recommend financial institutions to adopt immutable infrastructure practices with no human access to better meet their audit and compliance needs in a well architected framework.

Recommendation: To prevent misinterpretation of this provision, we suggest amending footnote 3 to *"Immutable infrastructure in the context of cloud computing refers to a paradigm where the cloud service subscribers' virtual infrastructure components (virtual servers, virtual network etc) are never modified after deployment. Should a new version of services or application require changes in the underlying infrastructure components, their instance(s) are replaced rather than changed."*

3. Part B, Section 4(c) – Change Management

A financial institution should establish a process to systematically manage releases by cloud service providers in relation to existing infrastructure, network, upstream and downstream systems to minimize the impact of any service disruption.

Comments

The wording of this clause may give the impression that financial institutions should vet and approve new releases of cloud infrastructure, network, services, and systems. As hyperscale cloud service providers offer a global standardized infrastructure to users ranging from thousands to millions of users globally, it is not practical or feasible for providers to manage releases on a per-customer basis. Conversely, because of this hyperscale approach, cloud service providers are able to implement patches and hotfixes to remediate vulnerabilities and counter cyberattacks at speed and scale for all users.

Recommendation:
To change Part B, Section 4(c) to add “manage possible consequences due to new” before releases, as follows –

“A financial institution should establish a process to systematically assess and take appropriate action on releases by cloud service providers in relation to existing infrastructure, network, upstream and downstream systems to minimize the impact of any service disruption”

4. Part B, Section 5(d)(v) – Cloud Backup and Recovery

A financial institution should assess the resilience requirements of the cloud services and identify appropriate measures that commensurate with the criticality of the system, to ensure service availability in the extreme adverse scenarios. To ensure service availability, financial institution should consider a risk-based approach and progressively adopt one or more of the redundancy approaches, including diversifying away from a single cloud service provider. Amongst the viable options are:

> *v) adopt multi-cloud strategy, with the use of services from different cloud service providers to mitigate concentration risks and geopolitical risks.*

Comments

Similar to security and compliance, cloud service providers approach resiliency as a shared responsibility. Cloud service providers are responsible for ensuring that the services used by customers—the building blocks for their applications—are designed to be continuously available, as well as ensuring that they are prepared to handle a wide range of events that could affect their infrastructure. We support BNM’s emphasis on a risk-based and progressive approach to resiliency. In that context, we wish to take the opportunity to share key principles and its implementations.

We believe that financial institutions should ensure that they—and the critical economic functions they perform—are resilient to disruption and failure, whatever the cause. In the design, building, and testing of their applications, customers are able to achieve their objectives for operational resilience.

CSP infrastructure guards against outages and incidents, and accounts for them in the design of their services—so when disruptions do occur, their impact on customers and the continuity of services is as minimal as possible. To avoid single points of failure, CSPs minimize interconnectedness within theie global infrastructure. Regions are isolated from each other, meaning that a disruption in one Region does not result in contagion in other Regions. Compared to global financial institutions' on-premises environments today, the locational diversity of cloud infrastructure greatly reduces geographic concentration risk.

CSPs also employ compartmentalization throughout our infrastructure and services. They have multiple constructs that provide different levels of independent, redundant components. Starting at a high level, consider our Regions. To minimize interconnectedness, CSPs deploy a dedicated stack of infrastructure and services to each Region. Regions are autonomous and isolated from each other, even though CSP allow customers to replicate data and perform other operations across Regions. To allow these cross-Region capabilities, CSPs take enormous care to ensure that the dependencies and calling patterns between Regions are asynchronous and ring-fenced with safety mechanisms.

In that context, CSPs build — and encourages its customers to build — for failure to occur, at any time. In that context, CSPs support financial institutions to architect resilient workload(s) that has the capability to recover when stressed by load (more requests for service), attacks (either accidental through a bug, or deliberate through intention), and failure of any component in the workload's components. Field teams, composed of technical managers, solution architects, and security experts, help financial institutions build their applications according to their design goals, security objectives, and other internal and regulatory requirements.

As mentioned in the shared responsibility model, financial institutions are responsible for deciding how to protect their data and systems in the cloud. There exist workbooks, guidance documents, and on-site consulting to assist in the process. Financial institutions not only benefit from controls and designs native to a cloud service provider's infrastructure, but also maintain full control of their deployment, security, and resiliency in the cloud.

Recommendation: Our members support their customers' choice to make decisions that best meet their needs, including using multiple IT environments across existing on-premises facilities and other cloud service providers. However, we caution against the recommendation of a multi-cloud strategy to mitigate risk. While multi-cloud could potentially reduce concentration risk to some extent, the technical, process, and resource complexity needed to support multiple cloud service providers can lead to decreased resilience overall. A multi-cloud approach can create challenges, such as increased costs, technical complexity, and additional specialist skillsets required to onboard and manage multiple cloud service providers.

If there is a business need to use multiple vendors (such as accessing the unique services of a second cloud service provider), and it outweighs the resulting complexity, financial institutions would typically deploy the majority of workloads in one primary cloud, and run a small percentage of their workloads with a second cloud. This would enable them to gain some experience and advantages without all the adverse consequences.

We further recommend that there should be minimum interference in a financial institution's third-party vendor management strategy (e.g., multi-cloud, hybrid, type of provider, etc.). Such decisions should instead be driven by distinct features that will be unique to each individual financial institution, e.g., their overall business strategy and risk profile.

5. Part B, Section 6 – Interoperability and Portability

Interoperability standards for cloud services continue to evolve such that porting data, related configuration and security logging across different cloud service providers may be challenging. To facilitate the smooth process of interoperability and portability between on-premise IT systems and alternate cloud service providers, financial institutions are encouraged to:

(a) ensure technical requirements for interoperability and portability are included in the contractual agreement with the cloud service provider to avoid vendor lock-in;
(b) maintain a list of cloud service providers and tools that are needed to facilitate a smooth transition;
(c) ensure usage of standardized network and communication protocols for ease of interoperability and portability with on- premise IT systems or alternate cloud platforms;
(d) ensure the use of common electronic data formats, where applicable, to ease the movement of data between cloud service providers or to on-premises IT system; and
(e) extend patch and EOL management to ensure technology solutions employed remain effective and protected against system vulnerabilities.

<u>Comments</u>

Data is a cornerstone of successful application deployments, analytics workflows, and machine learning innovations. A major benefit of the cloud is that it offers financial institutions with the ability move their data at any time to another vendor. For instance, cloud services are built to support both data migration into and out of the cloud. Many cloud service providers offer several tools to help move data between networks and technology partners. Their services are generally built on numerous open standards like SQL, Linux, and Xen. This flexible foundation enables customers to securely move information in and out of the cloud regardless of where that information is going, such as cloud-to-cloud or cloud-to-data center.

In addition, financial institutions can implement certain architectural techniques to improve the portability of their data and applications, such as:

- Use Docker containers that can be deployable virtually anywhere; build using microservices to reduce the "blast radius" of changes to parts of an application (enabling the testing of each one independently if changes are needed on a large scale).
- Have loosely coupled services, especially when using a service specific to a cloud service provider — building a façade for each service to swap it out as transparently as possible.
- Build applications on open standards like Xen, SQL, KVM, and Linux.

Recommendation: We do not recommend that the Guidelines propose including technical requirements for interoperability and portability as contractual terms between a cloud service provider and financial institution. As outlined above, there are existing and numerous technology pathways for financial institutions to architect their cloud workloads with portability as an objective. Applying such contractual terms as a regulatory requirement may also have the unintended consequence of raising the cost of cloud services. We propose BNM delete the contractual requirement and retain the recommendation for financial institutions to account for interoperability and portability in their risk management strategy. We believe this aligns with BNM's principal to empower financial institutions to choose their cloud provider that meets their business objectives and risk management strategy, including the ease of portability.

6. Part B, Section 7(b)(iii) – Exit Strategy

A financial institution's exit strategy should be supported by an exit plan that establishes the operational arrangements to facilitate an orderly exit from a cloud service provider, which include the following:

> *iii. obtain written confirmation from the cloud service provider or via an independent external service provider's attestation that all sensitive data has been completely removed and destroyed from the cloud service provider's facilities upon completion of the exit process;*

Comments

In general, whilst cloud service providers should be part of the exit planning process and supporting the customer in any actual exit activities, it is important that both the regulator and the financial institution acknowledge that the level of support expected from the cloud service provider in exit activities must exhibit a degree of reasonableness. The cloud service provider should make a baseline set of controls and support available to all customers, and any bespoke support should then be agreed separately with the customer. Cloud service providers operate at a hyperscale, meaning that requests to use specific software or technologies may not be feasible.

Specific to this clause, we agree on the importance for financial institutions to be able to demonstrate that upon completion of an exit process, all sensitive data has been completely removed and destroyed. To do so, financial institutions are responsible for deleting all sensitive data from its specific instance or environment, as cloud service providers do not have visibility of the data being stored in line with the shared responsibility model.

Financial institutions should seek out cloud service providers that have automated data deletion and data destruction procedures instead of seeking manual processes and or certifications of deletion. This can be demonstrated by the appropriate industry standards such as such as PCI DSS, ISO27001, ISO27017, ISO27018, ISO22301 and others. Financial institutions should ensure that they review a cloud service provider's independent third party audit reports where such controls and procedures are being tested and documented. With the implementation of scalable and automated programs individual attestations are not required and possibly redundant.

Recommendation: To change Section 7(b)(iii) to:
ensure that it deletes all sensitive data from the cloud service provider's services upon completion of the exit process and verify that the cloud service provider has processes and controls to completely remove and destroy the sensitive data from the cloud service provider's facilities ~~obtain written confirmation from the cloud service provider or via an independent external service provider's attestation that all sensitive data has been completely removed and destroyed from the cloud service provider's facilities upon completion of the exit process~~.

7. Part B, Section 8(c) - Cryptographic key management

For critical systems hosted on the cloud, financial institutions should retain ownership and control of the encryption key (themselves or with an independent key custodian), independent from the cloud service provider, to minimize the risk of unauthorised access to the data hosted on the cloud. As example, this could be achieved by deploying the hardware security module (HSM) on-premises or by utilising HSM-as-a-service from a different cloud service provider.

Comments

There is a general misunderstanding that physically separation requirements will provide better protection against unintended information or system disclosure, tampering, and unauthorized access compared to logically separated multi-tenant cloud environments.

However, when examining the most common attack vectors for unauthorized access — such as remote exploitation, human error, and insider threat — a physically separated environment does not reduce the risk profile. In fact, for any system that is accessible over a network or the internet, physical separation — such as placing them in a locked cage or a separate data center facility — does not inherently provide added security or control over the most important forms of access.

In addition, being prescriptive by mandating a separately hosted key management component(s) may have a number unintended consequence over and above costs. For example, if there is a high volume of key exchange traffic to a centralized HSM hosted remotely may lead to bottlenecks that impacts performance, or downtimes due to wide area network outages.

In the era of modern hyperscale cloud computing, logical security mechanisms meet and often exceed the security results of physical separation of resources and other on-premises security approaches. For instance, cloud service providers address the concerns driving physical separation requirements through the logical security capabilities they provide customers and the security controls they have in place to help protect customer data. The strength of that isolation combined with the automation and flexibility that it provides is on par with or better than the security controls seen in traditional, physically separated environments.

Recommendation: We recommend deleting physical separation requirements. To make this principles-based and let the financial institutions decide how best to address the risks, we recommend removing "(themselves or with an independent key custodian), independent from the cloud service provider" and to offer options by replacing "As example this could be achieved by deploying the" with "Examples include, but are not limited to," and removing "different" as follows –

For critical systems hosted on the cloud, financial institutions should retain ownership and control of the encryption key to minimize the risk of unauthorized access to the data hosted on the cloud. Examples include, but are not limited to, hardware security module (HSM) on-premises or by utilizing HSM-as-a-service from a cloud service provider.

8. Part B, Section 9(c) – Access Control

A financial institution should ensure access controls to all hypervisor management functions or administrative consoles for systems hosting virtualized systems are effectively implemented as per the requirements and guidance under the Access Control section of RMiT policy document. These controls should mitigate the risk of any unauthorised access to the hypervisor management functions and virtual machine.

Comments

The way this clause is worded may be interpreted as the need for financial institutions to be able to manage the virtualization layer of the cloud that is not accessible by its users to avoid introducing risks to other customers of cloud service providers. Such intervention would also be inconsistent with the shared responsibility model and compromise the security and integrity of other users' environment.

Recommendation: To add *"their tenant"* after ensure in the first sentence, as follows –

"A financial institution should ensure ***their tenant*** *access controls to all hypervisor management functions or administrative consoles for systems hosting virtualized systems are effectively implemented as per the requirements and guidance under the Access Control section of RMiT policy document. These controls should mitigate the risk of any unauthorised access to the hypervisor management functions and virtual machine"*

9. Part B, Section 11(b) – Distributed Denial of Service (DDoS)

The risk of a single point of failure (SPOF) may surface when a financial institution leverages solely on a cloud-based solution to mitigate DDoS attacks. As such, a financial institution is encouraged to engage alternative DDOS mitigation providers or establishing circuit breakers to avoid service disruption when the main DDOS mitigation provider is disrupted.

Comments

The encouragement for the use of an alternative DDoS provider to mitigate the risk could be interpreted as a requirement. In general, DDoS mitigation typically involves different layer. For example, layer 3 DDoS attacks involve network layer high volume attacks that is typically handled by the ISPs at the front-end, while layer 7 attacks involve low and slow attacks at application level that is handled by Intrusion Prevention Solution. A cloud-based solution already relies on multiple layers of DDoS protection. The use of an alternative DDoS provider to address single point failure across all the layers not only increases the cost and complexity, but could introduce unintended additional risks, technical and performance challenges unnecessarily. Further, ISPs and cloud-based DDoS mitigations are inherently fault tolerant and address single point failure.

Recommendation: To change wordings in the clause to be principle based that focus on the intent rather than being prescriptive. We propose to substitute 11 (b) with the following:

"The Financial Institutions should mitigate any single point of failure (SPOF) for solutions implemented to mitigate DDoS attacks."

10. Part B, Section 12 (a) 2 – Data Loss Prevention

ii) manage the expansion of the endpoint footprint if the financial institution allow staff to use their own devices to connect to cloud services.

Comments

We believe that the intent of this clause is to prevent employees, when using their own devices, from leaking the financial institution's data (knowing or unknowingly) into unauthorized storage locations and/or services. However, the risk would be the same if these devices allowed to connect to the on-premise services.

Recommendation: To address the issue of data leak prevention that reflects the intent, we propose to replace "connect to cloud services" to "access sensitive company data," as follows –

"manage the expansion of the endpoint footprint if the financial institution allow staff to use their own devices to access sensitive company data."