



afme/

asifma

sifma

Public Cloud Portability

GFMA White Paper

March 2025

Table of Contents

Executive Summary	3
What is the current cloud regulatory landscape?	3
What is portability?	3
What are the challenges and risks to its usage?	3
What is industry's recommendation?	3
The Regulatory Landscape	4
Defining Portability and Other Resilience Solutions	5
Methods to Achieve Portability and their Trade-Offs	7
Why Portability is Not a Comprehensive Resilience Strategy for Public Cloud	8
Conclusion	9

Executive Summary

Financial Institutions' (FIs') growing reliance on cloud services raises regulatory concerns about concentration risk and financial stability. To address this, regulators are mandating portability of data and services between cloud providers. However, this paper details why portability – or any other resiliency solutions – should not be prescribed, and that regulators should take a risk-based approach.

What is the current cloud regulatory landscape?

Regulatory Concerns: As FI cloud adoption increases, regulators are focused on enhancing security and resilience to mitigate risks like concentration risk and vendor lock-in.

Portability Mandates: Some regulators are already mandating solutions such as portability, which could pose technical and financial challenges for FIs.

What is portability?

Definition: Portability is an organization's ability to move applications, data, and systems across different hardware, software, and cloud environments.

What are the challenges and risks to its usage?

Challenges of Portability: While portability offers theoretical benefits, it introduces challenges such as technical infeasibility, potential capacity limitations during mass migrations, and increased operational risks, including technology and cybersecurity concerns.

Risks of Overreliance: Relying too heavily on portability as a resilience tool may compromise security and limit the benefits of cloud adoption.¹

Limitations of Portability:

- **Short timeframes are challenging:** In wide-scale outages, lack of access to CSP control planes can hinder quick migration.
- **Limitations in porting specialized services:** Proprietary capabilities of CSPs require extensive reconfiguration and testing and limit the ability to port these types of services.
- **Too much data:** Duplicating storage for large-scale data is impractical.
- **Legal impediments:** Jurisdictional differences and data localization laws impede portability.
- **Stifled innovation:** Imposing portability would limit CSPs' ability to differentiate themselves in an effort to homogenize services to ensure portability. This could limit innovation.
- **Bottlenecks:** In a crisis, there could be limits to compute liquidity and CSPs may lack capacity to accommodate mass migrations.
- **On-premise limitations:** Datacenter infrastructure cannot quickly scale to meet excess demand due to long investment cycles.
- **Expanded attack surface:** Duplicating environments increases vulnerability to cyberattacks.
- **Increased complexity:** Managing multiple cloud environments complicates processes and increases likelihood of errors.
- **Observability challenges:** Maintaining multiple cloud environments often requires third-party monitoring. These solutions tend to be hosted in the cloud and are susceptible to disruption, thereby increasing a firm's operational risk.

What is industry's recommendation?

Regulators should not mandate portability. Instead, they should adopt a risk-based approach, allowing FIs to use portability as one of several tools, tailored to FIs' specific needs and risk profiles.

¹ FIs move to public cloud because they can realize value for themselves and clients – increased business efficiencies, better services for customers, enhanced security and resilience, and ultimately, competitive advantage.

The Regulatory Landscape

Key Implication: There is evidence that some regulators may require financial institutions (FIs) to implement portability as a resilience tool to mitigate the risk of cloud concentration risk.

Regulatory Focus and Drivers

As cloud adoption increases, regulators are concerned about FIs' security and are developing guidance to enhance resilience and to mitigate cloud adoption risks such as concentration risk and vendor-lock-in. Some regulators believe current guidance falls short of safeguarding systemic risk, which is intensified by the limited number of CSPs in the market. This has led some regulators to mandate solutions such as portability, interoperability, and multi-cloud. Such mandates could pose technical, contractual, or business challenges to FIs.

Current Regulatory Landscape

The regulatory landscape for cloud adoption is evolving, with regulators increasingly focused on FIs' resilience strategies. This is reflected in their inquiries about FIs' cloud plans and public statements. Key examples include, the EBA's 2019 Guidelines on Outsourcing, ENISA's guidance on cloud security, and the Bank Negara Malaysia's Cloud Technology Risk Assessment Guidelines, etc. Globally, regulators like the HKMA², MAS³, and the UK PRA⁴ are setting similar expectations for portability and interoperability to mitigate vendor lock-in and enhance resilience.

Regulators have strongly suggested that FIs consider portability and interoperability to manage cloud adoption risks, however the trend of mandating these tools is concerning. Today, portability is a mandatory requirement for FIs in at least one jurisdiction. The Securities and Exchange Board of India (SEBI) explicitly requires firms to implement data portability and interoperability as part of their exit and data transfer strategy and to help avoid vendor lock-in and concentration risk.⁵ Similarly, the 2023 EU Data Act mandates CSPs to enable seamless data transfers and service switching within a short timeframe for their clients.⁶

² Hong Kong Monetary Authority (HKMA), [Guidance on Cloud Computing](#), August 31, 2022.

³ Monetary Authority of Singapore (MAS), [Financial Sector Cloud Resilience Forum](#).

⁴ Prudential Regulatory Authority (PRA), [SS2/21 Outsourcing and third party risk management](#), March 29, 2021 (updated November 15, 2024).

⁵ Securities and Exchange Board of India (SEBI), [Framework for Adoption of Cloud Services by SEBI Regulated Entities](#), March 6, 2023.

⁶ European Commission (EC), [European Data Act](#), December 13, 2023.

Defining Portability and Other Resilience Solutions

Before discussing the methods to achieve portability and their potential complications, it is critical to establish a baseline definition of portability and key public cloud resilience-related terms.

Portability vs. Interoperability

Portability is a firm's ability to move applications, data, and systems across different hardware, software, and cloud environments, with minimal modification and disruption.⁷

Interoperability is the ability of data, systems, or software to be processed by different services on different cloud environments.^{8,9} While interoperability facilitates the cooperative use of multiple environments, it does not inherently simplify the task of moving entire workloads from one infrastructure to another.

Key Implication: Interoperability does not automatically imply portability. For example an application might use common integration methods to interact with storage services across multiple cloud providers (interoperability) but may rely on cloud-specific control planes or configurations that prevent portability.



Multi-Cloud vs. Hybrid Cloud

Portability requires a **multi-cloud** strategy – (a firm using more than one CSP) which may also be **hybrid** (a combination of on-premise infrastructure with public cloud services). While many firms adopt multi-cloud strategies to access “best in class” services from different providers, true portability requires cloud environments to be suitable for general purpose use. Simply having multiple cloud environments does not guarantee the ability to port between them.

Key Implication: Having an appropriate multi-cloud strategy in place and designing applications to be portable are the choice of the FI. Providing services or data formats that can be interoperable is the responsibility of CSPs.

⁷ [FFIEC IT Examination Handbook – Architecture, Infrastructure, and Operations \(AIO\) IT Booklet](#), June 2021.

⁸ [NIST Cloud Computing Standards Roadmap](#), July 2013.

⁹ [FFIEC IT Examination Handbook – Architecture, Infrastructure, and Operations \(AIO\) IT Booklet](#), June 2021.

Recovery Capability vs. Exit Planning

During an operational outage, FIs will have defined recovery targets or set minimum service levels that must be maintained to avoid breaching impact tolerances for critical operations. However, portability is typically not a fast enough recovery capability to allow FIs to meet those targets or service levels. On the other hand, exit planning is the ability to leave a CSP, either in the short-term (stressed), or long-term (unstressed). Portability can play a role in both types.

Key Implication: In most cases, the timeframes required for portability will exceed the targets for recovering applications/products and services during an outage within impact tolerance.

Stressed Exit	Unstressed Exit
Mechanics: A forced exit to an alternative cloud platform. Applications are prioritized and accordingly redeployed to an alternative cloud platform	Mechanics: An unforced exit to an alternative cloud platform
Timeframe: Compressed timeframe to exit to an alternative cloud platform	Timeframe: Exit to alternative cloud platform based on the firm's controlled timeframe
Drivers: ○ Events such as regulatory, cybersecurity, contractual/relationship breakdown, service deterioration, cost inflation, or M&A ○ Business continuity/disaster recovery	Drivers: ○ Avoid vendor lock-in (reasons include pricing, features, performance) ○ Cost optimization ○ Innovation

Other Resilience Techniques

Other techniques to support resiliency that should be considered alongside portability include: a **placement strategy** of where to host applications in the public cloud to prevent concentration risk, **capacity management**, **regional failover** strategies, **backups**, alternative action plans for business processes and **resiliency testing**. All of these techniques can be included in resiliency plans when adopting public cloud. These other resilience techniques are separate to and distinct from wider tools aimed at addressing location risk, for example the EU's critical third party provider oversight regime.

Key Implication: By understanding these foundational concepts and tools, financial institutions can make informed decisions about when and how to leverage portability; and paves the way for more nuanced discussions of its merits and challenges.

Methods to Achieve Portability and their Trade-Offs

Key Implication: Portability for data, applications, and services can be achieved through various methods. One approach is to architect cloud-agnostic applications using third-party abstraction technologies. Another is using Infrastructure-as-a-Service (IaaS) for greater control over the stack to enable portable solutions. Both methods involve trade-offs in functionality, operations, security and resiliency.

1. Use of third-party abstraction techniques (e.g., Cloud Foundry, Heroku, or Kubernetes).

Abstraction technologies provide flexibility, scalability, and ease of use, but they also present challenges such as performance issues, increased security risks, and complex management requiring specialized expertise.

These challenges have limited their popularity and they now occupy a very niche part of the market:

- Abstraction layers struggle to keep up with the rapid evolution of cloud services.
- The limited number of providers of abstraction technologies could, ironically, lead to vendor lock-in of those suppliers.
- Limited abstraction providers could also lead to systemic concentration risk.

2. Develop the full technology stack as Infrastructure as a Service (IaaS) platform, etc.

In this model, the FI has the greatest control over its stack which, in some cases, allows them to design solutions with microservices, serverless compute clusters, standard APIs, cloud-agnostic data formats, and other open-source components. Techniques like containerization allow applications to run securely in virtualized environments across different cloud infrastructures.¹⁰ However, this is not always possible (e.g., if the container only authenticates through one control plane). Additionally, since containers do not offer as “clear and concrete of a security boundary as a [virtual machine]”¹¹, it would necessitate complex, container-specific security measures like additional vulnerability monitoring. More generally, any pure IaaS build has significant downsides in terms of development and operational effort, access to innovation, and cost implications. This approach can also reduce the ability to leverage unique capabilities offered by the CSP.

3. Implement an active-active resilient multi-cloud architecture

In this instance, applications are architected to run in active-active mode, allowing load sharing between two application instances and full transaction processing if one instance fails. Architecting this in a multi-cloud environment is possible, but introduces significant complexity as the underlying IT, security infrastructure, and application stacks differ and need to be customized to the cloud’s internal environment. This requires significant additional resources, bespoke software and security architectures for each cloud environment, IT and business process operations enhancements, complex BCP, and cyber incident management processes and contractual overhead—likely doubling implementation costs. Therefore, managing a multi-cloud portability solution may be impractical due to increased complexities and expense.

¹⁰ NIST Computer Security Resource Center Glossary – [Container](#).

¹¹ [NIST SP 800-190 Application Container Security Guide](#), September 2017.

Why Portability is Not a Comprehensive Resilience Strategy for Public Cloud

While portability offers theoretical benefits, it also introduces several complexities and challenges that limit how broadly it can be applied:

- Technical infeasibility, or the diminishment or full negation of inherent cloud benefits
- Potential capacity limitations of providers in the event of a mass migration
- Operational risk (including technology and cyber)

Firstly, portability is often not practically feasible without negating the benefits of public cloud:

- **Short Timeframes Are Difficult to Achieve:** In a wide-scale outage, the lack of access to CSP control planes could prevent an FI being able to port to a different CSP or back on-premise.¹²
- **Portability of Specialized Services is Limited:** FIs use different CSPs often for their unique capabilities, which are not comparable or easily substitutable. Migrating to another CSP could take months due to the need for configuring, testing, and redeveloping applications. Similarly, custom-off-the-shelf products used for innovation may be CSP-dependent and not portable.
- **Impractical Movement of Large Scale Data:** Duplicating storage is impractical because firms must manage consistency, recoverability (synchronous transactions make it impossible to achieve a recovery point objective in a distributed database), and third-party connections in both venues.
- **Legal Impediments:** Jurisdictional discrepancies exist between CSPs that could prevent portability due to a lack of comparable services in certain locations, data localization obligations, or conflicting regulations.¹³
- **Limited Flexibility/Choice:** Because portability would require CSPs to offer interoperable services at the 'lowest common denominator,' firms would be unable to choose the most innovative services or those best suited to their needs. Relatedly, imposing portability would reduce the incentive for CSPs to differentiate themselves, stifling innovation.

Secondly, there may be inadequate capacity in alternatives:

- **CSP Capacity Bottlenecks:** In a widespread loss of confidence in a CSP where many customers try to port data, applications, or services to a new CSP, adequate capacity could be unavailable.¹⁴ CSPs do not allow customers to reserve capacity in most of their services. This has been referred to as a compute liquidity problem or a stampeding herd scenario.
- **On-Premise Capacity is Limited:** The global supply chain of data center infrastructure is unlikely to cope with significant excess demand should firms rehouse significant workloads. Relevant investment cycles are multi-year; firms cannot scale up capacity even over months.
- **Limited potential for alternatives during live incident:** Even where capacity would normally be available, this may not be the case in the midst of a live incident, and could in any event risk contagion of contaminated data.

Thirdly, portability can increase the firm's operational risk, including cybersecurity risk:

- **Expanded Attack Surface:** Simultaneously running workloads on duplicate environments increases the attack surface. If identity and access management (IAM) services are designed to be interoperable, attackers may also find it easier to move laterally from CSP to CSP in a breach.
- **Process Complexity:** Stretching controls, observability tooling, human resources and processes (such as handling updates and patches) over multiple cloud environments increases complexity of a firm's management and raises the likelihood of error.
- **Observability:** Maintaining multiple cloud environments commonly requires third-party monitoring software rather than firm-build solutions. These solutions tend to be hosted in the cloud and therefore susceptible to disruption during incidents.

¹² AFME, "[Building Resilience in the Cloud](#)," September 22, 2021.

¹³ Ibid.

¹⁴ Ibid.

Conclusion

Many FIs have long relied on mainframes for critical workloads, and public cloud should be viewed similarly regarding concentration risk. FIs must deploy fit for purpose strategies to succeed in the marketplace and maintain resilience. While portability is part of risk management, it must be balanced with other resilience approaches. This balance should be based on the unique make-up of an FI's public cloud strategy and footprint.

An FI's resilience strategy for its data, applications, and systems in the public cloud should be guided by its risk appetite, not prescriptive regulations. Continuous evaluation of strategies are essential to address evolving technologies and emerging risks to maintain robust operational resilience.

We recommend:

- **Adopt a Principles-Based Regulatory Approach:** Allow FIs to tailor resilience strategies based on their specific operational needs and risk profiles, rather than prescribing uniform technical measures.
- **Consider Portability for Targeted Use Cases:** Portability may be considered where it adds demonstrable value, such as for high-risk or mission-critical applications.
- **Promote Collaborative Innovation:** Support industry collaboration to establish standards and best practices that enhance flexibility and resilience without stifling innovation.
- **Prioritize Transparency and Preparedness:** Emphasize the importance of transparent resiliency strategies and regular testing, ensuring institutions are well-prepared for disruptions without mandating specific technical solutions.