WHITE PAPER

anjuna

Financial Compliance Made Easy with Anjuna Seaglass™

Executive Summary

Cyber resilience, cloud data security, shared responsibility, and insider threat risks across IT supply chains have become top areas of focus for financial regulators worldwide. In the UK, the Prudential Regulation Authority (PRA), a regulatory body of the Bank of England, is responsible for supervising financial institutions. Recently, the PRA released its Supervisory Statement 2/21 (SS2/21) document, requiring financial firms to implement adequate governance and controls for outsourced systems, especially those of cloud providers. Chapter 7 of SS2/21, titled "Data Security," mandates that firms protect data-in-memory. This regulation applies to a wide range of financial firms, including UK branches of overseas banks and insurers, UK banks, building societies, PRA-designated investment firms, insurance firms, and reinsurance firms, including the Society of Lloyd's and managing agents.

Non-compliance with the PRA regulation can lead to severe penalties, making it critical for financial firms to understand and adhere to the regulation. One powerful control for protecting data-in-memory is Confidential Computing. However, implementing Confidential Computing can be challenging for financial firms due to several factors. Organizations often need to re-engineer their applications to take advantage of the technology, and Confidential Computing technologies are not standardized across cloud providers, making it difficult to support multi-cloud architectures (a PRA recommendation). Additionally, while Confidential Computing protects data in memory, organizations also need to protect data as it moves from memory to storage or to the network.

Anjuna Seaglass offers a straightforward solution for financial firms to implement Confidential Computing, providing robust protection for their data whether it's in memory, at rest, or in motion. Anjuna Seaglass has already proven its value by helping a multinational bank based in London meet compliance requirements rapidly, thus enabling the institution to pursue cloud transformation goals without any complications. By adopting Anjuna Seaglass, financial firms can meet PRA SS2/21 regulation, safeguard their data in all states, and securely leverage the cloud without extensive engineering efforts.

INTRODUCTION:

The Rise of Clouds and SS2/21

The increasing use of third-party cloud services providers, such as AWS, Microsoft Azure, and Google Cloud, in the financial industry has led to regulators worldwide expanding their controls. In March 2021, the Prudential Regulation Authority (PRA) of the United Kingdom issued Supervisory Statement 2/21 (SS2/21), titled "Outsourcing and Third-Party Risk Management,"[1] which reflects a significant shift in regulatory attitudes from a prevention approach to a harm reduction approach. Regulators have recognized that relying solely on preventative controls is inadequate to safeguard the banking system from potential risks. Instead, they must implement controls that can mitigate the impact of inevitable failures.

SS2/21 outlines several critical control objectives, including the need for continuity plans in the event of a stressed exit from an outsourcer (e.g. a third-party provider is acquired or becomes insolvent), the expansion of scope to encompass sub-outsourcers, and most importantly, the protection of data held in memory, also known as "data-in-use."[23] By placing a greater emphasis on data-in-memory protection, the PRA is aligning with other regulators, such as the Monetary Authority of Singapore[4] and the European Banking Authority[5], in developing controls against an increasingly popular attack vector.

Protecting Data-in-Memory

While most compliance and security professionals are familiar with data-in-transit and data-at-rest protections, fewer may be aware of what data-in-memory protection is and how to implement it. Data-in-memory refers to data that is currently being updated, processed, erased, accessed, or read by a system and is stored in a non-persistent digital state, usually in computer random access memory (RAM). Data-in-memory must be in cleartext for the system to operate on it. Attackers find memory-based attacks particularly attractive because even data that was originally encrypted at rest or in transit is in the clear in memory when

[1] Bank of England PRA 2022
[2] NCC Group 2023
[3] PWC 2021
[4] Monetary Authority of Singapore 2021
[5] European Banking Authority 2021

being processed and can be accessed with privileged access using simple built-in Linux commands. Furthermore, even if encrypted data isn't actively being processed, an attacker can find encryption keys in cleartext in memory that can be used to decrypt data.

Regulators have taken note of memory protection as an increasingly popular threat vector. For example, CircleCI suffered a breach when an attacker stole root credentials and accessed databases that were running plaintext in memory[6]. In another instance, Sysdig discovered a hack called SCARLETEEL[7] where the attacker used a privilege escalation attack in Kubernetes to steal proprietary intellectual property from an AWS customer.

The PRA has responded by including Section 7.11 of SS2/21, which specifically calls for data-in-memory protections: "The PRA expects firms to implement robust controls for data-in-transit, data-in-memory, and data-at-rest."

One way to safeguard data-in-memory is by using Confidential Computing technology. Confidential Computing employs Trusted Execution Environments (TEEs) built into manufacturers' CPU and memory architectures to isolate workloads from all other users and processes on a host. Intel's Software Guard Extensions (SGX) was the first to market in 2015, followed by AMD Secure Encrypted Virtualization (SEV) and AWS Nitro Enclaves in subsequent years. The major cloud providers have adopted these technologies and now offer various Confidential Computing services.

Challenges with Adopting Confidential Computing

Although cloud providers and chipset manufacturers are investing heavily, organizations face several challenges when adopting Confidential Computing. The first challenge is the complexity of implementing Confidential Computing to protect existing applications that need to be migrated to the cloud. Applications often require rewriting to meet specific Confidential Computing implementations, such as Intel SGX, which involves breaking the application into trusted and untrusted parts and implementing new interfaces and function calls that leverage SGX.

The second challenge is multi-cloud, which is required as a compensating control in SS2/21. According to Section 10.5, organizations must choose one or more cloud resiliency options, such as multiple or backup vendors, in material cloud outsourcing arrangements. While multi-cloud can increase resilience and reduce concentration risk, the implementation of Confidential Computing varies across chipsets and cloud providers, making it difficult and costly to run an application in a multi-cloud scenario. The application may need to be modified in unique ways for each chipset and cloud provider to support Confidential Computing.

The third challenge of adopting Confidential Computing is securing the chain of trust. While Confidential Computing can protect data in memory, data must be running inside a secure enclave to be protected. Safely getting data into and out of the secure enclave, such as to/from persistent storage or network, can be challenging. Some providers may require data-in-transit to be decrypted on the host before being sent into the enclave, while others may not provide a secure storage environment for persisting data-at-rest outside of an enclave. Additionally, some providers use a "VM model" for securing memory, where the trust boundary for the enclave is the entire VM. If multiple services run inside the VM, an attacker who penetrates the VM through the application layer can move laterally and attack other services as well, leading to further security risks.

Solving Confidential Computing Challenges and Achieving Compliance with Anjuna Seaglass

Until today, organizations looking to implement Confidential Computing had to tackle three significant challenges: rewriting existing applications, supporting multi-cloud, and securing the chain of trust. However, those who attempted to solve these issues through a DIY approach often failed due to a lack of viable solutions in the market to purchase, coupled with the specialized knowledge required and the significant investment of engineering time. Fortunately, Anjuna Seaglass provides a proven solution for financial firms to address each of these challenges.

Anjuna Seaglass helps enterprises create high-trust environments in the cloud where data is always encrypted and code is verified for authenticity. With Anjuna Seaglass, workloads remain confidential and trusted during execution, enabling financial institutions to embrace the cloud and migrate existing applications and regulated workloads without the threat of attackers or insiders accessing or altering code or data. Moreover, the platform obviates

[6] Anjuna Security 2023
[7] Sysdig 2023

the need to rewrite applications and instead provides a Confidential Runtime on which enterprises can simply run their applications as is, with no modification. Anjuna Seaglass works seamlessly with any major cloud provider, satisfying the PRA's multi-cloud requirement.

Anjuna Seaglass also ensures that code and data are protected with always-on encryption: in-use, at-rest, and in-transit, exceeding the PRA's requirements for protecting data-in-memory. This means that data processed in a secure enclave can be encrypted via a key available only to that enclave, securing the chain of trust and enabling the data to exit the enclave and be safely saved on disk or transmitted across the network. Taken together, Anjuna Seaglass addresses each of the challenges faced by financial firms looking to implement Confidential Computing, making it a valuable tool for any organization seeking to enhance its cloud security posture.

Anjuna Seaglass goes further to enhance compliance by offering cryptographic proof to "attest" that an organization's workload runs on secure Confidential Computing hardware and has not been modified since its creation. Anjuna Seaglass generates attestation reports that can be submitted to a regulator for cryptographic verification that workloads executed precisely as intended and only within an enclave. This hardware-based memory protection provides the most robust control to minimize risk and satisfy compliance requirements. Using Anjuna Seaglass, organizations can secure their applications and generate cryptographic proof in a matter of seconds, unlike the years of engineering effort required to develop a similar process independently.

Many Financial Firms are Subject to SS2/21 - Including Foreign Ones

The Prudential Regulation Authority (PRA) has a broad scope that affects many financial firms, including UK branches of overseas banks and insurers, UK banks, building societies, PRA-designated investment firms, insurance firms, and reinsurance firms. Failure to comply with SS2/21 may result in penalties, as demonstrated by the PRA's issuing its highest-ever fines in 2022[8]. In light of recent events, such as the second-largest bank failure in history at Silicon Valley Bank, its downstream effects on Signature Bank, and subsequent acquisition by HSBC[9], regulators are increasingly scrutinizing risks to the global financial system. Enforcement is expected to continue. Anjuna Seaglass reduces the cost of compliance significantly by lowering engineering investment and accelerating time-to-compliance.

Summary

Anjuna Seaglass offers a comprehensive solution for financial institutions aiming to comply with PRA SS2/21 by using Confidential Computing to secure their data-in-memory. With Anjuna Seaglass, organizations can surpass the PRA's requirements and fully protect their data in all states, enabling them to achieve compliance effortlessly.

To learn more:

- Read our Financial Services Case Study
- Read our Financial Services Two Pager

Get started with Anjuna Seaglass

[8] Bank of England PRA 2022

[9] HSBC 2023

APPENDIX:

Table of Controls Satisfied with Anjuna Seaglass

PRA SS2/21 Control	Anjuna satisfies by...	Anjuna provides evidence via...
7.10 The PRA expects firms to implement appropriate measures to protect outsourced data and set them out in their outsourcing policy... and, where appropriate, in their written agreements for material outsourcing...	Anjuna Seaglass enables Confidential Computing across all major cloud providers, enabling firms to comply with SS2/21 without worrying about whether their cloud provider of choice is supported.	Deployment of applications across multiple clouds
7.11 The PRA expects firms to implement robust controls for data-in-transit, data-in-memory, and data-at-rest. ... these controls may include a range of preventative and detective measures...	Anjuna Seaglass protects data-in-memory by facilitating easy implementation of Confidential Computing.	Attestation report proving applications running in secure enclaves
10.5 In material cloud outsourcing arrangements, the PRA expects firms to... decide on one or more available cloud resiliency options, which may include... multiple or back-up vendors	Anjuna Seaglass enables Confidential Computing across all major cloud providers, enabling firms to comply with SS2/21 without worrying about whether their cloud provider of choice is supported.	Deployment of applications across multiple clouds
8.6 The PRA expects firms to exercise their access, audit, and information rights in respect of material outsourcing arrangements in an outcomes-focused way, to assess whether the service provider is providing the relevant service effectively and in compliance with the firm's legal and regulatory obligations and expectations, including as regards operational resilience.	Anjuna Seaglass provides an easy-to-use "attestation report" cryptographically proving that memory is being protected in a secure enclave and that code running in the enclave has not been tampered with.	Attestation report proving applications running in secure enclaves
7.11 ...the ongoing monitoring of 'insider threats', (ie employees at the firm and at the third party who may misuse their legitimate access to firm data for unauthorised purposes maliciously or inadvertently). The term 'employee' should be construed broadly for these purposes and may include contractors, secondees, and sub-outsourced service providers (see Chapter 9);	Anjuna Seaglass facilitates the use of secure enclaves, which prevent unauthorized access to data-in-memory by cloud provider employees with root access.	Attestation report proving applications running in secure enclaves
10.13 The PRA does, however, expect firms to identify viable forms of exit in a stressed exit scenario, and give meaningful consideration to those that best safeguard their operational resilience, which may include but not be limited to: • bringing the data, function, or service back in-house/on-premises; • transferring the data, function, or service to an alternative or back-up service provider; or • any other viable methods.	Anjuna Seaglass enables firms to deploy Confidential Computing in any major cloud provider, eliminating the need to re-architect or re-build applications to work in a multi-cloud scenario. Stressed exits can be facilitated by simply deploying the Anjuna runtime with the workload in question in another cloud.	Deployment of applications across multiple clouds