digiDations

ATLAS Security Validation for ABS AASE 2.0

Mapping the ATLAS Security Validation Platform™ to the ABS Adversarial Attack Simulation Exercises 2.0

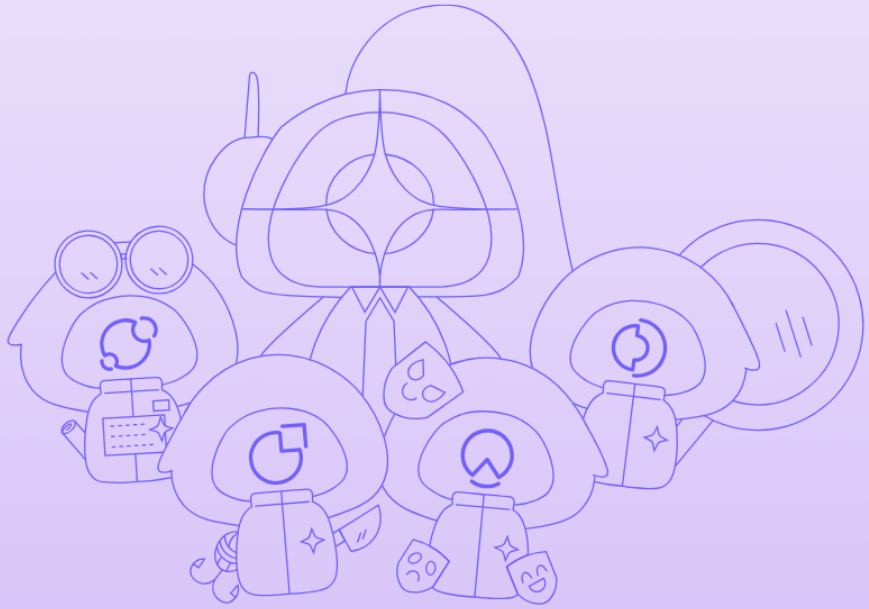

Table of Contents

Executive Summary **3**

Table 1: Protection Requirement (Clause 5) 5

Table 2: Methodology (Section 7) 11

Table 3: Documentation (Section 8) 17

Overview of ABS's Adversarial Attack Simulation Exercise 2.0 19

The Challenge of Demonstrating Security Control Effectiveness 20

ATLAS Security Validation Platform **21**

Overview of the ATLAS Security Validation Platform 21

ATLAS Security Validation Methodology 22

Capabilities Supporting Security Control Validation 23

Going Beyond Compliance **24**

Limitations of Traditional Security Assessments 24

Continuous Security Validation 25

Strengthening Cyber Resilience 25

About digiDations **27**

Executive Summary

Cyber security attacks against organization such as financial institutions (Fis) are evolving rapidly in scope, complexity, and sophistication. To address this risk, Fis deploy layers of defensive measures, solutions, and controls to reduce their exposure to attack and improve their response readiness.

AASE 2.0 are designed to challenge the FI's cyber security defences against real world adversary's Tactics, Techniques and Procedures (TTPs). The primary goal of the exercise is to assess the organisation's ability to prevent, detect and respond to cyber-attacks and identify potential gaps that may not be identified through traditional security assessment such as Vulnerability Assessment & Penetration Testing (VAPT); Penetration testing or vulnerability scans often provide only point-in-time insights and may not fully validate whether deployed security controls can detect, prevent, and respond to modern cyber threats. AASE 2.0 is intended to assist Fis in planning and execution of adversarial attack simulation but should not be relied on solely to achieve compliance and regulations.

This white paper explains how the **ATLAS Security Validation Platform**™ from digiDations support FIs in planning and execution of real world adversary simulation according to AASE 2.0. ATLAS uses adversarial simulation techniques to emulate real-world threat scenarios and evaluate the performance of security controls, detection rules, and incident response processes across the organization's environment.

The platform supports multiple areas highlighted in the AASE 2.0, including the **Guiding Principles, Methodology and Documentation.** Through continuous security validation, organizations can verify whether their defenses operate as intended, identify control gaps, and improve their ability to detect and respond to emerging threats.

By aligning continuous security validation with regulatory expectations, ATLAS enables FIs to move beyond theoretical compliance and gain measurable assurance of their cybersecurity posture. This approach strengthens operational resilience and helps organizations better defend against evolving threat actors targeting respective CIIs.

AASE 2.0 Guideline Mapping Table

The following tables provide a mapping between selected clauses from the Association of Banks in Singapore's Adversarial Attack Simulation Exercise (AASE 2.0) and the capabilities of the ATLAS Security Validation Platform.

The mapping illustrates how ATLAS can support FIs in the planning and execution of real-world adversary simulation exercises as outlined by ABS AASE 2.0. By simulating realistic adversary techniques and evaluating how security controls perform during these simulations, organizations can gain measurable insight into the operational effectiveness of their defences.

The mapping focuses on AASE 2.0 clauses most closely related to assisting in the planning and execution of adversarial attack simulations. This includes Section 6 (Guiding Principles), Section 7 (Methodology), and Section 8 (Documentation).

Table 1: Protection Requirement (Clause 5)

AASE 2.0 Section	AASE 2.0 Guideline	ATLAS Validation Capability	Validation Evidence
AASE 6.1	**Exercise Goals** **Exercise simulations should be based on the likely goals of real-world sophisticated adversaries.** **AASEs are goals driven. The intent of the exercise is not to identify and report vulnerabilities an FI may have. The goals represent what real-world adversaries may aim to achieve by compromising an FI.** **For example, an FI could be targeted by an adversary for financial gain and the objective could be to demonstrate an ability to perform a large unauthorised funds transfer. The scenarios will therefore be**	ATLAS contain list of scenarios/playbooks that simulates real-world adversary behaviors. These scenarios/playbooks are built based on latest threat intelligence of the adversary. On ATLAS, it is also possible to filter to show adversary that are targeting FIs.	N.A

	developed to include all the steps to reach that goal. It is not necessary to rotate or vary goals for every exercise as threat intelligence or other information may indicate that real-world adversaries could still have the same motivation. Scenarios and goals need only be changed if Threat Intelligence indicates that the threat landscape has changed, or controls around Critical Functions vary.		
AASE 6.3	Targeting of Live Critical Functions Exercise scenarios should be designed to target Critical Functions of the organisation and in a manner that is aligned with the motives of expected	The different components of ATLAS allow it to be deployed on a production environment and run simulation with no major impact on the current production assets.	N.A

	adversaries (in the production environment). Real-world adversaries are motivated to compromise Critical Functions to achieve maximum impact. Motives of adversaries can include business disruption, financial gain, and/or information theft. As such, the same functions should therefore be targeted by the Attackers in the exercise. Given the purpose of the exercise is to stress and enhance the FI's current resilience to a real-world adversary, the exercise should be conducted against the live environment. Assessments in staged or otherwise non-live environments may influence the outcome of the exercise and would not be representative of		

	organisational cyber resilience against real- world threat, unless the specific target is a development environment, or the team is testing for weaknesses in the protections between a test and production environment within the organization. **Assessment in a live environment would include, for example, the exploitation of staff cognitive biases, identification and exploitation of gaps within processes utilised by an organisation, and subversion of existing business process or technologies within an environment**		
AASE 6.4	**Exercise Frequency** **Organization should employ a risk-based approach to determine**	By utilizing ATLAS, FIs can launch simulation as and when required. This allows a continuous way of improving the Fis cyber resilience against latest threats.	N.A

	the appropriate frequency for AASE. **Factors such as the prevailing threat landscape, recent changes to FI's cybersecurity operations as well as results and lessons learned from past AASE may be considered in arriving at the suitable cadence.** **This guide recognizes that different FIs may be at different stages in their AASE journey and therefore recommends that AASE be performed once every 24 months minimally, but FIs have the flexibility to increase the frequency to every 12 months should this be deemed necessary.** **With each AASE iteration, the scenarios and sophistication of the TTPs may be adjusted in line**		

	with the improvements made in the organization's cybersecurity posture and changes in the threat landscape.		

Table 2: Methodology (Section 7)

AASE 2.0 Section	AASE 2.0 Guideline	ATLAS Validation Capability	Validation Evidence
AASE 7.1	**Planning Phase** **Prior to the start of the AASE, careful planning is required to ensure that the exercise can be conducted effectively and with minimal risks to the FI** **It is recommended that the planning stage includes the following steps:** - **Determine the scope and duration of the exercise** - **Define the key exercise parameters** - **Determine the participation model** - **Creation of the Exercise Working Group** - **Determine the escalation process**	ATLAS contain list of scenarios/playbooks that simulates real-world adversary behaviors. These scenarios/playbooks are built based on latest threat intelligence of the adversary. FIs may use this information to assist in their planning to determine the scope and goal of the exercise.	N.A

AASE 7.2	**Attack Preparation Phase** **During the attack preparation phase, the Exercise Working Group will develop attack scenarios that would impact Critical Functions or business continuity.** **Firstly, Threat Intelligence is gathered to determine the tooling and known capability of relevant threat actors of concern. Information is gathered, and the range of possibilities are expanded. With that information available, an attack model is created that helps define possible ways to attack the organization based on known TTPs, tooling and threat actor motivations. Finally, the possible attack paths from the attack model into attack**	ATLAS contain list of scenarios/playbooks that simulates real-world adversary behaviors. These scenarios/playbooks are built based on latest threat intelligence of the adversary. FIs may use this information to map it to their own gathered threat intelligence. In addition, if required, Fis may craft custom scenarios on ATLAS based on their own gathered threat intelligence.	N.A

	scenarios that broadly define how the exercise will be executed		
AASE 7.3	**Attack Execution Phase** **The Attack Execution phase involves the execution of the attack on the identified Critical Function based on the attack plan and scenarios that are formulated in the previous phase. During the attack execution phase, the Exercise Director should closely monitor the progress of the exercise and supervise the Attackers. While it is recommended for the Exercise Director to have full supervision at all stages, his involvement at certain stages where there is no impact such as reconnaissance can be reduced to a "keep informed" basis.**	ATLAS can run simulation as and when required. ATLAS also contains large amount of attack payloads that can be filter based on MITRE ATT&CK or Attack Kill Chain. Simulation on ATLAS would require selecting the source + destination assets for network-based attack and a single source asset for host-based attack. This allow simulation to be performed in a controlled and safe manner	N.A

	All steps taken by the Attacker and their observations should be continuously captured in the Exercise Log Report with the level of details determined in the requirements during the preparation phase. The Execution Log Report will also contain changes of plan, use of concessions and relevant approvals provided by the Exercise Director. As mentioned in the Risk Management Section, sharing the details of the attack with the Exercise Working Group carries a risk should the content be exposed unwittingly to either the Defenders or a real- world adversary. To reduce risks during the attack execution phase, the “test/halt engagement model” should		

	be applied. - **Test/Halt Engagement Model** - **External Reconnaissance and Perimeter Breach** - **Lateral Movement** - **Internal Information Gathering** - **Privileges Escalation** - **Persistence Access** - **Action on Objectives** - **Directing the Attack** - **Use of Concessions**		
AASE 7.4	**Exercise Closure Phase** **The exercise should be called to a close when either all outlined objectives have been met, or when the planned timeline has ended** - **Clean-up and Tactical Vulnerability Containment** - **Reconciliation** - **Attack/Defence Joint Replay**	Every attack payload in ATLAS would be paired with an undo command to for post simulation clean-up. For example, creation of Windows Registry Key will have a deletion of Windows Registry Key command. ATLAS also allows re-execution of any simulation if required.	

	- **Exercise Reports and Recommendations**	ATLAS can generate report to support FIs in crafting their exercise report.	

Table 3: Documentation (Section 8)

AASE 2.0 Clause	AASE 2.0 Requirement	ATLAS Validation Capability	Validation Evidence
AASE 8.1	**Attack Preparation Phase** - **Threat Intelligence Report** - **Attack Modeling Artifact** - **Exercise Preparation Report**	ATLAS contain list of scenarios/playbooks that simulates real-world adversary behaviors. These scenarios/playbooks are built based on latest threat intelligence of the adversary. FIs may use these information as the basis of the documentation required for the Attack Preparation Phase	N.A
AASE 8.2	**Attack Execution Phase** - **Execution Log Artifact**	Every simulation ran on ATLAS will be timestamped, this allows correlation with security product to identify potential gaps if required. In addition, ATLAS consist of a log gateway component which can retrieve logs from respective security product and perform correlation on simulated attacks. This allow FIs to have a quick overview of the effectiveness of the current security controls/products	Simulation result provides detailed timestamp of different payloads ran and the timestamp of it being logged by security product, if any.

AASE 8.3	**Exercise Closure Phase** - **Exercise Report** - **Clean-up Report** - **Defence Report** - **Remediation management Action Plan**	ATLAS can generate csv file of a given simulation to assist FIs in crafting the necessary report for this phase. This csv file will contain the payload ran, timestamp, block status, logged status, etc.	N.A

ABS Adversarial Attack Simulation Exercise 2.0

The **Association of Banks in Singapore's Adversarial Attack Simulation Exercise 2.0** serve as a guideline in supporting FIs in planning and executing of attack simulation exercise but should not be relied on solely to achieve compliance with regulations.

Cyber security attacks against organization such as financial institutions (FIs) are evolving rapidly in scope, complexity, and sophistication. To address this risk, FIs deploy layers of defensive measures, solutions, and controls to reduce their exposure to attack and improve their response readiness. The Adversarial Attack Simulation Exercise 2.0 issued by the Association of Banks in Singapore (ABS) underscores the need for Financial Institutes to conduct Adversarial Attack Simulation to assess and enhance the resilience of FIs against sophisticated attacks.

Overview of ABS's Adversarial Attack Simulation Exercise 2.0

The AASE 2.0 provides guidance for financial institutions (FIs) on conducting adversarial attack simulation exercises to assess and strengthen their resilience against sophisticated cyber threats.

AASEs are deigned to challenge an FI's cyber security defenses by modelling and then executing attacks based on real adversary Tactics, Techniques and Procedures (TTP). Scenarios are designed to be as realistic as possible, and may target the FI's People, Processes and Technology with the intent to compromise organization's Critical Functions.

The guideline is intended to support FIs within Singapore but may also be used by other organizations to plan and conduct such exercises. The document provides guidance on best practices and recommendations on how to adopt them partially or in full and depending on the maturity of the FI's capability in conducting these exercises.

The Challenge of Demonstrating Security Control Effectiveness

Financial institutions operating in Singapore's dynamic threat environment face a critical question that goes beyond deployment of security technologies: are these controls actually working?

Traditional security assessments such as vulnerability scans, penetration tests, and periodic audits have long formed the backbone of cybersecurity validation programmes. Yet these approaches share a fundamental limitation. They offer a snapshot of security posture at a single point in time, while adversaries operate continuously, refining their tactics, techniques, and procedures to bypass the very defenses organizations have invested in.

The Association of Banks in Singapore (ABS) recognized this gap in developing the Adversarial Attack Simulation Exercise (AASE) guidelines. By requiring financial institutions to validate their defenses through realistic, intelligence-led attack simulations, the guidelines shift the standard from passive control deployment to active proof of effectiveness, assessing not just whether controls exist, but whether they can prevent, detect, and respond to sophisticated real-world attacks. .

ATLAS Security Validation Platform

Organizations today operate in increasingly complex technology environments where traditional security assessments alone may not provide sufficient visibility into the effectiveness of deployed security controls. As cyber threats evolve, organizations require continuous methods to validate whether defensive technologies, monitoring mechanisms, and response processes operate effectively against modern adversary techniques.

The **ATLAS Security Validation Platform** from digiDations enables organizations to continuously test and validate the effectiveness of their cybersecurity defenses. By simulating realistic adversary techniques in controlled scenarios, the platform provides measurable insight into how security controls detect, prevent, and respond to potential attacks across the organization's environment.

Through continuous validation, organizations can identify gaps in detection coverage, evaluate the effectiveness of monitoring and response processes, and strengthen their overall cybersecurity posture.

Overview of the ATLAS Security Validation Platform

The **ATLAS Security Validation Platform** is designed to evaluate the operational effectiveness of security controls through adversarial simulation. The platform replicates techniques used by real-world threat actors and executes them safely within controlled testing scenarios.

These simulations enable organizations to observe how deployed security technologies — including endpoint protection platforms, network security controls, and security monitoring systems — respond to adversary techniques across different stages of the attack lifecycle.

Unlike traditional security assessments that provide only periodic insight, ATLAS supports **continuous validation**, allowing organizations to regularly test defensive capabilities as infrastructure changes and new threats emerge.

ATLAS Security Validation Methodology

ATLAS follows a structured validation methodology that helps organizations continuously evaluate and improve their cybersecurity defenses.

This validation process can be viewed as four stages:

Verify – Operational Design

Validate baseline configurations and confirm that foundational security controls are functioning correctly.

Evaluate – Security Effectiveness

Simulate realistic adversary techniques to determine whether deployed security technologies detect or prevent malicious activity.

Analyze – Threat Correlation and Security Operations Efficiency

Assess how effectively monitoring systems and analytics platforms correlate alerts and identify potential attack activity across the environment.

Optimize – Incident Response Readiness

Evaluate incident response processes and identify opportunities to improve detection accuracy and response speed.

Through this continuous validation cycle, organizations can measure the effectiveness of their security defenses and strengthen operational readiness.

Capabilities Supporting Security Control Validation

The ATLAS Security Validation Platform provides several capabilities that help organizations validate the effectiveness of their cybersecurity defenses.

First, the platform enables **realistic adversary simulation**, replicating techniques aligned with established attack frameworks such as MITRE ATT&CK. These simulations allow organizations to evaluate how security technologies detect and respond to specific adversary behaviors.

Second, ATLAS supports **continuous validation of security controls**, allowing organizations to regularly test detection mechanisms, monitoring systems, and defensive technologies against simulated attack scenarios.

Third, the platform provides **visibility into detection and response performance**, enabling security teams to identify gaps in monitoring coverage, detection logic, and incident response processes.

By providing continuous insight into how security controls operate under simulated attack conditions, ATLAS enables organizations to strengthen their defensive capabilities and maintain greater confidence in the effectiveness of their cybersecurity controls.

Going Beyond Compliance

Regulatory frameworks such as the **ABS Adversarial Attack Simulation Exercises 2.0** establish important expectations for managing technology risks and maintaining effective cybersecurity controls. However, compliance with regulatory requirements alone does not guarantee that an organization's defenses are capable of detecting or preventing modern cyber attacks.

As cyber threats evolve and adversaries develop increasingly sophisticated techniques, organizations must move beyond periodic assessments and compliance checks to ensure that their security controls, monitoring capabilities, and incident response processes remain effective.

Limitations of Traditional Security Assessments

Traditional security assessments such as penetration testing and red team exercises have long been used to evaluate an organization's cybersecurity posture. While these assessments can provide valuable insights into vulnerabilities and potential attack paths, they also have several inherent limitations.

Penetration testing engagements are typically conducted over limited timeframes and depend heavily on the expertise and experience of the testers involved. As a result, the range of techniques explored during an engagement may be constrained by time and available resources. In addition, testing activities conducted in production environments often require safeguards to prevent operational disruption, limiting the realism of the assessment.

Similarly, red team exercises simulate adversarial behavior to test detection and response capabilities, but these exercises are often performed infrequently—such as annually or quarterly. Because of their limited duration and scope, they may not provide continuous visibility into how defensive controls perform as systems evolve and new threats emerge.

As a result, these traditional approaches provide **point-in-time insights** rather than continuous assurance of security control effectiveness.

Continuous Security Validation

To address the limitations of traditional assessments, organizations are increasingly adopting **continuous security validation** practices. This approach involves regularly simulating adversary techniques to evaluate how security controls detect and respond to potential attacks across the organization's technology environment.

By replicating realistic attack techniques and observing how defensive technologies respond, security teams can identify gaps in detection coverage, validate the effectiveness of monitoring systems, and measure the readiness of incident response processes.

Continuous validation provides a more comprehensive understanding of an organization's security posture by testing defensive capabilities against a wide range of adversary techniques across multiple stages of the attack lifecycle.

Strengthening Cyber Resilience

Continuous security validation enables organizations to move beyond a purely compliance-driven approach to cybersecurity. Instead of focusing solely on meeting regulatory requirements, organizations can continuously measure and improve the effectiveness of their defensive capabilities.

The **ATLAS Security Validation Platform** supports this approach by enabling organizations to safely simulate real-world adversary techniques and evaluate how security controls perform during these simulations. By incorporating continuous validation into their cybersecurity programs, organizations can strengthen their ability to detect and respond to emerging threats while maintaining greater confidence in the effectiveness of their security defenses.

behavior across different stages of the attack lifecycle. These simulations enable security teams to observe how existing defensive tools and processes respond to real-world attack techniques.

Unlike traditional point-in-time assessments, ATLAS supports continuous validation of security controls across on-premises, cloud, and hybrid environments. The platform orchestrates controlled adversarial simulations and analyzes the resulting telemetry to determine whether security controls successfully detect, block, or respond to simulated attack activities.

The platform architecture includes multiple components that work together to execute simulations and collect validation results, including a centralized management console, simulation agents, telemetry collectors, and sandbox environments for executing high-risk attack behaviors safely. These components enable organizations to validate detection, response, and mitigation capabilities across their entire security architecture.

About digiDations

digiDations is an AI Native cybersecurity company helping organizations replace security assumptions with continuous, measurable validation. The question it answers is not whether an organization has vulnerabilities. It is whether their defenses can keep up with the threats targeting them right now.

Founded by cybersecurity veterans and AI researchers, digiDations built its product portfolio on Digital Mind, a proprietary AI foundation that continuously collects, correlates, and operationalizes real-world attack intelligence. The AI-Driven AEV portfolio includes ATLAS for Security Validation, APOLLO for Continuous Automated Red Teaming, and HELIOS for External Attack Surface Management. ORION for Threat Intelligence extends digiDations' CTEM coverage with adversary intelligence and visibility into emerging attack surfaces.

digiDations is building toward Autonomous Cyber Defense with TARA AI, the Autonomous CTEM Agent that self-directs across all four products: determining which scenarios to run, orchestrating execution, interpreting results, and generating reports, with teams approving decisions rather than defining the workflow.