Silverfort

Silverfort Identity Security for MAS Technology Risk Management

How Silverfort's identity security controls map to the MAS TRM Guidelines, the proposed TRM Notice amendments, and MAS guidance on AI agents



Executive summary

The Monetary Authority of Singapore (MAS) issued its Technology Risk Management (TRM) Guidelines to help financial institutions (FIs) build sound technology risk management and cybersecurity practices. MAS last revised the Guidelines on January 18, 2021. They apply to all FIs under MAS regulation, from banks, insurers, and exchanges to payment service providers and capital markets firms.

The Guidelines work alongside the TRM Notices, which set legally binding requirements. In June 2026, MAS proposed amendments that would extend the Notices to IT risk assessment, continuous monitoring, change control, and incident management. The amendments will take effect 12 months after MAS publishes the final Notices.

MAS has also extended its focus to artificial intelligence (AI). Its proposed Guidelines on Artificial Intelligence Risk Management (AIRG) cover generative AI and AI agents. In July 2026, the MAS BuildFin.ai program published Safeguards for Agentic Finance at Runtime (SAFR), an industry reference approach for controlling what AI agents do at runtime.

Addressing the Identity Security aspects of MAS TRM

Attacks on financial institutions increasingly start with an identity rather than a technical exploit. Attackers use compromised credentials to log in as legitimate users, then move laterally and escalate privileges until they reach critical systems and data. Non-human identities (NHIs), such as service accounts, API keys, and cloud roles, make this easier. They often hold broad privileges, use credentials that never rotate, and run without anyone reviewing their activity. AI agents now add a third group, and they often act through NHIs themselves. They authenticate, call tools, and act on behalf of users with delegated permissions that often exceed what their task needs.

The TRM Guidelines address this risk directly. Section 9 (Access Control) sets expectations for user access management, privileged access management, and remote access management. These include multi-factor authentication (MFA) for users with access to sensitive system functions and for all remote access, and monitoring of system and service accounts for suspicious or unauthorised activity. Identity also runs through other sections: risk assessment (Section 4), network security and lateral movement (Section 11), and cyber event monitoring and incident response (Section 12).

The proposed TRM Notice amendments would make several related practices mandatory. FIs would need to continuously monitor critical systems against defined alert thresholds, prevent unauthorised changes to any system, and preserve evidence for incident investigation. The existing requirement to protect customer information from unauthorised access remains in place.

MAS applies the same questions to AI agents. The proposed AIRG covers AI agents alongside other AI systems, and SAFR's runtime checkpoint starts with Agent Identity: every proposed action is bound to a recognised, registered agent before any other evaluation takes place. For Identity Security teams, AI agents now need the same discovery, ownership, least privilege, and monitoring as any other privileged identity.



Silverfort Identity Security Platform

Silverfort Identity Security Platform discovers human, non-human, and AI agent identities across on-prem and cloud environments, and analyses their exposures and activity. Its Runtime Access Protection (RAP) technology integrates with Active Directory (AD) to see every authentication, analyse its risk, and enforce access-based policies inline for workforce users, privileged users, third parties, and service accounts. This extends protection to resources that couldn't be protected before, such as legacy systems, command-line tools, and OT/IT infrastructure. For cloud NHIs, Silverfort builds geofencing policies and deploys them to Entra ID and AWS for enforcement. For AI agents, the platform applies runtime controls through native integrations with agentic platforms and the MCP Gateway.

Silverfort helps FIs address the identity-related expectations of the TRM Guidelines and the proposed TRM Notice amendments. It blocks the use of compromised credentials, stops lateral movement, and prevents unauthorised access to critical systems. It enforces least privilege for privileged users, AD service accounts, and AI agents, and gives security teams the continuous monitoring and incident evidence the proposed Notices call for.



Multi-factor authentication

Extend MFA protection to command-line access, legacy apps, IT infrastructure, and other critical resources that couldn't be protected before.



Detect and respond to identity threats

Detect common credential access, privilege escalation, and lateral movement attacks, and respond automatically with real-time blocking.



Securing privileged users and NHIs

Discover and classify all your privileged users and NHIs, enforce least privilege and adaptive access-based policies in AD, and geofence cloud NHIs in Entra ID and AWS.



AI agent security

Discover AI agents, map each one to its human owner and the NHIs it uses, and allow, flag, step up, or deny each tool call at runtime.

Mapping Silverfort capabilities to MAS TRM

The following tables map Silverfort's capabilities to the identity-related requirements in the MAS TRM Guidelines and the proposed TRM Notice amendments. Each table covers one section and lists the MAS requirement, quoted as published, alongside the Silverfort controls that support it.

TABLE 1:

TRM Guidelines Section 4, Technology risk management framework

Section 4 sets out how FIs should identify, assess, treat, and monitor technology risks. For identity, that means knowing which users, service accounts, and AI agents create risk, and tracking whether controls reduce it.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
4.2 Risk identification	
4.2.1 The FI should identify the threats and vulnerabilities applicable to its IT environment, including information assets that are maintained or supported by third party service providers. Examples of security threats that could have a severe impact on the FI and its stakeholders include internal sabotage, malware and data theft.	Silverfort continuously discovers human and non-human identities across AD, cloud identity providers (IdPs), SaaS applications, and cloud infrastructure. It detects identity exposures such as excessive privileges, misconfigurations, and dormant identities, and analyses the risk of every authentication and access attempt, including access by third parties. For AI agents, Silverfort identifies over-privileged and misconfigured agents and supply chain risks.
4.3 Risk assessment	
4.3.1 The FI should perform an analysis of the potential impact and consequences of the threats and vulnerabilities on the overall business and operations. The FI should take into consideration financial, operational, legal, reputational and regulatory factors in assessing technology risks. **4.3.2** To facilitate the prioritisation of technology risks, a set of criteria measuring and determining the likelihood and impact of the risk scenarios should be established.	Silverfort assigns a risk level to every authentication and to the users and resources involved. Each exposure gets a severity level that combines the likelihood that it leads to a breach with the impact of that breach. Exposures are mapped to MITRE ATT&CK tactics (all exposures in the cloud and most in AD), and storylines show the potential blast radius of each one. For AI agents, Silverfort prioritises risks by the sensitivity of what each agent can reach, its exposure level, and its potential blast radius. FIs can use these scores as identity inputs to their own likelihood and impact criteria.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
4.4 Risk treatment	
4.4.1 The FI should develop and implement risk mitigation and control measures that are consistent with the criticality of the information assets and the level of risk tolerance. The IT. control and risk mitigation approach should be subject to regular review and update, taking into account the changing threat landscape and variations in the FI's risk profile.	Silverfort enforces access-based policies inline, at the moment of authentication in AD. Each policy can trigger an alert, require MFA, or block access, based on the identity, the resource, and the risk of the access attempt. For AD service accounts, Virtual Fencing restricts each account to its approved sources, destinations, and protocols. For cloud NHIs in Entra ID and AWS, Silverfort deploys geofencing policies that the IdP enforces. For AI agents, FIs can set runtime policies by agent, user, application, and tool that allow, flag, or deny each tool call. FIs can apply stricter policies to their most critical systems and tighten them as their risk profile changes.
4.4.2 As there are residual risks from threats and vulnerabilities which cannot be fully eliminated, the FI should assess whether risks have been reduced to an acceptable level after applying the mitigating measures. The criteria and approving authorities for risk acceptance should be clearly defined and it should be commensurate with the FI's risk tolerance.	Silverfort continues to monitor identities and access after security controls are in place, so residual exposures stay visible. Its Authentication Logs record every authentication and access attempt, including the policy decisions and blocked access, so FIs can show whether controls are working between assessments.
4.5 Risk monitoring, review and reporting	
4.5.1 The FI should institute a process for assessing and monitoring the design and operating effectiveness of IT controls against identified risks. **4.5.2** A risk register should be maintained to facilitate the monitoring and reporting of technology risks. Significant risks should be monitored closely and reported to the board of directors and senior management. The frequency of monitoring and reporting should be commensurate with the level of risk. **4.5.3** To facilitate risk reporting to management, technology risk metrics should be developed to highlight information assets that have the highest risk exposure. In determining the technology risk metrics, the FI should take into account risk events and audit observations, as well as applicable regulatory requirements.	Silverfort Identity Security Posture Management (ISPM) gives centralised visibility into identity exposures across on-prem and cloud environments, with severity scores and guided remediation. Authentication Logs show the policy decision behind every authentication and access attempt, so FIs can check that controls operate as designed. To show whether remediation reduces exposure over time, the Executive Risk Report tracks trends in AD environments, and the new versus resolved exposures view does the same in the cloud. FIs can use these findings in their risk register and in reporting to the board and senior management.

TABLE 2:

TRM Guidelines Section 9, Access control

Section 9 sets out how FIs should manage user, privileged, and remote access, including least privilege, MFA, access reviews, and the monitoring of system and service accounts.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
9.1 User access management	
9.1.1 The principles of 'never alone', 'segregation of duties', and 'least privilege' should be applied when granting staff access to information assets so that no one person has access to perform sensitive system functions. Access rights and system privileges should be granted according to the roles and responsibilities of the staff, contractors and service providers.	Silverfort enforces access-based policies at the authentication layer in AD, so each user, contractor, and service provider can reach only the resources their role requires. Identity-based segmentation and deny policies limit access to sensitive systems to the identities that need it. For privileged users, Just-In-Time (JIT) access removes standing access by keeping admin accounts disabled until access is approved, and Virtual Fencing keeps admins within their assigned tier. For AI agents, which act with access delegated by staff, FIs can set runtime policies by agent, user, application, and tool, so each agent can reach only the tools and systems its task requires.
9.1.2 The FI should establish a user access management process to provision, change and revoke access rights to information assets. Access rights should be authorised and approved by appropriate parties, such as the information asset owner.	Silverfort evaluates every authentication in AD against the FI's access-based policies. These policies can trigger an alert, require MFA, or block access when an authentication to a protected system is insecure or risky. For privileged access, Silverfort's JIT workflow can validate each request against an approved ticket in the FI's IT service management (ITSM) platform (starting with ServiceNow). It confirms the approval, the requesting user, and the approved timeframe before granting access.
9.1.3 For proper accountability, the FI should ensure records of user access and user management activities are uniquely identified and logged for audit and investigation purposes.	Silverfort's Authentication Logs record every authentication and access attempt by each user and service account across on-prem and cloud environments, including the source, destination, protocol, risk level, and policy action. Policy changes are logged with the administrator who made them. The FI can send all of these logs to its SIEM over syslog or in OCSF format for audit and investigation.
9.1.4 The FI should establish a password policy and a process to enforce strong password controls for users' access to IT systems.	Silverfort detects stale accounts, passwords that never expire or are never rotated, admin accounts exposed to Kerberoasting, and weak Kerberos encryption that lets attackers crack password hashes offline. FIs can set exposure thresholds, such as maximum password age, to match their own password policy. Silverfort enforces MFA on AD authentications, so a compromised password alone isn't enough to access protected systems.
9.1.5 Multi-factor authentication should be implemented for users with access to sensitive system functions to safeguard the systems and data from unauthorised access.	Silverfort enforces MFA on authentications in AD for every user, from standard users to administrators. This extends MFA protection to resources and access methods that don't support MFA natively: legacy and homegrown applications, mainframes, databases, file shares, IT infrastructure, command-line tools such as PsExec, WMI, and PowerShell remoting, and local Windows logons. For FIs, this extends MFA to core banking, payment, and trading systems that authenticate through AD, with no changes to the systems themselves. FIs can use the MFA provider they already have, such as Microsoft Entra ID, Okta, Ping, or Duo, as the second factor. With Silverfort Identity Bridge, they can also apply their cloud IdP's existing access policies, including phishing-resistant methods such as FIDO2, to these AD authentications.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
CONTINUED: 9.1 User access management	
9.1.6 The FI should ensure appropriate parties such as information asset owners perform periodic user access review to verify the appropriateness of privileges that are granted to users. The user access review should be used to identify dormant and redundant user accounts, as well as inappropriate access rights. Exceptions noted from the user access review should be resolved as soon as practicable.	Silverfort continuously surfaces dormant and stale accounts, orphaned service accounts, and accounts with excessive privileges, ranked by severity. Security teams and asset owners get a prioritised list of what to fix first, with remediation guidance for each exposure. Silverfort also maps each identity's effective access in AD, including direct and inherited rights. Security teams can export privileged accounts and their protection policies for review. For AI agents, it identifies orphaned agents with no human owner and agents with excessive permissions.
9.1.7 Users should only be granted access rights on a need-to-use basis. Access rights that are no longer needed, as a result of a change in a user's job responsibilities or employment status (e.g. transfer or termination of employment), should be revoked or disabled promptly.	Silverfort ISPM detects shadow admins: accounts that hold admin-level privileges through access control list (ACL) inheritance or nested group membership and rarely appear in standard access reviews. When the FI connects its HR system, Silverfort also flags users who have likely left the organisation but still have active accounts in the IdP, along with NHIs and credentials they still own. JIT access keeps admin accounts disabled until access is approved, so privileged access exists only while it's needed. For AI agents, Silverfort flags orphaned agents with no human owner and agents whose underlying identities have gone dormant.
9.1.8 The FI should subject its service providers, who are given access to the FI's information assets, to the same monitoring and access restrictions on the FI's personnel.	Service providers and other third parties get the same security controls as the FI's own staff. Silverfort monitors every third-party authentication through AD, analyses its risk, and enforces the same access-based and MFA policies that apply to employees.
9.2 Privileged access management	
9.2.1 Users granted privileged system access have the ability to inflict severe damage on the stability and security of the FI's IT environment. Access to privileged accounts should only be granted on a need-to-use basis; activities of these accounts should be logged and reviewed as part of the FI's ongoing monitoring.	Silverfort discovers and classifies privileged identities based on their actual activity and effective permissions. In AD, JIT access removes standing access by keeping privileged accounts disabled outside approved access windows. Silverfort also enforces MFA on administrator tools and applies Virtual Fencing to keep each privileged account within its approved sources, destinations, and tier. For FIs that use CyberArk, Virtual Fencing or deny policies can block privileged access to protected servers unless it comes through the CyberArk Privileged Session Manager (PSM). Authentication Logs record every authentication by privileged accounts for ongoing review.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
CONTINUED: 9.2 Privileged access management	
9.2.2 System and service accounts are used by operating systems, applications and databases to interact or access other systems' resources. The FI should establish a process to manage and monitor the use of system and service accounts for suspicious or unauthorised activities.	Silverfort automatically discovers service accounts and other NHIs across on-prem and cloud environments. It maps each NHI to a human owner, either synced from existing attributes and tags, or assigned manually when ownership changes. For AD service accounts, it learns each account's normal behaviour and applies Virtual Fencing, restricting the account to its approved sources, destinations, and protocols without credential rotation or application changes. Smart Policies automatically bring new service accounts under protection after an adjustable learning period, and FIs can start in notify mode before moving to deny. Silverfort can also block service accounts from interactive logins. For cloud NHIs in Entra ID and AWS, Silverfort builds geofencing policies from each NHI's observed activity and deploys them to the IdP, which blocks authentication from outside approved locations. Silverfort extends this approach to AI agents. It maps each agent to its human owner and the NHIs it uses, and flags agents with excessive permissions, no human owner, dormant underlying identities, credentials approaching expiration, or an identity shared with other agents. At runtime, Silverfort evaluates each agent tool call before it runs and allows, flags, steps up, or denies it. Step-up is available through the MCP Gateway and supported native integrations.
9.3 Remote access management	
9.3.1 Remote access allows users to connect to the FI's internal network via an external network to access the FI's data and systems, such as emails and business applications. Remote connections should be encrypted to prevent data leakage through network sniffing and eavesdropping. Strong authentication, such as multi-factor authentication, should be implemented for users performing remote access to safeguard against unauthorised access to the FI's IT environment.	Silverfort enforces MFA on remote access to the FI's environment. For VPN connections, it applies MFA whether the VPN authenticates through AD (Kerberos, NTLM, or LDAPS) or through RADIUS, with Silverfort acting as the RADIUS server. It also protects Remote Desktop (RDP) connections to workstations, servers, and domain controllers, and restricts access to the jump hosts used for remote administration. Policies can trigger an alert, require MFA, or block access when Silverfort detects risk.

TABLE 3:

TRM Guidelines Section 11, Data and infrastructure security

Section 11 sets out how FIs should protect data, networks, systems, and virtualised environments. For identity, that means controlling who can reach systems that hold sensitive data, stopping attackers from moving between systems, and protecting administrative access to infrastructure.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
11.1 Data security	
11.1.2 The FI should implement appropriate measures to prevent and detect data theft, as well as unauthorised modification in systems and endpoint devices. The FI should ensure systems managed by the FI's service providers are accorded the same level of protection and subject to the same security standards. **11.1.3** Systems and endpoint devices are often targeted by cyber criminals to gain access or exfiltrate confidential data within an organisation. As such, confidential data stored in systems and endpoint devices should be encrypted and protected by strong access controls.	Silverfort analyses the risk of every authentication in AD, and enforces access-based policies on systems that hold sensitive data. These policies can trigger an alert, require MFA, or block access. The same policies apply to systems managed by service providers when access to them authenticates through AD.
11.2 Network security	
11.2.2 To minimise the risk of cyber threats, such as lateral movement and insider threat, the FI should deploy effective security mechanisms to protect information assets. Information assets could be grouped into network segments based on the criticality of systems, the system's functional role (e.g. database and application) or the sensitivity of the data.	Silverfort enforces MFA and access-based policies on the interfaces attackers use to move laterally in AD, including RDP, PsExec, PowerShell, and WMI. Identity-based segmentation adds deny policies by user, source, destination, and protocol, so FIs can enforce their network segments at the identity layer and block insecure authentication protocols such as NTLMv1. Together, these security controls stop attackers from using compromised credentials to move laterally between systems and limit the spread of ransomware.
11.4 System security	
11.3.1 The security standards for the FI's hardware and software (e.g. operating systems, databases, network devices and endpoint devices) should outline the configurations that will minimise their exposure to cyber threats. The standards should be reviewed periodically for relevance and effectiveness.	Silverfort ISPM continuously checks AD and cloud identity platforms for identity-related misconfigurations, such as weak Kerberos encryption, insecure authentication protocols, and passwords that never expire. Each finding includes remediation guidance, and FIs can set exposure thresholds to match their own security standards.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
CONTINUED: 11.3 System security	
11.3.5 To facilitate early detection and prompt remediation of suspicious or malicious systems activities, the FI should implement detection and response mechanisms to perform scanning of indicators of compromise (IOCs) in a timely manner, and proactively monitor systems', including endpoint systems', processes for anomalies and suspicious activities.	Silverfort Identity Threat Detection and Response (ITDR) detects identity-based attacks in real time, including Kerberoasting, AS-REP Roasting, pass-the-ticket, and MFA bombing. It groups related activity into incidents prioritised by severity, with supporting evidence for investigation and threat hunting. Silverfort can respond automatically by requiring MFA or blocking access, and shares identity risk with SIEM, SOAR, and XDR tools.
11.4 Virtualisation security	
11.4.1 Virtualisation is used by organisations to optimise the use of computing resources and to enhance resilience. The technology allows several virtual machines (VMs) that support different business applications to be hosted on a physical system. A system failure or security breach in one of the VMs could have contagion impact on other VMs. The FI should ensure security standards are established for all components of a virtualisation solution.	Silverfort enforces access-based policies on administrative access to virtualisation platforms, including VMware vCenter and ESXi, and Microsoft Hyper-V. These policies require MFA from approved administrators and deny access to every other account and source machine. A compromised account can't be used to reach the hypervisor, which limits the risk of a breach spreading between VMs on the same host.

TABLE 4:

TRM Guidelines Section 12, Cyber security operations

Section 12 sets out how FIs should monitor cyber events, detect anomalies, and respond to incidents. For identity, that means seeing every authentication, spotting behaviour that departs from normal, and containing compromised accounts quickly.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
12.2 Cyber event monitoring and detection	
12.2.1 To facilitate continuous monitoring and analysis of cyber events; as well as prompt detection and response to cyber incidents, the FI should establish a security operations centre or acquire managed security services. The processes, roles and responsibilities for security operations should be defined.	Silverfort gives the security operations centre (SOC) continuous visibility into identity activity across on-prem and cloud environments. It analyses the risk of every authentication and access attempt by users and NHIs, and by AI agents protected through a native integration or the MCP Gateway, and detects identity threats such as privilege escalation and unauthorised access in real time.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
CONTINUED: 12.2 Cyber event monitoring and detection	
12.2.2 A process to collect, process, review and retain system logs should be established to facilitate the FI's security monitoring operations. These logs should be protected against unauthorised access.	Silverfort's Authentication Logs record every authentication and access attempt. The FI can send them to its SIEM over syslog or in OCSF format, where its existing retention and protection controls apply.
12.2.3 To facilitate identification of anomalies, the FI should establish a baseline profile of each IT system's routine activities and analyse the system activities against the baseline profiles. The profiles should be regularly reviewed and updated. **12.2.4** The FI should consider applying user behavioural analytics to enhance the effectiveness of security monitoring. User behavioural analytics might include the use of machine learning algorithms in real time to analyse system logs, establish a baseline of normal user activities and identify suspicious or anomalous behaviours.	Silverfort builds a behavioural baseline for every user and AD service account, and analyses each authentication against it in real time. When a user's activity departs from the baseline, Silverfort raises their risk level, and policies can trigger an alert, require MFA, or block access. For service accounts, Virtual Fencing blocks any authentication outside the account's learned sources, destinations, and protocols. For cloud NHIs in Entra ID and AWS, Silverfort turns observed sign-in locations into geofencing policies that the IdP enforces. For AI agents, it compares each agent's permissions with its actual activity and flags risky or unusual behaviour.
12.2.5 Correlation of multiple events registered on system logs should be performed to identify suspicious or anomalous system activity patterns.	Silverfort ITDR correlates related authentications and identity events into incidents, prioritised by severity, revealing attack patterns that single events don't show. In cloud environments, ITDR correlates activity across IdPs, IaaS, and SaaS platforms to uncover cross-platform attack chains. The FI can also correlate Authentication Logs with other system logs in its SIEM.
12.2.6 A process should be established to ensure timely escalation to relevant stakeholders regarding suspicious or anomalous system activities or user behaviour.	Silverfort sends real-time alerts on suspicious authentications and identity incidents to the FI's SIEM and SOAR tools. When it detects a risk indicator, such as an abnormal location or a suspicious IP address, it can also email the user to confirm the activity, so likely compromises reach the response team quickly.
12.3 Cyber incident response and management	
12.3.1 The FI should establish a cyber incident response and management plan to swiftly isolate and neutralise a cyber threat and to securely resume affected services. The plan should describe communication, coordination and response procedures to address plausible cyber threat scenarios.	Silverfort contains identity threats in real time, without waiting for manual action. It requires MFA or blocks access for compromised user accounts, and Virtual Fencing blocks service accounts used outside their approved scope. Third-party security tools can also raise a user's or resource's risk level in Silverfort through its Risk API, which triggers risk-based policies that require MFA or block access.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
CONTINUED: 12.3 Cyber incident response and management	
12.3.2 As part of the plan, the FI should establish a process to investigate and identify the security or control deficiencies that resulted in the security breach. The investigation should also evaluate the full extent of the impact to the FI.	Silverfort ITDR presents each identity incident in a single view, prioritised by severity, so investigators can trace which accounts and systems an attacker used. Authentication Logs show how each access was granted and which policies allowed or blocked it. For cloud incidents, ITDR adds supporting evidence, an end-to-end timeline, remediation recommendations, and the blast radius of each identity involved, so teams can assess the full extent of the impact.
12.3.3 Information from cyber intelligence and lessons learnt from cyber incidents should be used to enhance the existing controls or improve the cyber incident management plan.	After an incident, FIs can act on ITDR's remediation recommendations and use ISPM to find and fix the same identity exposures across the rest of the environment, so lessons from one incident strengthen controls everywhere. Deny policies cut off the access paths the attacker used, and step-up MFA policies challenge affected users when they access sensitive systems or when their risk level rises.

TABLE 5:

TRM Guidelines Section 15, IT audit

Section 15 sets out how FIs should use IT audit to assess whether their technology risk controls are adequate and effective. For identity, that means evidence that access controls are working, across every identity in the environment.

MAS TRM GUIDELINES	SILVERFORT SECURITY CONTROLS
15.1 Audit function	
15.1.1 Audit plays an important role to assess the effectiveness of the controls, risk management and governance process in the FI. The FI should ensure IT audit is performed to provide the board of directors and senior management an independent and objective opinion of the adequacy and effectiveness of the FI's risk management, governance and internal controls relative to its existing and emerging technology risks. **15.1.2** A comprehensive set of auditable areas for technology risk should be identified so that an effective risk assessment could be performed during audit planning. Auditable areas should include all IT operations, functions and processes.	Silverfort's identity inventory covers users, NHIs, and AI agents across on-prem and cloud environments, so auditors can bring every identity into the scope of the audit. Authentication Logs then show every authentication and access attempt and the policy decision behind it, which is evidence that controls are operating, not just configured. ISPM adds the posture view, tracking identity exposures and their remediation over time. The Executive Risk Report brings these findings together for the board and senior management, with trend metrics that show whether identity controls in AD environments are improving.

TABLE 6:

Proposed TRM Notice amendments (Consultation Paper P012-2026)

MAS proposes to amend its 11 TRM Notices to strengthen the technology resilience of the financial services sector. In **Consultation Paper P012-2026**, MAS links the changes to major IT incidents caused by poor change management and by delayed detection and response. It also proposes that FIs consider threats from their use of AI in their IT risk assessments. The amendments will apply to every FI covered by the Notices and take effect 12 months after MAS publishes the final Notices.

PROPOSED TRM NOTICE	SILVERFORT SECURITY CONTROLS
7 A Bank must maintain a comprehensive and up-to-date inventory of all IT assets, including cryptographic assets, open-source components, and third-party components, in accordance with the requirements set out in Appendix A. *[Appendix A lists the minimum details to record for each IT asset, including its name, type, version, supported systems, criticality, and owner.]*	For AI agents, Silverfort maintains an up-to-date inventory that identifies each agent and its human owner, along with the systems it connects to and the NHIs it relies on. FIs that treat AI agents as IT assets can use this inventory to make sure every agent has an owner, as Appendix A requires for each asset.
8 A Bank must establish and maintain a framework and process to conduct regular IT risk assessments in respect of every system, the disruption of which may affect the Bank's operations or the services provided to its customers. The Bank must, in respect of each IT risk assessment it conducts: (a) identify the threats to and vulnerabilities within the system; (b) assess the risks that could arise from these threats and vulnerabilities, including the potential impact and likelihood of such risks affecting the Bank's operations or the services provided to its customers (the "identified risks"), in accordance with the Bank's established risk assessment criteria; and (c) implement risk mitigation measures that are commensurate with the identified risks.	Silverfort identifies identity threats and vulnerabilities across users, NHIs, and AI agents, including the risks AI agents create when they act with delegated access. To assess these risks, ISPM gives each exposure a severity level that combines likelihood and impact, which FIs can apply within their own risk assessment criteria. Mitigation then matches the risk: alerts, MFA, and deny policies for users in AD, Virtual Fencing for service accounts, geofencing for cloud NHIs, and runtime controls for AI agents.
9 A Bank must establish and maintain a register that records: (a) the identified risks referred to in paragraph 8(b) that are material ("material identified risks"); (b) the risk owners who will be accountable for managing the material identified risks effectively; and (c) the measures to mitigate the material identified risks. **10** A Bank must establish and maintain key risk indicators (KRIs) to effectively monitor the material identified risks referred to in paragraph 9(a) and assess the effectiveness of the measures referred to in paragraph 9(c).	ISPM ranks identity exposures by severity and provides remediation guidance for each, giving FIs both the identity risks and the mitigation measures to record in their risk register. To track KRIs and whether mitigation is working, the Executive Risk Report shows trend metrics across MFA, segmentation, service accounts, privileged access, posture, and threat detection in AD environments, and ISPM's new versus resolved exposures view does the same in cloud environments.

CONTINUED: Table 6 Proposed TRM Notice amendments (Consultation Paper P012-2026)

PROPOSED TRM NOTICE	SILVERFORT SECURITY CONTROLS
12 A Bank must establish and maintain a framework and process to continuously monitor all critical systems to detect and respond to issues affecting their performance or security in a timely manner. The Bank must ensure that the framework and process include, at a minimum: (a) defined indicators and thresholds that trigger alerts; and (b) response procedures and remedial actions that are commensurate with the nature and potential impact of the identified issues.	Silverfort continuously monitors authentication and access to critical systems, and ITDR detects identity threats against them in real time. FIs set their own thresholds, such as a risk level or a specific risk indicator, as policy triggers. Each trigger leads to a response matched to the risk: an alert, an MFA challenge, or blocked access in AD. Every alert also reaches the FI's SIEM and SOAR tools in real time, where the FI's own response procedures take over.
13 A Bank must implement effective controls to prevent unauthorised changes from being made to any system.	Silverfort controls who can reach the interfaces used to change systems. For administrators, it enforces MFA and deny policies on RDP, PowerShell, PsExec, and hypervisor management consoles, and JIT access enables admin accounts only for the timeframe of an approved ITSM ticket (starting with ServiceNow). For service accounts, Virtual Fencing and interactive login blocking prevent their use from unapproved sources or against unapproved systems. For AI agents, FIs can set runtime policies that deny an agent access to production systems or block the specific tool calls that make changes.
17 A Bank must establish and maintain an incident management framework and process which include, at a minimum: … (b) procedures to collect and preserve evidence for investigation purposes; (c) procedures to notify the Bank's senior management upon identification of the IT incident to enable informed decision-making… **19** A Bank must submit a root cause and impact analysis report to the Authority, within 14 days or such longer period as the Authority may allow, from the discovery of the relevant incident…	Silverfort captures the evidence an investigation needs. Authentication Logs record every authentication and access attempt, AI Agent Access Control Logs tie each agent action to the user who triggered it, and the FI can send Authentication Logs to its SIEM and export AI Agent Access Control Logs to its governance log to preserve them. When ITDR detects an incident, real-time alerts to SIEM and SOAR tools support prompt escalation to senior management. For the root cause and impact analysis, ITDR and Authentication Logs then show which identities and systems an attacker used and how each access was granted. For cloud incidents, ITDR also shows each identity's blast radius.
20 A Bank must implement effective IT controls to protect customer information from unauthorised access or disclosure.	Silverfort protects customer information at the point of access. For users, it enforces MFA and deny policies on access to core banking systems, databases, and legacy applications that authenticate through AD. For service accounts, Virtual Fencing keeps each account within its approved sources and destinations. For AI agents, FIs can set runtime policies by agent, user, application, and tool that deny access to the systems holding customer information, evaluated on every tool call. Together, these controls stop attackers from using compromised credentials or over-privileged agents to reach customer information or move laterally to the systems that hold it.

TABLE 7:

Securing AI agents: AIRG, MindForge, and SAFR

AI agents now act inside FIs. They call tools, query data, and take actions on behalf of users, often with broader access than their task needs. MAS addresses this risk in its proposed Guidelines on Artificial Intelligence Risk Management (AIRG), which set supervisory expectations for AI use and, as MAS has confirmed, cover AI agents. The Project MindForge Operationalisation Handbook, developed by MAS and an industry consortium, turns these expectations into practices, including least privilege for agent tool and data access, kill switches, and traceability through logging.

Safeguards for Agentic Finance at Runtime (SAFR) goes a step further and shows how to apply these practices at the moment an agent acts. Published under the MAS BuildFin.ai initiative as an industry reference approach rather than regulatory guidance, and updated to version 1.1 in October 2026, SAFR places a checkpoint between every agent decision and its execution. The checkpoint has four components: Agent Identity, Controls Repository, Disposition Engine, and Audit Log.

SAFR governs the financial action, and Silverfort secures the identity and access underneath it. Silverfort verifies which agent is acting and for whom, and decides before each tool call runs whether it may proceed. The FI's own financial controls then weigh the action's materiality and impact, such as exposure limits and payment mandates. The table below shows how Silverfort supports each SAFR component.

SAFR COMPONENT	SILVERFORT SECURITY CONTROLS
Agent Identity Binds each proposed action to a recognised, registered agent, verified against that agent's registry entry before any other evaluation proceeds.	Silverfort discovers AI agents across cloud and agentic platforms, SaaS applications, and local endpoints, and keeps them in one inventory that maps each agent to its human owner and the NHIs it uses. At runtime, each request carries the agent, the invoking user, and the target, supporting the closed-loop registry SAFR describes for agents that operate inside the FI's own infrastructure.
Controls Repository The institution's configurable rulebook: the controls… against which a proposed action is checked.	Silverfort supports SAFR's authorisation controls: which agents may act for which users, and which tools, applications, and actions each agent may use. To define them, FIs set policies by agent identity, user context, target application, and action, down to a specific MCP server or tool, and manage them from a central dashboard. New tools can be denied by default until someone reviews them.
Disposition Engine Evaluates each in-scope action against the controls retrieved from the Controls Repository and resolves it to one of the possible outcomes.	Silverfort evaluates each tool call before it runs and returns one of four outcomes: Allow, Flag (the call runs and is marked for review), Step-up (the invoking user approves or denies it), or Deny. These correspond to SAFR's Auto-Execute, Observe, Escalate, and Deny. Step-up is available through the MCP Gateway and supported native integrations.
Audit Log A tamper-evident, append-only record of every governance decision.	Silverfort evaluates each tool call before it runs and returns one of four outcomes: Allow, Flag (the call runs and is marked for review), Step-up (the invoking user approves or denies it), or Deny. These correspond to SAFR's Auto-Execute, Observe, Escalate, and Deny. Step-up is available through the MCP Gateway and supported native integrations.

CONTINUED: Table 7 SAFR Components

SAFR COMPONENT	SILVERFORT SECURITY CONTROLS
Human reviewer escalation Escalation with defined authority, turnaround and timeouts.	Silverfort holds tool calls that the FI's policy marks as sensitive until the invoking user approves or denies them, and AI Agent Access Control Logs record the approver and the outcome. Step-up is available through the MCP Gateway and supported native integrations.
Gateway integration (SAFR deployment pattern) A SAFR gateway intercepts outbound API calls at the infrastructure layer... evaluating it without any changes to the agent code itself.	Silverfort's MCP Gateway sits in the same position: inline between the agent and the MCP servers it calls. It's hosted by Silverfort, and agents are pointed to it through configuration, with no changes to agent code. This brings existing and third-party agents under runtime control, including local agents on employee machines. FIs can start in monitor mode, with full visibility into tool calls, before turning on enforcement, in line with SAFR's advice to establish coverage first. For agents built in Microsoft Copilot Studio or running through Google Cloud Agent Gateway, native integrations apply Silverfort's runtime policies without the MCP Gateway.

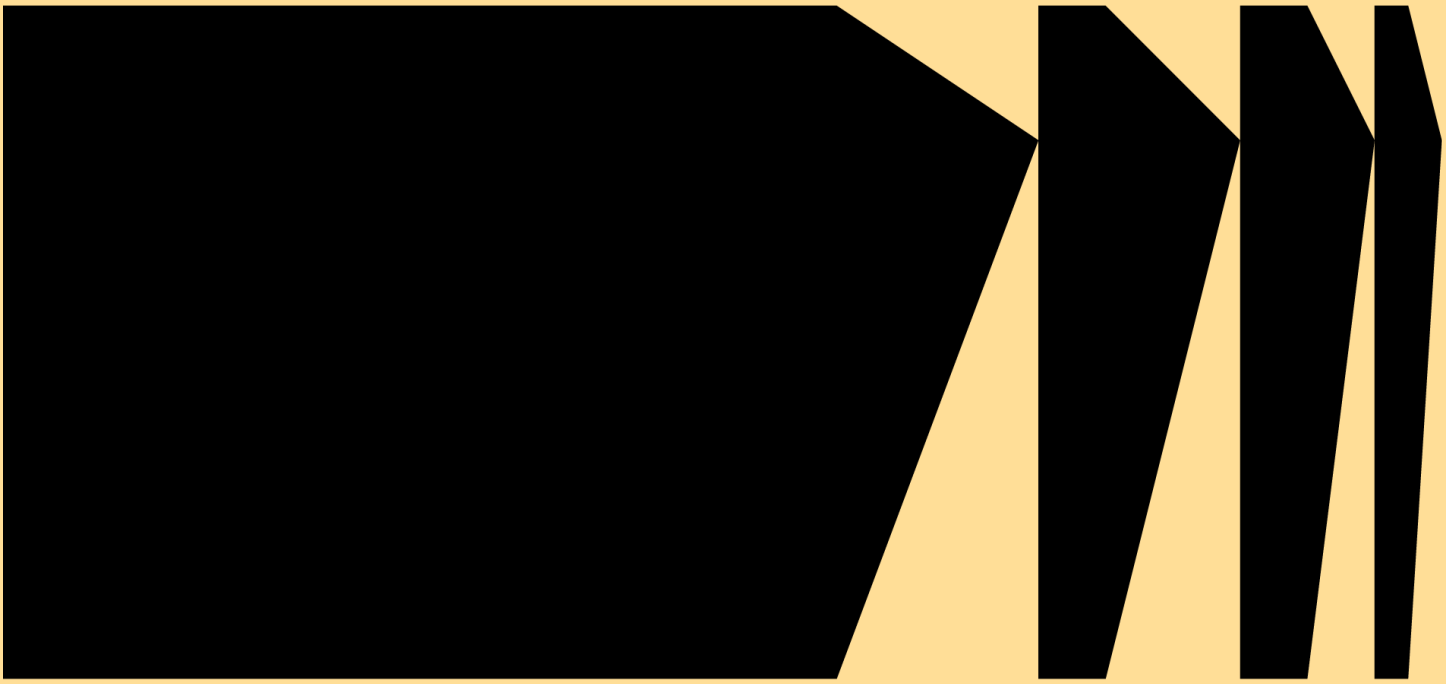

About Silverfort

Silverfort secures every dimension of identity. We deliver end-to-end identity security that is easy to deploy and won't disrupt business operations, resulting in better security outcomes with less work. Discover every identity, analyze exposures, and enforce protection inline to stop lateral movement, ransomware, and other identity threats.