AWS INCIDENT RESPONSE

MAS TRM and AWS Incident Response: Hitting the 1-Hour Notification Window

Singapore's MAS requires regulated financial institutions to notify within one hour of a relevant IT incident. This is the practical playbook for AWS workloads in ap-southeast-1, with the runbook, the evidence stack, and the common failure modes.

SINGAPORE

By Matt Gurr

12 May 2026

ghostvector.ai/guides

GHOST VECTOR — GUIDE

MAS TRM and AWS Incident Response: Hitting the 1-Hour Notification Window

TL;DR

The Monetary Authority of Singapore requires regulated financial institutions to notify within **one hour** of discovering a relevant IT or cyber security incident, followed by a root cause and impact report within 14 days. The 1-hour clock is the tightest cyber incident notification deadline in major jurisdictions and is non-negotiable. Hitting it on AWS requires three things in place before an incident: an unambiguous internal severity model that maps to MAS's "severe and widespread" / "material customer impact" tests, a pre-staged notification path to your MAS relationship manager and after-hours channel, and an AWS-native detection stack in ap-southeast-1 that can confirm scope in minutes rather than hours.

What is MAS's 1-hour notification rule?

The 1-hour rule does not come from the **Technology Risk Management Guidelines** themselves — the TRM Guidelines are principles-based and not legally binding in the strict sense. The actual deadline lives in the **MAS Notices** that bind specific categories of regulated FI:

- **Notice 644** for banks
- **Notice 1108** for merchant banks
- **Notice SFA04-N06** for capital markets services
- **Notice 127** for direct insurers
- Equivalent notices for finance companies, payment service providers, and other supervised categories

All of these impose substantively the same obligation: notify MAS as soon as possible and **not later than one hour** after discovery of a "relevant incident", followed by a comprehensive root cause and impact analysis within **14 days**.

Three details determine how the rule applies in practice:

1. **Discovery, not detection.** The clock starts when the financial institution becomes aware that a relevant incident has occurred. An unconfirmed alert is not discovery. An L1 analyst confirming an alert is real, or a customer-impact report reaching the incident commander, is discovery.
2. **Severity test is yours to apply, consistently.** MAS does not enumerate which incidents are "severe and widespread" or "materially impact customers." Your IT Risk Management framework — which MAS expects you to have, document, and follow — defines the criteria. Inconsistent application is what triggers supervisory concern more than any individual incident.
3. **The first notification can be brief.** MAS expects a short factual notification at the 1-hour mark, not a forensic report. The 14-day follow-up is where depth is expected. Do not delay the first call while you investigate.

Which incidents qualify under MAS TRM?

MAS TRM does not publish a 20-category list like CERT-In. Instead it relies on the dual test of "severe and widespread operational impact" or "material customer service impact." Mapping that to AWS detection sources:

MAS-relevant incident category	Examples	Primary AWS detection sources
Unauthorised access to critical systems	Compromised IAM credentials, privilege escalation, root account access	GuardDuty IAM findings, CloudTrail anomaly detection, IAM Access Analyzer
Cyber attack disrupting services	DDoS, ransomware, web application compromise	Shield Advanced, WAF, GuardDuty, custom CloudWatch alarms on availability metrics
Significant data theft or leakage	S3 exfiltration, RDS dump, customer data exposure	Macie, GuardDuty S3 findings, CloudTrail S3 data events, VPC flow logs
Material customer service disruption	Customer-facing outage, transaction processing failure, login system unavailability	CloudWatch synthetic canaries, Health Dashboard, ALB/CloudFront 5xx rates
Compromise of critical IT systems	Core banking system intrusion, payments platform compromise	EDR on EC2, EKS audit logs, Security Hub aggregation
Misuse of legitimate access	Privileged user abuse, third-party access incident	CloudTrail, IAM Access Analyzer, Detective

The cluster framing matters because a single AWS-side event often hits more than one row. A compromised IAM key that is used to exfiltrate customer data triggers both the unauthorised access category and the data theft category. Your runbook should drive a single notification, not two.

The 60-minute MAS notification runbook

The MAS 1-hour clock is the tightest mainstream incident notification window in the world. Every minute must be designed for in advance:

Minute 0 — Discovery. An incident commander or designated senior IT/security responder confirms a relevant incident. Record the exact timestamp in the incident ticket. *This is your t=0 for MAS reporting.*

Minutes 0–10 — Severity classification. Apply the FI's documented severity model. Confirm the incident meets the "severe and widespread" or "material customer impact" tests. If unclear, escalate up rather than down — MAS will not

penalise an over-cautious notification, but they will penalise a missed one.

Minutes 10–20 — Mobilise. Page the chief information security officer, head of IT, and the CRO or MAS relationship lead. Open a war-room channel. Begin parallel workstreams: containment, customer impact assessment, MAS notification preparation.

Minutes 20–35 — Initial containment and evidence. Run pre-staged automations. Snapshot affected EBS volumes, isolate compromised resources via security group swap, copy the relevant CloudTrail window to a forensics S3 bucket in ap-southeast-1, freeze any IAM credentials suspected of compromise.

Minutes 35–50 — Draft notification and contact MAS. Populate the pre-staged notification template with:

- Time of discovery and time of incident (best estimate)
- Type of incident
- Affected systems and customer-facing services
- Initial estimate of customer impact
- Containment status
- Named accountable contact

Call your MAS relationship manager during business hours. After hours, use the MAS 24/7 contact channel agreed with your supervisor in advance — having this number written down somewhere reachable is itself an audit item.

Minutes 50–60 — Submit written notification. Send the formal notification email to the MAS contact established for your institution. Keep proof of delivery.

The remaining nine hours of your working day go to deeper containment, customer impact assessment, drafting the next regulatory update, and starting the 14-day root cause and impact report.

Pre-incident hardening — what to set up now

Hitting one hour is mostly a question of how much you decided in advance. Four AWS-specific items make the difference between meeting and missing the clock.

1. A pre-tested MAS contact path

Most 1-hour failures are not technical. They are organisational. The pattern is: an incident occurs, the team confirms it is reportable, then someone spends forty minutes finding the right MAS contact and getting authority to make the

call. Avoid this by:

- Maintaining the named MAS supervisor's direct line in a runbook accessible without VPN
- Maintaining the after-hours MAS contact procedure agreed for your institution
- Documenting standing authority for the CISO, head of IT, or CRO to make the initial notification without further approval
- Running this contact path through every tabletop

2. A documented severity model

MAS expects you to apply your severity model consistently. The model should explicitly include criteria that map to "severe and widespread" and "material customer impact" — for cloud workloads on AWS, useful proxies are:

- Customer-facing API availability below 95% for more than 5 minutes
- Successful authentication to a production account by a non-approved IAM principal
- Confirmed exfiltration of any production database snapshot or S3 object containing customer data
- Loss of integrity of any system processing customer transactions

Severity criteria belong in your IT Risk Management policy, signed off by the IT steering committee, and exercised in tabletops. Inventing the criteria during an incident is the surest way to fail consistency.

3. Logging and evidence baseline in ap-southeast-1

For the 14-day follow-up report you will need substantially more evidence than the initial notification. Pre-stage:

- **CloudTrail** organisation trail with S3 destination in ap-southeast-1, ≥12 months retention with Object Lock
- **CloudWatch Logs** for application and system logs, ≥12 months retention
- **VPC Flow Logs** delivered to S3 in ap-southeast-1
- **GuardDuty** enabled in every active region with findings forwarded to Security Hub
- **Macie** enabled on any account holding customer data
- **AWS Config** recording configuration history — invaluable when the 14-day report asks what changed and when

MAS auditors will ask for chain-of-custody on every piece of evidence. Object Lock with compliance mode is the cheapest mechanism for demonstrating it.

4. Pre-staged notification templates

Have the initial notification email and the 14-day report template loaded into your incident management platform. Fields should mirror what MAS expects — incident type, discovery time, affected services, customer impact estimate, containment status, contact person. Drafting from a blank document inside 60 minutes is a poor use of the window.

AWS region selection and Singapore-specific considerations

Singapore has one AWS region — **ap-southeast-1 (Singapore)** — which has been operating since 2010 and supports the full AWS service catalogue including all security services, SES, Bedrock, and a deep set of AWS Local Zones. Three IR-relevant points:

1. **MAS does not impose strict data localisation but expects practical demonstrability.** Keeping production data and logs in ap-southeast-1 makes the right-of-audit and exit-plan requirements of the **MAS Outsourcing Guidelines** straightforward to satisfy. Multi-region active-active across ap-southeast-1 and a non-Singapore region (Sydney is the common choice) is acceptable provided the outsourcing register documents it.
2. **AWS Local Zones in Singapore extend ap-southeast-1 with lower-latency edge sites.** Latency-sensitive trading workloads sometimes use these. Evidence collection from a Local Zone works the same way as the parent region.
3. **AWS Direct Connect locations in Singapore are well established.** Most MAS-regulated FIs operate hybrid architectures with Direct Connect. Your IR runbook must explicitly cover incidents that span on-premises and cloud — a compromise that starts on a workstation and pivots to AWS via Direct Connect is one of the more common scenarios in this market.

Common mistakes

Five patterns we see repeatedly when reviewing MAS readiness for AWS workloads:

- **Treating "discovery" as "investigation complete."** The 1-hour clock starts when the responsible person becomes aware. Wanting to "make sure" before calling MAS is the single most common 1-hour failure mode.

- **No pre-agreed after-hours MAS path.** Discovery at 2am on a public holiday is exactly when this matters and exactly when teams fall over. Agree the path with your relationship manager *before* you need it.
- **Severity model that's a slide deck, not a policy.** If your severity definitions are not in a board-approved policy that operations teams can apply at 2am, they are decorative.
- **Logs centralised outside Singapore.** Routing CloudTrail and security logs to a regional observability hub in Tokyo or Sydney creates audit friction and slows evidence collection. Keep the Singapore copy authoritative.
- **Treating the 1-hour notification as the report.** It is a notification, not a report. Plan the 14-day report from minute one and dedicate someone to drafting it in parallel with the response.

What to do this week

If your institution is MAS-regulated and runs production workloads on AWS, three actions return the most value in the first week:

1. **Test the MAS contact path.** Walk through a notional 2am discovery on a Saturday. How does the call to MAS actually happen? Document and rehearse.
2. **Audit your severity model.** Open your IT Risk Management policy. Identify the criteria that map to MAS's "severe and widespread" and "material customer impact" tests. If they are not there, add them.
3. **Run a 1-hour tabletop.** Use a realistic scenario — compromised privileged IAM credentials in production, or a confirmed S3 data exfiltration. Time yourselves from discovery to MAS notification. The first time you do this it will not be one hour. The point is to find out what slows you down.

The 1-hour clock is unforgiving in absolute terms, but the supervisory expectation underneath it is essentially "do you know what you are doing." Institutions that maintain a tested runbook, a documented severity model, and a working contact path consistently hit the window. The ones that improvise consistently miss it.

Related guides

- AWS Incident Response in APAC — Pillar guide — the regulatory and technical overview across Australia, Singapore, India, and ASEAN.

- Meeting CERT-In's 6-hour incident reporting rule on AWS — the equivalent rule for Indian workloads, with substantial overlap on AWS-side hardening.
- APRA CPS 234 notification obligations for AWS — the Australian FSI parallel, with a more forgiving 72-hour window.
- OAIC Notifiable Data Breach scheme on AWS — the Australian privacy regulator parallel.
- Ghost Vector managed detection — our managed detection service includes MAS-aligned 1-hour notification workflows.