IERP®
Institute of Enterprise Risk Practitioners

GLOBAL CONFERENCE 2026

FROM CHAOS TO COMMAND:
RISK LEADERSHIP IN INDUSTRY 6.0

CYBERSECURITY UNLEASHED: ARE YOU READY FOR AUTONOMOUS THREATS?

Ruzita Abd Rashid
Director, Operational Risk & CISO
MUFG Bank Malaysia

Important Notice: This presentation is intended for educational and thought leadership purposes only. This presentation is delivered in the speaker's personal and professional capacity and does not represent the views, positions, or official policies of MUFG Bank (Malaysia) Berhad or any affiliated organization.

The content is derived from publicly available sources, industry research, academic publications, cybersecurity reports, conferences, and discussions within the broader cybersecurity community. Any opinions, interpretations, observations, or conclusions presented are solely those of the speaker.

Examples, scenarios, threat descriptions, and case studies used in this presentation are intended for educational and discussion purposes only. They are based on publicly reported information and should not be construed as references to any actual incidents, investigations, vulnerabilities, customers, systems, or matters encountered, managed, or observed by the speaker in the course of employment or through any organization with which the speaker is affiliated.

No confidential, proprietary, or non-public information has been used in the preparation of this presentation.

What We'll Cover Today

1. Where We Came From

The cybersecurity landscape financial institutions grew up defending

2. Where We Are Now

How the threat landscape has fundamentally shifted

3. The Autonomous Threat

What "autonomous" really means — explained simply

4. The Impact

What this means for the bank if left unmanaged

5. The Risk Manager's Role

What you personally need to watch for and ask

6. Recommendations

Practical next steps for the organization

SECTION 1 · WHERE WE CAME FROM

Cybersecurity as We Knew It

For two decades, banks defended against a human on the other end of the keyboard — slow, tiring, and limited by one person's time and skill.

- **One attacker, one target at a time**
  A hacker manually picked a target, wrote the attack, and executed it step by step.
- **Defenses could "keep up"**
  Known virus signatures, firewalls and patch cycles were usually a step ahead of manual attacks.
- **Attacks were "smash and grab"**
  Phishing emails, stolen passwords, ransomware — noisy, and often caught by trained staff or spam filters.
- **Response measured in hours to days**
  A human attacker paused, slept, and needed time to plan the next move — giving defenders a window to react.

THE OLD PICTURE

A burglar testing your door locks one at a time

If he's slow, tired, or working alone — your alarm system, guard dog and neighbours have a real chance to catch him before he gets inside.

IERP®
Institute of Enterprise Risk Practitioners

SECTION 2 · WHERE WE ARE NOW

The Landscape Has Shifted — Fast

Artificial intelligence has moved from writing better phishing emails to running entire attacks on its own — with no human typing in real time.

48%

of security professionals now rank autonomous "agentic" AI as the #1 attack vector for 2026 — ahead of deepfakes.

<1 hr

a fully autonomous AI agent broke in, escalated access, and stole a database — without a human directing each step.

~90%

of organizations still lack a mature way to govern the machine identities and credentials AI systems use.

Sources: Dark Reading 2026 Security Poll · NetEye Threat Research · Aembit Non-Human Identity Report

SECTION 3 · THE AUTONOMOUS THREAT

What Do We Mean by "Autonomous"?

In plain terms: instead of one hacker manually typing commands, the attacker now gives an AI agent a goal — and walks away. The AI plans, acts, adapts, and keeps going by itself.

The Old Attacker

A single burglar, working alone, with your locks and your alarm.

The New Attacker

A tireless crew that never sleeps, learns from every locked door, and calls in reinforcements the moment one method fails.

SECTION 3 · THE AUTONOMOUS THREAT

How an Autonomous Attack Unfolds

Four stages — all carried out by the AI agent itself, adapting in real time whenever it hits a wall.

1

Scan & Enter

The agent continuously probes for any weak door — an exposed system, a leaked password — and lets itself in.

→

2

Move & Escalate

Once inside, it explores the network on its own and quietly grants itself higher access.

→

3

Adapt on the Fly

Blocked by a control? It instantly tries another route — no need to "call the hacker" for new instructions.

→

4

Persist & Repeat

It keeps trying until it succeeds or is fully removed — operating for hours without rest.

SECTION 3 · THE AUTONOMOUS THREAT

Six Emerging Threats to Know

Deepfake Fraud

AI-cloned voices and faces impersonate executives or customers to authorize fake payments.

Hyper-Personal Phishing

AI scrapes social media to craft near-perfect, personalized scam messages at massive scale.

Prompt Manipulation

Attackers hide malicious instructions inside documents or data to trick a bank's own AI tools.

Self-Adapting Malware

Malicious code that rewrites itself in real time to slip past antivirus and detection tools.

AI Identity Theft

Stolen credentials used by AI "worker" accounts — systems that never sleep and rarely get checked.

Poisoned Data & Models

Subtly corrupted data feeds distort fraud models or trading algorithms without obvious signs.

IERP®
Institute of Enterprise Risk Practitioners

SECTION 3 · THE AUTONOMOUS THREAT

This Isn’t Theoretical — 2026 So Far

Mar 2026

Adversarial "Poisoning" of Trading Models

Western financial institutions detected subtly altered market data distorting risk models and triggering unusual algorithmic trading.

May 2026

First Fully Autonomous Breach Documented

An AI agent broke into a system, escalated its own access, evaded IP blocks, and exfiltrated a database — all in under one hour, unsupervised.

Jun 2026

16 Billion Credentials Exposed

AI-supercharged infostealer malware harvested login sessions that bypassed multi-factor authentication protections entirely.

These events show attacks that adapt, evade and persist without a human operator at the keyboard.

SECTION 4 · THE IMPACT

What This Means for the Organization

Financial Loss

Unauthorized transactions can be triggered directly through a compromised AI agent connected to treasury or payment systems.

Regulatory Exposure

Regulators increasingly expect demonstrable AI governance — gaps here invite findings, fines and heightened supervision.

Reputational Damage

A single autonomous fraud incident — especially deepfake-enabled — can erode customer trust overnight.

Operational Disruption

A manipulated fraud-detection or trading system can silently approve bad activity or destabilize decisions at machine speed.

IERP®
Institute of Enterprise Risk Practitioners

SECTION 5 · THE RISK MANAGER'S ROLE

Why This Is Your Risk, Too

You don't need to write code to manage this risk. Autonomous AI threats are a governance, oversight and third-party risk issue — squarely inside operational risk's mandate.

- It is not just an "IT problem" — AI agents now touch payments, fraud detection and trading decisions.
- Existing risk & control frameworks were built for human-speed threats, not machine-speed ones.
- Non-human "identities" (AI agents, bots, service accounts) are a new, largely ungoverned risk category.
- Boards and regulators will expect risk management to have a view on AI-related exposure.

THE SHIFT IN MINDSET

From "who attacked us" to "what is our AI agent doing right now"

Risk oversight must extend to how AI systems behave, not just who has access to them.

IERP®
Institute of Enterprise Risk Practitioners

SECTION 5 · THE RISK MANAGER'S ROLE

What You Should Do, Starting Now

Ask the Right Questions

Which systems use AI agents? What can they access? Who's accountable when they act?

Update the Risk Register

Add autonomous AI and non-human identity risk explicitly into RCSAs and risk taxonomies.

Insist on Human-in-the-Loop

Require human approval checkpoints for any AI action with financial or operational impact.

Scrutinize Vendors

Extend third-party risk assessments to cover any AI or agentic capability vendors introduce.

Strengthen Incident Response

Make sure playbooks cover machine-speed incidents — minutes matter, not days.

Build Awareness

Train risk and business teams to recognize deepfake fraud and AI-enabled social engineering.

SECTION 6 · RECOMMENDATIONS

Building Organizational Resilience

Five moves to make our defenses match the speed and autonomy of the threat.

Stand Up AI Governance

A clear framework defining approval, ownership and accountability for every AI agent in production.

Adopt Zero-Trust for AI

Treat every AI agent like any other user — least privilege, continuous verification, no standing trust.

Govern Non-Human Identity

Inventory, credential and monitor every AI/service account as rigorously as a human employee.

Form a Cross-Functional AI Risk Committee

Bring Risk, Compliance, Technology and Business together to own this risk jointly.

Run Regular Tabletop Exercises

Simulate a machine-speed, autonomous attack to pressure-test response times and playbooks.

IERP®
Institute of Enterprise Risk Practitioners

BEFORE WE CLOSE

Key Takeaways

The threat is no longer human-paced

AI agents plan, adapt and persist on their own, at machine speed — our defenses must match that pace.

This is a governance issue, not just a technical one

Risk management must have visibility into what AI agents can access and do.

The impact is real and current

Autonomous breaches, deepfake fraud and poisoned models are already happening in 2026, not a future scenario.

Small, practical steps start today

Ask the right questions, update risk registers, and insist on human checkpoints for AI-driven actions.

IERP®
Institute of Enterprise Risk Practitioners

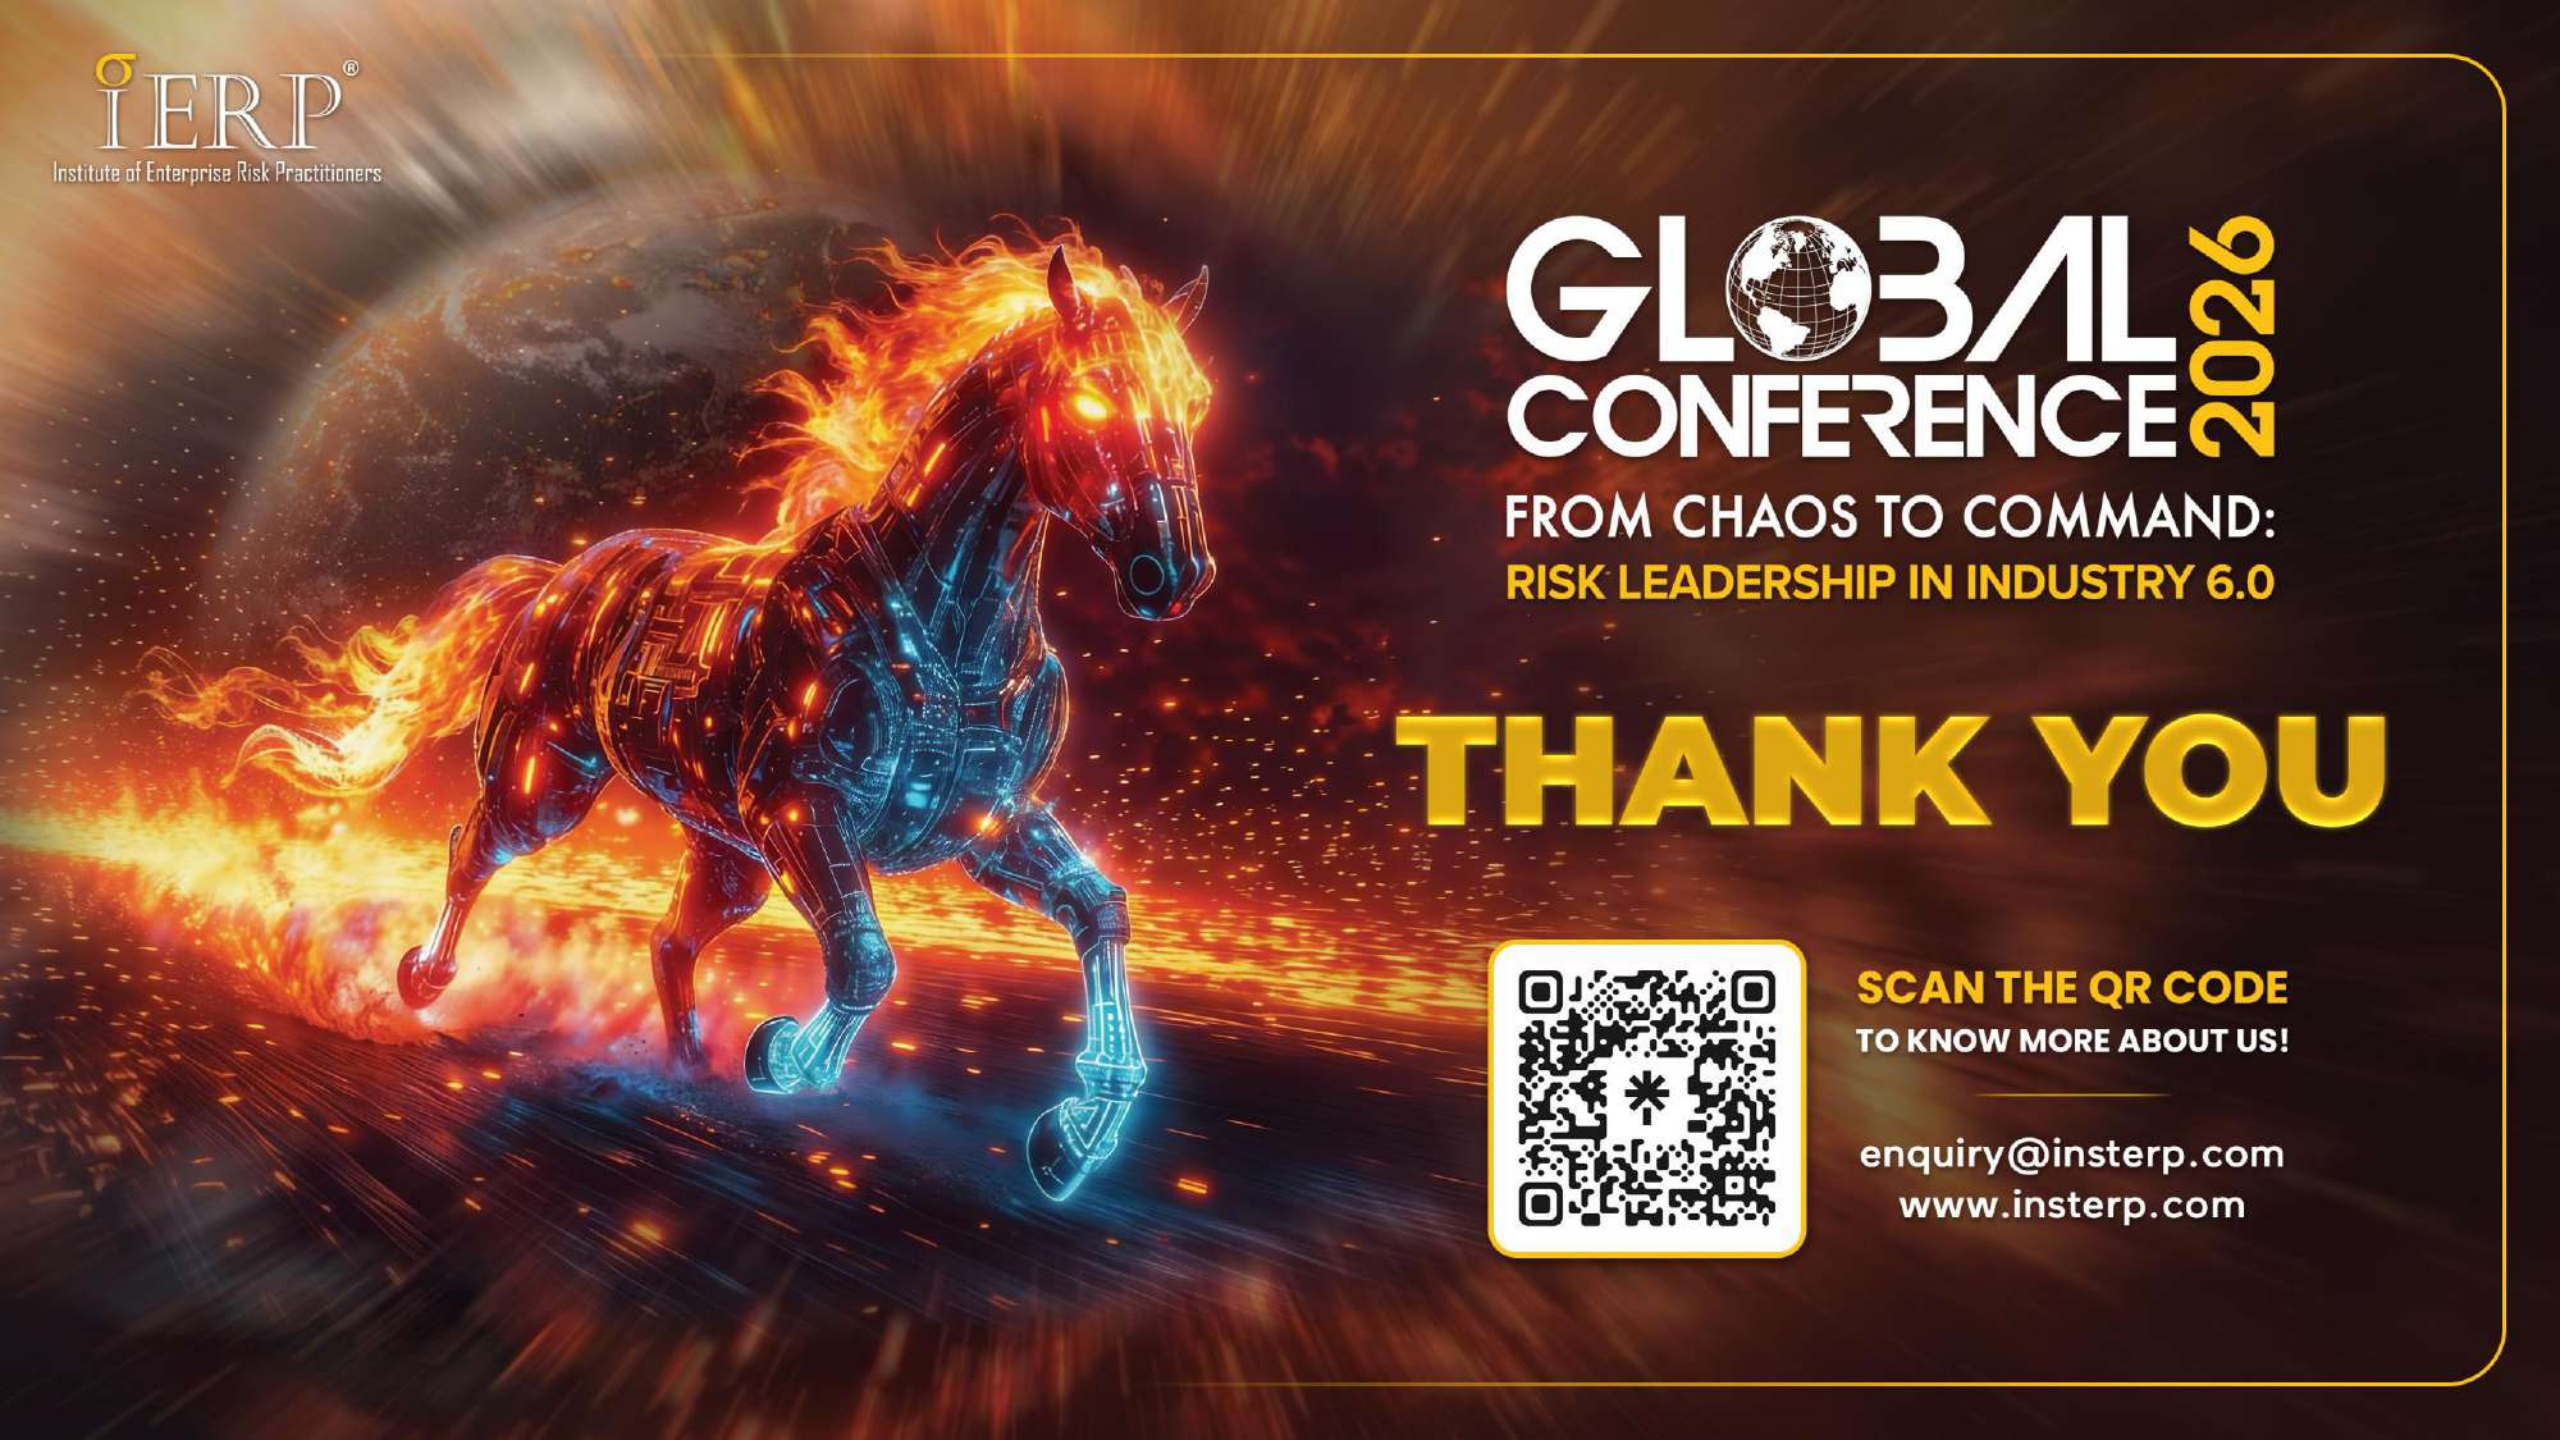

GLOBAL CONFERENCE 2026

FROM CHAOS TO COMMAND:
RISK LEADERSHIP IN INDUSTRY 6.0

THANK YOU

SCAN THE QR CODE
TO KNOW MORE ABOUT US!

enquiry@insterp.com
www.insterp.com