RCE WORKSHOP · CYBER-ENABLED FRAUD · 10 NOVEMBER 2026

Before the First Click

Preemptive intelligence against AI-enabled financial crime

Fin O’Leary

1 · THE THREAT

AI is a threat. It is also an opportunity.

€95m

moved from Fideuram, February 2026 (reported)

A fake WhatsApp message from the CEO, then a cloned voice.

The largest deepfake fraud on public record, by amount moved.

About four times the Arup case in 2024.

How early can we see these threats coming?

Sources: press reports, September 2026. Amount moved, not net loss.

1 · THE THREAT

What has changed

Cheap

A few seconds of audio can clone a voice.

Convincing

A perfect email or voice no longer proves anything.

Fast

Thousands of versions are tried while we review one.

So we cannot judge fraud by how it looks.

Where most of the effort goes

Detect and respond

Most organisations invest heavily here.

Most tools activate once compromise has already begun.

The click has happened. The password is entered. The payment is sent.

Then the clock starts.

What can we see before the break-in?

Where today's tools start working

Earlier: before the first click

Where most tools start

1 Build

Sites, servers, certificates, toolkits, inc. malware

2 Launch

First email, call or message

3 Compromise

A click, a login, a payment

4 Respond

Detect, investigate, block

5 Recover

Clean-up, reporting, lessons

The attacker is most exposed on the left. Almost nobody is looking there.

Earlier sight means more options: block it, take it down, warn customers.

2 · OBSERVABLE EARLY

AI threats and speed

Websites | Servers | Certificates | Toolkits | Timing

The message is fake. The set-up behind it is real, and it leaves traces.

AI-enabled attacks still rely on observable infrastructure.

AI-Agents and ML find the pattern. People decide what to do.

2 · OBSERVABLE EARLY

From reactive to preemptive: five moves

1. Look outside your own network for your name, brand and executives
2. Map the connections, not just single alerts
3. Set a clock: time to sight, time to act
4. Agree in advance who acts: block, take down, warn, report
5. Include suppliers, and share what you learn

A layer in front of detect and respond, not a replacement.

2 · OBSERVABLE EARLY

Defenders are going agentic

Agent teams in a loop

Seen in a major vendor's new security platform

- **Red agents**
  Probe for weaknesses and attack paths
- **Blue agents**
  Detect, investigate and assess
- **Green agents**
  Fix and harden

Mythos

A frontier AI model with restricted access.

Reported to have found thousands of serious flaws in major operating systems and browsers.

US and UK authorities briefed banks.

Machine-speed attacks need machine-speed defence.

AI agents and agentic AI mean the same thing. This talk uses them interchangeably.

3 · CASE STUDY · A LARGE INTERNATIONAL BANK

One connected network

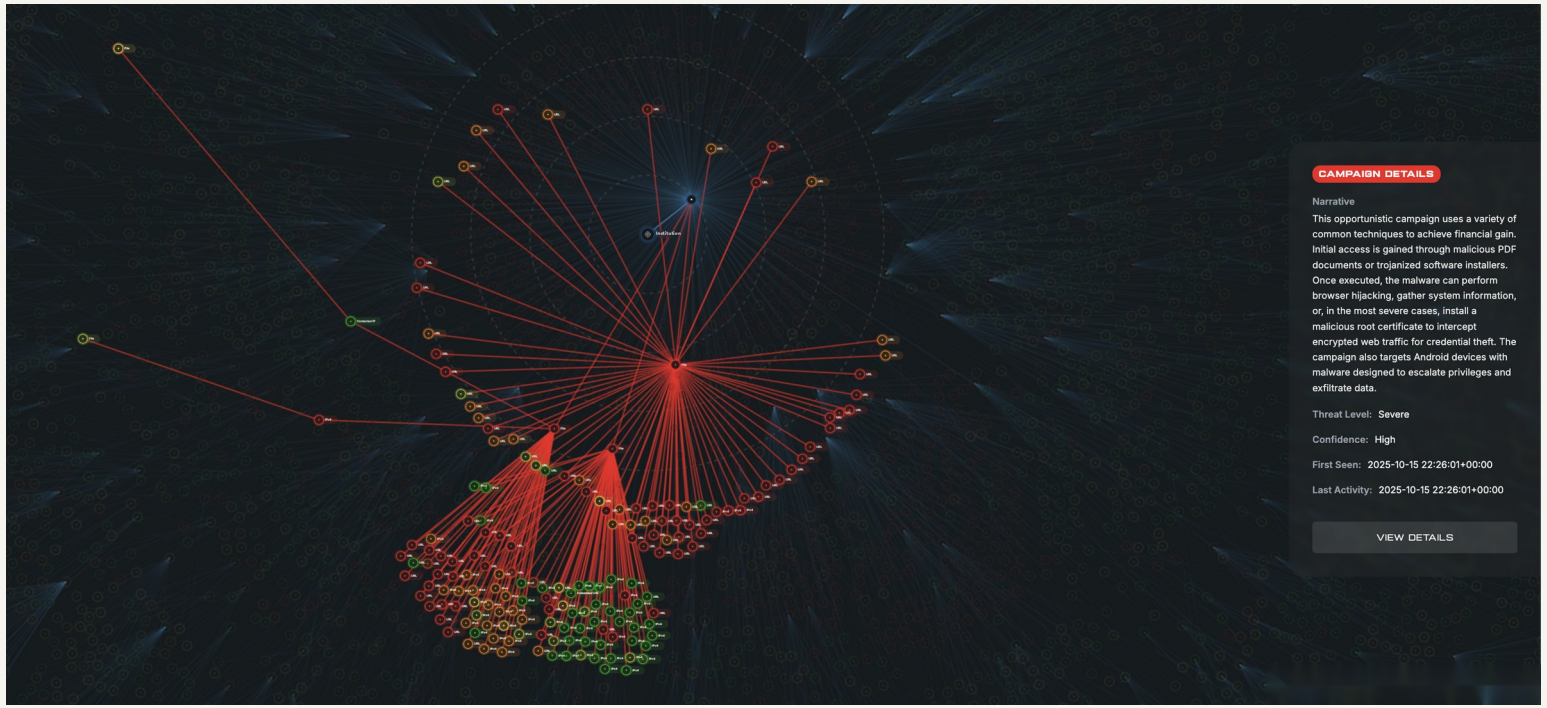

Each dot is a website, server or file. Each line is a link between them.

Drawn before it went live – 5 campaigns.

3 · CASE STUDY

Five campaigns, in plain English

1. Stealing passwords from infected computers
2. Fake software installers that redirect web traffic
3. A network of infected computers
4. Cracked software carrying banking malware
5. Phishing emails sending customers to a fake login page

This case is malware and phishing, not deepfakes. Deepfakes would use the same set-up.

3 · CASE STUDY

Mapped before it went live

Then someone has to act.

The bank

Blocks it

Website hosts

Take it down

National CERT

Warns others

FIU and regulator

Share and supervise

The earlier the map exists, the more time everyone has.

Your suppliers are part of the risk

3 of 5

campaigns arrived through fake installers and cracked software

~~A yearly checklist~~

Ongoing visibility

EU DORA | Kuwait CBK | Mauritius BoM and FSC | South Africa | Nigeria | Kenya

Regulators are moving this way. How do you know, this week?

Five questions supervisors can ask

1. How early do you see a threat?
2. How fast do you act?
3. Which suppliers matter most, and how are they watched?
4. Are your controls tested against fake voices and faces?
5. Who do you tell, and how quickly?

Suggestions for discussion, in proportion to the size of the institution.

Why threat intelligence is strategic

AI attacks move at machine speed.

Attackers hide behind legitimate cloud hosting and free certificates.

Direct threat intelligence tells legitimate from malicious use, before launch.

Regulators: ask for outcomes

- How early do you see a threat?
- Explore continuous outside-in visibility, including suppliers.
- Enable early-warning sharing between banks, CERTs and FIUs.

Industry: build the capability

- Put a preemptive layer in front of detect and respond.
- Mean time to detect and mean time to respond (MTTD & MTTR)
- Share what you see - early.

CLOSE

Three takeaways

1. Fakes are cheap, but the set-up behind them is real and can be seen.
2. Speed is the control: Pre-emptive AI and ML systems find it, people decide.
3. Your suppliers are part of your risk: watch them continuously.

Questions?

fin@blackwired.com