MOORE

Issue 001

TECHNOLOGY ALERT

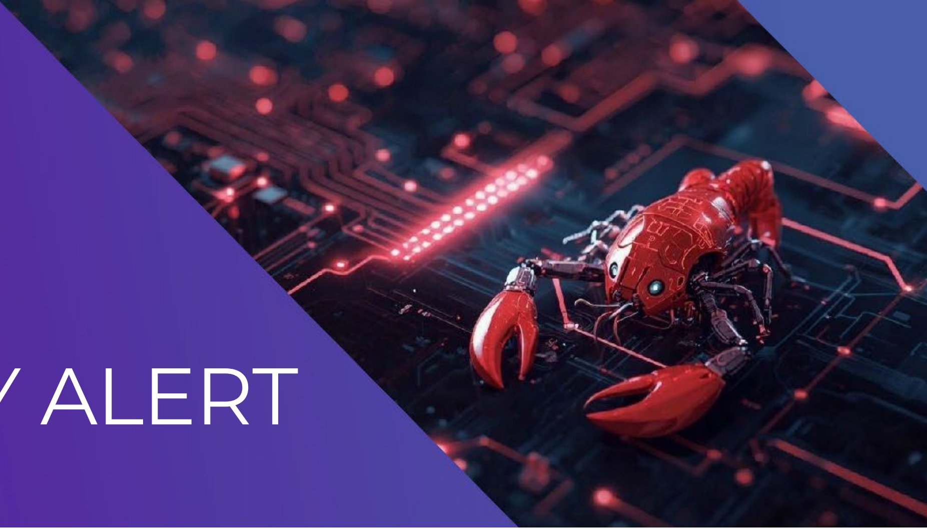

OpenClaw - Managing the Risks of AI-Driven Automation Tools

In recent months, “OpenClaw” has attracted growing attention as a new type of automation tool powered by artificial intelligence. It is an open-source application which can be installed on a computer or server and connected to major AI models. Once set up, users can interact with OpenClaw via chat to help perform routine tasks, such as triaging emails, drafting documents, browsing websites, and handling simple online processes.

Unlike traditional web-based chatbots, OpenClaw is designed to take actions on behalf of the user. It may read files on the host machine, log in to websites, and invoke tools or scripts. In many common configurations, it can perform a range of operations similar to those available to the person who installed it.

In this alert, we highlight the main issues businesses should be aware of when staff are considering or already using OpenClaw, and set out a practical approach for evaluating such tools in a controlled manner.

KEY RISK AREAS

1. Level of access

OpenClaw typically runs under the user account that installs it. As a result, it may inherit access to corporate email, shared folders, collaboration platforms, and internal applications available to that user. From a risk perspective, this places OpenClaw closer to a powerful local application than to a lightweight productivity add-on.

2. Influence from external content

OpenClaw follows instructions contained in the information it processes, including web pages, emails and documents. Malicious or specially crafted content can be used to steer its behaviour (often referred to as “prompt injection”), potentially leading to unintended data transmission or system changes from the organisation’s perspective.

3. Third-party “skills”

The software can be extended with publicly available “skills” which add extra functions. Independent reviews have found that a material portion of popular skills contain security vulnerabilities, including some rated high or critical, such as possible credential theft or unauthorised network access. These skills operate effectively as small software modules with the same level of access as the main tool.

4. Deployment on standard endpoints

Given its ability to execute code and store long-lived credentials, several expert groups and vendors recommend treating OpenClaw as a high-risk or untrusted execution environment. It may not be appropriate to deploy it directly on standard user workstations that store sensitive or business-critical information without additional safeguards.

PRACTICAL IMPLICATIONS FOR ORGANISATIONS

From a governance and control perspective, OpenClaw raises questions in a number of areas.

First, there is a possibility of informal adoption by staff seeking efficiency gains, sometimes on corporate devices, without going through normal IT approval processes. This may result in OpenClaw having access to systems and data which have not been formally risk-assessed.

Second, there is an impact on accountability and regulatory compliance. When a tool can act using a user's access rights, any misconfiguration or compromise may give rise to questions regarding responsibility, incident management, and adherence to contractual, legal, or regulatory requirements, particularly in regulated sectors.

Third, the skills ecosystem introduces a further layer of third-party risk. Each additional skill may have its own security posture and data handling practices. Without appropriate review, the combined risk profile may be difficult to understand and manage.

SUGGESTED APPROACH TO EVALUATION

For organisations that wish to explore OpenClaw in a structured way, current industry guidance indicates several key measures:

- Conduct any initial testing in a segregated environment, such as a dedicated virtual machine or test server, rather than on everyday user laptops.
- Configure OpenClaw to use dedicated, low-privilege test accounts, and restrict its access to non-critical systems and non-sensitive data where feasible.
- Maintain an inventory of instances and document which systems each instance can reach, together with a record of any installed skills or extensions.
- Establish a review process for skills, focusing on their source, requested permissions and any external network connections, before they are approved for use.
- Communicate clearly with staff regarding the firm's policy on installing and using tools of this nature on corporate devices.

Moore IT & Cybersecurity Services

How can we help

We are closely monitoring technical developments and market practices relating to OpenClaw and similar agent-based tools. Our team is available to assist you with:

- reviewing and updating internal policies and guidelines to cover AI-driven automation tools
- assessing the extent to which such tools may already be in use within your organisation
- designing secure pilot environments that align with your risk appetite and regulatory obligations

If you would like to discuss any of the matters outlined in this alert, please contact our advisory team directly.

If you would like to discuss any of the matters outlined in this alert, please contact our advisory team directly.

PATRICK ROZARIO
Advisory Services
Managing Director
T +852 2738 7769
E patrickrozario@moore.hk

KENNETH MA
Transaction Services
Managing Director
T +852 2738 4633
E kennethma@moore.hk

KEVIN LAU
IT & Cybersecurity
Principal
T +852 2738 4631
E kevinlau@moore.hk

The information provided is for general guidance only and should not be treated as professional advice. While we believe the content is correct at the time of publication, we cannot accept responsibility for any loss or inconvenience caused by actions taken based on the information herein. We recommend consulting a qualified advisor to ensure your specific circumstances are properly evaluated before making any decisions.

© 2026 Moore Advisory Services Limited