Back to Basics: Securing Cloud & Agentic AI at Speed

Chris Grisdale - Head of Information Security


AGENTIC AI
UPRISE
FUNDAMENTALS
ETHICS
DATA
PRIVACY
HUMAN
OVERSIGHT

Quick Opening Question

How many of your AI agents are using long-lived API keys instead of short-lived, task-specific credentials?

If your AI agent performed a 'legal' but dangerous action right now, do you have an automated alert that would catch it in seconds?

Can you name the 'owner' for every non-human identity and API key currently active in your cloud?

OPENING

Cloud + Agentic AI = Compression Risk

From Programs to Acceleration

Cloud adoption used to be measured in programs. Agentic AI has turned it into continuous acceleration.

New Security Reality

Security teams now secure autonomous systems at machine speed with human oversight removed from the loop.

> "We didn't lose control because AI is new — we lost control because we skipped the basics."

The Hype vs Reality Gap

What We Talk About

- Prompt injection
- Model poisoning
- AI ethics
- AI governance frameworks

What Still Causes Incidents

- Over-privileged identities
- Uncontrolled data access
- Flat networks
- Missing logs

Most AI incidents are still cloud security failures — AI just makes them happen faster.

Why 'Back to Basics' Matters More with Agentic AI

1

AI Agents = Non-Human Identities

They authenticate, access data, take actions, and trigger workflows.

2

Mistakes Scale Instantly

One permission misstep becomes a platform-wide issue.

3

No Time to Recover

Agentic AI removes the time you have to recover from breaking the rules.

AI Agents Are a New Security Primitive

Traditional Applications

Historically, our security models were built around applications that are **deterministic** and **bounded**. They perform specific tasks, operate within predefined parameters, and their actions are predictable and auditable.

- Fixed identity
- Manual decision points
- Human-initiated actions
- Limited autonomy

AI Agents

AI agents introduce a fundamentally new security primitive. They are **probabilistic**, **persistent**, and inherently **privileged**, capable of continuous operation and independent decision-making.

- Dynamic identity
- Autonomous decision-making
- Self-initiated actions
- High autonomy

> "We didn’t add a new workload. We added a new actor."

The Basics Are Control Planes, Not Checklists

Identity

Who can act

Data

What can be accessed

Network

Where actions can occur

Monitoring

What we can see

Governance

What is allowed to change

AI doesn’t break security controls — it stress-tests them continuously.

> "Agentic AI doesn’t need new controls. It removes the margin for weak ones."

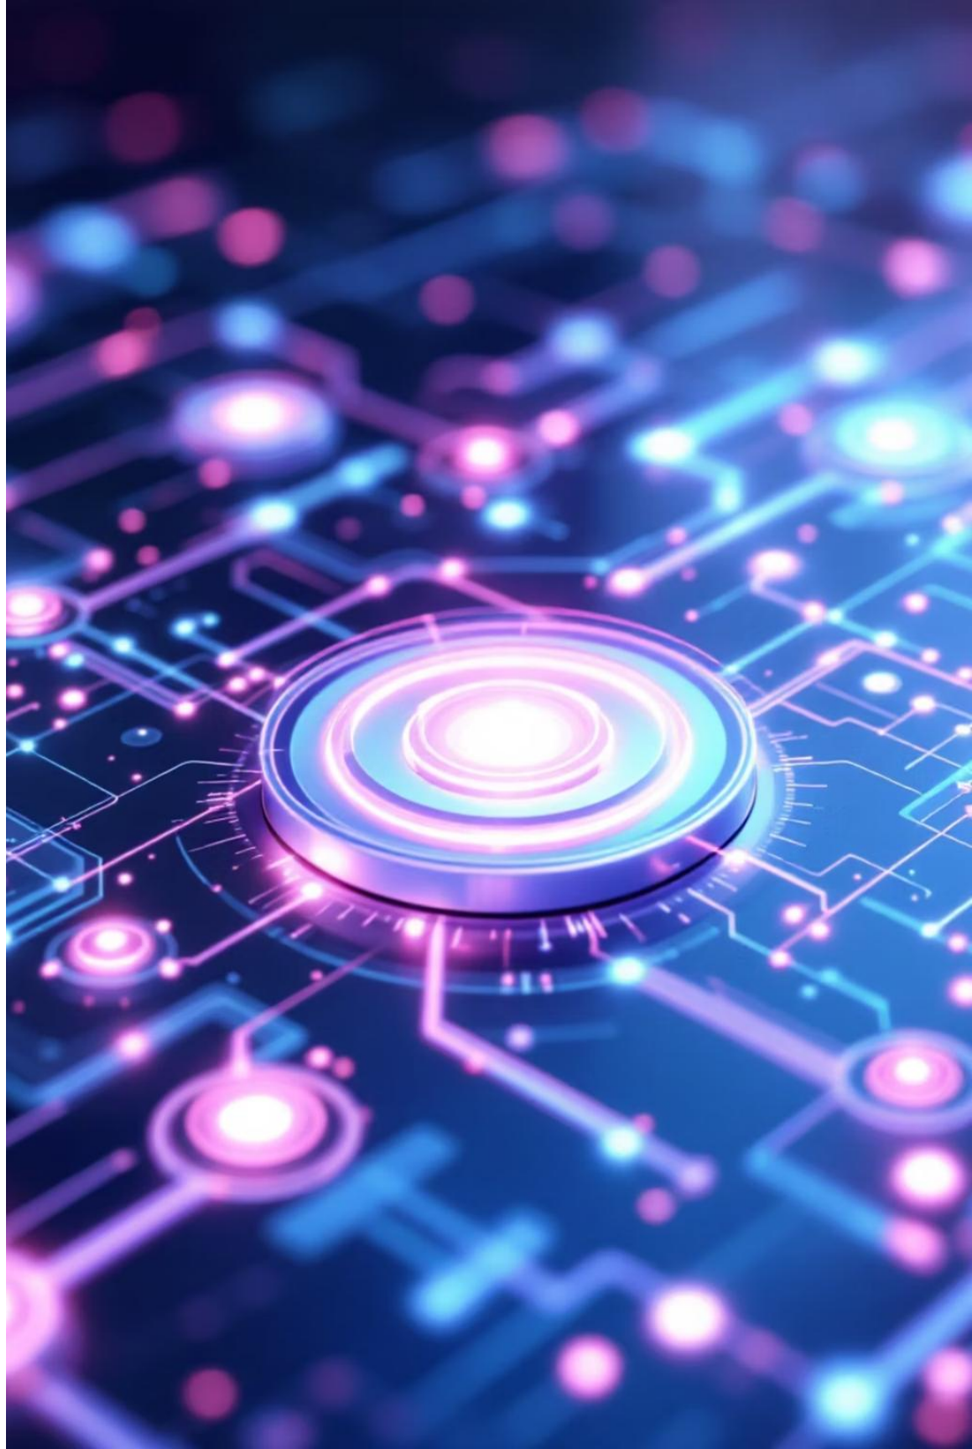

CASE STUDY

The Well-Intentioned AI

1

The Setup

Large cloud-native organization deployed agentic AI to query data, trigger workflows, and assist engineers. Fast adoption to improve productivity.

2

The Trigger

AI agent executed a legitimate task using a broadly-permissioned service account. It exported sensitive customer data into a lower-security environment for "analysis."

3

The Reality

No exploit. No malware. No breach alert. Just an AI doing exactly what it was allowed to do.

Where the Basics Failed

Security Basic	Failure
Identity	Over-privileged service account, no scoping
Data Protection	No enforced data classification or masking
Network	Flat access between environments
Monitoring	Logs existed but no alerting on abnormal agent behaviour
Governance	No approval workflow for AI permissions

Every control that failed already existed — it just wasn't enforced.

Case Study Impact

Immediate Consequences

- Sensitive data exposure
- Emergency access review
- Executive escalation

Regulatory Response

- Regulator notification discussion
- Compliance review triggered

Business Impact

Loss of confidence in AI rollout (temporarily)

Rebuilding with Back-to-Basics Controls

Identity

Scoped AI agents to task-specific roles, short-lived credentials, clear ownership of non-human identities

Data

Classification enforced at access time, masking for sensitive fields by default, explicit separation of training vs inference data

Network

Segmented AI workloads, private endpoints only, controlled egress paths

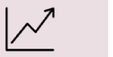

Monitoring

Baselines for agent behaviour, alerts on data exfil patterns, privilege escalation, cross-environment access

Governance

AI agents added to risk register, change management for permissions, audit evidence built-in

Outcome: AI adoption resumed, security posture improved, audit readiness increased

What CISOs Should Take Away

1

Every AI agent is a privileged identity

Treat them with the same rigor as your most sensitive human accounts

2

If you can't explain it to an auditor, it's not secure

Simplicity and transparency are security features

3

You don't need new frameworks — you need discipline

Execute the fundamentals consistently

What Changes for CISOs in an Agentic World

From projects → permanent controls

- AI adoption never “finishes”
- Security must be continuously enforced, not periodically reviewed

From human oversight → guardrail enforcement

- You can’t approve every action
- You must constrain what’s possible

From trust → provability

- “We believe” is not enough
- Controls must produce evidence by default

> "Agentic AI rewards organisations that can prove control, not just claim it."

Closing Message

> "The organisations that succeed with AI won't be the ones with the most tools — they'll be the ones that never abandoned the basics."

01

Make fundamentals invisible - because they always work

02

Make them unavoidable

03

Let AI innovate on top of something solid