# Protecting Cloud Access in the Era of Autonomous AI

**George Roberts**

Founder, Mod Digital LLC

Principal IAM Architect, McDonald's

MVP
Microsoft
Most Valuable Professional

# Why Agents Are Different

Agents are actors

Reason, plan, act

Operate without GUIs

Assistive → Agentic → Autonomous

# The Cloud Problem We Created

- IAM assumed humans
- RBAC assumed static apps
- Logging assumed request/response
- Over-permissioned agents
- Shadow agents

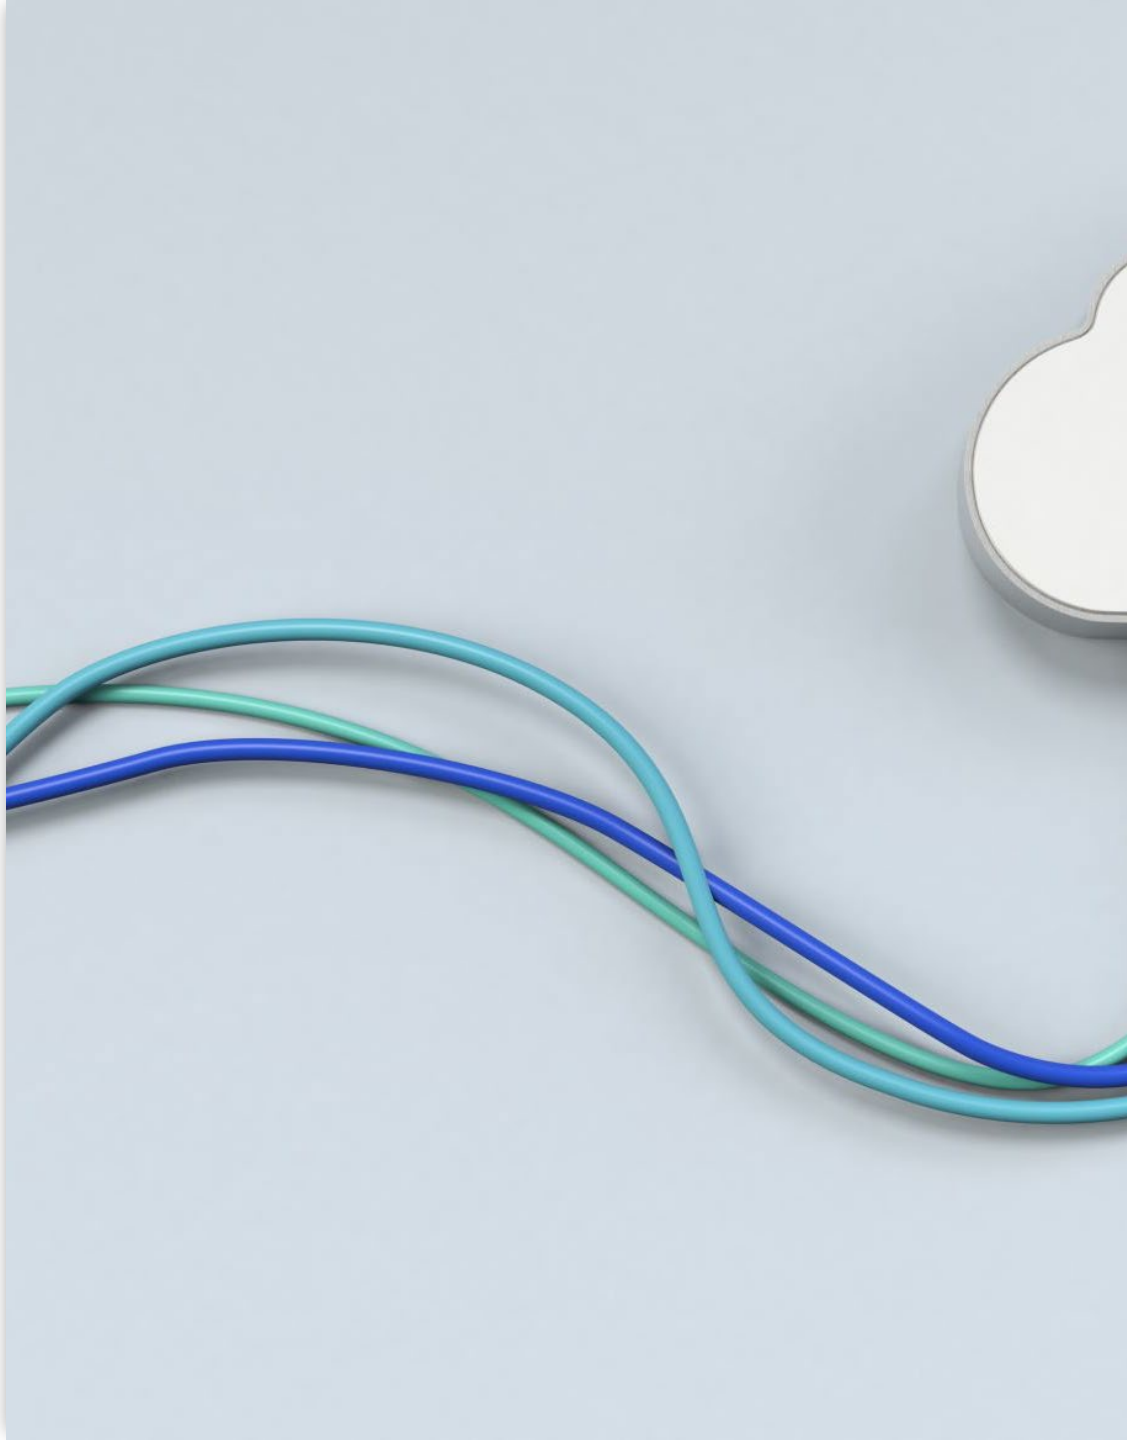

# Are Agents Workloads or Virtual Users?

# Yes

# Common Agent Types

Single-task

Tool-calling

Multi-agent

Event-driven

# Typical Flow

How do we govern these agents?

- Inventory
- Non-human identities
- Least privilege
- Lifecycle policies

# Identity for Non-Human Workers

- Each agent gets an identity
- Scoped permissions
- Conditional access

# Observability

- Who triggered
- What data
- What tools
- What actions
- Traceability

# Treat Agent Lifecycle like SDLC

DEV/TEST/PROD

CHANGE CONTROL

VERSIONING

RETIREMENT

# The Risk is Real

Prompt injection

Data exposure

Loops

Amplification

Mitigations

# Lessons for Any Cloud

- Separate build and control planes
- Inventory first – you need to know what exists
- Instrument before autonomy
- Agents will outnumber humans

# Where This Is Heading

AGENTS MANAGING AGENTS

NEW AUDIT REQUIREMENTS

NON-HUMAN IDENTITY & ACCESS MANAGEMENT

PLATFORM OWNERSHIP

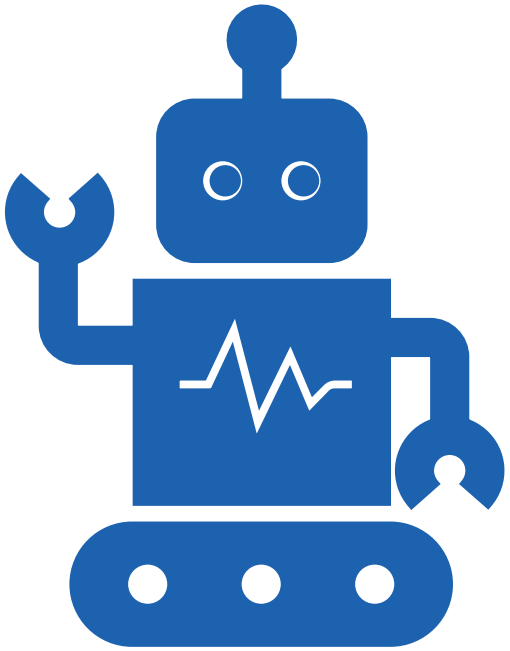

# The Agents are Already Here

- Agents are workloads with autonomy
- Identity and governance is required
- Need to be treated as virtual users
- No autonomy without trust

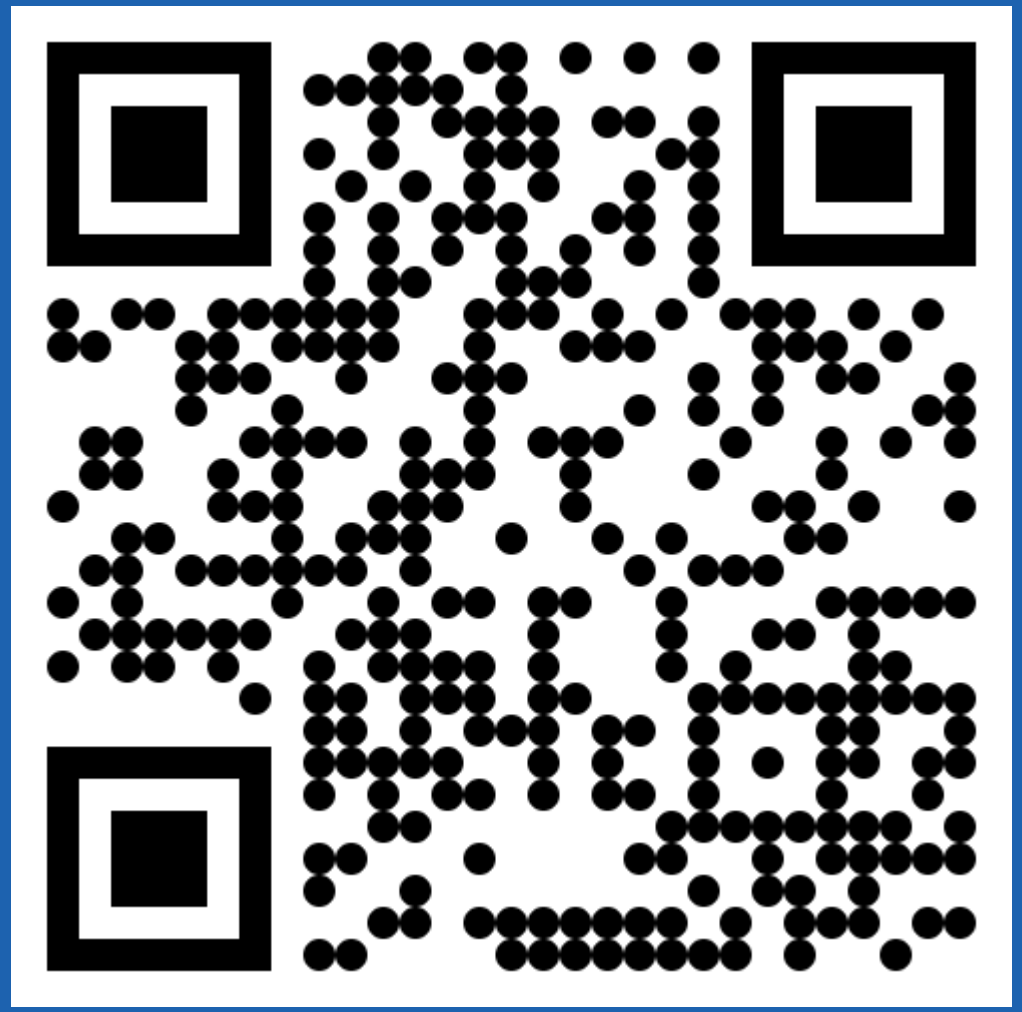

# Thank you!

**Connect with me!**

linkedin.com/in/**sirtwist**