THE LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

DATA GOVERNANCE AND SYSTEM FOUNDATIONS

Foundations for the Agentic Era

The new agent products, OpenAI's Dots, Meta's Muse and NVIDIA's Safe Agent platform, each ship impressive controls. But those controls only hold if the data and systems beneath them are sound. This session is about the foundations that decide whether an agent is safe to let loose.

EXECUTIVE SESSION · OCTOBER 2026 · INDUSTRY FACILITATOR SESSION

WHY THIS IS SUDDENLY URGENT

Foundations used to be hygiene. Now they are load-bearing.

Good data governance was always sensible. What has changed is the cost of getting it wrong. When a person reads a report, a data error is caught by a human. When an autonomous agent acts on that data, the error becomes an action, taken at machine speed, with credentials, before anyone looks.

The old world

Humans sat between the data and the action. A messy field, a stale record or a vague permission was annoying, but a person caught it before it did harm.

The agentic world

The agent reads the data and acts on it directly. Your governance is no longer a filing standard; it is the guardrail that stands between a bad record and a real action.

Eight in ten firms name data limitations as the roadblock to scaling agents. The bottleneck is not the model; it is the ground it stands on.

THE ARGUMENT IN ONE LINE

The agent is only as safe as the floor beneath it

Picture the stack. The product's own controls sit on top. They are real, but they govern the agent. Underneath sits your data and your systems, and that is what the agent actually acts on. A strong lid on a weak floor is still a weak system.

The agent's own controls	*Dots, Muse, NVIDIA: sandboxes, approvals, watchdogs*	VENDOR
Your access & identity layer	*Who, and which agent, may touch what*	YOU
Your governance & lineage	*Definitions, quality, provenance, audit*	
Your data	*The records the agent reads and acts on*	

Vendors are racing to secure the top layer. The three layers below it are yours, and no product can fix them for you.

READING THE SIGNAL

What Dots, Muse and NVIDIA are telling us

Three of the most capable agent products of 2026 all shipped with elaborate controls. Read what they chose to build, and what they left for you.

THE SAME LESSON, THREE TIMES

Everyone is building guardrails now

The striking thing about this year's launches is not capability; it is how much of each product is control rather than intelligence. Each secures a different layer, and each proves the same point: capability is assumed, safety is the differentiator.

OPENAI

Dots

A permission ladder in software: each agent on its own cloud computer, hidden credentials, configurable permissions, auto-review and monitoring that can pause it.

Secures: the agent's own conduct.

META

Muse

An isolated Secure VM, a separate Sentinel that approves every outbound action, credentials the model never sees, and mandatory approval at checkout.

Secures: a person's digital life.

NVIDIA

Safe Agent platform

OpenShell sandboxes the agent on the host; Sentry, an out-of-band watchdog on separate silicon, quarantines it in milliseconds if it crosses a boundary.

Secures: enterprise execution.

Three different companies, three different layers, one message: in 2026 the product is the guardrail. That should tell you where your own effort belongs.

THE GAP EVERY PRODUCT LEAVES OPEN

What none of them can do for you

Look closely and all three make the same assumption. They control what the agent does, but they take the data, the definitions and the permissions it works from as given. That part is yours, and it is exactly where things go wrong.

They sandbox the agent	but cannot tell if the record it read was wrong, stale or duplicated.
They hide credentials	but cannot decide which data a given agent should be allowed to reach in the first place.
They require approval	but a human approving an action built on bad data still approves a bad action.
They log everything	but a log is only useful if your data has the lineage to make the entry meaningful.

Every one of these controls governs the agent's behaviour. None of them governs your data. That is the unglamorous work no vendor will do for you.

THE REAL WORK

The foundations that make an agent safe

If the products secure the lid, these are the floor. Five foundations that decide whether an autonomous agent helps or harms.

THE FLOOR BENEATH THE AGENT

Five foundations, in plain terms

None of these is new. What is new is that an agent now depends on them to act correctly, continuously, without a human in between. Each is a question you should be able to answer before an agent touches your systems.

1	**Quality**	Is the data right, fresh and free of duplicates, continuously, not once a quarter?
2	**Meaning**	Does everyone, and every agent, read a term the same way? One shared definition.
3	**Lineage**	Can you trace where a number came from and what it touched downstream?
4	**Access**	Does each agent reach only the data its task needs, and nothing more?
5	**Identity**	Can you tell which agent did what, when and why, after the fact?

Think of these as the pre-flight check. An agent that acts on data failing any one of them is an accident waiting for a credential.

Right, understood, and traceable

The first three foundations are about the data an agent reads. Get them wrong and even a perfectly behaved agent does the wrong thing confidently.

Quality, continuously

Agents act in real time, so quality has to be monitored in real time. Automated validation and anomaly detection catch an error before it becomes an action, not in next month's audit.

A shared meaning layer

An ontology or knowledge graph so “customer”, “revenue” or “active” means one thing. Without it, two agents act on conflicting readings of the same word.

Lineage you can follow

Every figure should carry its origin and its downstream reach. When an agent acts, you can see what it relied on and what its action will ripple into.

Clean, consistently-defined, traceable data is the difference between an agent that assists and one that amplifies your worst record at scale.

Least privilege, and a name on every action

The last two foundations are about control and accountability. The products assume you have these; most organisations do not, at least not for non-human actors.

Access, scoped to the task

An agent should reach only what its job requires, and no more. That means treating agents as first-class identities with their own, narrow permissions, not letting them inherit a human's broad access. Explicit policy on what each agent can touch, and when a human must approve.

Identity and audit

Every action needs an accountable name attached: which agent, acting for whom, on what data, when, and why. As autonomy rises this is what makes an incident investigable and a regulator answerable. Logging is cheap; meaningful logging depends on the foundations above.

Give an agent a human's access and you have handed its credentials to software that never sleeps. Scope it tightly and name every action.

WHERE YOUR WORK MEETS THEIRS

Your foundations plus their controls

The point is not to distrust the products. It is to see them as the top layer of a system whose lower layers you own. Done well, the two halves reinforce each other; done badly, the vendor's guardrail is the only thing standing, on sand.

What the vendor brings

Isolation, approval gates, credential hiding, out-of-band enforcement and activity logs. The runtime safety of the agent itself. Buy this; do not rebuild it.

What you must bring

Clean, defined, traceable data; agents as scoped identities; and an audit trail that is meaningful because the data under it is. The vendor's controls act on whatever you give them.

A good product on good foundations is a safe agent. A good product on poor foundations is a confident one. The difference is entirely your side of the line.

WHERE ARE YOU NOW

A maturity ladder for agent-ready foundations

Most organisations are not where they think. Place yourself honestly: an agent deployed a rung too high for your foundations is the common, expensive mistake of this year.

4 · Agent-ready	Real-time quality, shared semantics, agent identities, scoped access, full lineage. Safe to delegate.
3 · Governed	Clear ownership, policies and definitions, but still built around human users and periodic checks.
2 · Managed	Data is catalogued and mostly trusted, governance is partial, access is coarse and human-shaped.
1 · Ad hoc	Quality is unknown, definitions vary by team, access is broad. An agent here is a liability.

Deploy the agent at the rung your foundations support, not the rung the demo suggests. Most of 2026's agent failures are a rung-three firm buying a rung-four product.

A TEST BEFORE YOU DEPLOY

Six questions to ask before an agent goes live

A short diligence test for any team about to point an agent at real systems. If you cannot answer these, the foundation is not ready, whatever the product promises.

1. **Is the data monitored in real time?**
   Or will the first sign of a bad record be the agent's action on it?

2. **Is there one definition per term?**
   Would two agents agree what “active customer” means?

3. **Can you trace any figure's lineage?**
   Origin upstream, and reach downstream, for anything an agent relies on.

4. **Is each agent a scoped identity?**
   Narrow permissions of its own, not a borrowed human login.

5. **Is every agent action attributable?**
   Which agent, for whom, on what, when and why, after the fact.

6. **Who is accountable when it errs?**
   A named owner and an escalation path, agreed before deployment.

WHAT TO CARRY FORWARD

The foundation is the strategy

1. Agents turn data errors into actions. Foundations stopped being hygiene; they are the guardrail.
2. Dots, Muse and NVIDIA all secure the agent. None of them governs your data, definitions or access.
3. Five foundations decide safety: quality, meaning, lineage, scoped access, agent identity.
4. Buy the vendor's runtime controls; build the data and governance layer yourself.
5. Deploy at the rung your foundations support, not the one the demo promises.

TO DISCUSS

Which rung of the ladder is your organisation really on? · *Which foundation is weakest, and what would it take to close it before you deploy an agent?*