

AWS State, Local, and Education Learning Days

Washington D.C



Generative AI/ML and AI governance for the public sector

Sergio Ortega

AI/ML BD and Sales Lead
State and Local Governments
sergioai@amazon.com



AI/Machine learning (ML) is at an inflection point

Key drivers: Compute capacity increase | Data growth | Model sophistication

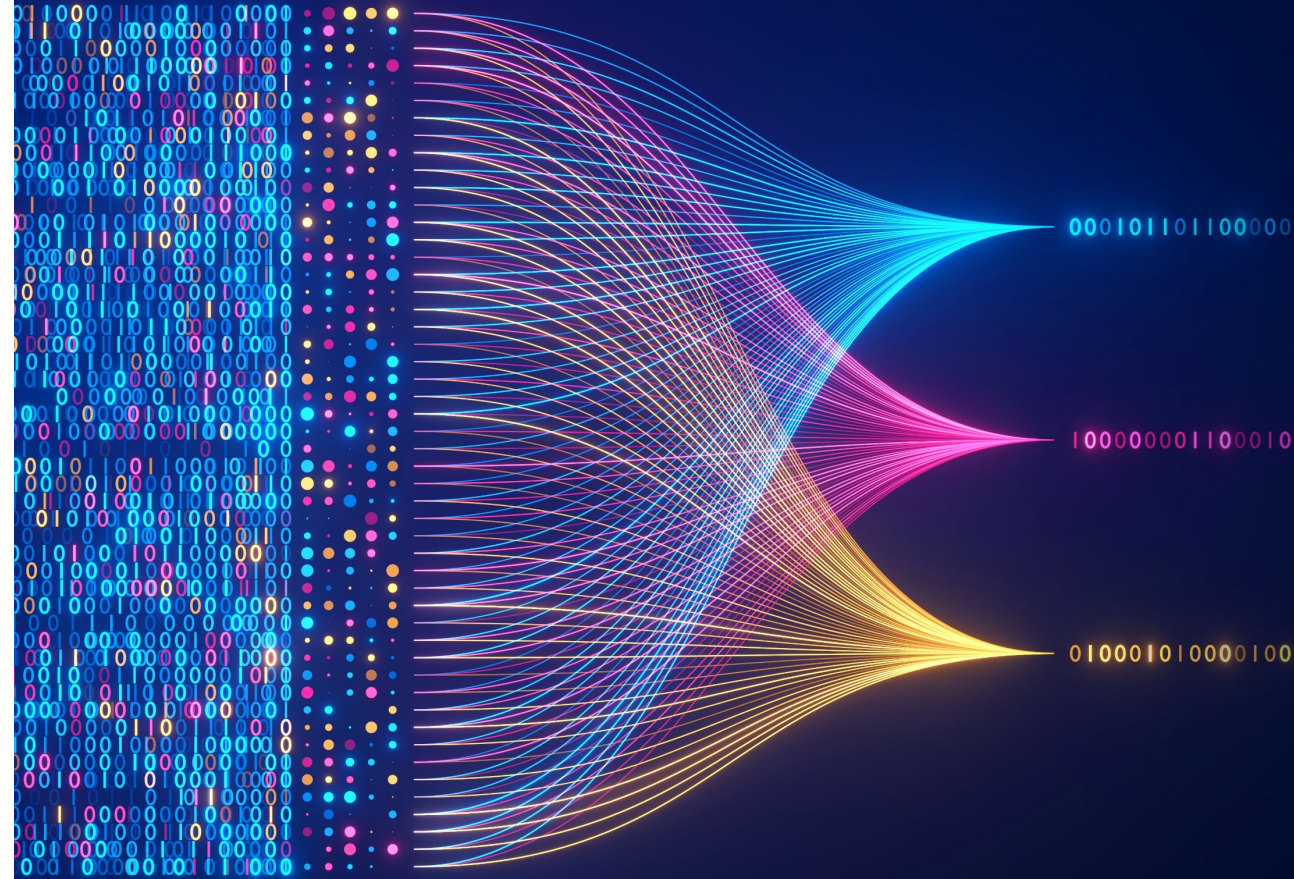
Generative AI is powered by foundation models

Pretrained on vast amounts of unstructured data

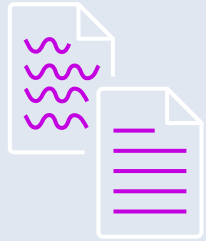
Contain large number of parameters that make them capable of learning complex concepts

Can be applied in a wide range of contexts

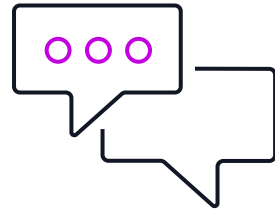
Customize FMs using your data for domain specific tasks



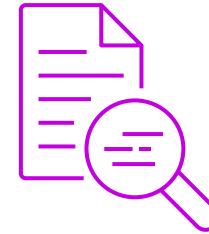
Foundation model use cases



Productivity
Text generation



Chat
Virtual assistant



Summarization
Text extraction



Search



Code generation



Image generation




Image classification

Current challenges

- **Consumer facing applications for enterprise problems**
- **Relying on foundation models alone to solve problems**
- **Inaccessible, unintegrated and poor-quality data**
- **Not aligning use cases to strategic priorities**
- **Shadow AI**

TOP *of* MIND

GEN AI: TOO MUCH SPEND, TOO LITTLE BENEFIT?



Tech giants and beyond are set to spend over \$1tn on AI capex in coming years, with so far little to show for it. So, will this large spend ever pay off? MIT's Daron Acemoglu and GS' Jim Covello are skeptical, with Acemoglu seeing only limited US economic upside from AI over the next decade and Covello arguing that the technology isn't designed to solve the complex problems that would justify the costs, which may not decline as many expect. But GS' Joseph Briggs, Kash Rangan, and Eric Sheridan remain more optimistic about AI's economic potential and its ability to ultimately generate returns beyond the current "picks and shovels" phase, even if AI's "killer application" has yet to emerge. And even if it does, we explore whether the current chips shortage (with GS' Toshiya Hari) and looming power shortage (with Cloverleaf Infrastructure's Brian Janous) will constrain AI growth. But despite these concerns and constraints, we still see room for the AI theme to



Generative AI Application



Generative AI
Application

Data Foundation

STORAGE

**GOVERNANCE
& COMPLIANCE**

**DATABASES,
ANALYTICS,
& DATA LAKES**

**DATA
INTEGRATION**

Your data is the **differentiator**



Generic
generative AI



Generative AI that
knows your business
and your customers

What could go wrong?



Inaccuracies

Answers that are factually incorrect, irrelevant, or nonsensical, because of limitations in their training data and architecture

“The world record for crossing the English channel on foot is 15 hours”



Bias

Answers that display discriminatory behaviour resulting in prejudiced or unequal treatment of a particular group or groups

“Generate a picture of a person cleaning” returns overwhelmingly women



Copyright and IP

The rights of content creators from whom training data is collected remains uncertain and is currently being challenged

Artists suing creators of foundation models alleging the improper use of its photos



Security and privacy

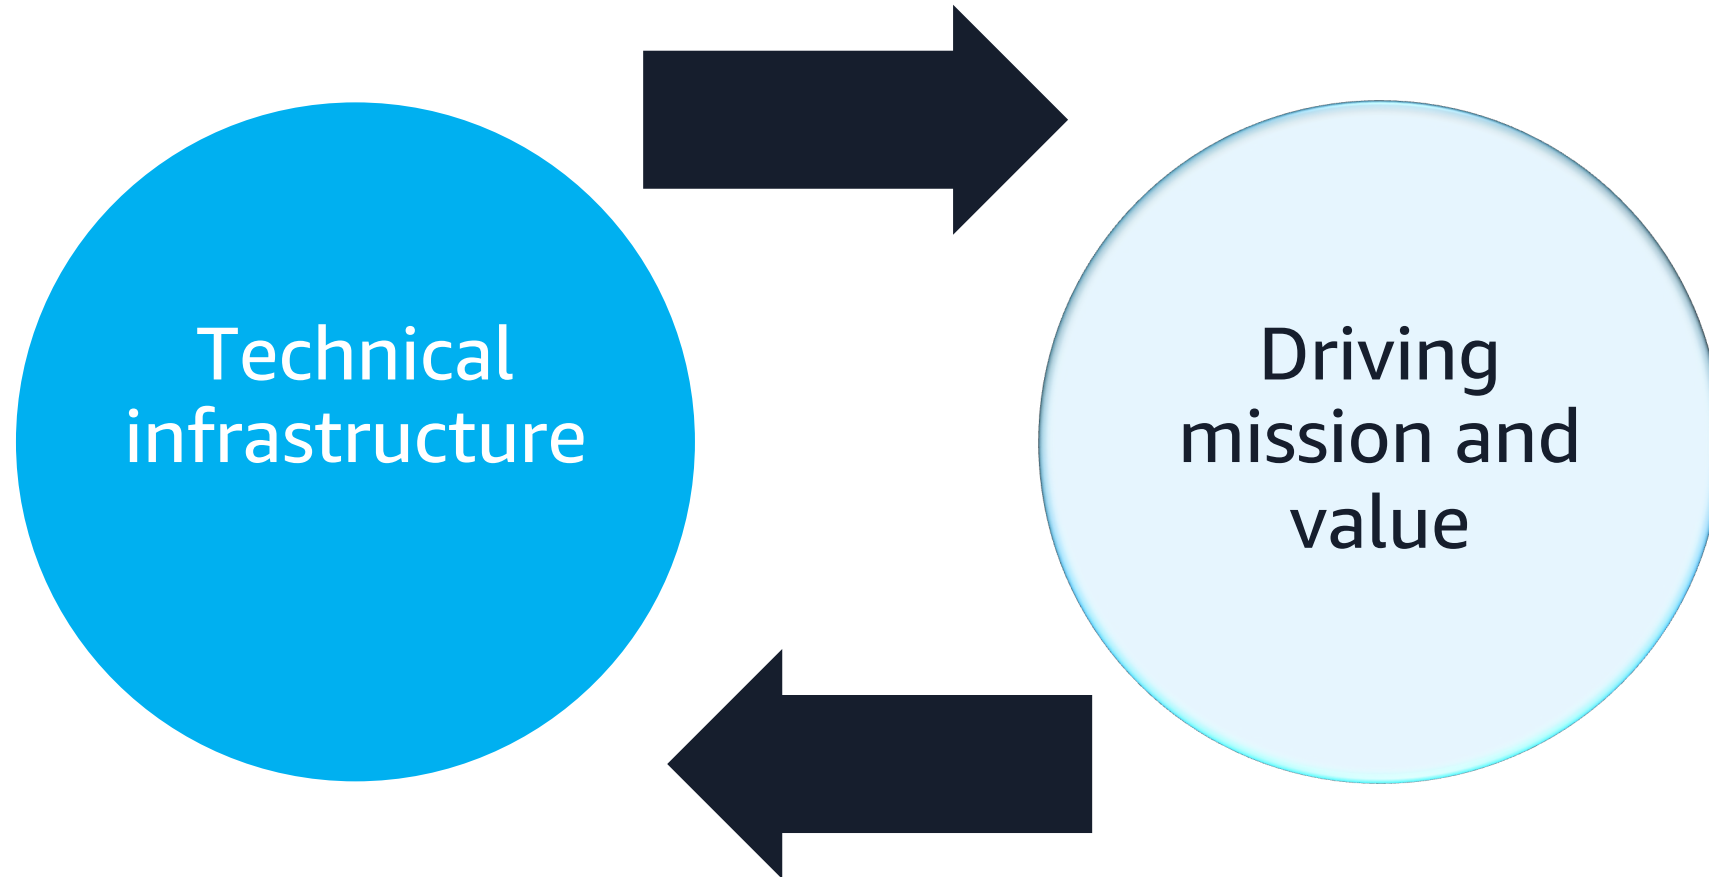
Some model providers use and store data for training purposes. Entire end-to-end data pipelines require security and data privacy controls.

Engineers accidentally releasing source code by putting into a foundation model for debugging

The importance of using AI responsibly

Consider how critical it is to use AI responsibly for reducing risks and deliver value comprehensively, at scale, while keeping the AI logic equitable and unbiased

AWS: Supporting Generative AI in the public interest



Security & Compliance considerations for generative AI

COMPLIANCE & GOVERNANCE

The policies, procedures, and reporting needed to empower the business while minimizing risk

Create generative AI usage guidelines

Establish process for output validation

Develop monitoring & reporting processes

LEGAL & PRIVACY

The specific regulatory, legal, and privacy requirements for using or creating generative AI solutions.

Retain control of your data

Encrypt data in transit and at rest

Support regulatory standards

CONTROLS

The implementation of security controls that are used to mitigate risk.

Human-in-the-loop

Explainability & auditability

Testing strategy

Identity and access management

RISK MANAGEMENT

Identification of potential threats to generative AI solutions and recommended mitigations.

Threat modeling

Third-party risk assessments

Ownership of data, including prompts and responses

RESILIENCE

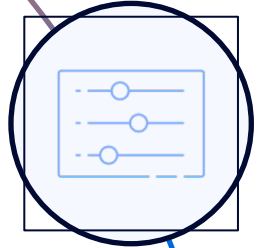
How to architect generative AI solutions to maintain availability and meet business SLAs.

Data management strategy

Availability

High Availability and Disaster Recovery strategy


Guardrails for Amazon Bedrock




Apply guardrails to multiple foundation models and Agents for Amazon Bedrock



Configure harmful content filtering based on your responsible AI policies

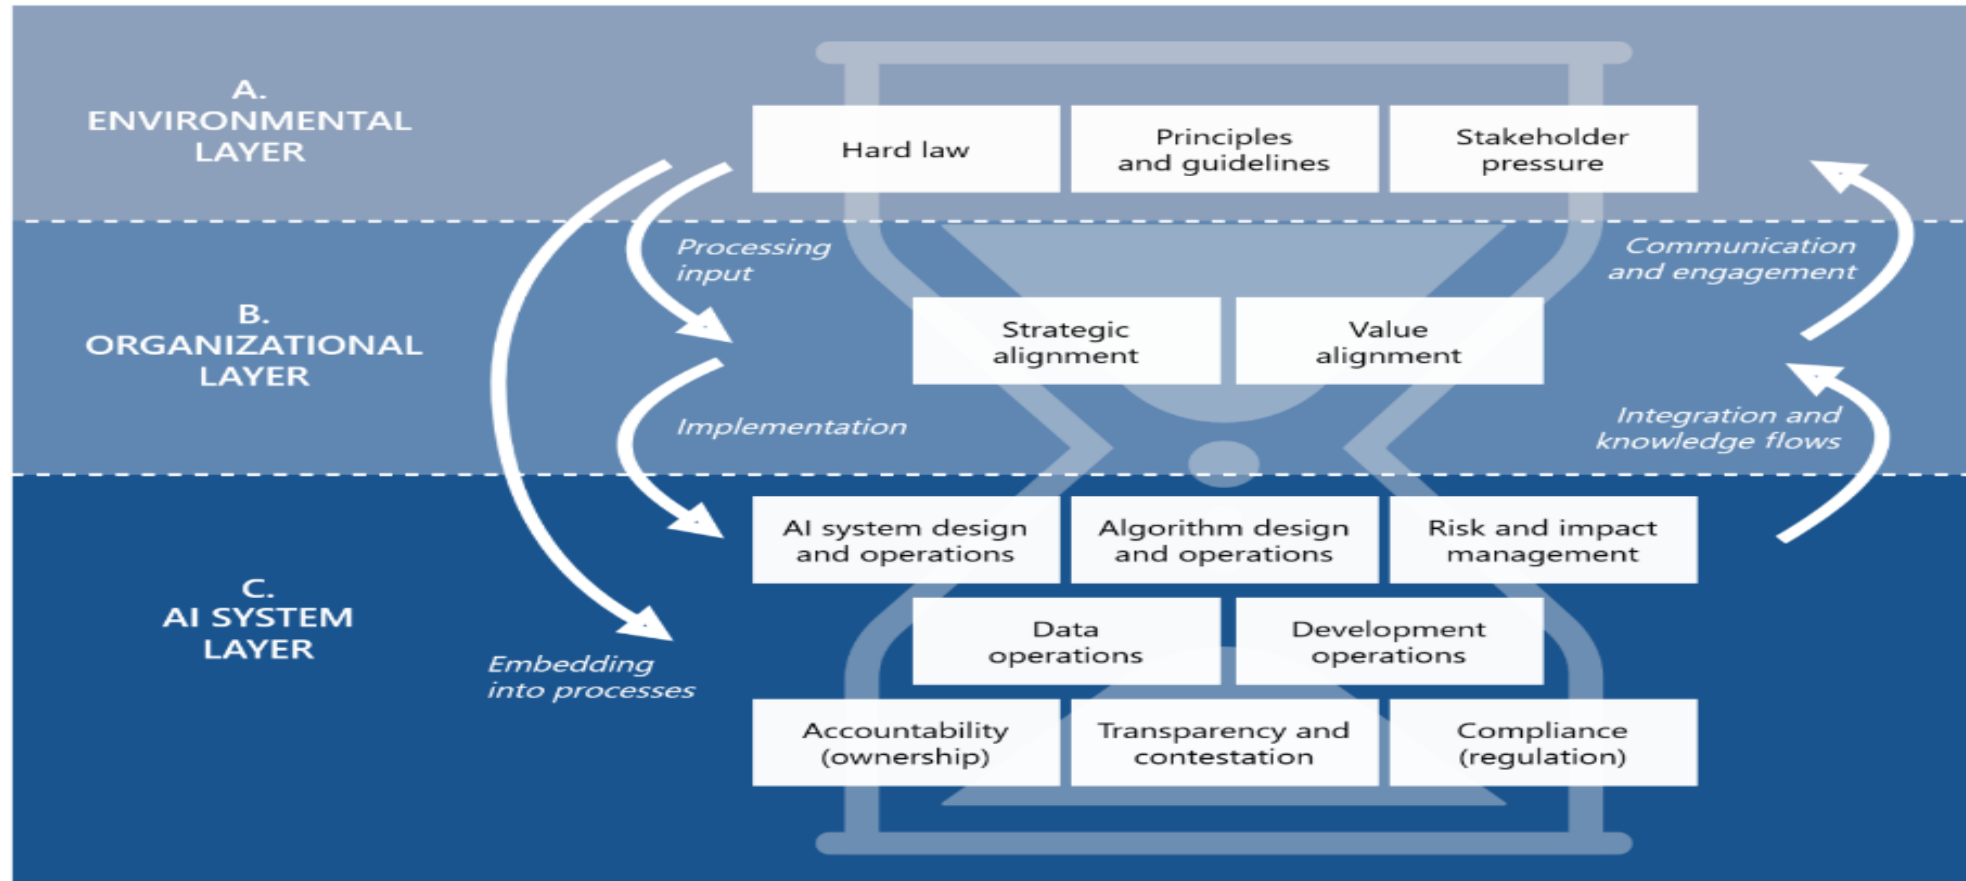


Define and disallow denied topics with short natural language descriptions



Redact sensitive PII information in FM responses

AIGA: EU AI Governance Framework



Citation: Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). *Putting AI Ethics into Practice: The Hourglass Model of Organizational AI Governance* (arXiv:2206.00335). arXiv. <https://doi.org/10.48550/arXiv.2206.00335>

Risks impacting organizations

Reputational impact

Poor organization perception; erodes customer base and hinders sales

Revenue loss

Diminished credibility and trust

Regulatory repercussions

Legal penalty or restrictions resulting from failure to adhere to laws or infringing on rights


“[Organizations] fail to focus on ethical, social, and regulatory implications, leaving themselves vulnerable to potential missteps when it comes to data acquisition and use, algorithmic bias, and other risks, and exposing themselves to social and legal consequences.”

HBR's Year in Business and Technology: 2021
referencing McKinsey & Company article “Ten Red Flags Signaling Your Analytics Program Will Fail”


A multi-disciplinary problem



Economics



**Moral
philosophy**



Technology



Law



**Social
science**

- Responsible AI is a complex, multi-disciplinary problem, blending requirements across a range of specialist fields
- Although some organizations have begun to establish a basic awareness of the problems associated with responsible AI, few have access to the requisite skills or experience to tackle this problem in a comprehensive manner

Pillars for the responsible use of AI

Value alignment

Systems should be designed and used in ways that align with company mission, social norms, and legal compliance

Inclusion

Inclusion of unique skills, experiences, perspectives, and cultural backgrounds

Training & education

Appropriate knowledge sharing and education to understand purpose, use, and impact

Accountability

Structured maintaining human involvement and responsibility for design, development, decision processes, and outcomes

Privacy & security


Protects the quality and integrity of data used, its relevance, access, and processing

Fairness

Systems must be designed to minimize bias and promote inclusive representation

Transparency & explainability

Understanding how data is used, how decisions and outcomes are made in a human understandable way



Governance: Educause Assessment

- Institutional generative AI readiness
- Multiple domains: data, training, governance, security procurement
- Recommendations:
 - Documentation of mitigation for algorithmic biases
 - Human review of Gen AI content
 - Transparency for content moderation to prevent toxic, bias or inaccurate information
 - Data security

Strategy

1. Does your institution provide these types of access to generative AI tools?


	Yes	No	Don't know
Licenses for students for a publicly available AI tool			
Licenses for faculty for a publicly available AI tool			
Licenses for staff for a publicly available AI tool			
An existing AI tool trained with your institution's data			
An integrated suite of AI products			

2. If your institution has a strategy for guiding future investments in AI, does it include the following elements?

	Yes	No	Don't know
High-level guiding principles or values			
A roadmap detailing specific actions or steps, milestones, and metrics for achieving the strategy			
Resources (funding, personnel, technology) allocated to support AI strategic initiatives			



Generative AI Stack




APPLICATIONS THAT LEVERAGE LLMs AND FMs




TOOLS TO BUILD WITH LLMs AND OTHER FMs



INFRASTRUCTURE FOR FM TRAINING AND INFERENCE



Benefits of building responsibly



Accelerate adoption



Institute appropriate governance structure



Align AI risk management with broader risk efforts



Develop people resources and skills



Build operational capability



Drive inclusive innovation



**Technological advancement must respect
the rule of law, human rights, and
dignity, as well as our shared values of
inclusivity, privacy, and fairness**

Accelerate your impact

Rapid AI prototyping and innovation engagements

Noah Eden

Global Lead, Cloud Innovation Centers
Worldwide Public Sector
Amazon Web Services

Cloud Innovation Centers: Solving public sector challenges

What is a CIC?

Cloud Innovation Centers enable public sector and AWS to **collaborate and solve real-world challenges**, through the collaboration with students and higher education.

Where are the CICs located?



How can CICs help?

Industry-Leading Innovation Processes

Cloud Technology Expertise

Prototype
Development

Innovation
Events



Innovating together with university students

Public sector organizations bring organizational or mission-related challenges...

We deliver technical proof of concept solutions...



My eCISO

A cybersecurity chatbot to help assess entities' compliance with NIST Cybersecurity Framework



ScopeBuilder

Uses LLMs to build high quality scopes of work for campus procurement



Educational Simulation Tool

Helps pharmacy students practice in real-world patient interactions



PDF Accessibility Project

“ With the introduction of new accessibility standards, AI and machine learning offer what may be the most viable path to success, given our resources and scope.

Cory Tressler

Assistant Dean for Technology and Digital Programs


The Ohio State University Libraries



THE OHIO STATE
UNIVERSITY

UNIVERSITY LIBRARIES

Engagement lifecycle



What is your top innovation priority?

- ...to more easily communicate with constituents?
- ...to index your websites and enable smarter, multi-lingual search?
- ...to quickly analyze or aggregate thousands of paper records?
- ...to enable citizens to access the resources and guidance they need?
- ...to improve graduation rates via support of faculty?

Or something else entirely?

Let's innovate and build together.



Thank you!

Sergio Ortega

AI/ML BD and Sales Lead
State and Local Governments
sergioai@amazon.com

Noah Eden

Global Lead, Cloud Innovation Centers
Worldwide Public Sector
noaheden@amazon.com

**Please complete the survey
for this session**



AI/ML track

**Generative AI/ML and AI
governance for the public
sector**