

AWS State, Local, and Education Learning Days

Salt Lake City, UT



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

[Linkedin.com/sergioortegaai](https://www.linkedin.com/company/amazon-ai/ml/?viewAsMemberOf=sergioortegaai)



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

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

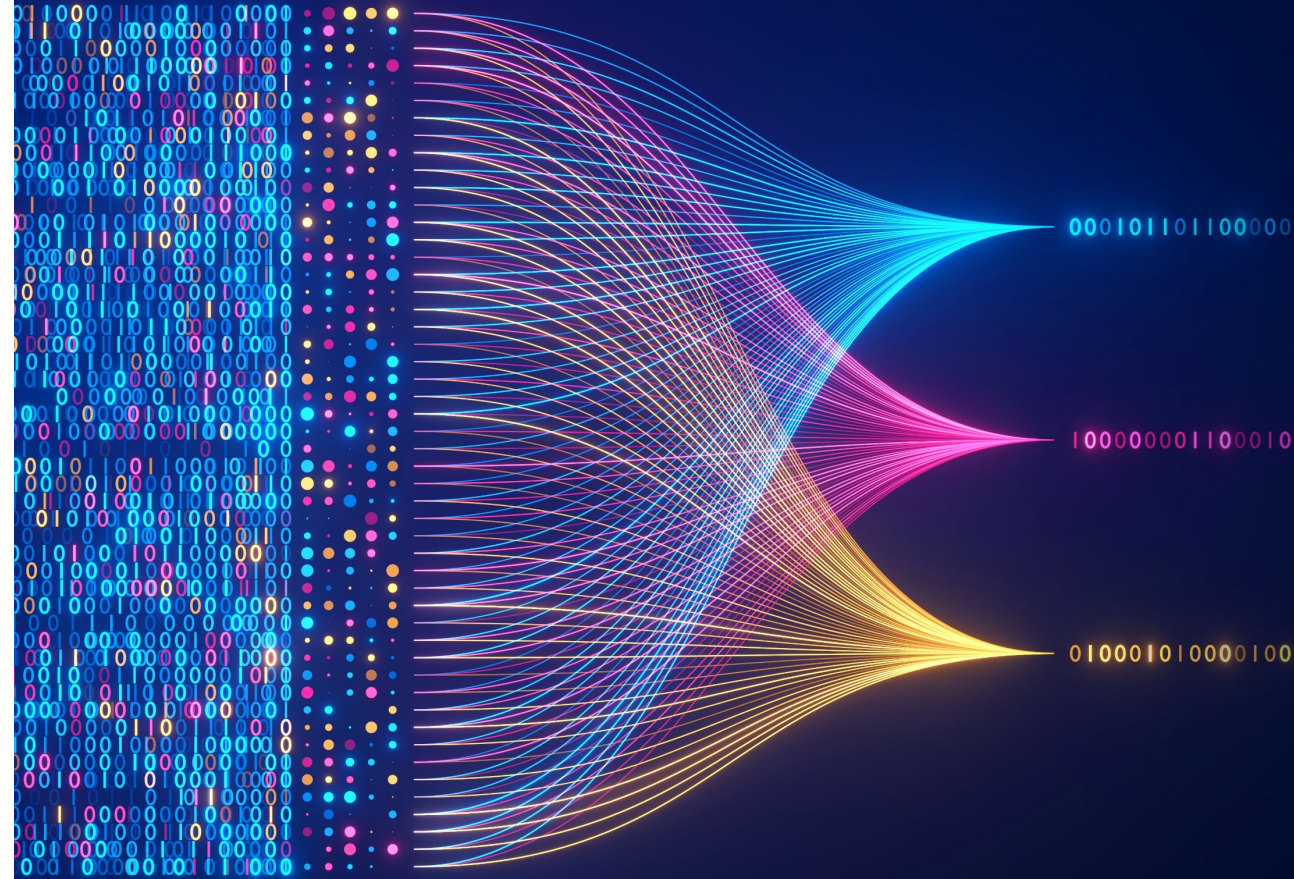
Generative AI is powered by foundation models

Pretrained on vast amounts of unstructured data

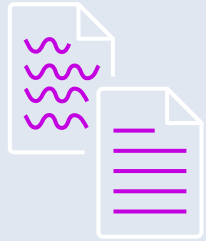
Contain large number of parameters that make them capable of learning complex concepts

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

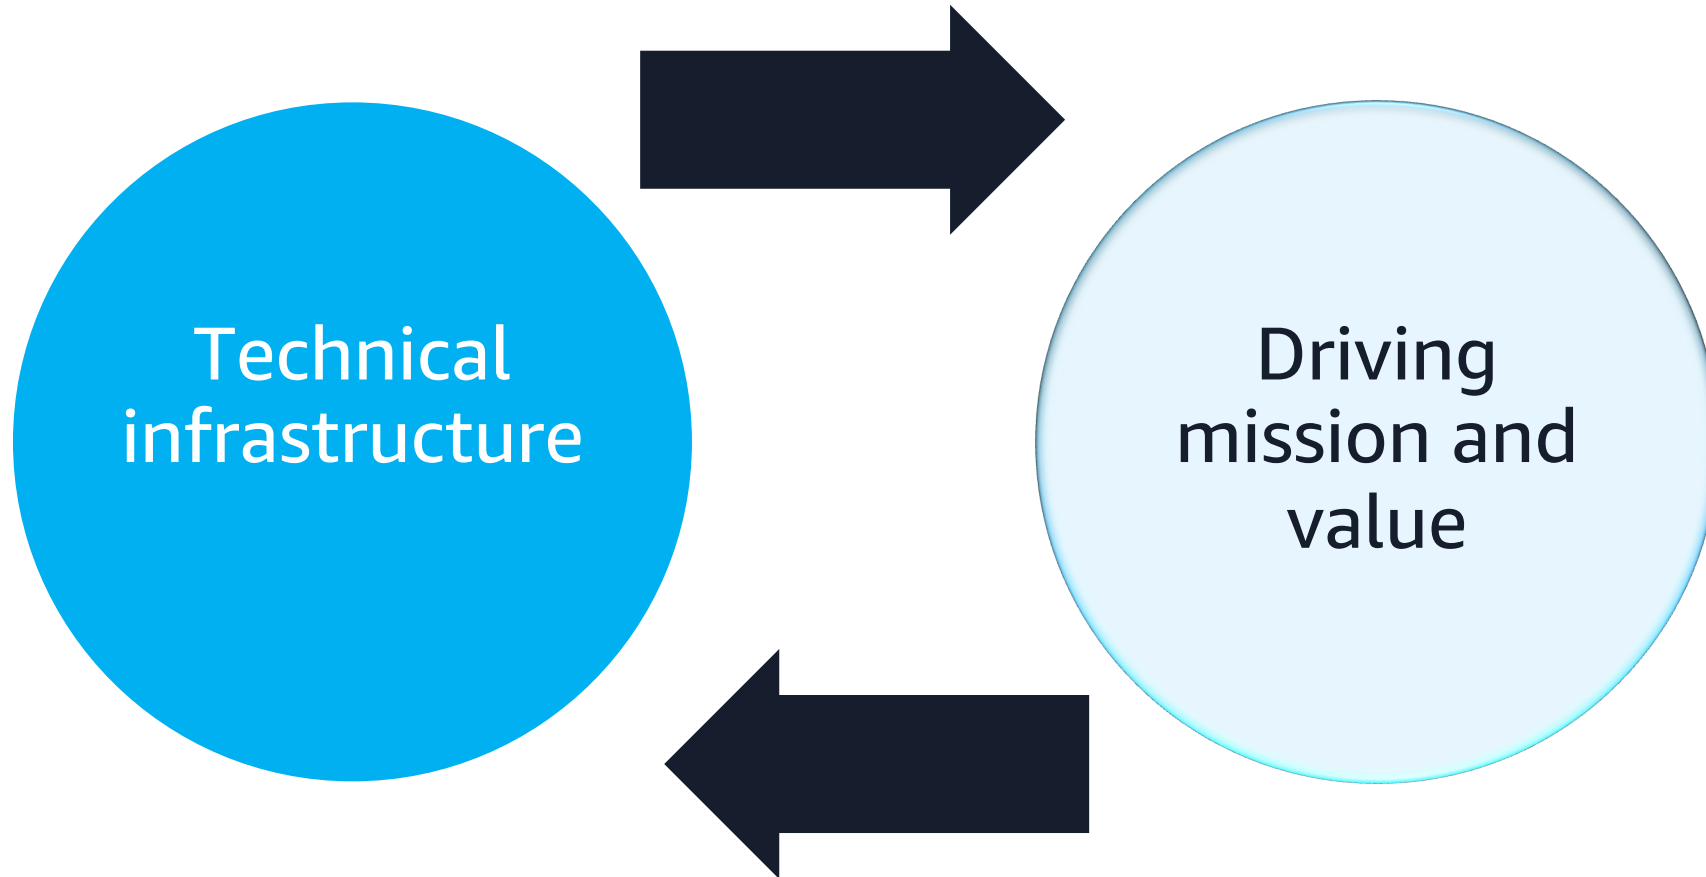
Some model providers use and store data for training purposes. Entire end-to-end data pipelines require security and data privacy controls.

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

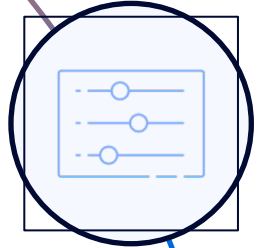
How to architect generative AI solutions to maintain availability and meet business SLAs.

Data management strategy

Availability

High Availability and Disaster Recovery strategy

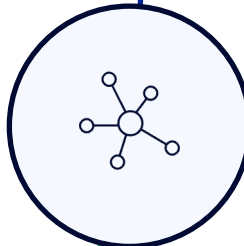
Guardrails for Amazon Bedrock




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

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

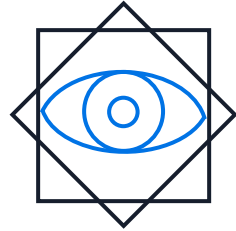
“[Organizations] fail to focus on ethical, social, and regulatory implications, leaving themselves vulnerable to potential missteps when it comes to data acquisition and use, algorithmic bias, and other risks, and exposing themselves to social and legal consequences.”

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

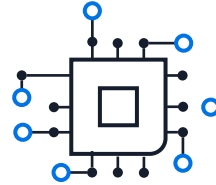
A multi-disciplinary problem



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

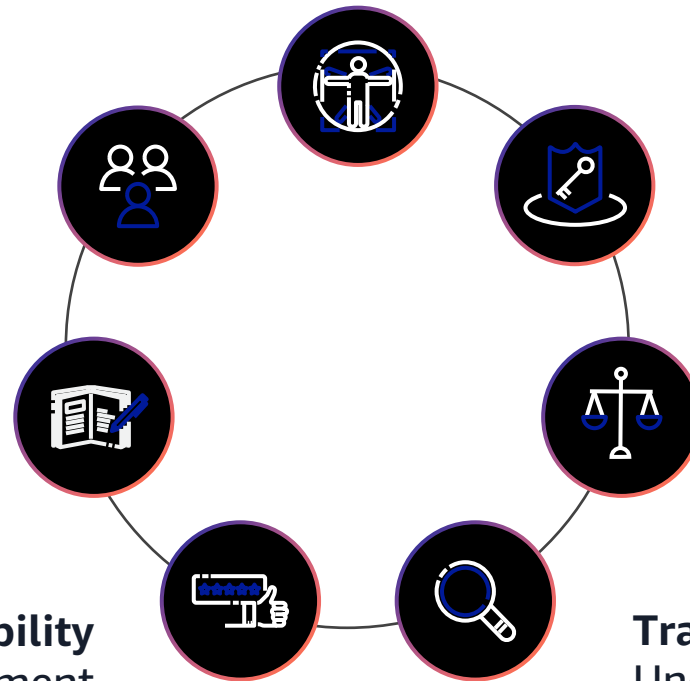
Protects the quality and integrity of data used, its relevance, access, and processing

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



Benefits of building responsibly



Accelerate adoption



Institute appropriate governance structure



Align AI risk management with broader risk efforts



Develop people resources and skills



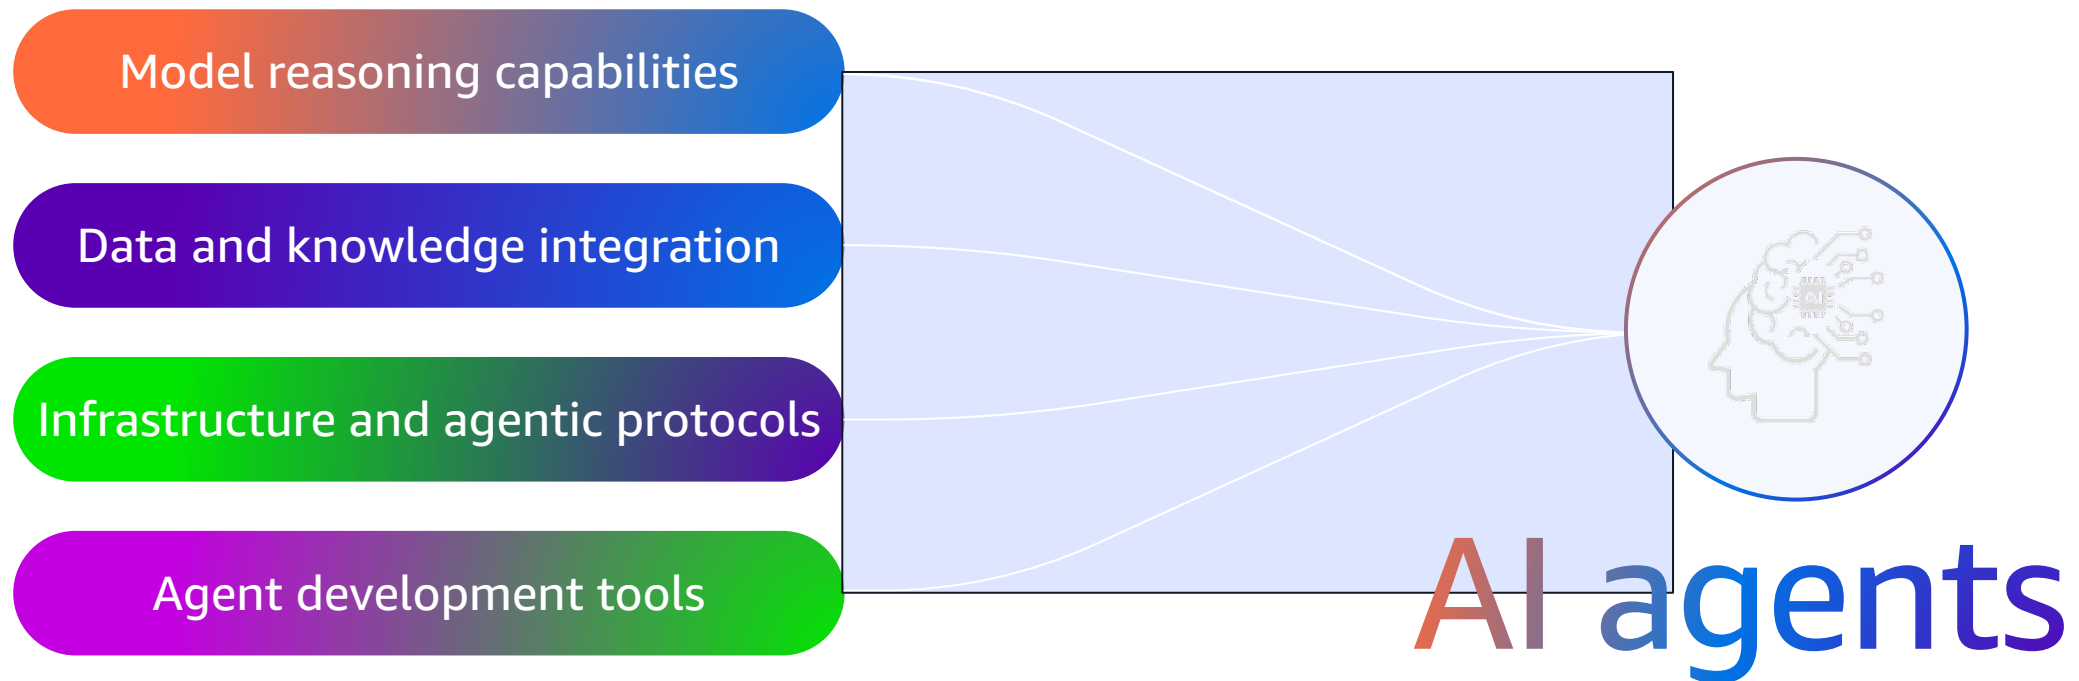
Build operational capability



Drive inclusive innovation



Technology advancements make AI agents practical for enterprise use



Software systems that leverage AI to reason, plan, and complete tasks on behalf of humans or systems

Enterprises are doubling down on agents

33%

of enterprise software apps will include agentic AI by 2028, up from less than 1% in 2024.

Gartner, "Top strategic Technology Trends for 2025," October 2024

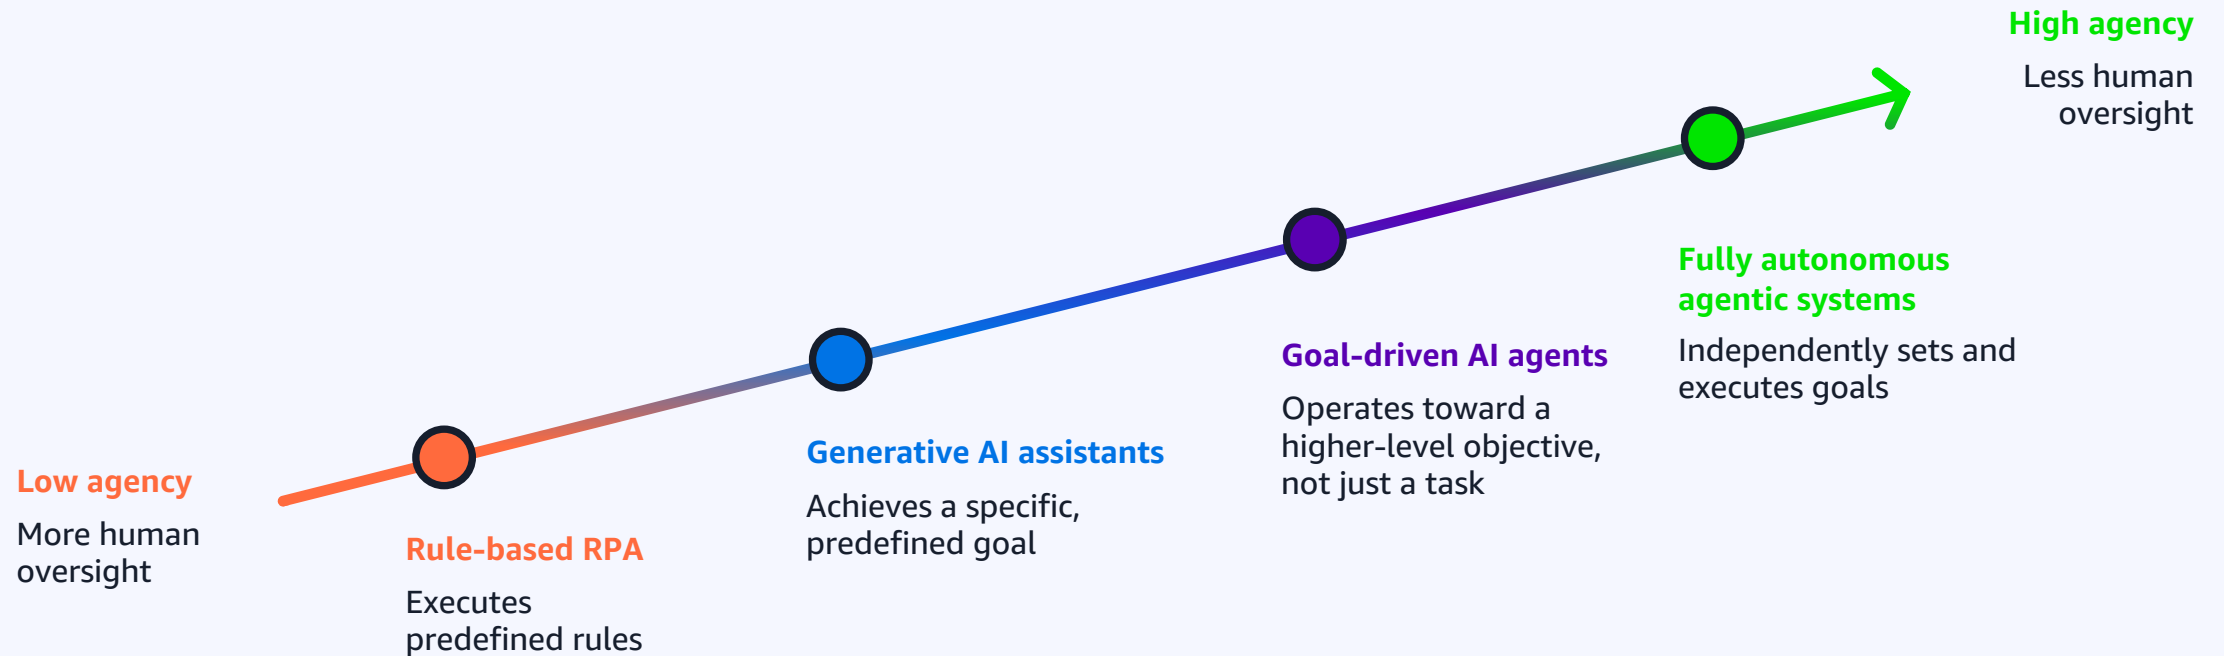
15%

of day-to-day work decisions will be made autonomously through agentic AI by 2028.

Gartner, "Top Strategic Technology Trends: Agentic AI—the Evolution of Experience" February 2025



Agentic AI maturity scale – from rules to reasoning



Where does your use case fit?

Scale requires operational excellence

Excitement
and potential



POC

Challenges on the path to production



Performance



Scalability



Security

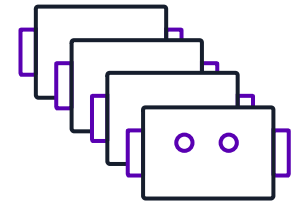


Context



Governance

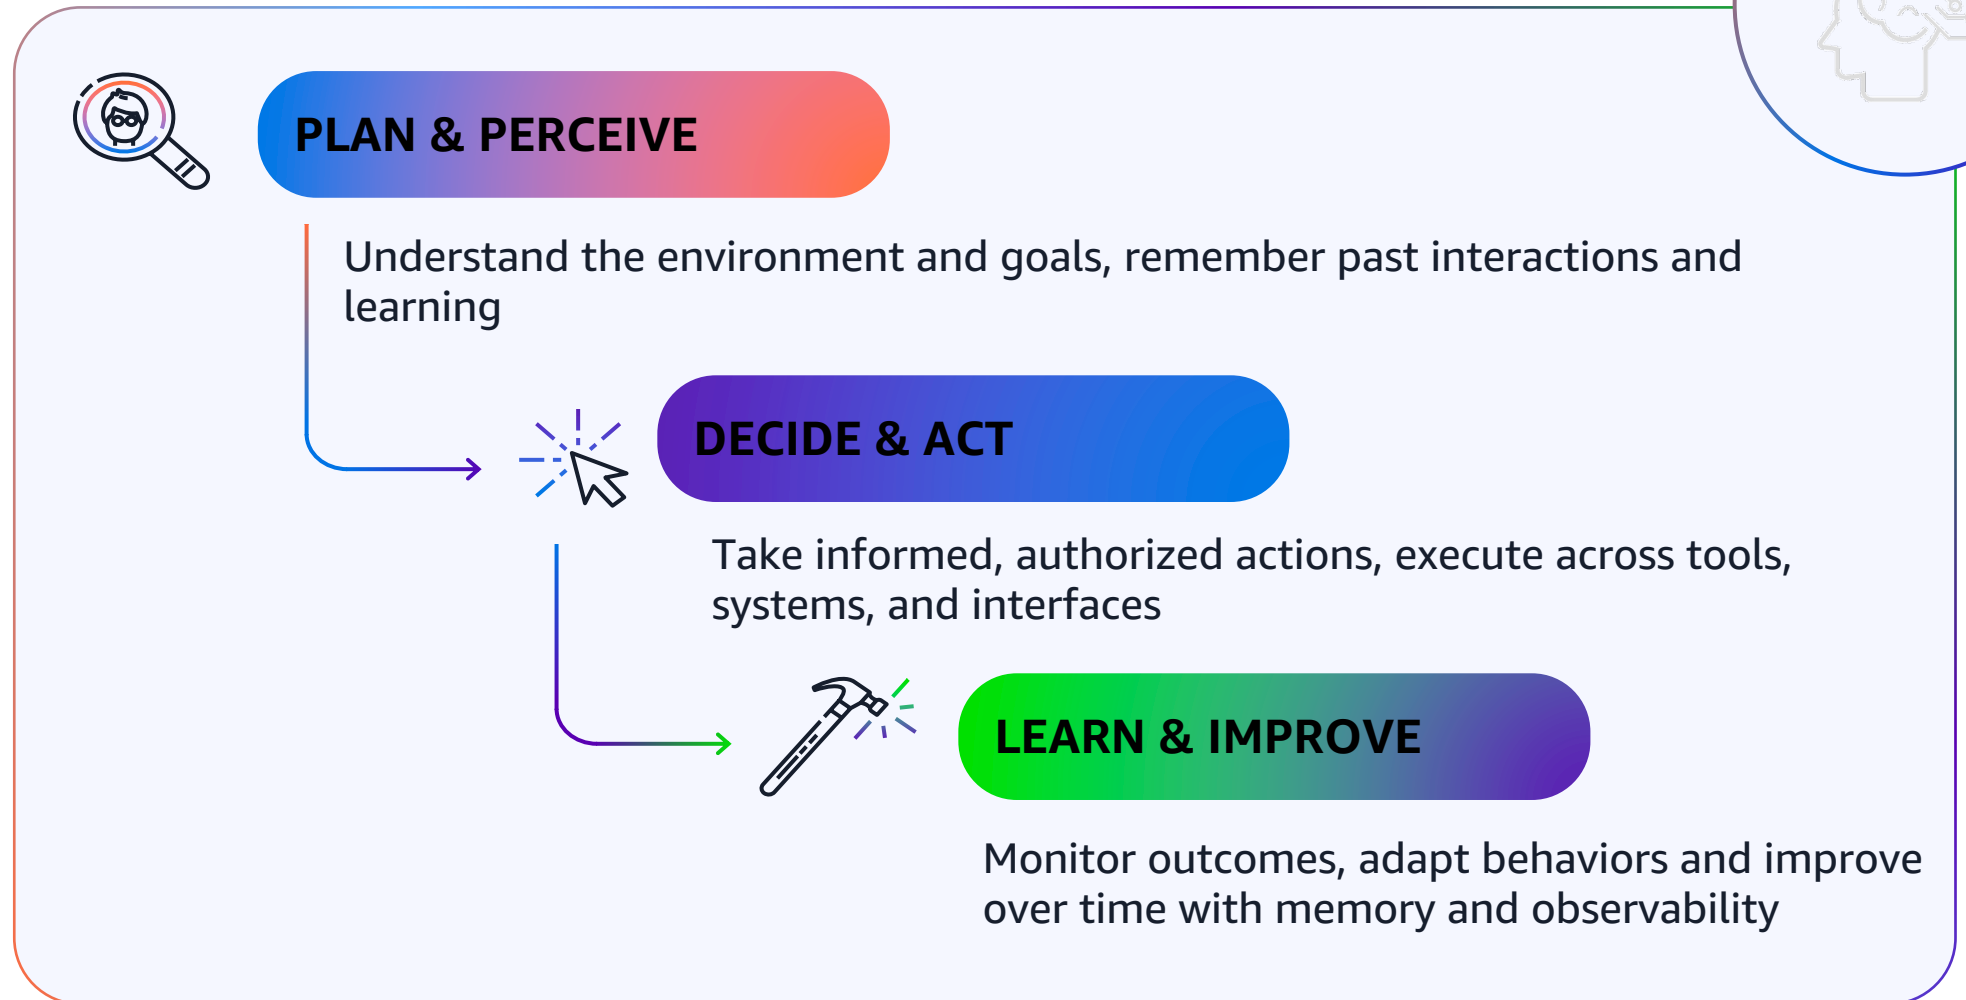
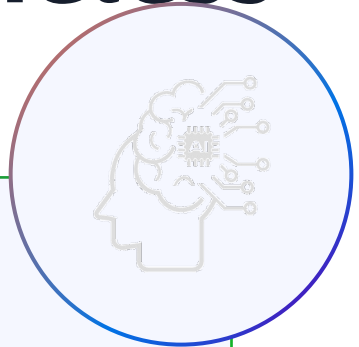
Meaningful
business value



AI production
agents

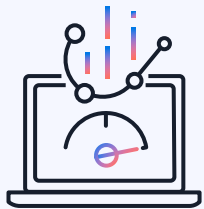


Agentic AI is where innovation meets timeless fundamentals



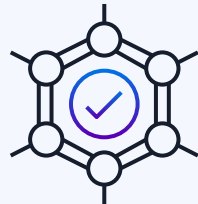
Our guiding principles for scaling agentic AI

EMBRACE AGILITY



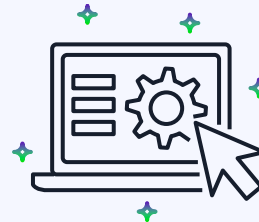
Stay flexible to outpace change. Build open, adaptable systems that evolve with your needs

EVOLVE FUNDAMENTALS




Build future-proof foundation for agents that enables secure, scalable autonomy

INNOVATE WITH MODELS & DATA



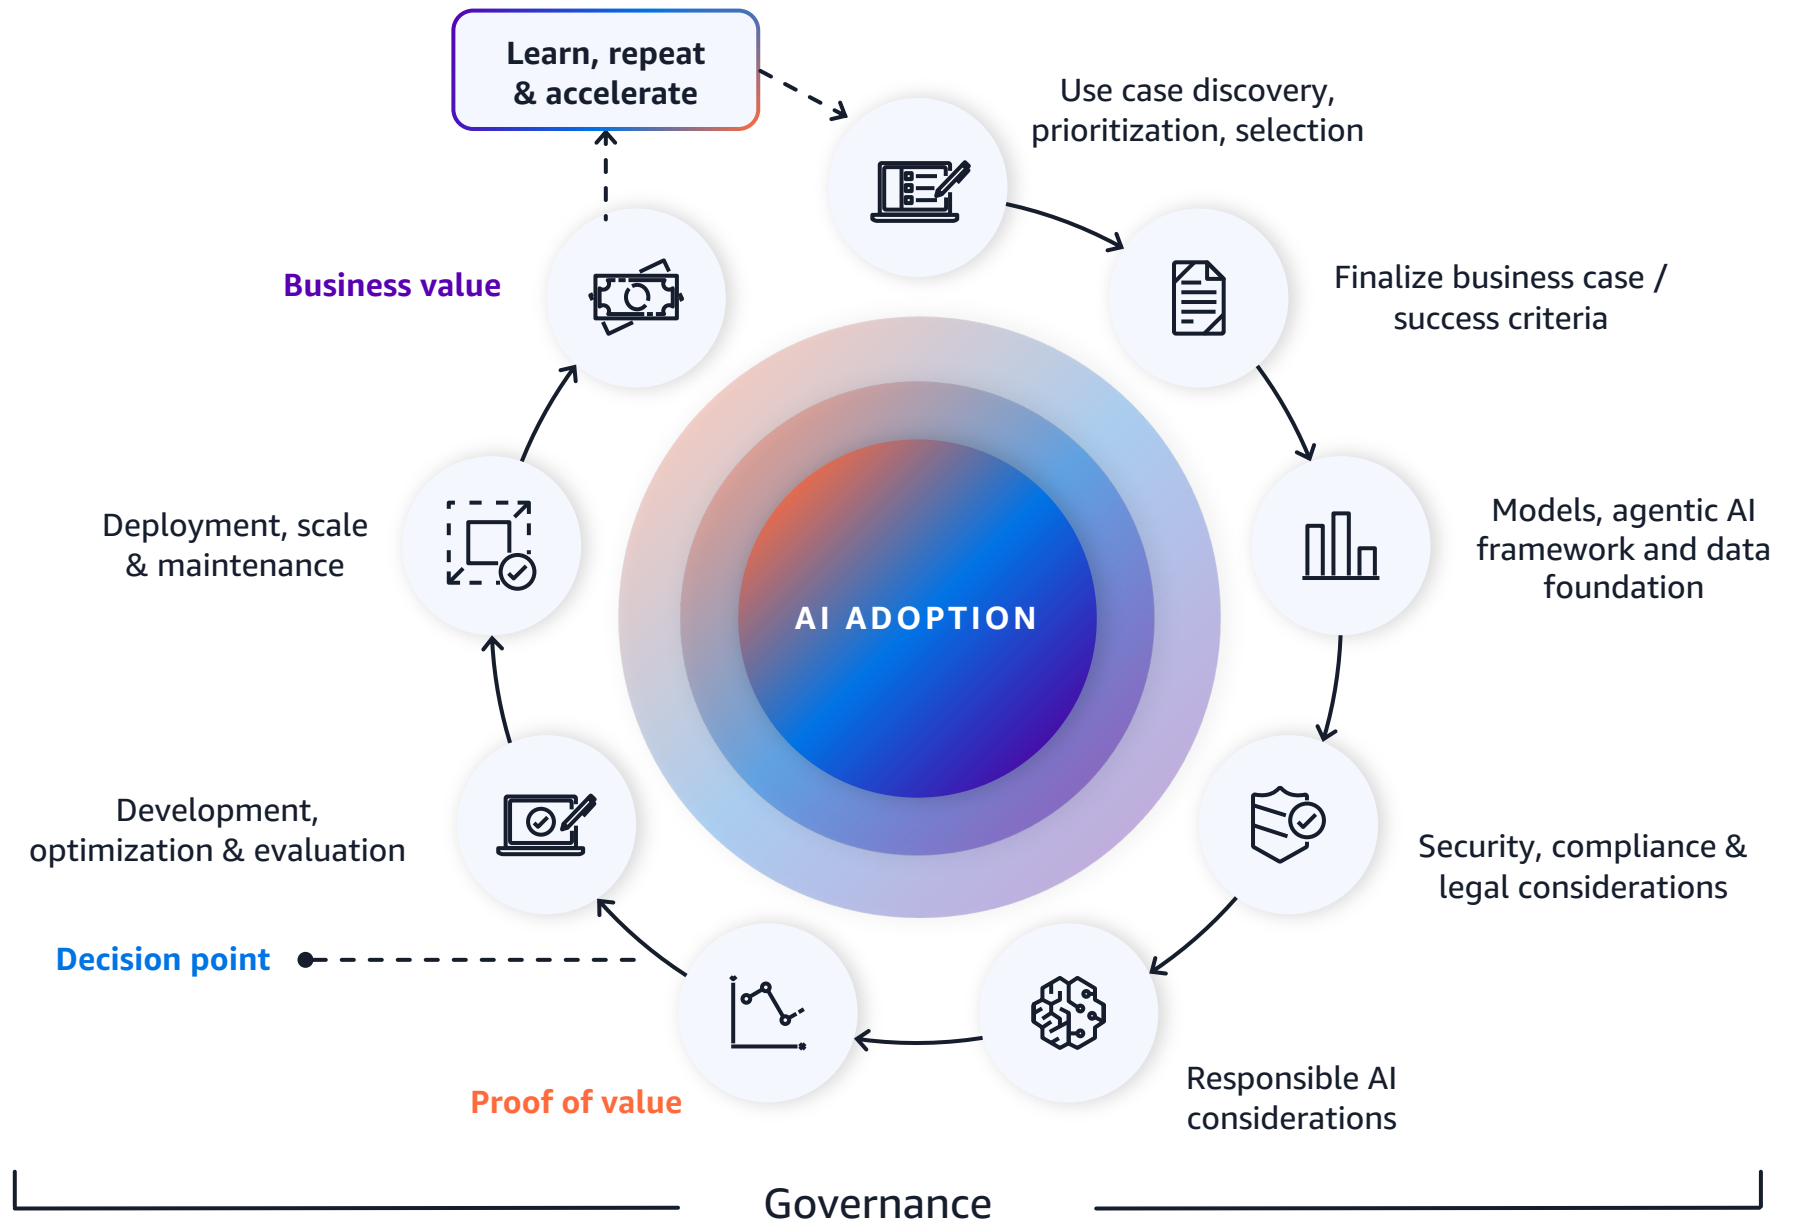
Pair the best models with your data. Deliver smarter decisions by combining choice and context

DEPLOY READ- TO-USE AGENTS

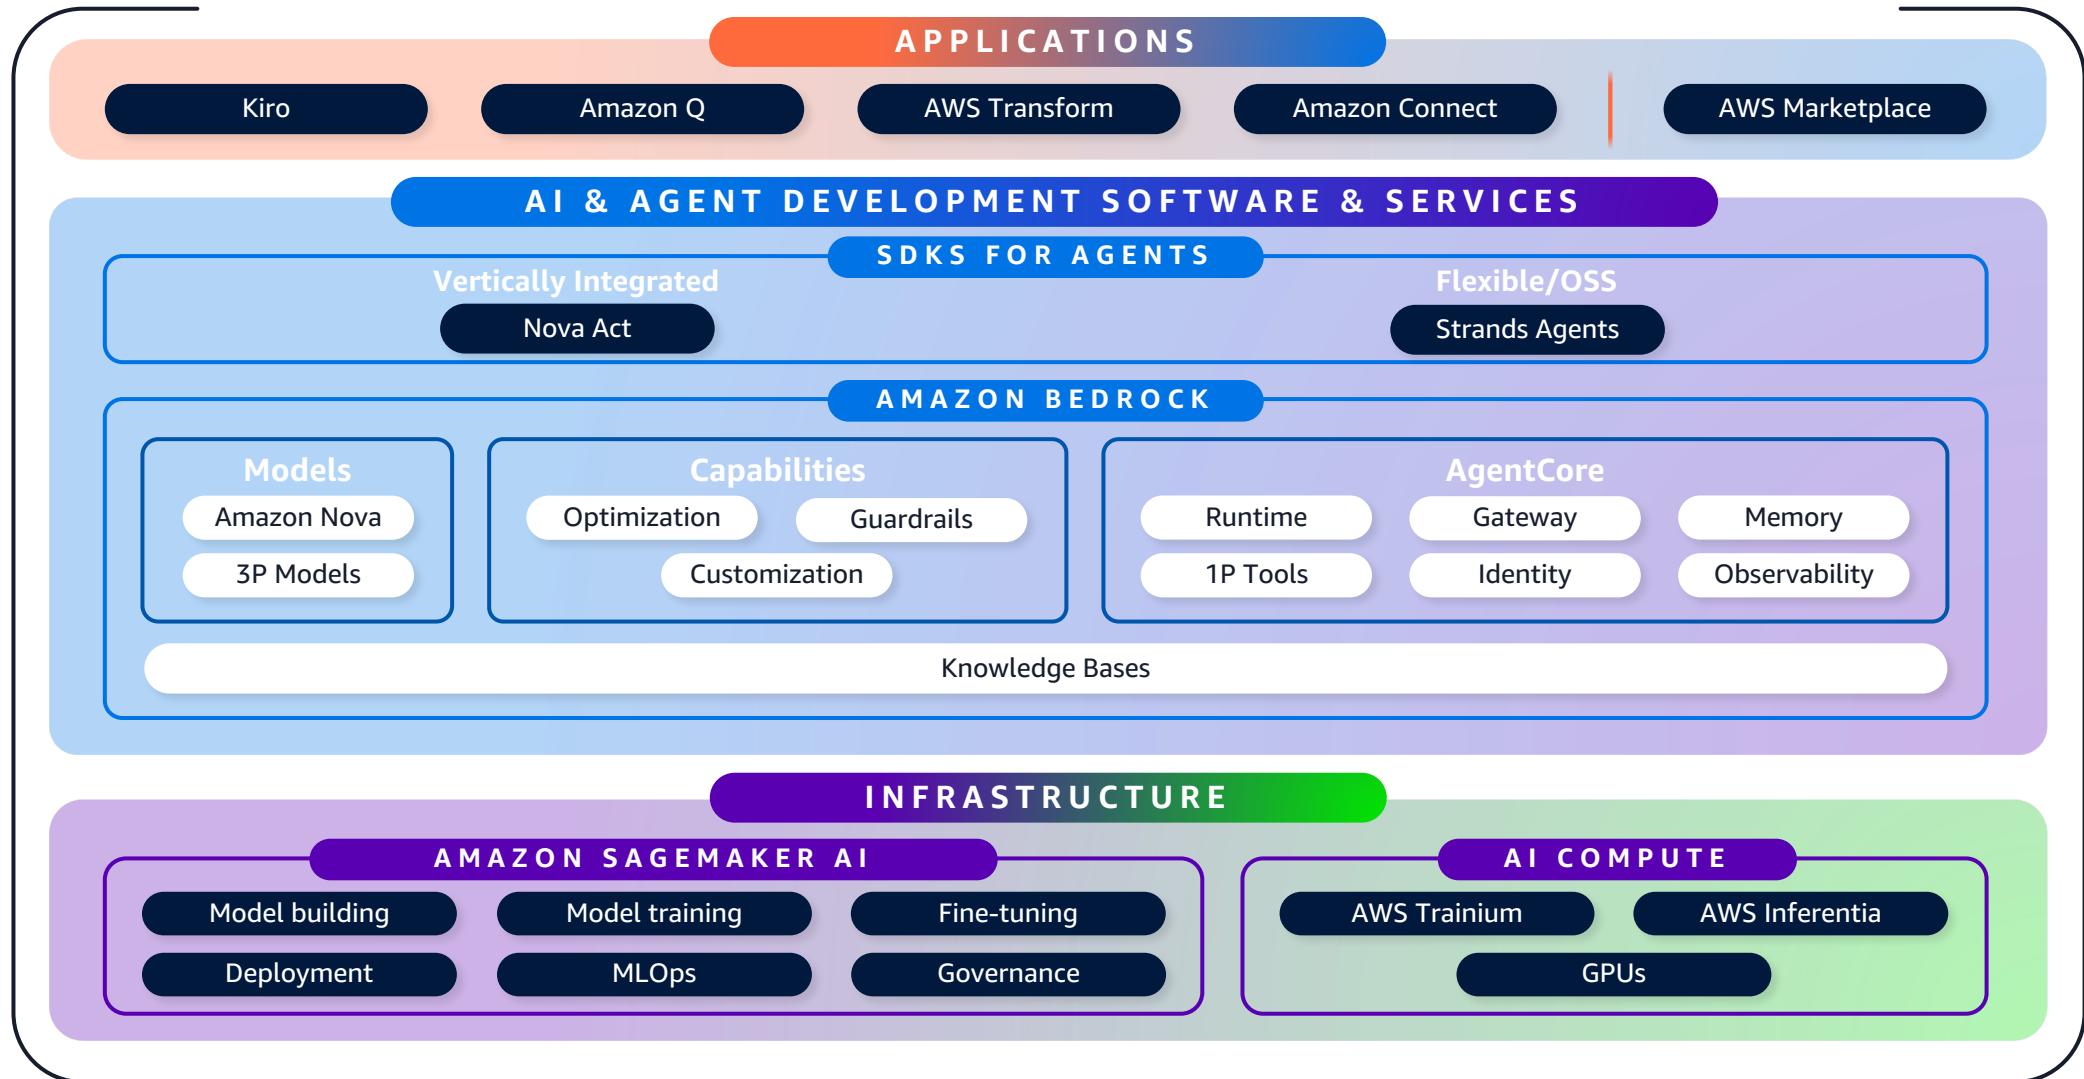


Transform work with intelligent solutions. Automate tasks and unlock human creativity at scale

Start now



AWS for Agentic AI



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



Innovating for public sector

No-cost, no-contract rapid POC + open-source solution development


Colleen Schwab

Global Innovation Strategist, Head of AI
Cloud Innovation Center

AWS

cmschwab@amazon.com

Organizing for innovation



Cloud Innovation Centers (CICs)



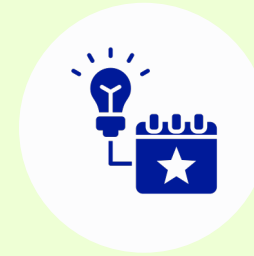
Leading innovation
processes



Cloud
expertise




Prototype
development



Innovation
events

ASU Artificial Intelligence
Cloud Innovation Center
Arizona State University

powered by 




CAL POLY
Digital Transformation Hub

powered
by 



University of
Pittsburgh

Health Sciences and
Sport Analytics

Cloud Innovation Center
powered by 

Student workers

- Undergraduate and graduate students
- Paid hourly
- Work with CIC for 1+ year
- **Roles**
 - Cloud developer
 - UI/UX
 - Marketing



Sample of CIC customers



Oregon State University



Alameda County



BreastCancer.org



Brightpoint



City of San Diego



City of Chandler Police Department



CalFire



Orthodox Union



Disability Rights Texas



City of San Luis Obispo



Colorado State University



Phoenix Children's Hospital



St. Louis Zoo



University of Pittsburgh



The Ohio State University



WorldBank



U.S. Department of Agriculture (USDA)



Wisconsin State Laboratory of Hygiene




Catholic Charities



Virginia Department of General Services
Division of Consolidated Laboratory Services



Developing prototypes in 6-10 weeks



PDF Remediation: The Ohio State University

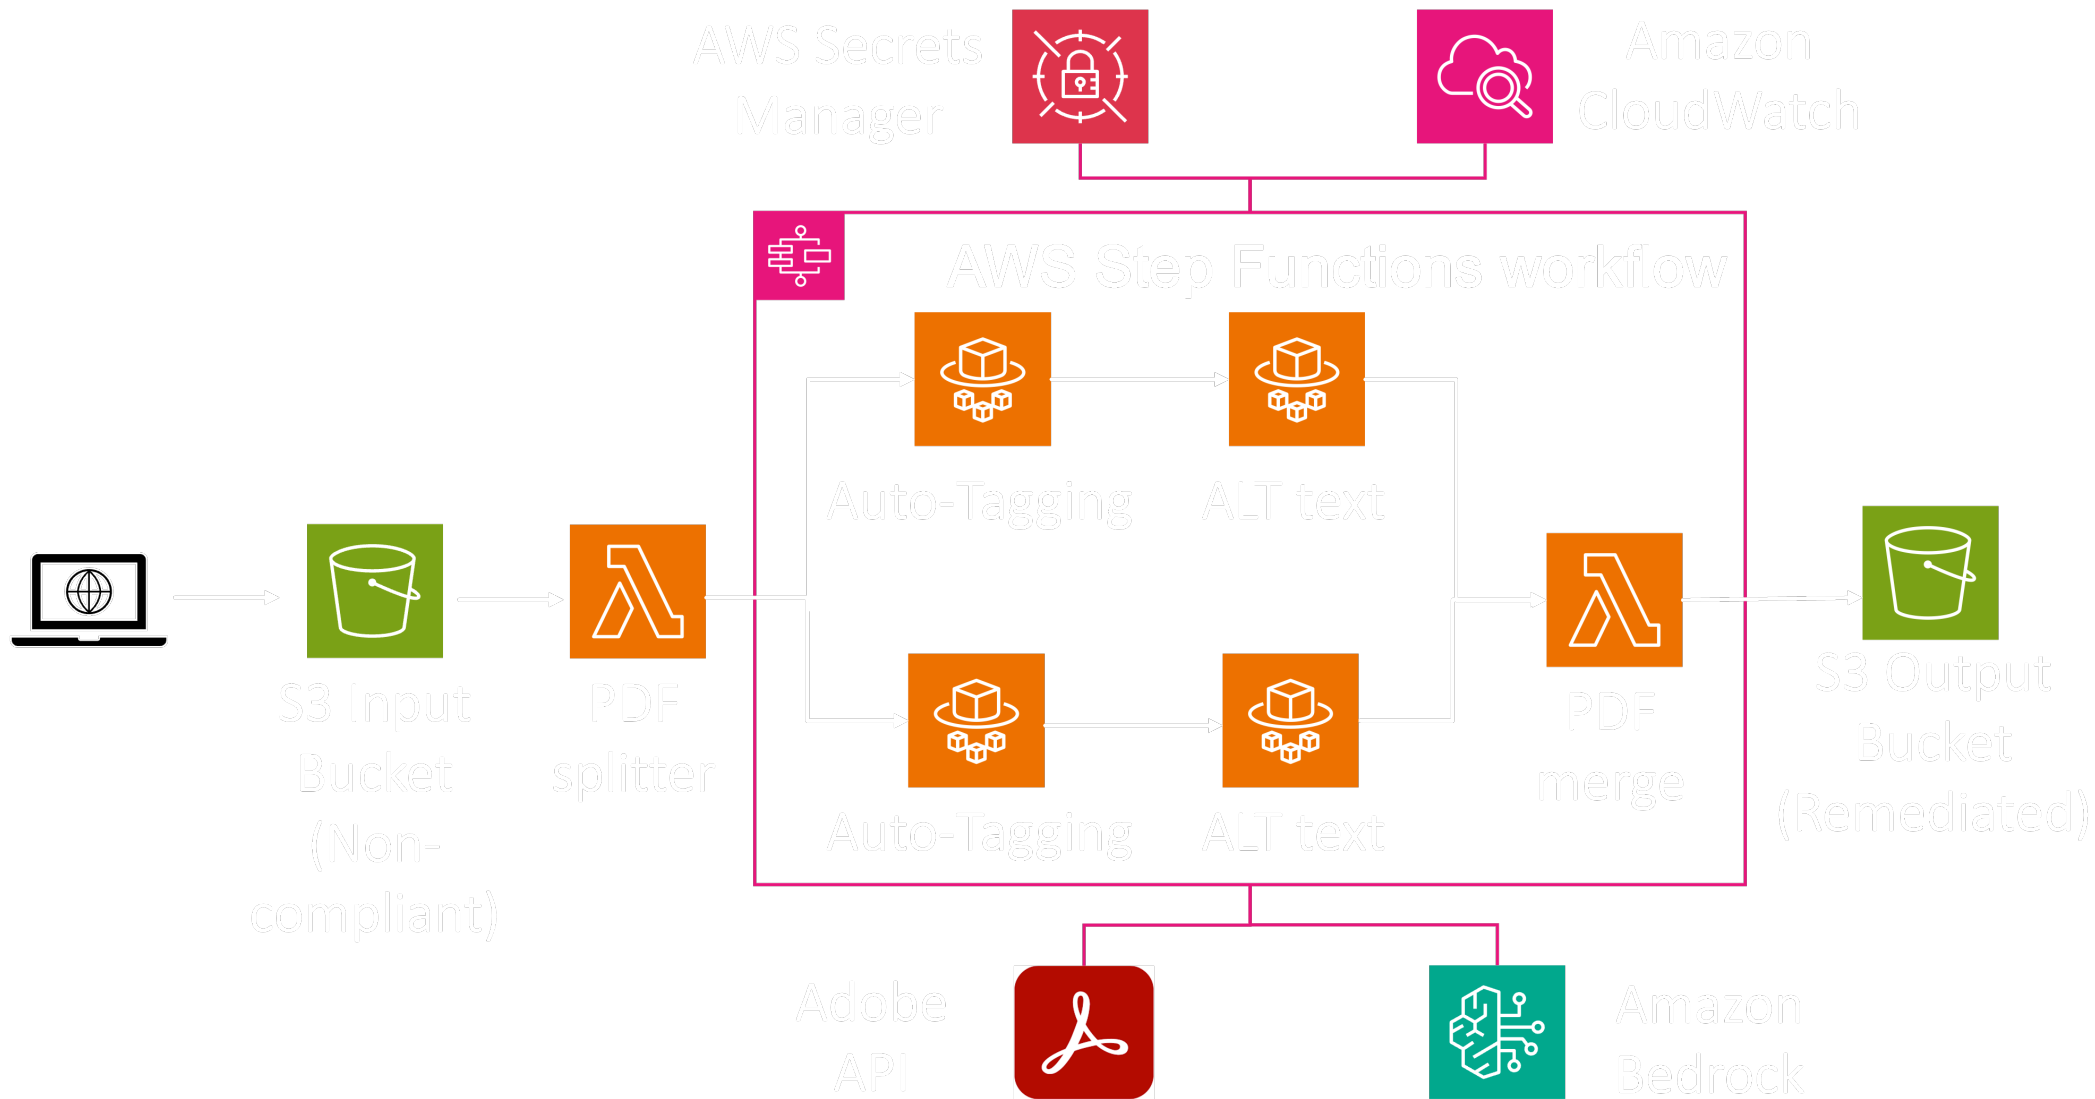
Problem statement: Many public sector organizations have document collections containing hundreds of thousands of PDF documents, many of which do not meet the Web Content Accessibility Guidelines (WCAG) 2.1 Level AA standards, making it difficult or impossible for individuals relying on assistive technologies to access those documents.

Solution

Our team combined AWS services + Adobe Auto Tag API to create a comprehensive solution to remediate PDF documents.

Other applications

AWS also has an open-source PDF-to-HTML solution.



Assessor Assistant: State of CA

Problem statement: Understanding property tax regulations in California is complicated and confusing for most homeowners. This conversational AI assistant uses natural language processing and semantic search to provide accurate, cited answers from sources like the California Assessor's Handbook, Tax Codes, and Letters to Assessor documents. The chatbot also features an intelligent exemption form finder that guides citizens through tailored yes/no questions, recommending appropriate exemption forms.

Solution

This solution significantly enhances citizen interactions, reduces assessor workloads, and streamlines training, transforming regulatory navigation.

Other applications

Complex information ecosystems that are hard or nuanced for people to navigate.

Chat

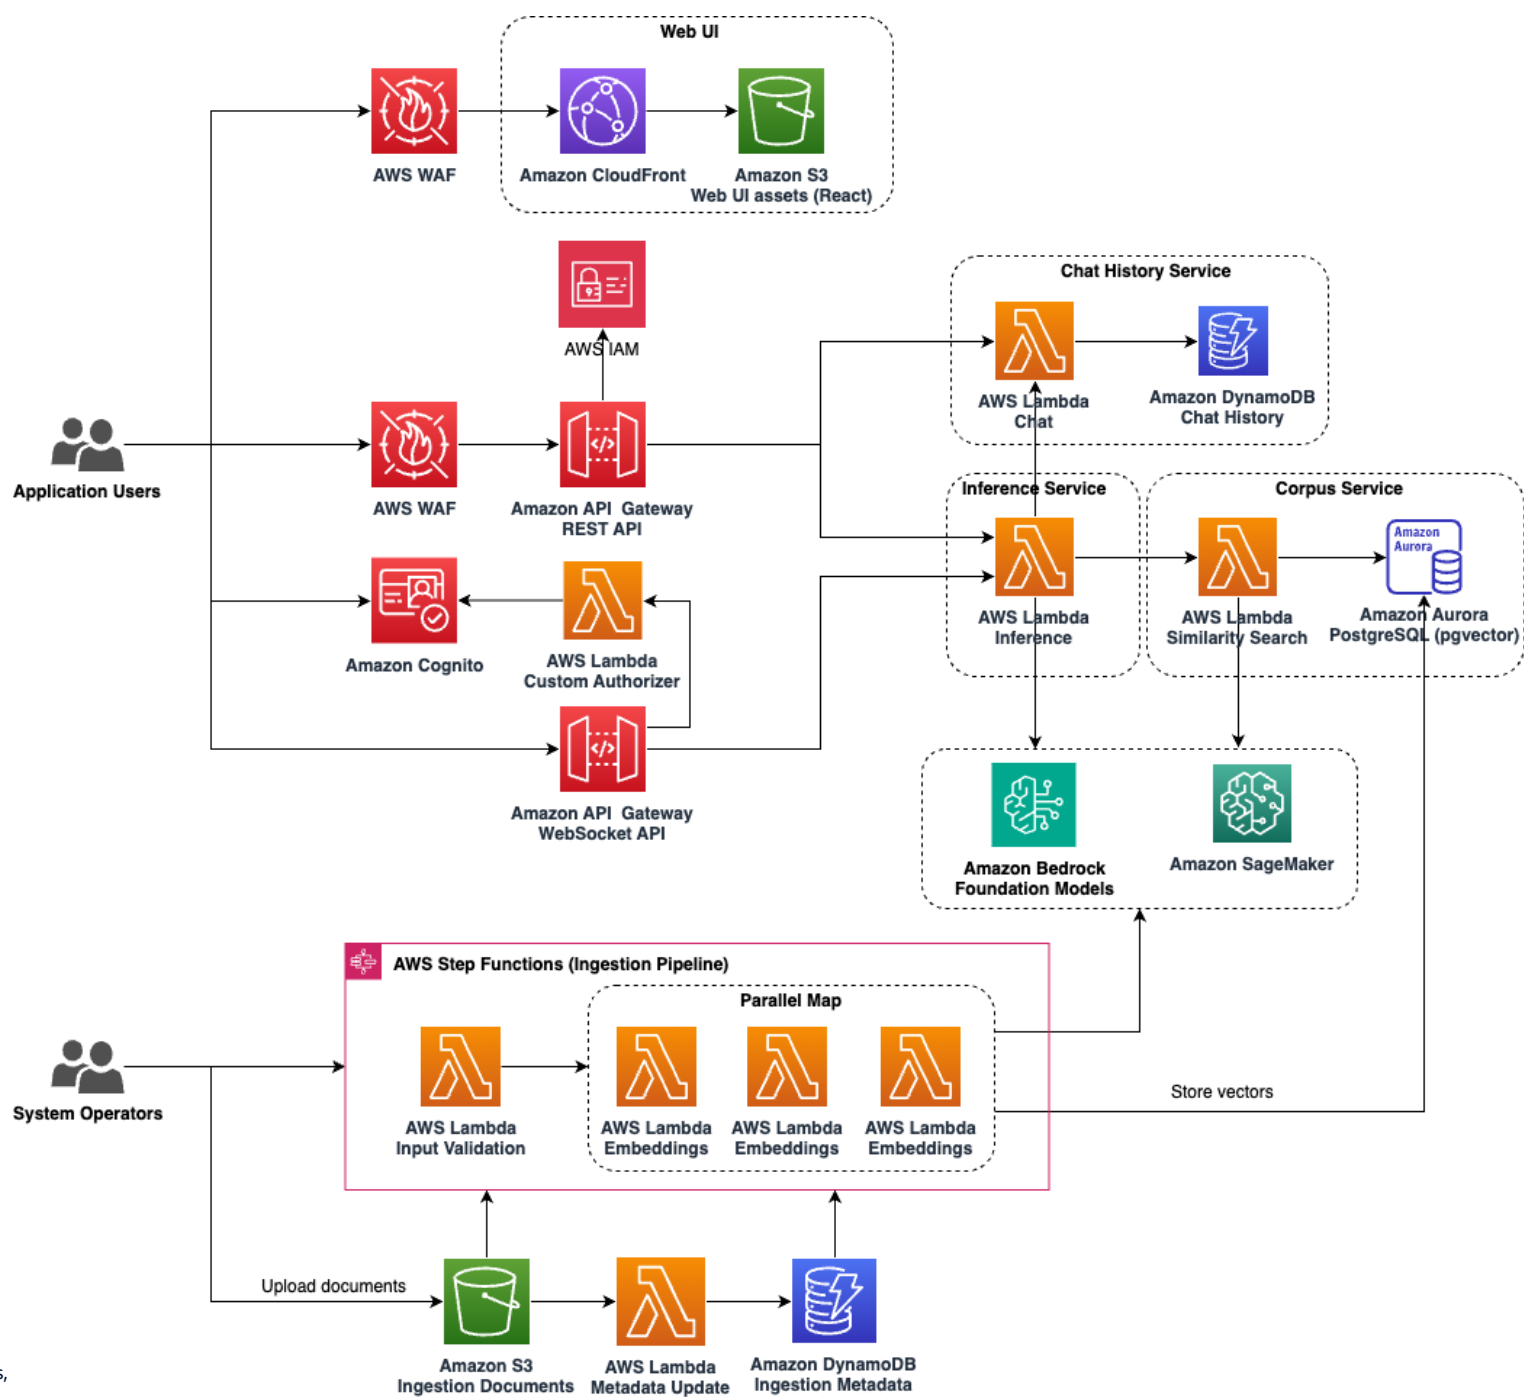
Chat [Info](#)

Content creation chats

Chats [+ Create Chat](#)

< 1 >

No chats
No chats to display.



Blueberry Bot: Oregon State University

Problem statement: Additionally, there is often a natural lag between the completion of research and the publication or communication of results. As a result, the most up-to-date insights may not always be immediately available in responses to growers. The current system also lacks a feedback loop to track common questions and integrate answers into a shared knowledge base, making it harder to streamline support and ensure consistent dissemination of new information.

Solution

Seamlessly connect people to the information they need w/ a commitment to accuracy. When the bot encounters questions beyond its current knowledge base, it seamlessly transitions to a smart email routing system, connecting growers directly with Dr. Yang for personalized expertise.

Other applications

Knowledge capture of a long-time employee; healthcare professionals accessing the latest medical protocols, legal teams querying vast case law databases, or engineers referencing technical specifications. Anywhere that disparate bodies of information exist that need to be tapped to answer questions



About us

Powered by AI-driven insights, Blueberry Bot provides real-time guidance on best farming practices, pest control, and sustainable harvesting techniques.

FAQs

What is meant by Maximum Residue Limit?

Tell me about Postharvest Handling

Give me the best practice for pest management.

What is the best way to grow blueberries in Central Oregon?

Blueberry AI Assistant



Welcome user! In order to provide the most accurate responses, can you please tell me where you are growing blueberries?

Oregon



Thank you for sharing that information! How can I help you today?


What is the blueberry shock virus?



The blueberry shock virus (BIShV) is a viral disease that affects highbush blueberry plants. It was first identified in 2004 in a single 'Bluetta' plant. The virus causes shock symptoms such as the death of flowers and foliage. Plants may recover from the virus in subsequent years and produce a normal fruit load, even though they remain infected. The virus spreads gradually through susceptible plants, as depicted in an animation showing its nine-year spread through a half-acre of blueberry plants.

Type a Query...







Thank you!

Colleen Schwab

Global Innovation Strategist
AWS

cmschwab@amazon.com

Sergio Ortega

AI/ML BD and Sales Lead
State and Local Governments

sergioai@amazon.com

[Linkedin.com/sergioortegaai](https://www.linkedin.com/company/sergioortegaai)

**Please complete the survey
for this session**



Executive track

Generative AI/ML and AI governance
for the public sector