



INFRASTRUCTURE & ARCHITECTURE SUMMIT

November 4 – 5 2025
London Hilton Olympia, UK

EUROPE'S LEADING EVENT FOR BUILDING, SCALING AND SECURING ENTERPRISE AI SYSTEMS



Renato Paco
Global Director Cloud
& Infrastructure
Sky



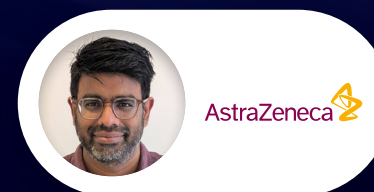
Bryan Carmichael
Director Digital Innovation
& Data Science
ThermoFisher Scientific



Lucas Vinh Tran
Director of AI/ML & Engineering
JPMorganChase



Jie Zheng
Technology Team Lead, Innovation
TUI



Vijay Vammi
Enterprise MLOps Services
& Technology Owner
Astrazeneca

#AIAS 

[Visit Website](#)

[Book for 2025](#)

[Partner with Us](#)

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

WELCOME TO THE #AIIA SUMMIT 2025

The AI Infrastructure & Architecture Summit returns – Europe's Leading Enterprise Event for AI Leaders Across Technology, Engineering, Data and Infrastructure

Dear AI & Data Leaders,

As we gather again for the AI Infrastructure & Architecture Summit, we meet at a critical inflection point. The acceleration of AI adoption has outpaced the maturity of enterprise infrastructure – exposing new challenges, new demands, and a clear need for resilient, scalable, and adaptable platforms and architecture.

This summit is designed to meet that moment. Built specifically for enterprise leaders working across AI, Data, Engineering, and Infrastructure, the program brings together the architects of modern AI systems within enterprise, those responsible for turning innovation into durable, production-grade value.

✓ The **evolving global landscape of AI infrastructure** and Europe's strategic position within it

✓ Core architectural decisions for **building and scaling Generative AI applications**

✓ What it takes to construct **infrastructure built for exponential scale**, spanning hardware, cloud, orchestration, and platform engineering

✓ **Engineering for reusability, portability & performance** to go beyond isolated pilots and enable scalable production-grade AI applications

✓ Emerging strategies to monitor AI models in production, and ensure **reliability through observability**

✓ The trade-offs between **cloud, on-prem, and hybrid environments** for enterprise AI deployment

The AI Infrastructure & Architecture Summit is where the strategic meets the technical – giving you the insights, peer connection, and clarity needed to scale AI with confidence.

Hope to see you there!

Warm regards,



Sam Lehmann

Event Director

AI Infrastructure & Architecture Summit



#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

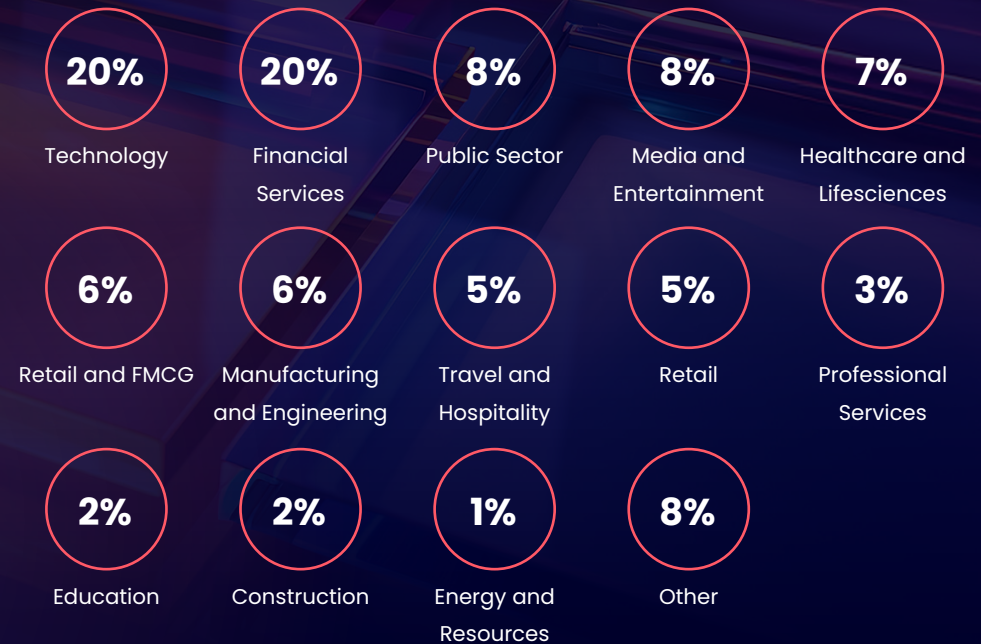
AUDIENCE OVERVIEW

Designed to focus on the technical side of AI projects, our last summit brought together 145+ enterprise leaders in AI, Data, Analytics, Infrastructure and Architecture.

FUNCTION



FUNCTION



SENIORITY



#AIIAS 

[Visit Website](#)

[Book for 2025](#)

[Partner with Us](#)

WELCOME























SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

EXPERT SPEAKERS

  Lucas Vinh Tran Director of AI/ML & Engineering JPMorganChase	  Renato Paco Global Director Cloud & Infrastructure Sky	  Bryan Carmichael Director Digital Innovation & Data Science ThermoFisher Scientific	  Jie Zheng Technology Team Lead @ TUI Innovation Hub – Sourcing TUI	  Vijay Vammi Enterprise MLOps Services & Technology Owner Astrazeneca
  Neil Thewlis Head of AI and Model Transformation HSBC	  Amit Nandi VP Solutions & Data Architect Barclays	  Deepak Paramanand Director of AI JPMorgan Chase & Co	  Ahmed Elahi Senior Director of Data & Technology SkyShowtime	  Matthew Sherriff AI Platform Manager Telegraph
  Ako Sabir Chief Data Officer	  Peter Rees Lead Architect – Enterprise Data & AI/ML AP Moller–Maersk	  Robin Mobasseri Head of AI and Analytics Services, Support and Enablement Wells Fargo	  Dhiraj Saxena Senior Vice President – Data JPMorganChase	  Ajwad Ali Head of AI Engineering AXA UK

“The summit was incredibly insightful, bringing together top minds to discuss scalable AI systems and real-world applications. The sessions offered practical solutions and fresh ideas, leaving me inspired and better equipped to tackle AI challenges. A big thank you to the organisers for an excellent event!”

Executive Director, OMD

#AIIAS 

Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

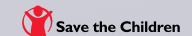
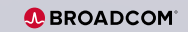
REGISTRATION

2024 BRANDS IN ATTENDANCE

"A fully enriching experience for
anyone shaping the future of AI."

Head of AI and Analytics Services and Enablement, Wells Fargo

WITH LEADERS FROM



#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY ONE

Tuesday 4th November

REGISTRATION AND PASTRIES

8.00

OPENING REMARKS FROM THE CHAIR

8.45

KEYNOTE PRESENTATION

Europe's AI Infrastructure Landscape – Global Trends and the Role Europe Will Play

9.00

As AI becomes core to business competitiveness, the infrastructure powering it has moved from a technical concern to a geopolitical and strategic priority. This keynote offers senior enterprise technology leaders a macro view of the evolving global AI infrastructure race – and Europe's position within it. From compute and networking to data governance and sovereign cloud initiatives, the session will explore how Europe's policy frameworks, investment strategies, and technical ecosystems are shaping the region's ability to compete in a compute-constrained, AI-driven world.

- The global landscape of AI infrastructure: investment flows, supply chain risks, and regional advantages
- What "AI sovereignty" means in practice – and why it matters to European enterprises
- How European initiatives (e.g. EuroHPC, Gaia-X, Chips Act) will impact infrastructure decisions
- Balancing global interoperability with regional regulation and data protection
- Strategic considerations for enterprises building long-term AI capabilities in Europe

Speaker To Be Announced

PANEL DISCUSSION

Architecting Scalable AI Ecosystems: Engineering for Reusability, Portability & Performance

9.30

How do leading enterprises build AI platforms that go beyond isolated pilots and enable scalable, production-grade AI applications? In this panel, senior platform and infrastructure leaders share practical insights into architecting AI platforms that deliver reusability, governance, and speed, at scale.

- Designing reusable, composable AI platforms for enterprise-wide adoption
- Striking the balance between centralized control and developer autonomy
- Tooling and orchestration strategies for managing the end-to-end model lifecycle
- Infrastructure choices for supporting hybrid and multi-cloud AI workloads
- Embedding governance, performance, and cost-efficiency into platform design

Vijay Vammi, Enterprise MLOps Services & Technology Owner, AstraZeneca

Matthew Sherriff, AI Platform Manager, Telegraph

Dhiraj Saxena, Senior Vice President – Data, JPMorganChase

Ajwad Ali, Head of AI Engineering, AXA UK



The Telegraph



PRESENTATION

Full-Stack AI Infrastructure Built To Scale Exponentially

10.10

Infrastructure teams must move beyond patchwork solutions and design integrated, full-stack platforms purpose built for scale. This session explores how leading enterprises are seeking to architect a comprehensive, total infrastructure ecosystem that aligns hardware, software, and operational layers – from GPUs and networking to orchestration tools and model lifecycle management. Building high-performance, scalable foundations that enable AI innovation without sacrificing reliability or control.

- Core components of a scalable AI infrastructure stack: compute, storage, and orchestration
- Integrating ML frameworks, pipelines, and developer tooling into infrastructure design
- Strategies to support flexibility, performance, and cost-efficiency at scale
- Avoiding bottlenecks and rework: aligning architecture to long-term AI roadmaps

Session Available For Sponsorship

MORNING COFFEE BREAK

10.40

PRESENTATION

Infrastructure Design for High-Performance Model Deployment

11.10

As AI workloads evolve from experimentation to enterprise-scale deployment, infrastructure leaders face new challenges in performance, cost, latency and accuracy. This session, led by Renato Paco, Global Director of Cloud & Infrastructure at Sky, explores the infrastructure, resource and financial considerations behind Sky's strategy to build a scalable, optimized environments for AI transformation at scale.

- Building hybrid environments that integrate on-prem and cloud-based GPU clusters to support elastic, high-throughput AI workloads.
- Capacity Planning for AI Workloads: Forecasting compute demand across model types, including considerations for batch size, concurrency, and inference latency.
- Balancing Cost, Latency & Accuracy: Frameworks for evaluating trade-offs and optimizing infrastructure for business outcomes.
- Applying non-functional testing, load simulation, and observability to validate infrastructure under real-world AI traffic patterns.
- Futureproofing Compute Tooling: Building adaptable, shared environments that support innovation, multi-tenancy, and evolving AI algorithms.

Renato Paco, Global Director of Cloud & Infrastructure, Sky



#AIAS in

Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY ONE

Tuesday 4th November

PRESENTATION

Building Trachyon: An AI Platform for Governance, Interoperability & Scalability

11.40

Wells Fargo's internal AI platform—named Trachyon—serves as the critical middleware between internal users and external LLMs to drive everything from an internal chatbot, personalization at scale and regulatory compliance. It strengthens governance, enables seamless model swapping, and ensures robust guardrails. In this session, Robin will explore:

- Governance-first orchestration – how Trachyon enforces data sanitisation and privacy
- Poly model & poly cloud architecture – the flexibility to plug in different LLMs and clouds without breaking workflows
- Scaling secure high volume use cases – powering Fargo's massive interaction volume using Trachyon's secure pipeline

Robin Mobasseri, Head: AI and Analytics Services, Support and Enablement, Wells Fargo



PRESENTATION

Scaling MLOps Foundations to Deliver Transformational Platform Impact

12.10

Vijay Vammi, Enterprise MLOps Services & Technology Owner at AstraZeneca join us to share his team's journey to build scalable, observable, and developer-friendly platforms that balance innovation with control. Moving the business towards MLOps to better align engineers with your AI platform infrastructure.

- Enabling product velocity while protecting innovation through scalable ML pipelines
- Embedding observability and performance monitoring into every stage of the model lifecycle
- Unifying workflows through a single internal developer platform to reduce silos and increase adoption
- Driving adoption of GenAI and platform standards through internal consulting, prompt design, and user-centric tooling
- What does success look like? Measuring impact for your AI and ML investments.

Vijay Vammi, Enterprise MLOps Services & Technology Owner, AstraZeneca



NETWORKING LUNCH

12.40

PRESENTATION

Reimagining Personalisation in Banking: From Classical Pipelines to Multi-Agent LLMs and Emerging Technologies

1.40

This session explores the evolution of personalisation and recommendation systems in the banking and finance sector – from traditional machine learning pipelines to cutting-edge multi-agent LLM architectures and emerging technologies. Drawing on real-world leadership experience, Lucas will share insights into the technical and strategic challenges of designing intelligent, compliant, and scalable personalisation systems in a complex regulatory environment.

Building end-to-end ML pipelines for personalisation in banking and finance: A look into the foundations of personalised systems – from data and feature engineering to model development, deployment, monitoring, and feedback loops – within a highly regulated, data-rich environment. Evolving to Generative AI: Architecting multi-agent LLM systems for personalisation: Designing agent-based LLM architectures to support personalisation at scale across financial services. Unlocking customer intelligence with GenAI: Using generative AI to power customer understanding and enable hyper-personalised user journeys.

Applying quantum-inspired algorithms for enhanced recommendations: Exploring how quantum-inspired techniques can optimise recommendation efficiency and handle complex challenges. Immersive personalisation: Extended Reality (XR) in next-generation customer experience: Examining the potential of augmented and virtual reality to reshape customer interaction models and deepen personal engagement in digital finance.

Lucas Vinh Tranh, Director of AI/ML & Engineering, JPMorganChase



PRESENTATION

Developing & Running Large-Scale Enterprise AI Applications For Business Success

2.20

To deliver real transformational change, enterprise leaders must understand how they can effectively scale to meet billions of constantly changing data items, thousands of queries per second, with minimal latencies. During this session we explore how you can use big data, RAG, vector search, machine learning and LLMs to deliver fast, precise decisions that drive real business success.

- Query, organize and make inferences in vectors, tensors, text and structured data at enterprise level
- Building data-driven application that deliver unbeatable performance and infinite automated scalability
- Efficiently retrieving and processing vast amounts of information with RAG architecture

Session Available For Sponsorship

#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY ONE

Tuesday 4th November

PRESENTATION

Repeatable By Design: Building AI Infrastructure for Regulated Environments

2.50

Highly regulated sectors – from life science to financial services – demand zero tolerance for failure when deploying AI. In these environments, “close enough” is never good enough. During this session, Bryan Carmichael, Director of Digital Innovation & Data Science explores how Thermo Fisher are building robust AI infrastructures that adhere to regulatory issues, ensure auditability, and guarantee repeatability.

- Designing AI systems where precision, not probability, is the standard
- Choosing tools and data pipelines that ensure regulatory compliance
- Integrating human-in-the-loop mechanisms for validation and control
- Collaborating effectively with regulators, auditors, and legal advisors

Bryan Carmichael, Director of Digital Innovation & Data Science,
Thermo Fisher Scientific



AFTERNOON COFFEE BREAK

3.20

PRESENTATION

Building AI-Ready Data Infrastructure: Quality, Integration, and Performance at Scale

3.50

Crap in, crap out – You hear it all the time, but how can enterprise leaders manage and prepare high-quality data to power AI models effectively, ensuring optimal results in dynamic and data-rich environments. Seamlessly incorporating the latest AI advancements into existing infrastructure, ensuring smooth integration with your data stack to maximise the performance of your AI applications.

- Scalability Strategies for AI Models: Explore the best practices for deploying AI models that can scale across large enterprise environments, ensuring performance and efficiency.
- Scalable AI Infrastructure Solutions: Explore vendor-provided solutions that deliver specialized infrastructure for scaling AI workloads, enhancing flexibility, and maintaining cost efficiency.
- Optimising AI for Enterprise Deployment: Gain insights into real-world AI deployment strategies that tackle infrastructure, model complexity, and scalability challenges, ensuring flexibility and long-term success.
- Ensuring Security and Compliance in AI Deployments: Discover how to secure AI solutions at scale, while meeting regulatory and compliance requirements in various industries.

Session Available For Sponsorship

PANEL DISCUSSION

The Real ROI of AI

4.20

AI investment is accelerating but proving a tangible near-term business value remains a challenge. This panel brings together enterprise AI leaders to unpack how to garner continuous internal buying, collaborate across non-technical functions, and position AI as a core business enabler – not just a technical experiment.

- Defining and measuring the real impact of AI investments: cost, productivity, revenue, and beyond
- Securing executive support and funding for AI products in competitive prioritisation environments
- Communicating value to non-technical stakeholders and navigating organisational AI illiteracy
- Embedding AI thinking across product, operations, and business strategy
- Evolving your role from technical lead to enterprise AI change agent

Jie Zheng, Technology Team Lead, TUI Innovation Hub – Sourcing, TUI



Deepak Paramanand, Director of AI, JPMorgan Chase & Co



Ahmed Elahi, Senior Director of Data & Technology, SkyShowtime



END OF CONFERENCE DAY ONE

5.00

“An expertly run event, bringing together a diverse group of industry leaders and experts. I was able to learn from every speaker and take actionable insights that I can deploy in my work and organisation.”

VP – Global Head of Software Asset Management, Marsh McLennan

#AIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY TWO

Wednesday 5th November

REGISTRATION AND PASTRIES

8.00

KEYNOTE PRESENTATION

How to Scale Your AI Tooling, Design & Architecture: People, Platform, Process

9.00

Led by Neil Thewlis, HSBC's Head of AI and Model Transformation, this session explores how enterprises can build AI architectures that scale both technically and organizationally. As AI projects move from proofs-of-concept to production, IT leaders must balance quick wins with long-term vision. Neil will share HSBC's journey to build long-term AI initiatives, focusing on key pillars of people, platform, and process to drive re-useability, portability and scale.

- People, Platform, Process: Leveraging the right talent and culture, robust AI models and infrastructure, and streamlined workflows/governance to support AI growth.
- Building Beyond Short-Term: Assessing how you can avoid falling into the trap of focussing on near-term investments, in favour of longer-term goals and objectives.
- Developing AI solutions with re-useability and portability front of mind and avoiding vendor lock-in utilizing open standards and cloud-agnostic tools will ensure your models and data can move freely as the landscape changes.
- Bridging the gap between technical implementation and business functionality by mapping AI use cases to the right architectural components so that technology choices directly serve business goals.
- Accommodate emerging requirements around scale, control, and agentic by designing with adaptability in mind

Neil Thewlis, Head of AI and Model Transformation, HSBC

PANEL DISCUSSION

How to Architect Your Generative AI Applications

9.30

Our first panel explores the core architectural decisions behind successful Gen AI products and systems, from model selection and customisation to platform engineering, orchestration, and infrastructure optimisation. Attendees will gain actionable insight into building and deploying GenAI applications that meet enterprise-grade requirements for performance, scalability, and reliability.

- Selecting and customising foundation models: fine-tuning, prompt engineering, and managing centralised model repositories
- Leveraging MLOps platforms to orchestrate model lifecycle management within application pipelines
- Infrastructure design for training and inference: optimising GPU/TPU workloads, storage, and compute pipelines
- Deployment and distribution strategies for GenAI models in production environments
- Setting and meeting system-level performance benchmarks to ensure end-to-end reliability and scalability

PRESENTATION

LLMOps: Deploying and Managing Large Language Models in Production

10.10

Operationalizing large language models (LLMs) brings a new set of infrastructure, monitoring, and lifecycle challenges that go far beyond traditional MLOps. As enterprises move LLMs into production environments, ensuring stability, scalability, and governance becomes critical. This session introduces the emerging discipline of LLMOps; the tooling, practices, and architectures needed to manage foundation models at scale, providing actionable guidance for technical leaders overseeing production deployments.

- How LLMOps differs from traditional MLOps and why new patterns are needed for foundation models
- Managing model versioning, updates, and rollback strategies for large, stateful LLMs
- Monitoring model performance, latency, and drift in real time – including prompt-level observability
- Balancing cost, reliability, and performance at scale

Session Available For Sponsorship

MORNING COFFEE BREAK

10.40

PRESENTATION

Driving Value With AI-Powered Software Engineering At TUI

11.10

AI-assisted coding and software development tools are reshaping how engineering teams build and ship products. But beyond productivity gains, how do you turn AI into a strategic advantage and reframe software engineering as a commercial growth engine, not a cost center?

- Practical strategies for using AI in code generation to deliver business impact
- Mitigating risk: balancing autonomy and control with clear guardrails
- Overcoming internal resistance: addressing the learning curve and shifting team culture
- Communicating the value of AI-augmented engineering to business stakeholders

Jie Zheng, Technology Team Lead, TUI Innovation Hub – Sourcing, TUI



#AIIAS in

Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY TWO

Wednesday 5th November

PRESENTATION

AI x Blockchain x Compute: Building Trustworthy, Scalable Intelligence

11.40

As AI models grow more powerful and compute-intensive, blockchain technologies are emerging as tools to help manage trust, transparency, and decentralised access to compute resources. This session explores how the convergence of AI, blockchain, and compute is set to reshape infrastructure strategies.

Using blockchain to verify AI model provenance, usage, and audit trails
Decentralised compute networks: unlocking distributed, low-cost GPU resources
Tokenized incentives for data sharing, model training, and resource contribution
Challenges in integrating blockchain with AI pipelines: latency, scale, and governance
Emerging architectures enabling secure, transparent, and decentralized AI execution

Deepak Paramanand, Director of AI, JPMorgan Chase & Co

JPMORGAN
CHASE & CO.

NETWORKING LUNCH

12.20

PANEL DISCUSSION

Enabling AI-Ready Platforms and Building Valuable AI Products

1.20

AI foundations lie in robust, secure, and intelligent data platforms that can fuel innovation while preserving trust and control. In this panel, senior data and AI leaders will explore what it takes to architect enterprise-grade data environments that support high-impact AI products. From governance to business outcomes, this session will unpack the intersection of data architecture and product strategy in the age of AI.

- The critical role of high-quality, actionable data in powering GenAI and AI applications
- Strategies to secure, manage, and maximize the value of enterprise data while preserving privacy and compliance
- Building data platforms that make information findable, accessible, trustworthy, interoperable, and reusable (FAIR)
- Operationalising data governance at scale through automation, policy enforcement, and platform design
- How AI is reshaping product development and business models – and what that means for data architecture going forward

Bryan Carmichael, Director of Digital Innovation & Data Science
Thermo Fisher Scientific

ThermoFisher
SCIENTIFIC

PRESENTATION

Connecting Your LLMs to the Outside World with the Model Context Protocol

1.50

There is a critical need to build AI systems that enable LLMs and AI agents to interact with external data, tools, and services in a secure, scalable, and reusable way. During this session, Peter Rees will share lessons learned on their journey to embed Model Context Protocol across AI systems and agents.

- Architecting a scalable enterprise service, built for reusability and portability across workflows and domains
- Secure-by-design approaches for managing model input/output, data access boundaries, and traceability
- Sharing lessons learned from A.P Moller's Maersk journey to mature AI agents through MCPs

Peter Rees, Lead Architect – Enterprise Data & AI/ML, AP Moller



PRESENTATION

Operationalising AI: Infrastructure to Agents – Unlocking Business Value At Scale

2.20

As enterprises push beyond traditional Machine Learning towards autonomous AI agents, today's MLOps platforms are being stretched to their limits. New demands – like managing prompt versioning, orchestrating agent behaviour, and ensuring governance across complex AI systems – signal the need for a new operational paradigm. In this session, Amit will share a practical vision for the next generation of AI Infrastructure, evolving enterprise architecture for the Agentic AI era. Drawing from his work building systems where humans interact with autonomous agents through LLM interfaces – and where domain-specific LQMs and specialized SLMs are coordinated by agents. Rethinking the five foundational layers of AI (Infrastructure, Data, Code, Models, and UI) to introduce AgentOps as a new discipline focused on orchestration, memory, and real-time feedback loops, offering a blueprint for scalable, human-aligned AI systems.

- Aligning technical foundations for scalable, agent-driven AI adoption
- Understanding where MLOps ends and LLOps or AgentOps begins – and why it matters for your organisation
- What's needed to support multi-agent systems with real-time collaboration and adaptation
- How can you future-proof your AI stack for human-in-the-loop and autonomous use cases

Amit Nandi, VP Solutions & Data Architect, Barclays



#AIAS in

Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE DAY TWO

Wednesday 5th November

AFTERNOON COFFEE BREAK

2.50

PANEL DISCUSSION

Cloud, On-Prem, or Hybrid? Infrastructure Strategies for Enterprise AI at Scale

3.20

This panel brings together senior tech leaders to discuss frameworks for evaluating and architecting the right environment for your AI initiatives, covering:

- Trade-offs between cloud, on-premises, and hybrid deployments for training and inference workloads
- When to invest in dedicated GPU infrastructure vs. leveraging cloud-based AI services
- Real-world examples of enterprises shifting from cloud to on-prem to control GenAI costs
- Architectures and tooling for managing hybrid environments seamlessly (compute, data, orchestration)
- Financial modelling and cost control: balancing CapEx and OpEx while ensuring performance

KEYNOTE PRESENTATION

Next-Generation AI Hardware Accelerators: Powering Performance at Scale

3.50

AI innovation is increasingly bottlenecked by compute. This session provides a deep dive into the latest AI accelerators reshaping enterprise infrastructure – Attendees will gain a practical understanding of:

- The performance characteristics and trade-offs between GPUs, TPUs, FPGAs, and custom silicon
- Aligning accelerator selection with specific AI workloads (training vs. inference, LLMs vs. vision models, etc.)
- Memory, bandwidth, and interconnect considerations for high-throughput AI pipelines
- Infrastructure implications – power, cooling, and rack design for high-density accelerator deployment
- Future trends in AI compute and what enterprise architects need to plan for now

END OF CONFERENCE DAY TWO

4.20



“The excellent breadth of topics and styles helped keep my interest, variety is key to sustaining interest especially in a long and technically oriented event. The more normal, everyday business language that can be used to articulate AI the better, and some of the speakers were really able to do this.”

Global Head of Data Trust & Integrity, Entain PLC

#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

POSITION YOURSELF AS A KEY SOLUTION PROVIDER IN THIS EMERGING FIELD

Partner with AIIA Summit

AI Infrastructure & Architecture Summit is the leading conference for solution providers looking to meet prospects, grow their brand, and position themselves as thought leaders in this transformative field.

Find out how we can help you win business, generate pipeline and build your brand.

- ✓ Network and Interact with AI & data decision makers who are actively looking to invest in Generative AI solutions
- ✓ Grab the opportunity to demonstrate thought leadership in this emerging field
- ✓ Ensure you are front of mind for those looking to invest in the latest AI tooling through onsite branding



What is included?



Presentations / panel discussions

Demonstrate your thought leadership by directly addressing our senior level audience.



Your own dedicated booth

Build your brand, network with attendees and showcase your capabilities by having your own dedicated space on site.



Coffee breaks / lunch / drink receptions

Want to leave an impression with your prospects? Leverage one of our networking packages to generate networking opportunities with your key personas.

“It was well organised. The lineup of speakers and topics were well thought through. Meeting like-minded people from various industries was valuable.”

VP Architecture, Barclays

#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

PREVIOUS EVENT SPONSORS



#AIIAS 

[Visit Website](#)

[Book for 2025](#)

[Partner with Us](#)

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

CONFERENCE PACKAGES & REGISTRATION

Please choose your
Conference Package:

Event Pass

Access to main conference days
Day 1 Evening Networking Drinks Reception
Membership to AI, Data & Analytics Network
Online Community

For purchase by
Solution Providers,
Service Organisations,
and Consultancies:

Solution Provider Pass

2 Conference Days Only

Early Bird Discount:

Ends 29th August

SAVE £600

£799 +20% VAT

£1,899 +20% VAT

Standard Rate

£1,399 +20% VAT

£2,499 +20% VAT

Group Discounts



Booking for
2 delegates?
5% off



Booking for
3-4 delegates?
10% off



Booking for
5-7 delegates?
15% off



Booking for
8+ delegates?
20% off



Email
enquire@iqpc.co.uk

#AIIAS



Visit Website

Book for 2025

Partner with Us

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION

ABOUT THE VENUE

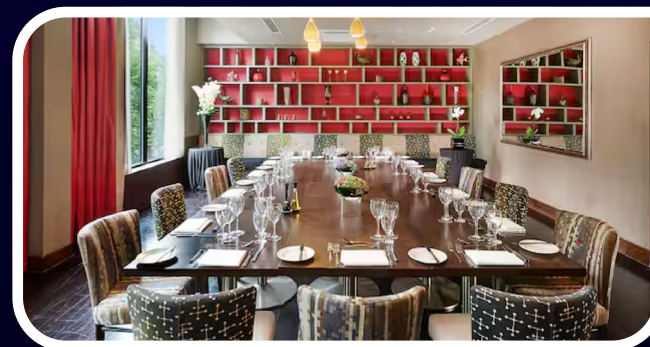
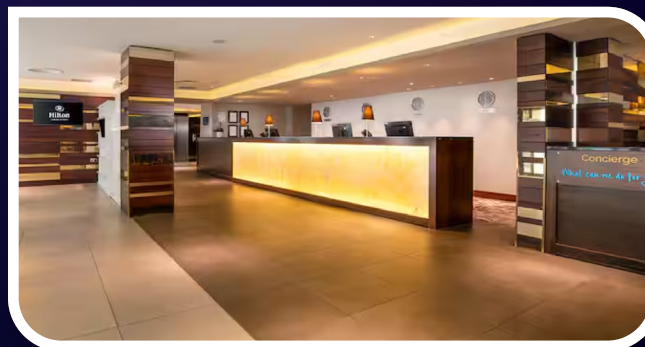
Our venue is steps from the shops of Kensington High Street, passing the gardens of Holland Park along the way. We're a three-minute walk from Kensington (Olympia) Station, and a short trip to visit Royal Albert Hall, Hyde Park, and other city center sights.

Hilton London Olympia
380 Kensington High St, London, UK
W14 8NL
Tel: +44 20 7603 3333

www.hilton.com

[Book here](#)

Please note: IQPC will never recommend, approve or appoint any third-party rooming service to act on our behalf. Please be extremely wary if you are approached by any such companies. We will always endeavor to negotiate the best available rates for you so please use the Hotel's website link provided.



#AIIAS 

[Visit Website](#)

[Book for 2025](#)

[Partner with Us](#)

WELCOME

SPEAKERS

AGENDA

WHY SPONSOR?

REGISTRATION