HPC + AI WALL STREET

QC / HPC / BD / AI / CXL

2023

SHOW GUIDE

21 years of delivering innovative thought leadership and providing world-class technology from industry pioneers for FinTech and Capital Markets.

Showcasing the fastest, smartest systems on the planet for FinServ.

TUESDAY & WEDNESDAY, SEPTEMBER 26 + 27, 2023

8:30AM-6:00PM ET
INTERCONTINENTAL, TIMES SQUARE, NYC

WELCOME

Welcome to the 2023 HPC + AI Wall Street conference and exhibition.

This event may very well be the only conference you'll attend this year where world-changing technologies like Quantum Computing will be unpacked and the convergence of HPC, Big Data, and AI will be explored, specifically for Financial Services.

Financial Services workloads have historically been compute-insatiable, and the industry has often engaged in the early adoption of bleeding edge systems. The need to handle data collection, analysis, and application-to-market at ever-faster speeds has produced transformative shifts across the entire ecosystem.

Organizations investing in HPC drive market innovation, speed-to-insight, and competitive advantage. And 2023 marks an exceptional year in AI as we enter the realm of Artificial General Intelligence and develop capabilities considered only a dream just a few years ago.

Our conference kicks off with a panel of analysts who will provide insights into the level of investments being made in Quantum Computing, the development of its vendor landscape, and where these systems are currently being deployed. Those who do not plan for Quantum now will face a chasm of competitive disadvantage.

After lunch, we'll hear from a cast of FSI tech luminaries unraveling the secrets of FinTech in the cloud. PaaS puts the task of optimizing the hardware on the shoulders of the cloud provider rather than the FSI technologist. This subtle distinction emphasizes optimizing the technology to the workloads vs. optimizing the workloads to the technology.

The promise of these compute technologies is the entry point to advanced analytics and AGI. None of this is possible without high-performance compute systems, coupled with complex big data management capabilities that align with customized AI algorithms and platforms.

Day Two will focus on these mission-critical challenges with sessions in Data Management and AI. Data has always been the North Star, and this is even more evident as we enter the era of AI / LLMs. FSI is laden with regulations and protocols governing data fabric, and selecting the tools — particularly the AI platform — screams of compliance challenges. Finally, it's imperative and responsible to end the program with a thought-provoking session on AI ethics.

Advanced scale systems for Financial Services are in a constant state of change and development. HPC + AI Wall Street is building a specialized community of users and vendors driven to utilize the fastest and smartest systems on the planet to meet the challenges and demands of this change.

We encourage you to engage with our luminary speakers, your peers, and our partner vendors at their exhibit tables as we explore the advancements in QC, HPC, Big Data, and AI together.

— Tom Tabor, *CEO & Founder, Tabor Communications Inc.*

AGENDA

Tracks: Quantum Computing | HPC | Data Management | Artificial Intelligence

TUESDAY, SEPTEMBER 26, 2023

8:30-8:45	Welcome + Opening Remarks
QUANTUM COMPUTING SESSIONS	
8:45-9:00am	Quantum Computing Overview: The State of QC in 2023
9:00-9:45am	Quantum Computing Analyst Panel: One Year Later
9:45-10:30am	**QUANTUM KEYNOTE: The Time to Start Your Journey to Hybrid (Quantum+HPC) Computing is Now**
10:30-11:00am	Break + Exhibits
11:00am-12:00pm	Ready and Willing: Quantum Computing Today and Tomorrow
HIGH PERFORMANCE COMPUTING SESSIONS	
12:00-12:15pm	High Performance Computing Overview: The State of HPC in 2023
12:15-1:00pm	**Lunch + HPC KEYNOTE: Optimizing HPC to Your FSI Workloads with AI**
1:00-2:00pm	HPC + AI Resources in the Great GPU Squeeze
2:00-3:00pm	Break + Exhibits + 4 Lightning Presentations
3:00-4:00pm	Legendary HPC Pioneers From Citigroup, Goldman Sachs, JPMorgan Chase, and Morgan Stanley Share Past Experiences and Provide Future Insights for HPC, AI, Big Data, and Cloud
4:00-5:00pm	Beating the Markets with HPC + AI Convergence: Generating Alpha and Excess Returns with Alternative Data
5:00-6:00pm	Vendor Showdown with Intersect360 Research
6:00-7:30pm	Cocktail Reception

WEDNESDAY, SEPTEMBER 27, 2023

DATA MANAGEMENT SESSIONS	
8:45-9:00am	Data Management Overview: The State of Big Data in 2023
9:00-9:45am	The Data Management Challenges with Generative AI
9:45-10:30am	**DATA MANAGEMENT KEYNOTE: Unify and Automate Data Across Any Storage, Any Data Center, Any Cloud, Anywhere**
10:30-11:00am	Break + Exhibits
11:00am-12:00pm	AI/ML in the Cloud: Data Management, Ops Stack, and Ecosystem Tooling
ARTIFICIAL INTELLIGENCE SESSIONS	
12:00-12:15pm	**Artificial Intelligence Overview: The State of AI in 2023**
12:15-1:00pm	**Lunch + AI KEYNOTE: What Is An AI Factory?**
1:00-2:00pm	Monitoring ML Models in a Financially Regulated Environment
2:00-3:00pm	Break + Exhibits + 4 Lightning Presentations
3:00-4:00pm	From Open Source to Third Party: Capturing the Early Potential of Generative AI for FinServ
4:00-5:00pm	**CLOSING KEYNOTE: Responsible AI: Building Value While Managing Risks and Regulation**
5:00-6:00pm	Analyst Crossfire with Intersect360 Research
6:00-6:15pm	Closing Remarks + Raffle Drawings

SESSIONS

TUESDAY, SEPTEMBER 26

QUANTUM COMPUTING

QUANTUM COMPUTING OVERVIEW
The State of QC in 2023 (8:45am)

PANEL DISCUSSION: Quantum Computing Analyst Panel: One Year Later (9:00am)

Moderator: John Russell, *QCwire, HPCwire*
Bob Sorensen, *Hyperion Research*
Heather West, PhD, *IDC*
Jay Boisseau, *Vizias (formerly Dell)*

As quantum computing moves from experiment to early deployment, what's needed to build a QC-readiness kit? Panelists will discuss how to separate hype from reality; identify the top questions you should ask potential QC suppliers; explore whether (or not) there are penalties for taking a wait-and-see attitude; and look at current and forecast spending plans in the quantum market. You'll walk away with a better idea of what constitutes a thoughtful QC on-ramp in financial services.

QUANTUM KEYNOTE: The Time to Start Your Journey to Hybrid (Quantum+HPC) Computing is Now (9:45am)

Philip Farah, *IonQ*
Bob Fletcher, *IonQ*

GenAI and LLMs have caught many Financial Services Institutions off guard. In this session we will share the necessary steps required to "set the Quantum foundations" and reveal why, given the speed of evolution of our Quantum hardware, we believe the time to start is now. We will cover specific use cases relevant to the Financial Services industry and explain why all Quantum will be hybrid (a combination of Classical and Quantum) for the foreseeable future. Lastly, we'll review the implications of preparing for the upcoming Quantum era.

Ready and Willing: Quantum Computing Today and Tomorrow (11:00am)

Travis Humble, *Oak Ridge National Laboratory*

Along with the dramatic shift in compute capabilities that quantum computing promises, comes the concern of being ready and willing to embrace this new technology. What will quantum computing actually change? And how close are we to seeing those changes occur? This talk will summarize the field, including its current successes and failures today, as well as measure the pace of progress toward new computing capabilities for tomorrow. We will also highlight the national ecosystem for research and development, including activities at the Department of

Energy's Quantum Science Center that are expanding opportunities to test and evaluate quantum technologies and transition them to production applications.

HIGH PERFORMANCE COMPUTING

HIGH PERFORMANCE COMPUTING OVERVIEW

The State of HPC in 2023 (12:00pm)

HPC KEYNOTE: Optimizing HPC to Your FSI Workloads with AI (12:15pm)

Ryan Quick, *Providentia Worldwide*
Craig Yamasaki, *Supermicro*

PANEL DISCUSSION: HPC + AI Resources in the Great GPU Squeeze (1:00pm)

Moderator: Doug Eadline, PhD, *HPCwire*

Prabhu Ramamoorthy, *NVIDIA*
Wyatt Gorman, *Google Cloud*
Thomas Jorgensen, *Supermicro*

The rapid and continuing rise of Generative AI (ChatGPT, Bard, and others) has pushed the demand for GPUs to unprecedented levels and created a "squeeze" on GPU availability. The prospect of extremely long lead times and prohibitive cost for GPUs may invite HPC users to seek more performance from the traditional (and available) CPUs and utilize GPUs available in the cloud. This panel will bring together industry leaders to discuss how HPC systems and the cloud can continue to provide users with high performance in the coming age of GPU scarcity.

Legendary HPC Pioneers From Citigroup, Goldman Sachs, JPMorgan Chase, and Morgan Stanley Share Past Experiences and Provide Future Insights for HPC, AI, Big Data, and Cloud (3:00pm)

Moderator: Roman Chwyl, *IBM*

Ty Panagoplos, formerly at *Santander US*, *TD Securities*, and *JPMorgan Chase Grid*
J Ram, formerly *at Goldman Sachs Grid* and *Morgan Stanley Grid*
Dino Vitale, *TD Bank Group (formerly Citigroup Grid and Morgan Stanley Grid)*

In this unique fireside chat, Wall Street's HPC legends, representing Citigroup, Goldman Sachs, JPMorgan Chase, and Morgan Stanley, come together to share their collective wisdom and experiences in large-scale Grid and HPC management. Delving into the past and present, they provide invaluable insights for navigating the evolving landscape of HPC, AI, Big Data, and Cloud technologies and the operational models required to support them at scale. Join us as they illuminate the path forward in a future where massive computational power, data, and the convergence of AI and Quantum computing promise to reshape the industry.

Beating the Markets with HPC + AI Convergence: Generating Alpha and Excess Returns with Alternative Data (4:00pm)

Prabhu Ramamoorthy, *NVIDIA*

The Holy Grail of Investing is to stay ahead of markets with real-time calculations as well as tapping into org sources, such as alternative data. With quantitative HPC + AI convergence, learn how financial firms are accelerating price discovery/risk for real-time exposures, as well as alternative data, with Gen AI NLP/LLM models from unstructured data to serve internal needs as well as those of their customers.

Vendor Showdown with Intersect360 Research (5:00pm)

Moderator: Dan Olds, *Chief Research Officer, Intersect360 Research*

Molly Presley, *SVP, Marketing, Hammerspace*

Doug Norton, *Chief Marketing Officer, InspireSemi*

Rob Glanzman, *Global Strategic Alliances Principal Architect, Financial Services, Pure Storage*

Nick Ihli, *Director of Solutions Engineering and Cloud, SchedMD*

WEDNESDAY, SEPTEMBER 27

DATA MANAGEMENT

DATA MANAGEMENT OVERVIEW
The State of Big Data in 2023 (8:45am)

The Data Management Challenges with Generative AI (9:00am)

Alex Woodie, *Datanami*

Generative AI has taken the world by storm in 2023, and for good reason: the technology has the potential to be truly transformative for an enterprise. But no matter how advanced GenAI is, it suffers from the same data management challenges as any other form of advanced analytic or machine learning. In this session, we'll take a look at the biggest data management challenges presenting themselves in 2023, as well as some of the technology, tools, and techniques that leading-edge companies are using to get beyond those challenges and realize the true potential of transformative GenAI.

DATA MANAGEMENT KEYNOTE: Unify and Automate Data Across Any Storage, Any Data Center, Any Cloud, Anywhere (9:45am)

David Flynn, *Hammerspace*

High-performance computing applications, web-scale storage systems, and modern enterprises increasingly have the need for a data architecture that will unify data at the edge, as well as in data centers and clouds. These organizations with massive-scale data requirements need automated data orchestration coupled with the extreme high performance of a parallel file system and a standards-based solution that will be easy to deploy on machines with diverse security and build environments. In this session you will learn:

- How to unify data created in different clusters and locations into a single namespace, and place locally to applications and compute for processing and AI
- The latest technologies available to deliver parallel file system performance from data sets stored in a hybrid cloud environment
- The latest in standards-based technologies available for data orchestration and storage at mass scale
- How to secure data at a global level to ensure protection, governance, and access rights are maintained no matter where data is being used or stored

AI/ML in the Cloud: Data Management, Ops Stack, and Ecosystem Tooling (11:00am)

Dino Vitale, *TD Bank Group*

One of the highest priorities for many financial service institutions centers around digital modernization and a strategy/approach to enable machine learning and AI at enterprise level. This talk will help identify key challenges and considerations for leveraging the cloud when running AI/ML workloads, including AI vs. ML ecosystems, data management, ML/AI ops, and a discussion of the challenges and considerations of adopting generative AI Large Language Models (LLMs).

ARTIFICIAL INTELLIGENCE

ARTIFICIAL INTELLIGENCE OVERVIEW
The State of AI in 2023 (12:00pm)

AI KEYNOTE: What Is An AI Factory? (12:15pm)

Troy Kaster, *Penguin Solutions*

There are many new challenges in delivering AI Factories at scale. The demand for large-scale GPU clusters is building, but a knowledge gap exists between what large enterprises are familiar with and what this new technology requires in order to design, build, deploy, and manage. The end result is usually poor performance and availability for a massive financial investment. This talk will discuss how to avoid some of those mistakes and how to get the maximum return on your AI Factory investment.

Monitoring ML Models in a Financially Regulated Environment (1:00pm)

Sagar Gaikwad, *Bucket Theorem (formerly Capital One)*

In the fintech domain, the integration of advanced Machine Learning (ML) methods encounters specific challenges due to regulatory requirements. It's vital for ML implementations in this sector to be both transparent and ethically compliant. This presentation will cover:

- **Explainability of ML Models:** Why simply having a powerful model isn't enough, and the vital role of transparency and explainability in today's finance landscape.
- **MLOps — Beyond the Buzzword**: A deep dive into the methodology that's revolutionizing model lifecycle management, making sense of the whirlwind from training to real-time monitoring.
- **Federated Learning:** Delve into the world of decentralized ML, a beacon for data privacy, and its growing relevance in a world clamoring for data rights.
- **Case Studies:** Engaging stories from the frontlines detailing the wins, pitfalls, and unexpected turns of AI/ML in regulatory scenarios.
- **Through the Regulator's Lens:** Gaining insights into the regulatory psyche, understanding their reservations, and aspirations for big data, AI, and ML in shaping the future of market regulation.
- **The Art & Science of Model Operationalization:** Navigating the journey of models from the lab to the limelight, focusing on deployment nuances, version tales, and the balancing act between development and production stages.

Join us to gain comprehensive insights into the best practices for ML model management in the regulated FinTech landscape.

PANEL DISCUSSION: From Open Source to Third Party: Capturing the Early Potential of Generative AI for FinServ (3:00pm)

Moderator: Sagar Gaikwad, *Bucket Theorem (formerly Capital One)*

Ricardo Portilla, *Databricks*
Christopher Parisien, *NVIDIA Guardrail*
Frank Liu, *Zilliz*

Join industry leaders from Zilliz, NVIDIA Guardrail, and Databricks MosaicML as we delve deep into the evolving world of Large Language Models (LLMs) and Generative AI. While LLMs like ChatGPT have made headlines and are a hot topic of discussion, they're just the tip of the iceberg. We'll discuss:

- What challenges are currently limiting the broader adoption of LLMs in enterprises?
- Are there inherent risks or pitfalls in over-relying on AI-driven decision-making?
- What infrastructural, data, and security considerations should technology leadership prioritize to ensure a mature AI strategy?
- With AI continually advancing, how can organizations strike a balance between leveraging open-source solutions and third-party managed services?

Whether you're an AI enthusiast or a FinServ professional grappling with AI's potential in your domain, this discussion promises to provide valuable insights into the state and future of AI. Join us in unraveling the complexities and charting the course forward.

CLOSING KEYNOTE: Responsible AI: Building Value While Managing Risks and Regulation (4:00pm)

Traci Gusher, *EY (Ernst & Young)*

This discussion will focus on how the emergence of Generative AI has created new areas of focus for organizations to develop and utilize AI responsibly. Topics will include traditional risk in AI, new risks unique to Generative AI, emerging legislation, and how organizations are beginning to address these new challenges.

Analyst Crossfire with Intersect360 (5:00pm)

Moderator: Dan Olds, Chief Research Officer, *Intersect360 Research*

Matt Jacobs, *Chief Commercial Officer, Cornelis Networks*

Alan Benjamin, *President & CEO,GigaIO*

Philip Farah, *VP Sales, Industries and Strategic Relationships, IonQ*

Troy Kaster, *Vice President of Generative AI, Penguin Solutions*

ENDLESS MEMORY
CXL™ FORUM at HPC + AI Wall Street

The CXL Forum
at HPC + AI Wall Street

Wednesday, 9:30 – 2:00
Central Park Room

GIGAIO LIGHTELLIGENCE LIQID

Massachusetts Institute of Technology MemVerge

SUPERMICRO SAMSUNG XCONN TECH

SPECIAL EVENTS

TUESDAY, SEPTEMBER 26

VENDOR SHOWDOWN

with Intersect360 Research (5:00pm)

Moderator: Dan Olds, *Chief Research Officer, Intersect360 Research*

Molly Presley, *SVP, Marketing, Hammerspace*

Doug Norton, *Chief Marketing Officer, InspireSemi*

Rob Glanzman, *Global Strategic Alliances Principal Architect, Financial Services, Pure Storage*

Nick Ihli, *Director of Solutions Engineering and Cloud, SchedMD*

Can a vendor's value proposition be presented with only three slides in five minutes? Can they stand up to tough questioning from a seasoned industry analyst? Find out in this session as Dan Olds from Intersect360 Research moderates a select group of vendors giving short presentations and being made to answer three tough questions — one of which will be a complete surprise. This is a fast-moving and fun session which will get past the marketing-speak and uncover the real story behind the story. At its completion, the audience will vote on which technology will have the biggest impact on Financial Services over the coming year.

WEDNESDAY, SEPTEMBER 27

ANALYST CROSSFIRE

with Intersect360 Research (5:00pm)

Moderator: Dan Olds, *Chief Research Officer, Intersect360 Research*

Matt Jacobs, *Chief Commercial Officer, Cornelis Networks*

Alan Benjamin, *President & CEO,GigaIO*

Philip Farah, *VP Sales, Industries and Strategic Relationships, IonQ*

Troy Kaster, *Vice President of Generative AI, Penguin Solutions*

The Analyst Crossfire is a compelling and quick-moving panel discussion where some of the best and brightest minds in the industry discuss some of the most interesting and provocative topics in the industry today. Broken into four 15-minute segments, Dan Olds from Intersect360 Research will moderate and pose provocative questions covering the state of Quantum computing, the links between HPC, AI, and Wall Street, and how customers can tackle the enormous data management challenges posed by the advent of AI. The result will be an unscripted, spontaneous conversation that will dive into some of the most interesting topics today. You'll hear some compelling arguments, and maybe even have a little fun along the way.

TUESDAY LIGHTNING TALKS

2:00: Boosting Your Intelligent Transformation with Energy Optimized Computing
Marc Fisher, *Senior Solution Specialist – HPC and AI, Lenovo*

2:15: Parallel Works: The Future of Secure, Hybrid, Multi-Cloud HPC
Matthew Shaxted, *President & Co-Founder, Parallel Works*

2:30: CXL and The Road to Endless Memory
Bernie Wu, *Vice President, Strategic Partnerships and Business Development, MemVerge*

2:45: Accelerating Everything AI
Thomas Jorgensen, *Senior Director, Technology Enablement, Supermicro*

WEDNESDAY LIGHTNING TALKS

2:00: WEKA Data Platform for AI, Machine Learning, and Deep Learning
Brian Mitchell, *Senior Solution Architect, WEKA*

2:15: Ease Into AI With the World's First 32 GPU Single-node Supercomputer
Alan Benjamin, *President & CEO, GigaIO*

2:30: AI Co-pilots: Empowering AI with Context for Better Decision-making
Zayd Sukhun, *Lead Solution Engineer, Banking, Quantexa*

2:45: Accelerating HPC and AI with Open Source Tools on Google Cloud
Wyatt Gorman, *Solutions Manager, HPC & AI Infrastructure, Google Cloud*

Lightning Talks will take place daily in the exhibits area.

Breakfast: A light breakfast will be served in the exhibits area from 8:15-10:15 each morning.

Lunch: Lunch will be served both days at 12:00pm in the main ballroom.

Coffee Break: A coffee break with snacks will be served in the exhibits area from 2:30-4:30 each afternoon.

Cocktail Party: All are welcome to attend a cocktail party sponsored by Dell and AMD in the exhibits area from 6:00-7:30 Tuesday evening.

ADVISORY BOARD

Kristin Boggiano

Co-Founder and President

CrossTower

Jay Boisseau, Ph.D

CEO and Co-Founder

Vizias

Former HPC & AI Technology Strategist, Dell

Roman Chwyl

Managing Director
Grid as a Service

IBM

Sagar Gaikwad

Managing Director

Bucket Theorem

Former Head of Machine Learning Foundations, Capital One

Ritesh Jain

Founder

Infynit

Former COO, Global Head of Digital Technology Delivery, HSBC

Ryan Quick

Principal

Providentia Worldwide

Harvey Stein

Senior VP, Research Analytics and Methodology

Two Sigma Labs

Ankit Manoj Vasa

Data Specialist, Privacy Lead GDPR – Cyber Threat Management

Nasdaq

Dino Vitale

Distinguished Engineer, Infrastructure Technology Solutions Cloud Engineering Team

TD Bank Group

Steve Yatko

CEO

Oktay Technology

SPEAKERS

Alan Benjamin

President and CEO, GigaIO

Alan Benjamin is one of the visionaries behind GigaIO's innovative solution. He was most recently COO of Pulse Electronics ($800M communication components and subsystem supplier) and previously, CEO of Excelsus Technologies. Earlier, Alan helped lead PDP, a true start-up, to a successful acquisition by a strategic buyer for $80M in year three. He started his career at Hewlett-Packard in Sales Support and quickly moved into management positions in Product Marketing and R&D. Alan graduated from Duke University with a BSEE in Electrical Engineering and attended Harvard Business School AMP program, as well as UCSD LAMP program.

Jay Boisseau

CEO & Co-founder, Vizias, formerly HPC & AI Technology Strategist, Dell Technologies

Jay Boisseau is an experienced, recognized leader and strategist in advanced computing technologies, with over 25 years in the field. Jay is CEO and co-founder of Vizias, a small consultancy with expertise in HPC, AI, technology community building, and technology event support. Jay also serves as the executive director and founder of The Austin Forum on Technology & Society. Recent major work included serving as the HPC & AI Technology Strategist for Dell Technologies and leading and greatly expanding the global Dell HPC Community. Formerly, Jay created and led the Texas Advanced Computing Center (TACC) at The University of Texas at Austin and worked at the San Diego Supercomputer Center and the Arctic Region Supercomputing Center. He received his doctorate in astronomy from UT Austin, and his undergraduate degree in astronomy and physics from the University of Virginia.

Roman Chwyl

Managing Director Grid as a Service, IBM

Roman Chwyl is a seasoned Wall Street technology executive with over 25 years of experience, specializing in the fusion of High Performance Computing/Grid workloads and Public Cloud solutions. His proficiency encompasses onboard and operate models crucial for supporting large-scale computing environments and their requisite business frameworks. Throughout his career, Roman has held pivotal roles at major public cloud providers. He led the Azure FinTech division at Microsoft, contributed to the incubation of the AWS FinTech Market at Amazon, and established and managed the Financial Services vertical at Google Cloud. Currently, Roman directs strategy for large-scale computing and data-intensive workloads at IBM within the Wall Street sector. Roman is also the co-founder of Concourse Labs, a venture-backed SaaS company specializing in automated cloud governance solutions. These solutions facilitate the secure and efficient adoption of Cloud Computing for the Financial Services vertical. Roman's decade-long tenure at Platform Computing solidified his reputation as a business architect of Grid Computing for Wall Street, particularly with Symphony, recognized as the industry standard Grid Software solution. He effectively managed IBM's 2012 acquisition and its integration into the Financial Services Vertical, serving top-tier financial institutions. Roman holds a B.Sc. in Physics and Mathematics from the University of Toronto and currently resides in New York.

Doug Eadline, PhD

Managing Editor, HPCwire

Douglas Eadline, PhD, began his career as Analytical Chemist with an interest in high performance computer methods. Starting with the first Beowulf "How-To" document, Doug has written hundreds of articles, white papers, and instructional documents covering many aspects of Linux HPC, Hadoop, and Data Analytics computing.

Prior to launching and editing the popular ClusterMonkey.net website in 2005, he served as editor-in-chief for *ClusterWorld Magazine* and was senior HPC Editor for *Linux Magazine*. Doug is author of several books and instructional videos covering the practical aspects of HPC and scalable tools for data science.

Philip Farah

VP Sales, Industries and Strategic Relationships, IonQ

As VP Sales, Industries and Strategic Partnerships at IonQ, Philip Farah is responsible for developing IonQ's strategic partnerships with a focus on revenue growth and on strengthening IonQ's position at the center of the Quantum Computing Ecosystem. Prior to joining IonQ, Philip ran Services and Engineering Sales, Global Accounts for Worldwide Technology, a leading systems integrator and IT infrastructure value-added reseller. Philip started his career in the US as a consultant with McKinsey & Co, and has held several leadership roles at Gartner (Senior Managing Partner, Financial Services), Cisco Systems (head of Global FSI Innovation strategy practice), and Capital One Financial (head of Recoveries Analytics).

Bob Fletcher

Director of Enterprise Sales, IonQ

Bob Fletcher brings 20 years of new product sales and strategy experience from the Internet Backbone, Data Center, HPC, and AI markets to drive IonQ's enterprise market development. Over the last eighteen months, he's worked with some of the largest early adopters of quantum computing to drive machine learning, chemistry, and object detection value.

David Flynn

Co-founder and CEO, Hammerspace

David Flynn is a recognized leader in IT innovation who has been architecting disruptive computing platforms since his early work in supercomputing and Linux systems. David pioneered the use of flash for enterprise application acceleration as founder and former CEO of Fusion-io, which was acquired by SanDisk in 2014. Previously, David served as Chief Architect at Linux Networx, where he was instrumental in the creation of the OpenFabrics stack and designed several of the world's largest supercomputers leveraging Linux clustering, InfiniBand, and RDMA-based technologies. David has also served as Project BlackDog's Chief Scientist and Vice President of Engineering and held positions at Network Computer Inc. and Liberate Technologies, a spin-off of Oracle Corporation that developed thin-client solutions. David holds more than 100 patents in areas across web browser technologies, mobile device management, network switching, and protocols to distributed storage systems.

Sagar Gaikwad

Managing Director, Bucket Theorem, former Head of Machine Learning Foundations, Capital One

A passionate leader with a track record of building high performing engineering teams, Sagar Gaikwad demonstrates daily that hard work can be fun. At Capital One, Sagar led teams to build a real-time cybersecurity platform processing petabytes of data, several big data initiatives, infrastructure-as-a-service to modernize data science for the cloud, and an AI-based recruitment sourcing tool. Sagar is currently scaling the cloud engineering team to handle the rapid adoption and growth of the data science infra-as-service platform. Before Capital One, Sagar spent seven years building data engineering products at Micron Technology. Sagar earned his BS in Electrical Engineering from the University of Mumbai and earned his Masters from George Mason University, where he studied Information Systems and Software Engineering. Sagar is an occasional speaker at technology and leadership conferences as well as a startup advisor.

Rob Glanzman

Global Strategic Alliances Principal Architect, Financial Services, Pure Storage

Rob Glanzman is a 25+ year veteran of the cloud and data industries, and currently leads Pure Storage's Financial Services vertical worldwide. Previously at VMware, Rob was a founder of VxRail, the world's leading hyper-converged Infrastructure Appliance and multi-billion-dollar run rate business for Dell | EMC. He also co-founded and led VMware's Global Innovation Incubator xLabs in the Office of the CTO. Prior to that Rob managed the VMware / EMC relationship for all Global Financial Accounts and architected some of the largest storage and virtualized environments in existence at the time. Rob also worked in various engineering and executive positions at AT&T and Morgan Stanley. He currently leads Pure's Financial Services Vertical world-wide.

Wyatt Gorman

Solutions Manager, HPC and AI Infrastructure, Google Cloud

Wyatt Gorman is the Solutions Manager for High Performance Computing (HPC) and Artificial Intelligence (AI) Infrastructure at Google Cloud. Wyatt focuses on enabling our customers' toughest workloads and requirements on Google Cloud, along with our engineers, architects, and partners. Wyatt's background is in distributed computing and HPC. He was an HPC & ML Specialist at Google Cloud before becoming a Solutions Manager. Before joining Google Cloud, Wyatt held roles at Seagate and Xyratex as Senior Systems Engineer. Wyatt's first exposure to HPC was helping run the HPC lab while a student at the State University of New York at Geneseo.

Traci Gusher

Americas Data and Analytics Leader, EY (Ernst & Young)

Traci Gusher brings extensive experience leading programs focused on artificial intelligence, advanced analytics, and emerging technologies. She is responsible for leading data and analytics teams in the Americas across Consulting & Strategy and Transactions, focused on building innovative approaches and platforms to enable digital transformation journeys. Traci teams with leaders to build sustainable data, analytics, and artificial intelligence capabilities, focusing on moving beyond experimentation and prototyping. She helps organizations build processes, teams, and technology that follow modern delivery approaches for production-scale analytics and AI, and product-centered data strategies with a relentless focus on generating tangible business benefits.

Travis Humble

Director, Quantum Science Center, Oak Ridge National Laboratory

Travis Humble is director of the US Department of Energy's Quantum Science Center, a Distinguished Scientist at Oak Ridge National Laboratory, and director of the lab's Quantum Computing Institute. Travis leads the development of quantum technologies and infrastructure to impact the DOE mission. Travis is also editor-in-chief for ACM Transactions on Quantum Computing and co-chair of the IEEE Quantum Initiative. Travis holds a joint faculty appointment with the University of Tennessee Bredesen Center to work with students in developing energy-efficient computing solutions.

Nick Ihli

Director of Solutions Engineering and Cloud, SchedMD

Nick Ihli is the Director of Solutions Engineering and Cloud at SchedMD. He is a 16-year veteran of the HPC scheduling space. In that time, he's helped hundreds of customers from all industries and workloads find ways to optimize their HPC scheduling environment. At SchedMD, Nick manages the solutions engineering organization, consulting with customers as they migrate to Slurm from other schedulers or to achieve more from their current

Slurm deployment. Nick also spent some time in public cloud computing and currently spearheads Slurm's cloud projects.

Matt Jacobs

Chief Commercial Officer, Cornelis Networks

As Chief Commercial Officer at Cornelis Networks, Matt is responsible for all go-to-market elements of the company, including revenue generation, partnerships, marketing, and corporate strategy. Prior to joining Cornelis Networks, Matt was Chief Strategy Officer for Penguin Computing, where he was responsible for driving the company's future direction, corporate development, and acquisitions. He also held executive roles in sales, marketing, and product development, where he architected Penguin's transition to the HPC market, developed the company's cloud strategy, and more recently led transformation efforts in the areas of solutions, software infrastructure, and as-a-service products. A 22-year veteran of the technical computing market, Matt has deep expertise in HPC and AI platforms, storage, interconnects, software infrastructure, workloads, cloud, and hybrid and alternative delivery and consumption models. Matt holds a BA in Spanish Linguistics and Literature from the University of Georgia.

Thomas Jorgensen

Senior Director, Technology Enablement, Supermicro

Thomas Jorgensen heads up the AI/ML efforts at Supermicro's Technology Enablement team. Before joining Supermicro, Thomas was VP Operations at Reniac, an Intel Capital-funded Database Acceleration company that used FPGAs and GPUs for hardware acceleration of NoSQL Databases.

Prior to Reniac, Thomas was co-founder of Napatech and grew the company from garage to IPO. Napatech is the world's leading provider of programmable Smart Network Interface Cards (SmartNICs) used for Data Processing Unit (DPU) and Infrastructure Processing Unit (IPU) services in telecom, cloud, enterprise, cybersecurity, and financial applications. Thomas holds several patents in Network Packet Processing and Packet Capture.

Troy Kaster

Vice President of Generative AI, Penguin Solutions

Troy Kaster is Vice President of Generative AI at Penguin Solutions. He has sold and delivered some of the largest AI factories in the world, with almost 50,000 GPUs under management. Troy prides himself on his ability to apply technology to solve a wide array of customer problems, leveraging AI/ML solutions to solve real-word problems and provide business value at scale. Prior to joining Penguin, Troy enjoyed success at Motorola, Emerson Network Power, Artesyn Embedded Technologies, and SMART Embedded Computing. His broad background includes positions in product management, sales, partner management, go-to-market strategy, and professional services. Troy is also a 5TONIC steering committee member and has collaborated with companies like Telefonica, Intel, and Ericsson to accelerate 5G technology adoption. A graduate of Algonquin College, Troy resides in Gilbert, Arizona with his family.

Frank Liu

Head of AI & ML, Zilliz

Frank Liu is the Head of AI & ML at Zilliz, where he serves as a maintainer for the Towhee open-source project. Prior to Zilliz, Frank co-founded Orion Innovations, an ML-powered indoor positioning startup and worked as an ML engineer at Yahoo. His passion for ML extends beyond the workplace; in his free time, he trains ML models and experiments with unique architectures. Frank holds MS and BS degrees in Electrical Engineering from Stanford University.

Doug Norton

Chief Marketing Officer, InspireSemi

Doug Norton is the CMO for Inspire Semiconductor, an Austin-based startup leveraging the open RISC-V CPU instruction set architecture (ISA). He began his career in product development at IBM and moved into the sales and marketing organization to expand IBM's reach into engineering/scientific markets. He then went on to various senior leadership positions at Cadence Design Systems and early stage companies such as CoWare, Newisys, Virtana, and Nimbix. Doug is also President of the Society of HPC Professionals, serves on multiple RISC-V International committees, and is a Mentor at the Austin Technology Incubator (ATI) at The University of Texas. Doug earned a BS in Electrical Engineering cum laude from Missouri University of Science & Technology.

Dan Olds

Chief Research Officer, Intersect 360 Research

Dan Olds is a veteran of the High Performance Computing industry with more than 25 years of experience in the high-end server market and as an industry analyst. As Intersect360 Research's Chief Research Officer, Dan leads the demand side and supplier-driven data analysis practice for Intersect360 Research's forward-looking market intelligence subscription service. In addition, he supports a range of client-specific services, including custom research studies and strategic consulting. An authority on technology trends and customer sentiment, Dan is a frequently quoted expert in industry and business publications.

Ty Panagoplos

EVP & CIO Transformation, Santander US, former CTO/ Executive Director, JPMorgan Chase Grid

Ty Panagoplos has held very senior C-level roles at large global firms over the past 25 years, including JP Morgan, Lehman Brothers, Barclays, TD, and Santander. Ty has extensive global experience having led and built out teams in North America, Europe, Asia, and Africa, spanning many financial services business lines including both Retail and Capital Markets. He also has a detailed understanding of the regulatory environment and requirements in many global jurisdictions. Formally educated as an electrical engineer with extensive hands-on software development experience, Ty has a deep technical background, which is critical for making multi-million dollar architecture, design, and technology decisions, and for ensuring delivery of the most effective technology in support of the business. Ty has been responsible for technology vision and strategy, and has a strong track record of helping FinTech firms to align their strategy and capabilities to current needs and gaps in the market.

Christopher Parisien, PhD

Senior Manager of Applied Research, NVIDIA Guardrail

Christopher Parisien is a Senior Manager of Applied Research at NVIDIA, leading the development of NeMo Guardrails, a toolkit for safety and security in Large Language Models. Chris holds a PhD in Computational Linguistics from the University of Toronto, where he used AI models to explain the strange ways that children learn language. During his time in industry, Chris helped build the first generation of mainstream chatbots, developed systems to understand medical records, and served as Chief Technology Officer at NexJ Health, a patient-centered health platform. His current focus at NVIDIA is to bring trustworthy language models to large enterprises.

Ricardo Portilla

Industry Principal, Databricks

Ricardo has 12+ years of experience working with large financial services customers to bring business value to data teams in the cloud. He has designed and consulted on solutions architecture with hundreds of customers in sub-industries from Capital Markets to Banking and Wealth Management. His previous experience was at FINRA, where he moved mission-critical workloads from on-prem data warehouses to the cloud and established a machine learning practice for financial fraud.

Molly Presley

SVP, Marketing, Hammerspace

Molly brings over 15 years of product and growth marketing leadership experience to the Hammerspace team. Molly has led the marketing organization and strategy at fast growth innovators such as Pantheon Platform, Qumulo, Quantum Corporation, DataDirect Networks (DDN), and SpectraLogic. In these companies she was responsible for the go-to-market strategy for SaaS, hybrid cloud, and data center solutions across a range of data-intensive verticals and use cases. At Hammerspace, Molly leads the marketing organization and is responsible for inspiring data creators and data users to take full advantage of a truly global data environment.

Ryan Quick

Principal, Providentia Worldwide

Ryan Quick has been focused on distributed systems for the last 25 years, with special attention placed on the interaction between applications, operating systems, and the hardware and networks underlying them. He has consulted on high-throughput transactional platforms, security, and systems integration for banking, holds patents for messaging middleware systems, and is a pioneer in bridging HPC technologies with enterprise best-practice infrastructure. Ryan is an expert at scale-

out systems, UNIX kernel design, and profiling, and has been recognized for innovation in hardware and application design, as well as in messaging ontology and distributed event-driven systems. His current efforts bring machine learning, real-time streaming, set-selection, and digital signal processing technologies to bear on predictive analytics to provide self-healing for command and control systems.

J Ram

Co-Founder, Concourse Labs, and former Managing Director, Goldman Sachs Grid and CTO/ Executive Director, Morgan Stanley Grid

J Ram serves as the VP/Co-Founder and CISO at Concourse Labs, a company specializing in automating and managing cloud governance and posture for multi-cloud environments. Before joining Concourse Labs, J Ram held the position of Head of Cloud at Goldman Sachs, where he oversaw cloud for RISK and Pricing (secBD). During his tenure, he successfully streamlined operations, facilitating significant growth in the secDB cloud, which in turn enabled new business opportunities. In his earlier career as the CTO of Morgan Stanley Infrastructure, J Ram played a pivotal role in developing the Aurora Platform and Compute Grid. His innovations and automation efforts resulted in the creation of a cutting-edge compute grid, ushering in a new era of financial products. J Ram's contributions have been instrumental in shaping the landscape of cloud technology and financial services.

Prabhu Ramamoorthy

CFA, FRM, CAIA, Global Partner Success Manager, NVIDIA

Prabhu Ramamoorthy is the financial ecosystem partner manager at NVIDIA, where he focuses on Quant and AI acceleration for financial services. Previously, he was Head of Technology of a margin software firm, Dash Regtech, that catered to leading investment banks, and he served as a director at Big Four firm KPMG/EY, where he helped 100+ financial institutions over the last 10 years. Ramamoorthy holds an MBA from the University of Wisconsin-Madison and an undergraduate engineering degree from BITS-Pilani, one of the top engineering institutes in India. He is a CFA charterholder, financial risk manager, and chartered alternative investment analyst specializing in financial transformation use cases.

John Russell

Managing Editor, QCwire, HPCwire

A previous winner of the Jesse H. Neal National Business Journalism Award (which has been called "the Pulitzer Prize of the business media,"), John Russell leads *HPCwire*'s coverage of quantum computing and helped spearhead development of *QCwire*, Tabor Communications' recently-launched monthly newsletter covering quantum information sciences. John has broad science and business journalism expertise with extensive knowledge of the biopharmaceutical and advanced IT industries. He was formerly a founding executive editor for *Bio-IT World*, which focused on the intersection of advanced computing technology and life sciences.

Bob Sorensen

Senior Vice President of Research, Chief Analyst for Quantum Computing, Hyperion Research

Bob Sorensen is Hyperion Research's Chief Analyst for Quantum Computing. Bob's areas of expertise include analysis of advanced computing hardware, architectures, interconnects, and performance metrics for both classical and quantum systems. Before joining Hyperion Research, Bob worked 30+ years for the U.S. Federal Government as a Senior Science and Technology analyst covering global advanced computing developments to support senior-level U.S. policy makers, including those in the White House, Department of Defense, and Treasury. Bob has a BSEE from the University of Rochester and a MSCS from George Washington University. He strongly prefers C to Python.

Dino Vitale

Distinguished Engineer, Infrastructure Technology Solutions Cloud Engineering Team, TD Bank Group

Dino Vitale is a Distinguished Engineer in the Cloud Engineering group of TD Bank Group, leading the adoption of ML and AI across the enterprise. He is currently focusing on enabling various ML/AI cloud-based services and capabilities, including Generative AI, AI/ML ops, model risk governance/compliance, data engineering/pipelines, and helping applications teams adopt AI services (speech, language, translation). In previous roles, Dino helped revolutionize the adoption of multi-tenant HPC compute grids at various large financial services and designed/built large-scale shared compute environments and automation that supported risk management and analytics workloads. Dino has extensive architecture experience working with a wide range of technologies in designing/implementing data lakes, micro-services, HPC clusters, event streaming, data platforms, and implementing automation tooling for operational scalability and dev/sec ops.

Heather West, PhD

Research Manager, Infrastructure Systems, Platforms, and Technologies, IDC

Heather West, PhD, is a research manager within IDC's Enterprise Infrastructure Practice. In this role, she leads IDC's research on quantum computing. Other areas of Heather's research coverage include AI and enterprise infrastructure workloads. Heather also manages primary research projects focused on end-user purchasing plans for infrastructure products and adoption of technologies shaping the infrastructure market.

Alex Woodie

Managing Editor, Datanami

Alex Woodie has written about IT as a technology journalist for more than a decade. He brings extensive experience from the IBM midrange marketplace, including topics such as servers, ERP applications, programming, databases, security, high availability, storage, business intelligence, cloud, and mobile enablement. Alex resides in the San Diego area.

Craig Yamasaki

Senior Director of Product Management, Supermicro

Craig Yamasaki has a broad background from Field Services, 3rd line technical support, writing/testing firmware for processors and servers, designing ASIC simulation environments, system and solution level program management, and most recently, product management and product marketing for some of the most technical (HPC) customers of servers, IaaS and HPC software in the marketplace. In addition to the tactical business and technical aspects, he has personally driven and subsequently built a broad team to drive a comprehensive roadmap and business strategy working with key value chain partners (such as engineering, supply chain, sales and services) to bring 'whole products' from 'investigation' to 'end of life' based on customer/market data to drive customer value to the top line and bottom line metrics. Craig prides himself on building world class team who are empowered and accountable to drive towards the goals and metrics that enable the business to succeed and the team members to grow to deliver their full potential while developing their career objectives.

CXL FORUM

WEDNESDAY, SEPT. 27 • 9:30AM – 3:30PM

9:30am: **Introduction to CXL and Big Memory Computing**
Dr. Charles Fan, *CEO and co-founder of MemVerge*

9:50am: **Applications and Stretch Goals in CXL-Accelerated Academic Research**
Kurt Keville, *Systems Engineer, MIT*

10:10am: **Memory in the AI/ML and Data Era**
Julie Choi, *Head of New Biz PM, Samsung*

10:30am: **Accelerating Everything AI**
Thomas Jorgensen, *Senior Director, Technology Enablement, Supermicro*

11:00am: **CXL Switches Accelerate the Next Generation AI Computing**
Jianping (JP) Jiang, *VP of Product Marketing and Business Operations, Xconn Technologies*

11:20am: **Advantages of Optical CXL for Disaggregated Compute Architectures**
Ron Swartzentruber, *Director of Engineering, Lightelligence*

11:40am: **Accelerating AI Clusters with Gismo (Global IO-free Shared Memory Objects)**
Yong Tian, *Field CTO, MemVerge*

1:00pm: **Supermicro CXL Storage Solutions**
Thomas Jorgensen, *Senior Director, Technology Enablement, Supermicro*

1:20pm: **The March of Composability Onwards to Memory with CXL**
Alan Benjamin, *President and CEO, GigaIO*

1:40pm: **Elevating Data Infrastructure with Liqid: A Look into Composable CXL**
Bryan Davies, *Solutions Architect, Liqid*

2:30pm: **CXL Birds-of-a-Feather Session**

SPONSORS

PLATINUM

Dell Technologies

Dell Technologies is helping financial services organizations improve investment returns and attract and retain more customers with scalable, flexible solutions designed to help you solve complex problems faster than ever. Powerful engineering-validated designs for HPC that are optimized for Monte Carlo simulations are capable of performing risk assessments across millions of scenarios. We offer a robust portfolio of products and services optimized for HPC and AI in financial services, providing performance and efficiency from a company invested in your future. In fact, we're one of the only companies in the world with a portfolio for HPC, AI, and analytics that spans workstations, servers, networking, storage, rack systems, and services. As an industry leader in advanced computing, Dell Technologies offers proven products, solutions, and expertise that reduce complexity and help you capitalize on the promise of using technology to increase profits and reduce losses by better managing risks.

Hammerspace

Hammerspace is the data orchestration system that unlocks innovation and opportunity within unstructured data. It orchestrates the data to build new products, uncover new insights, and accelerate time to revenue across industries like AI, scientific discovery, machine learning, extended reality, autonomy, corporate video, and more. Hammerspace delivers the world's first and only solution to connect global users with their data and applications on any vendor's data center storage or public cloud services, including AWS, Google Cloud, Microsoft Azure, and Seagate Lyve Cloud.

Penguin Solutions

With more than 20 years of experience in HPC platforms, Penguin Solutions is at the forefront of the generative AI revolution. In fact, we have designed, built, and managed some of the largest and most powerful AI supercomputers in the world. And today we are managing over 50,000 NVIDIA® GPUs for AI training. Penguin unites smart people, open technologies, and value-add services so your business can harness the power of emerging technologies like AI and edge computing. We consistently deliver new capabilities before customers realize they are available or possible. This helps you maintain your competitive advantage and releases substantial time and resources. We're the people to talk to when you're facing the most complex technology project of your career. Penguin Solutions is part of the SGH family of brands.

GOLD

AMD

For more than 50 years, AMD has driven innovation in high-performance computing, graphics, and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world, rely on AMD technology daily to improve how they live, work, and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, LinkedIn, and Twitter pages.

Lenovo

Producing three devices per second in 180 global markets, Lenovo delivers a full portfolio of solutions, from smartphones to smart data center infrastructure. Supercomputers are used across a wide range of industries, including finance, for compute workloads such as financial risk assessments. Lenovo is the #1 provider of high-performance computers and

the manufacturer behind the world's most energy efficient supercomputer according to the TOP500 and Green500, respectively. Additionally, as artificial intelligence changes the way organizations do business, Lenovo takes AI from the realm of research to add real value to any organization by simplifying the adoption process, with optimized infrastructure, proven expertise, and pre-validated solutions. Learn more at lenovo.com/HPC and lenovo.com/AI.

MemVerge

Generative AI has taken the world by a storm, creating a boom for GPUs, high bandwidth memory, and CXL memory. At this pivotal moment, MemVerge has built a Big Memory Computing Platform that continuously rightsizes computing and memory resources for AI workloads. For the public cloud, Memory Machine™ Cloud intelligently manages workloads on public cloud for users such as TGen, EBI, and Seekgene. For CXL, MemVerge's award-winning software enables applications to access pools of CXL resources without modifications, while providing optimal performance and cost. Learn more at www.memverge.com.

Supermicro

Supermicro (NASDAQ: SMCI) is a global leader in Application-Optimized Total IT Solutions. Founded and operating in San Jose, California, Supermicro is committed to delivering first to market innovation for Enterprise, Cloud, AI, and 5G Telco/Edge IT Infrastructure. We are transforming into a Total IT Solutions provider with server, AI, storage, IoT, and switch systems, software, and services, while delivering advanced high-volume motherboard, power, and chassis products. The products are designed and manufactured in-house (in the US, Taiwan, and the Netherlands), leveraging global operations for scale and efficiency and optimized to improve TCO and reduce environmental impact (Green Computing). Our award-winning portfolio of Server Building Block Solutions® allows customers to optimize for their exact workload and application by selecting from a broad family of systems built from our flexible and reusable building blocks that support a comprehensive set of form factors, processors, memory, GPUs, storage, and cooling solutions.

SILVER

CDW Corporation

CDW Corporation, a Fortune 500 company and member of the S&P 500 Index, provides a broad array of products and services ranging from hardware and software to integrated IT solutions such as security, cloud, hybrid infrastructure, and digital experience. CDW services more than 250,000 business, government, education, and healthcare customers in the United States, the United Kingdom, and Canada to help navigate an increasingly complex IT market and maximize return on their technology investments.

GigaIO

GigaIO is the leading provider of workload-defined infrastructure for AI and accelerated computing, featuring the world's only open rack-scale computing platform. The company's flagship product, the GigaIO SuperNODE™, is a 32-GPU single node server that utilizes their transformative, ultra-low latency PCIe memory fabric, FabreX™. This 32-GPU engineered solution offers a simplified system capable of scaling multiple accelerator technologies such as GPUs and FPGAs without the latency, cost, and power overhead required for multi-CPU systems. The innovative open architecture of FabreX uses industry-standard PCI Express/soon CXL technology, allowing GigaIO solutions to break the constraints of the server box and liberate resources to shorten time to results. Get the elasticity of the cloud at a fraction of cloud TCO (Total Cost of Ownership). Contact info@gigaio.com, visit www.gigaio.com, or follow on Twitter and LinkedIn.

JetCool

Explore the advancements in financial computing with JetCool's self-contained liquid cooling system, a proud part of Dell OEM Solutions, tailored for the demands of AI, high-frequency trading, and other applications. Designed for data centers, hyperscalers, and colocation facilities, our state-of-the-art system amplifies compute density, supports next-gen servers, and slashes fan power in half — all with zero facility modifications. Experience unparalleled efficiency and reliability while saving 50% on upfront costs compared to other liquid cooling approaches.

Pure Storage

Pure Storage uncomplicates data storage, forever. Our innovative, cloud-ready solutions enable financial firms to turn bottlenecks into breakthroughs and data into powerful outcomes. Pure accelerates modern AI/ML workloads with massively parallel, scale-out storage solutions that drive time to insight for data-driven businesses. In fact, Pure was recently recognized as the Best AI Solution for Big Data. Pure's storage as-a-service empowers customers, providing the real time agility required to meet ever evolving data needs, and Pure now offers the first and only energy efficiency guarantee in the enterprise storage as-a-service market. With a certified customer satisfaction score in the top 1% of B2B companies, Pure's customers are among the happiest in the world. Find out how our enterprise-class data solutions accelerate financial services and why Meta chose to partner with Pure for its AI Research SuperCluster.

BRONZE

BDO Digital

BDO Digital's comprehensive suite of AI services is designed to assist organizations at every stage of their AI journey, providing end-to-end support for successful AI implementation and adoption. Whether organizations are just beginning to explore the possibilities of AI or are already leveraging it for business transformation, we provide tailored assistance and help them mature their AI adoption overtime. We help our clients create competitive advantage and uncover new avenues of growth and profitability through digital strategies. Your strategies need to be bold, but they also need to deliver near-term results. That's why we take a practical approach that marries long-term vision with quick wins and measurable ROI milestones.

BittWare

BittWare, a Molex company, provides acceleration hardware, including solutions for the financial services industry (FSI). Our portfolio of solutions includes ultra-low latency PCIe cards, network interface (NIC) cards, and AI accelerators. Each is well-suited for low- and ultra-low latency using FPGAs and custom ASICs. In 2023 we have new solutions for low-latency network processing with Liquid-Markets-Solutions (LMS) and their UberNIC, built on BittWare hardware. Our AI partners include Groq, with our FPGA hardware complementing their fully deterministic, low-latency GroqChip AI processor.

Cornelis Networks

Cornelis Networks is the leading independent provider of purpose-built, scale-out interconnects for high-performance computing, artificial intelligence, and high-performance data analytics. Leveraging 20 years of experience dating back to the genesis of InfiniBand, Cornelis Networks designs leadership-class fabrics that deliver the lowest latency and highest message rate at extreme scale. The superior price-performance of Cornelis Omni-Path enables OEMs and system integrators to deliver highly cost-effective network solutions, freeing cluster budgets to deploy additional compute and accelerator resources, and resulting in greater overall system performance and faster time to insight for customers' mission-critical workloads.

Google Cloud

Google Cloud's HPC and AI services help financial institutions around the world reduce risk, improve trading, and manage their portfolios. Google Cloud's powerful, award-winning HPC and AI services and tools are secure, scalable, and cost effective, and have proven themselves at scale with financial services industry customers that include HSBC, Goldman Sachs, and CME Group. Come visit us at our table to learn more, and check out https://cloud.google.com/hpc and https://cloud.google.com/ai.

InspireSemi

InspireSemi is an Austin-based chip design company that has built a technology foundation that delivers best-in-class performance, energy efficiency, versatility, and a thriving open software ecosystem. This enables it to address the multiple, diversified markets of High Performance Computing (HPC), AI, graph analytics, and blockchain. It has a disruptive and strongly differentiated accelerated computing solution compared to existing approaches for these markets. InspireSemi is led by an accomplished team with a proven track record and is leveraging an efficient operating model with world-class outsourced manufacturing partners (e.g., TSMC, ASE). Leading industrial companies, national labs, computer OEMs, and crypto mining companies value InspireSemi's technology, innovative North America design team, and predictable supply chain.

IonQ

IonQ, Inc. is a leader in quantum computing, with a proven track record of innovation and deployment. IonQ's current generation quantum computer, IonQ Forte, is the latest in a line of cutting-edge systems, boasting an industry-leading 29 algorithmic qubits. Along with record performance, IonQ has defined what it believes is the best path forward to scale. IonQ is the only company with its quantum systems available through the cloud on Amazon Braket, Microsoft Azure, and Google Cloud, as well as through direct API access. IonQ was founded in 2015 by Christopher Monroe and Jungsang Kim, based on 25 years of pioneering research. To learn more, visit www.ionq.com.

Liqid

Liqid is revolutionizing IT resource utilization by unlocking the true potential of GPUs. As pioneers in composable infrastructure, we have eliminated traditional GPU deployment constraints. This results in a better total cost of ownership (TCO), unprecedented performance, and flexible on-demand allocation of GPUs. Using Liqid Matrix software, customers can pool and deploy GPUs, Liqid's NVMe drives, and soon, CXL, accelerating results, enhancing performance, and optimizing infrastructure efficiency.

Parallel Works

Parallel Works is a unified control plane and single-pane-of-glass interface that delivers a uniform practitioner experience across diverse HPC resources on-premises and across clouds. The platform facilitates collaborative research through the sharing of workflow tools, resource usage recipes, large datasets, and computational research methods. It also enables push-button creation of customizable and elastic HPC cloud clusters and makes it easy for practitioners to scale to tens of thousands of cores with near real-time cost control features and award-winning ML capabilities. The Parallel Works platform provides flexible, software-defined HPC as a service to customers in industry, government, and academia. It allows scientists, researchers, and engineers to make better use of their on-premises HPC resources while providing a uniform experience on any public cloud service provider. In CSP environments, the platform provides near real-time cost reporting and budget enforcement features. To see a demo, please contact Phyllis Rhodes at prhodes@parallelworks.com.

Quantexa

Quantexa is a global data and analytics software company that helps organizations make more informed operational decisions through meaningful data. Quantexa's platform uses the latest big tech and AI to power Contextual Decision Intelligence; a new approach that uncovers hidden risks and reveals new opportunities by providing a holistic, connected view of internal and external data, all in one place. These insights help companies solve major challenges across data management, KYC, customer intelligence, financial crime, risk, fraud, and security throughout the customer lifecycle. The Quantexa platform boosts operational performance with over 90% more accuracy, and offers 60 times faster analytical model resolution than traditional approaches. Learn more at Quantexa.com or follow us on https://www.linkedin.com/company/quantexa/mycompany.

SchedMD

SchedMD is the core company behind Slurm distribution and maintenance. We are also the sole provider for Slurm support, development, training, and configuration. SchedMD services eliminate security vulnerabilities, offer access to real-time support from Slurm experts, and accelerate Slurm scheduling results with proven best practices. Slurm is a policy-driven, open-source workload manager designed to satisfy the demanding needs of high-performance computing (HPC), high-throughput computing (HTC), and AI. Slurm's automation capabilities simplify administration, accelerate job execution, and improve end user productivity, all while reducing cost and error margins.

WEKA

WEKA is the data platform provider for AI and other performance-intensive workloads leading a paradigm shift in how data is stored, processed, and managed. The WEKA® Data Platform is a software solution that transforms legacy data architectures and stagnant data silos into dynamic data pipelines that power GPUs efficiently and fuel generative AI, ML, and HPC workloads seamlessly and sustainably. Its advanced cloud-native architecture is optimized to solve complex data challenges at scale, delivering 10-100x performance improvements across edge, core, cloud, hybrid, and multicloud environments. WEKA helps the world's leading data-driven organizations accelerate research and discovery breakthroughs and business outcomes.

WhereScape

WhereScape Data Automation software accelerates the design, build, documentation, and management of complex data ecosystems and ensures the delivery of data — wherever it is stored and however it is presented — to the people who need it. Thousands of users worldwide leverage WhereScape automation to remove reliance upon repetitive aspects of data infrastructure projects to deliver data warehouses, vaults, lakes, and marts in days or weeks rather than in months or years. WhereScape sits on top of and orchestrates your existing data infrastructure using automated ELT. This Data Automation enables developers to produce data architectures in a fraction of the time it would take with hand coding, without human error. It enables individual staff members to be more productive and innovative by eliminating repetitive manual tasks.

WEDNESDAY, 6:00PM — RAFFLE ON THE MAIN STAGE

Make sure to be in your seats at 6pm on Wednesday for closing remarks and drawings for tablets, Air Pods, and more fabulous prizes. Must be present to win – good luck!

EXHIBITOR TABLE GUIDE

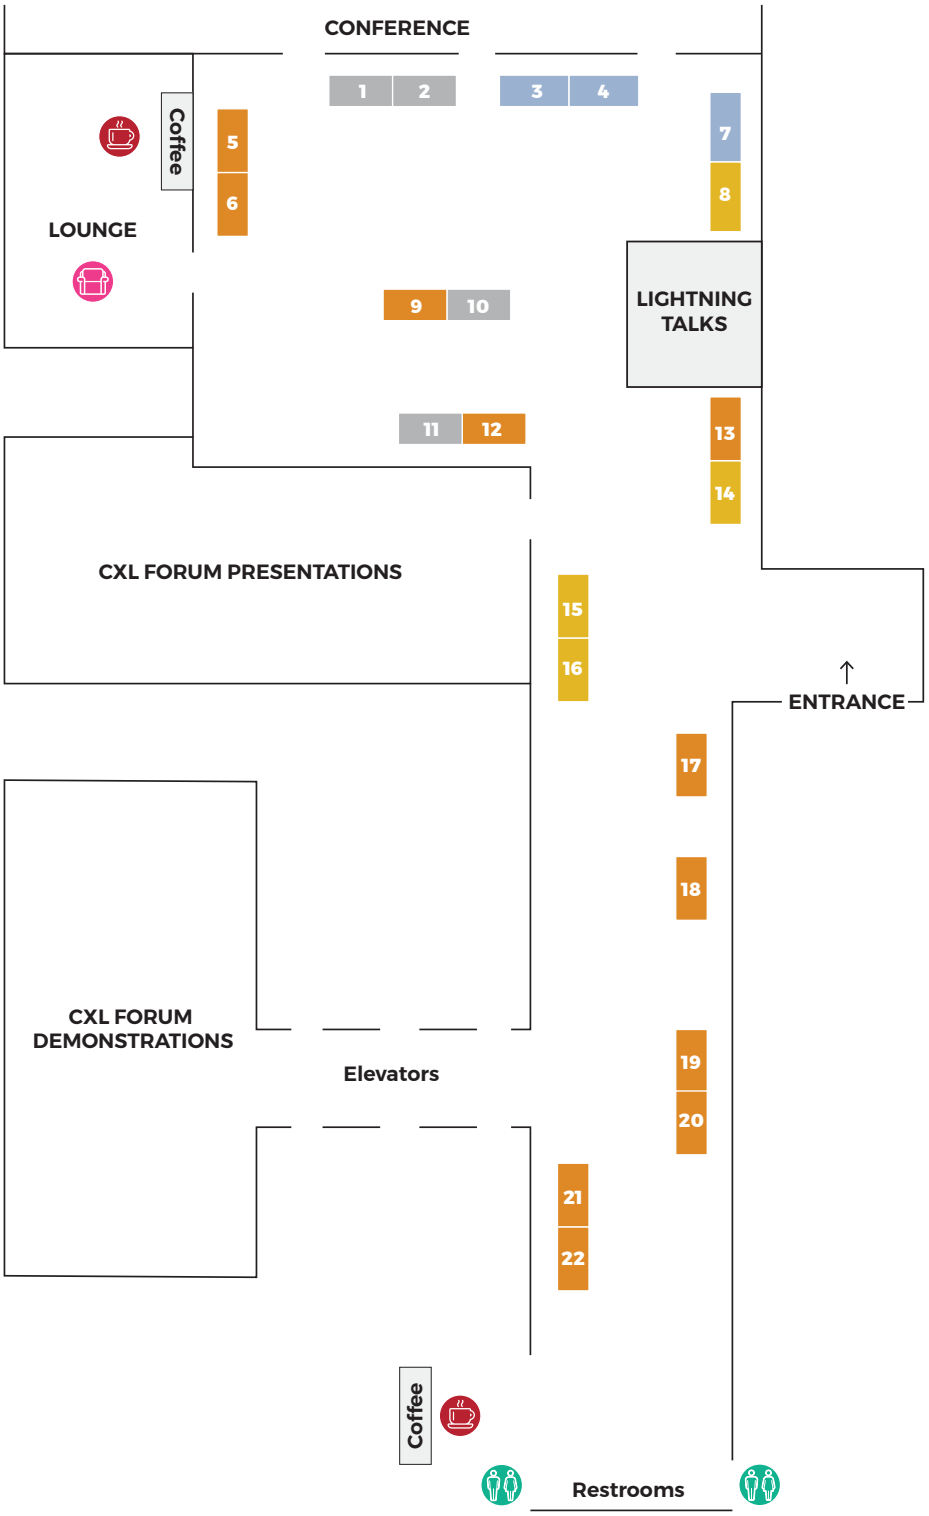

PLATINUM SPONSORS

7 Dell Technologies
4 Hammerspace
3 Penguin Solutions

GOLD SPONSORS

8 AMD
14 Lenovo
16 MemVerge
15 Supermicro

SILVER SPONSORS

10 CDW Corporation
11 GigaIO
2 JetCool
1 Pure Storage

BRONZE SPONSORS

20 BittWare
13 Cornelis Networks
9 Google Cloud
5 InspireSemi
17 Liqid
18 Parallel Works
19 Quantexa
6 SchedMD
12 WEKA
21 WhereScape

MEDIA PARTNER

22 Tabor Communications

Breakfast: A light breakfast will be served in the exhibits area from 8:15-10:15 each morning.
Lunch: Lunch will be served both days at 12:00pm in the main ballroom.
Coffee Break: A coffee break with snacks will be served in the exhibits area from 2:30-4:30 each afternoon.
Cocktail Party: All are welcome to attend a cocktail party sponsored by Dell and AMD in the exhibits area from 6:00-7:30 Tuesday evening.