HPC + AI
WALL STREET
INTELLIGENT GLOBAL FINTECH

SHOW GUIDE

2022

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

Showcasing the fastest systems on the planet for FinServ.

TUESDAY & WEDNESDAY, SEPTEMBER 20 + 21, 2022

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

WELCOME

Welcome to the 2022 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 organizations have historically been compute-insatiable and have engaged in early adoption of bleeding edge systems. The need to handle data collection, analysis, and application-to-market at ever-faster speeds have produced transformative shifts across the ecosystem.

Organizations investing in HPC drive market innovation, speed-to-insight, and competitive advantage. And 2022 marks an exceptional year in HPC as we enter the realm of *spooky science* and develop capabilities considered only a dream just a few years ago.

This year, we will delve into the world of Quantum Mechanics, more specifically Quantum Computing (QC) for Financial Services. QC skeptics abound, but a QC universe exists, and it is quickly shaping and building tomorrow's systems.

QUANTUM COMPUTING IS "A REVOLUTION FOR HUMANITY BIGGER THAN FIRE, BIGGER THAN THE WHEEL."

— Haim Israel, *Head of Global Thematic Investing Research, Bank of America*

Day One will lead off with a panel of analysts and a futurist who will provide insight into the level of investments being made in QC, 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.

Over lunch, we'll hear how HPC just broke into the Exascale Era over the summer. The Frontier Supercomputer at ORNL is solving grand challenge problems with compute speeds of a billion, billion operations per second... today! This compute capability is at the fingertips of FinTech, and our HPC sessions will cover onboarding these breakthrough speeds for today's workloads.

The promise of these compute technologies is the entry point to advanced analytics and AGI. None of this is possible without complex big data management capabilities that align with customized AI algorithms and models. Day Two will focus on these mission-critical challenges with sessions in Data Management and AI.

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 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

TUESDAY, SEPTEMBER 20, 2022

Time	Session
8:30-8:45	Welcome + Opening Remarks

QUANTUM COMPUTING SESSIONS

Time	Session
8:45-9:45am	Three Analysts and a Futurist: Lean In to the Realities of Quantum Computing
9:45-10:30am	The Best Roadmap for Quantum Computing Into a Classical Computing Environment
10:30-11:00am	Break + Exhibits
11:00am-12:00pm	The Future of Quantum Technology at JPMorgan Chase
12:00-1:00pm	Quantum is Coming! A Conversation About its Many Implications for Finance

HIGH PERFORMANCE COMPUTING SESSIONS

Time	Session
1:00-2:00pm	LUNCH: Together We Advance FSI with HPC & AI Solutions
2:00-3:00pm	Break + Exhibits + 4 Lightning Presentations
3:00-4:00pm	Endless Memory: An Overview of CXL for FSI
4:00-5:00pm	Accelerators for FinServ: What's Available and Which One is Right for Your Workload?
5:00-6:00pm	Vendor Showdown with Intersect360 Research
6:00-7:30pm	Cocktail Reception

WEDNESDAY, SEPTEMBER 21, 2022

DATA MANAGEMENT SESSIONS

Time	Session
8:45-9:45am	Top 10 Resources for Unlocking Big Data Value
9:45-10:30am	Analytics Everywhere: How to Scale and Accelerate Your Data Science and AI in Capital Markets
10:30-11:00am	Break + Exhibits
11:00am-12:00pm	Intelligent Searching Using Computation Storage and Neuromorphic Processing
12:00-1:00pm	Migrating Datalakes and Machine Learning to the Cloud

ARTIFICIAL INTELLIGENCE SESSIONS

Time	Session
1:00-2:00pm	LUNCH: Building and Training Large-Scale AI/ ML Models for Operational Efficiency
2:00-3:00pm	Break + Exhibits + 4 Lightning Presentations
3:00-4:00pm	The Tao of Open Source AI
4:00-5:00pm	Planning for AI Regulation and Ethics
5:00-6:00pm	Analyst Crossfire with Intersect360 Research
6:00-6:15pm	Closing Remarks + Raffle Drawings

SESSIONS

TUESDAY, SEPTEMBER 20

QUANTUM COMPUTING

PANEL DISCUSSION: Three Analysts and a Futurist: Lean In to the Realities of Quantum Computing (8:45am)

Moderator: John Russell, *QCwire, HPCwire*

Alex Challans, *The Quantum Insider*
Bob Sorensen, *Hyperion Research*
Dr. Heather West, *IDC*
Eric Hunter, *Bradford & Barthel, LLP*

What do you need to know to take advantage of the quantum revolution? Our panel of experts examines how the quantum computing landscape is evolving and which opportunities (and perils) it presents for financial services. Where are the investments going? What does the quantum vendor community look like? Who's taking the quantum plunge now? What are the likely FS use cases? With these topics on the agenda, you'll get a snapshot of quantum's FS potential that will guide your organization's directional thinking on developing a quantum computing strategy.

QUANTUM KEYNOTE: The Best Roadmap for Quantum Computing Into a Classical Computing Environment (9:45)

Sam Mugel, PhD, *Multiverse Computing*

Now that quantum is here, does it mean that HPC providers should tear down and replace their entire infrastructure? Absolutely not! This talk will review some high-impact industry use cases where current quantum computers can bring value today, discuss the best path to adoption for HPC leaders, and look at some of the wins that quantum computing offers — without the need to upgrade hardware.

The Future of Quantum Technology at JPMorgan Chase (11:00am)

Marco Pistoia, *JPMorgan Chase*

Although quantum has been around for quite some time, the technology is finally maturing at a pace that makes it relevant for experimentation in the finance industry. Finance is the first industry sector to benefit from the impact of quantum computing, thanks to the high-complexity use cases it offers and the need to be timely in financial transactions. This session will cover how JPMorgan Chase is proactively building its own algorithms and applications to be scaled across the firm as well as preparing, with the financial industry's best security defenses, against quantum computers that will be used by fraudsters and hackers to break modern public-key encryption.

LEADERSHIP COMPUTING ACROSS DATA CENTER WORKLOADS

AMD
together we advance_

Quantum is Coming! A Conversation About its Many Implications for Finance (12:00pm)

Jay Boisseau, *Dell Technologies*
Aparna Prabhakar, *IBM*
William 'whurley' Hurley, *Strangeworks*
Robert Hormuth, *AMD*

The hype around quantum computing is exploding, and for good reason: global investment is greater than $10B, global technology giants view it as the next big thing, and the hardware ecosystem alone now has dozens of companies vying to provide reliable, capable, quantum computing systems first — for potentially massive impact and profit. In this conversation we will discuss what is 'real' now, which challenges remain, realistic timelines for quantum computing and industry successes, and especially how quantum computing will transform financial Monte Carlo simulation, portfolio optimization, security techniques, and more in financial services.

HIGH PERFORMANCE COMPUTING

HPC KEYNOTE: Together We Advance FSI with HPC & AI Solutions (1:00pm)

Robert Hormuth, *AMD*

Capital markets, traders, hedge funds, banks, and investment firms depend on power and highly-performant solutions to drive the applications and analytics they need to succeed in their fast-moving and ultra-competitive industry. This discussion will describe the Exascale-informed AMD solutions that power those solutions for the FSI industry.

Endless Memory: An Overview of CXL™ for FSI (3:00pm)

Steve Scargall, *MemVerge*

For applications with a growing appetite for memory, Compute Express Link (CXL) technology promises to dramatically increase memory capacity and utilization. Along the way it will make memory a first-class citizen in the data center, with open standards for pooling and composing petabytes of memory. Adoption of the new technology starting in 2023 requires the involvement of a broad ecosystem of vendors that provide: a) processor, memory, re-timer, and switch chips, b) memory, storage, SmartNIC and DPU cards, c) server, memory array, and composable center infrastructure systems, and d) software to provision the pool of memory spread across the CXL bus. In this presentation, software architect Steve Scargall will provide an overview of CXL standards, technologies, products, and never-seen-before test data showing how CXL will impact FSI applications.

PANEL DISCUSSION: Accelerators for FinServ: What's Available and Which One is Right for Your Workload? (4:00pm)

Moderator: Dan Olds, *Intersect360 Research*
Steve Yatko, *Oktay Technology*
Harvey Stein, *Two Sigma*
Ryan Quick, *Providentia Worldwide*

Just a handful of years ago, there weren't many choices when it came to compute accelerators — you had GPUs,

and that was about it. Today, the landscape has radically changed with a large number of new devices and even entirely new architectures being offered by veteran industry players and new entrants alike. In this panel session, we're going to lay out the alternatives and discuss which accelerators best lend themselves to the unique needs of the FinServ industry. There's no free lunch when it comes to adding accelerators to compute infrastructure. Users have to consider learning curves, potential gotchas, and, above all, the total ROI of accelerating workflows and applications.

Vendor Showdown with Intersect360 Research (5:00pm)

WEDNESDAY, SEPTEMBER 21

DATA MANAGEMENT

Top 10 Resources for Unlocking Big Data Value (8:45am)

Alex Woodie, *Datanami*

The desire is there to leverage advanced analytics and AI for competitive gain in FinServ. What's missing is the capacity. The big holdup at this point: the data is a mess. In this wide-ranging talk, you'll gain insight into today's big data management pain points and learn how the software industry and open source community are responding with better technologies and techniques to truly unlock the value in big data.

DATA MANAGEMENT KEYNOTE: Analytics Everywhere: How to Scale and Accelerate Your Data Science and AI in Capital Markets (9:45am)

Monica Livingston, *Intel*

While companies use AI in different ways, they all face a common challenge — how to quickly move from concept to real-world scale, minimizing cost and mitigating risk while complying with industry regulations. The key is to utilize built-in hardware acceleration and optimizations for popular software tools to streamline AI workflow from data ingest to deployment, clearing the path to scale AI everywhere. This session will cover Data Management and AI tools as they apply to use cases in capital markets.

Intelligent Searching Using Computation Storage and Neuromorphic Processing (11:00am)

Earl Dodd, *World Wide Technology*

This presentation highlights successful use cases of next-generation search technology that handles real-time and historical data queries. We will demonstrate intelligent search with neuromorphic processors that compose computational storage capabilities with 100X search performance and no ETL or indexing required. WWT and Lewis Rhodes Labs (LRL) have combined their Extreme Search™ ability with leadership infrastructure to create the Search++ offering. For data-rich environments (e.g., Cyber, GitHub, and more), Search++ scales to many petabytes of data with guaranteed query completion times in <30 minutes. Leverage your existing data infrastructure investment and add computational storage for consistent results as data grows.

Migrating Datalakes and Machine Learning to the Cloud (12:00pm)

Dino Vitale, *TD Bank Group*

One of the highest priorities for many financial service institutions is creating a cloud strategy and migration approach for supporting various digital modernization initiatives. The traditional "Lift n Shift" and HPC workloads were among the initial workloads migrated to the cloud. As cloud adoption matures, companies are starting to look at creating innovative, data-driven products and services in the cloud, so more choices and challenges are introduced for migrating data assets and platforms for data analytics. This talk will help identify key challenges and considerations for data migration and cover next-generation data/machine learning platforms in the cloud (i.e.,. lakehouse, data mesh/fabric). As machine learning adoption increases, we will also cover key considerations/challenges and the technology stack for implementing MLOps capabilities.

ARTIFICIAL INTELLIGENCE

AI KEYNOTE: Building and Training Large Scale AI/ML Models for Operational Efficiency — Challenges and a Prescriptive Roadmap (1:00pm)

Arti Garg, PhD, *HPE*

Faster, easier, contextual and local language access to financial data for better decision making. Train AI to read vast amounts of financial data and reports (10-k etc.) from all over the world and answer financial queries in a natural language.

PANEL DISCUSSION: The Tao of Open Source AI (3:00pm)

Moderator: Alex Woodie, *Datanami*

Alexei Zhukov, *EPAM*
Arti Garg, PhD, *HPE*
Bob Gaines, *Intel*

What is fueling the meteoric rise in AI? Why, the data, of course! But beyond that, it's the widespread availability of deep learning development tools. But putting open source frameworks like PyTorch and TensorFlow to good use in FinServ is easier said than done. In this informative panel, we'll explore the secrets behind putting it all together.
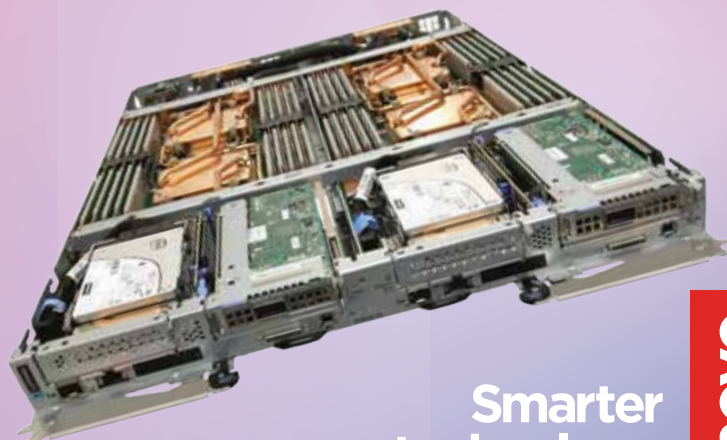
Planning for AI Regulation and Ethics (4:00pm)

Traci Gusher, *Americas Data and Analytics Leader, Ernst & Young*

For years, there has been discussion on ethical considerations for AI, including topics like bias, explainability, and use-case application. Debate and recommendations are now turning to regulation, changing the way organizations will be able to use AI, the steps that are necessary to meet regulatory guidelines, and, of course, transparency. In this discussion, we will examine some of the recent regulations to be enacted, break down some of the ethical concerns that exist in the AI landscape today, and look at how organizations are making plans to mitigate the risks associated with using AI at scale.

Analyst Crossfire with Intersect360 Research (5:00pm)

SPECIAL EVENTS

TUESDAY, SEPTEMBER 20

VENDOR SHOWDOWN
with Intersect360 Research (5:00pm)

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

William Cunningham, PhD, *Head of Quantum Software, Agnostiq*

Matt Jacobs, *Chief Commercial Officer, Cornelis Networks*

Eric Oberhofer, *Vice President, North American Sales, GigaIO*

Jeff Whitaker, *VP, Marketing and Product Management, Panasas*

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.

TUESDAY LIGHTNING TALKS

2:00: **The Future of Back Testing and How to Manage Market Data Workloads Accordingly**

Kurt Kuckein, *Vice President of Marketing, DDN*

2:15: **Sustainability in HPC**

Mark Fischer, *Solution Architect, Lenovo*

2:30: **Building AI Infrastructure Right**

Lior Balan, *Director of Sales & Cloud, Run:ai*

2:45: **Taming the Petabyte Monster**

Jeff Denworth, *CMO and Co-Founder, VAST Data*

Lightning Talks will take place daily in the exhibits area.

WEDNESDAY, SEPTEMBER 21

ANALYST CROSSFIRE
with Intersect360 Research (5:00pm)

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

Ryan Quick, *Principal, Providentia Worldwide*

Arti Garg, PhD, *AI Chief Strategist, Head of Advanced AI Solutions & Technologies, HPE*

Kurt Kuckein, *Vice President of Marketing, DDN*

Thierry Pellegrino, *SVP and President, 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.

WEDNESDAY LIGHTNING TALKS

2:00: **Harnessing the Power of HPC for Finserv**

Tony Parkinson, *VP of Commercial Sales, Penguin Solutions*

2:15: **Get a Grip on File Storage**

Jacob Farmer, *Founder / Chief Evangelist, Starfish Storage*

2:30: **The Future of Big Memory Computing: Cloud and CXL**

Charlie Yu, *Director of Systems Engineering, MemVerge*

2:45: **How to Become an HPC Thought Leader: A Follow-Up to Building AI Infrastructure Right**

Jose Alvarez, *Director of Life Sciences & Research Computing, Cambridge Computer*

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

Lunch: Lunch will be served at 1pm in the main ballroom for those with colored dot badges.

Coffee Break: A coffee break with snacks will be served in the exhibits area from 3:00-5:00 each afternoon.

Cocktail Party: All are welcome to attend a cocktail party sponsored by Dell and AMD in the exhibits area from 6:00-8:00 Tuesday evening.

ADVISORY BOARD

Kristin Boggiano

Co-Founder and President

CrossTower

Chris Brummer

Co-Founder

Fintech Beat

Ritesh Jain

(Former) COO, Global Head of Digital Technology Foundation

HSBC

Shaunt Sarkissian

Chief Innovation Officer, Global Head of Payments, IP, and Corporate Development

TBOL

Phyllis Lampell

Executive Director

Wall Street Technology Association

Harvey Stein

Senior VP, Research Analytics and Methodology

Two Sigma

Dino Vitale

Distinguished Engineer, Infrastructure Technology Solutions Cloud Engineering Team

TD Bank Group

Sagar Gaikwad

Dir. of Engineering – Cloud & Machine Learning Platforms

Capital One

Ankit Manoj Vasa

Data Specialist, Privacy Lead GDPR – Cyber Threat Management

Nasdaq

Tassos Sarbanes

Mathematician – Data Scientist – Data Architect

Credit Suisse AG

Darshak Gosalia

Technology Management

Steve Yatko

CEO

Oktay Technology

Ryan Quick

Principal

Providentia Worldwide

Arno Kolster

Principal & Co-Founder

Providentia Worldwide

SPEAKERS

Jay Boisseau

HPC & AI Technology Strategist, Dell Technologies

Jay Boisseau is an experienced, recognized leader and strategist in advanced computing technologies, with over 25 years of experience in building and leading computing programs, departments, and organizations. As the AI & HPC Technology Strategist at Dell Technologies, Jay evaluates new technologies and approaches for HPC and machine/deep learning, assesses customers' workloads needs and plans, and helps design new solutions to increase customers' innovation, productivity, and efficiency. His current primary focus areas are cloud-enabled HPC for simulation, data analytics (HPC on demand), contributing to the Omnia cluster software project, working with strategic HPC/AI customers, and developing strategies for new HPC/AI solutions in key emerging verticals. Additional projects include evaluations of several new AI & HPC technologies, assessing quantum computing prospects for HPC/AI, leading the Dell Technologies HPC Community, and some confidential (NDA required) projects.

Alex Challans

CEO, The Quantum Insider

Alex Challans is CEO of The Quantum Insider, the leading provider of news and market intelligence on the quantum technology industry. He previously worked as an Investment Director in a London-based private equity fund. Alex is passionate about all things quantum and is focused on how the quantum market can commercialize in a sustainable fashion.

Earl Dodd

HPC & ESG Architect, Alchemist & Business Practice Leader, World Wide Technology

Earl J. Dodd is World Wide Technology's Global HPC Business Practice Leader. He provides HPC/HPDA/ Supercomputing strategy, technology enablement, business development, and ESG/Sustainability support to WWT's global enterprise and government customers. Earl helps enterprises achieve desired ROI/ROC by leveraging HPC+AI technology, extreme data in motion, data center sustainability, and the cloud for secure ultra-scale architectures and collaboration environments. This effort drives the next generation of computationally-steered workflows in decision support environments for real-time situational awareness, institutional learning, and ESG compliance.

Bob Gaines

Senior AI, Advanced Computing, and HPC Solutions Director, Intel

Bob Gaines is the AI machine-learning/deep-learning thought leader for North America.
His specialty is Big Data / High-Performance Data Analytics, leveraging software-optimized heterogeneous computing environments to simplify, automate, and accelerate a client's vision of an integrated data pipeline to improve the degree of confidence and predictability of models, while reducing the time and cost for model development. Bob has a deep history in information architecture, analytics, and AI. He has led the financial services solutions practice for many industry leaders, including Oracle, DDN, Seagate, Cray, and IBM, before joining Intel's Accelerated Computing & Graphics Group (AXG) in 2022. Bob has spent the last 12 years helping financial firms build smarter algorithmic trading models leveraging HPC and AI technologies to address modeling in quantitative research, algo-trading strategy development, risk, compliance, payments, credit, and P&C insurance. He has an MBA (finance + computer science) from Loyola University and an advanced data science certificate from MIT School of Engineering.

Arti Garg, PhD

AI Chief Strategist, Head of Advanced AI Solutions & Technologies, HPE

Arti Garg is Chief Strategist for the AI Solutions business at Hewlett

Packard Enterprise. As part of her role, she oversees HPE's advanced AI technology programs, including the AI Advance start-up program and advanced AI processor partnerships. Over her career, Arti has held several data science leadership roles in the computing and industrial sectors. In the past, she has also worked in the U.S. government at the White House Budget Office and as a legislative adviser in the U.S. Congress. Arti holds a PhD in Physics and an MS in Aeronautical & Astronautical Engineering.

Traci Gusher

Americas Data and Analytics Leader, 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.

Robert Hormuth

Corporate Vice President, Architecture and Strategy of the Data Center Solutions Group, AMD

Robert Hormuth is Corporate Vice President, Architecture and Strategy of the Data Center Solutions Group at AMD. Robert has 33+ years in the computer industry, joining AMD in 2020 after 13 years with Dell, 8 years with Intel, and 11 years at National Instruments. At AMD Robert is charged with creating a long-term system-level vision for DSG and identifying the technical requirements/implications to the DSG portfolio. Robert's

experience includes IO peripheral designs, x86 system design, BIOS, firmware, application software, and FPGA/ASIC design. Robert has a B.S. in Electrical and Computer Engineering from The University of Texas at Austin and currently holds 30 patents.

Eric Hunter

Global Futurist, Head of Innovation Futures, Chief Technology Innovation Officer, Bradford & Barthel, LLP

Eric is a global futurist, speaker, author, and advisory member to multiple boards across the globe. He is the Chief Technology Innovation Officer with Bradford & Barthel, LLP. Eric leads the Innovation Futures Division, an industry-agnostic think tank focusing on future disruptions: how to best leverage the future of technology and human behavior in affecting change. An international keynote speaker, Eric speaks and conducts masterclasses on what is coming next: how to recognize, create, and leverage new business models in leading organizational change through the ongoing reality of global business model disruption. In this vein, Eric has worked with organizations such as Google, IBM, Shapoorji Pallonji Group in India, Polytechnic University in Hong Kong, KM Asia, KM Europe, KM Congress in India, Social Now in Portugal, Global Legal Forum in The Hague, TEDx and many more. Eric expanded his focus within the future of technology and human behavior in his TEDx: "The Future of Human Behavior." His latest book, Travels in Time, published in May of 2022.

William 'whurley' Hurley

Founder and CEO, Strangeworks

whurley is founder and CEO of Strangeworks, a quantum computing startup that makes the power of quantum computing easily accessible and available to all. He is an Eisenhower Fellow, Innovator in Residence for the Legatum Center for Development and Entrepreneurship at the Sloan School of Management at MIT, a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE), Founder

of the Quantum Computing Standards Workgroup at the IEEE, the first Ambassador to CERN and Society, a regular contributor to TechCrunch on the topic of Quantum Computing, and the co-author of *Quantum Computing For Babies*.

Monica Livingston

*Senior Director, Artificial Intelligence and Graphics Sales, Intel*

Monica Livingston leads AI and Graphics technical enablement at Intel Corporation, driving implementation and deployment of AI solutions for end customers, as well enabling workloads for data center graphics. She has previously held roles in large account management, field applications engineering, and hardware design. Monica is passionate about tech, accessible STEM education, and career development. She earned a Masters of Engineering in Electrical Engineering from the University of Florida and an MBA from the University of North Carolina at Chapel Hill.

Sam Mugel, PhD

CTO, Multiverse Computing

Sam Mugel, PhD is an expert in quantum computing and quantum machine learning. His PhD was jointly awarded by ICFO (Spain) and the University of Southampton (UK). Sam was formerly a computational physicist at Cortirio (UK), Technical Director at The Quantum Revolution Fund, and Founder and CTO at Groundstate Consulting. He is an advisor for the McKinsey Tech Council and the Forbes Tech Council, and a Mentor at the Creative Destruction Lab.
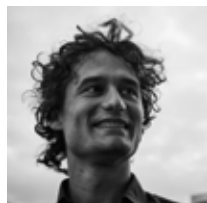
Dan Olds

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.

Marco Pistoia, PhD

Managing Director, Distinguished Engineer, Head of Research & Engineering, Head of Quantum Technology, JPMorgan Chase

Marco Pistoia, PhD, is Managing Director, Distinguished Engineer, and Head of JPMorgan Chase's Global Technology Applied Research. He joined JPMorgan Chase in January 2020. Formerly, he was a Senior Manager, Distinguished Research Staff Member and Master Inventor at the IBM Thomas J. Watson Research Center in New York, where he managed an international team of researchers responsible for Quantum Computing Algorithms and Applications. Marco is the inventor of over 250 patents granted by the U.S. Patent and Trademark Office, and over 300 patent-pending applications. He is also the author of over 400 scholarly papers published in international journals and conference.

Aparna Prabhakar

Vice President Quantum Partner & Alliances, IBM

Aparna Prabhakar is a transformational senior business executive and change catalyst with deep experience in building and scaling strategies for long-term success. Currently, she is a vice president at IBM where she leads the IBM Quantum Partner Ecosystem and Alliances, driving adoption, business development and enablement of the world's largest quantum computing network. Previously, Aparna was an executive leader at IBM corporate strategy, where she led and developed the strategic outlook and recommendations at the C-Suite level

for AI, Quantum Computing, High Performance Computing systems, and Cloud Computing. Aparna was inducted into IBM's Industry Academy in 2021 as one of the top 300 IBM industry leaders worldwide, and has multiple patents and publications. She earned an MBA from the MIT Sloan School of Management and an MS from Louisiana State University.

Ryan Quick

Principal, Providentia Worldwide

Ryan Quick designs and builds scalable architectures for special-purpose computing in blockchain and web3, finance, academia, healthcare, retail, manufacturing, and government. As a software architect with Data Vortex, he accents the revolutionary characteristics of Data Vortex for some of the most difficult problems in distributed computing, workload orchestration, and web3 safety and security. Ryan is an expert at scale-out systems, and has been recognized for innovation in hardware and application design, messaging, and event-driven systems. He has designed high performance platforms with Data Vortex, eBay, PayPal, Qualcomm, Ampere, HPE, Texas Instruments, IBM, the DoE, and others.

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.

Steve Scargall

CXL Consortium Member and Senior Product Manager & Software Architect, MemVerge

Steve Scargall is responsible for leading product management for CXL Fabric Management software at MemVerge. Steve recently joined MemVerge from Intel, where he was a member of the CXL Consortium and architected innovative memory tiering solutions on CXL for Linux, Kubernetes, and Virtualized Environments (KVM and VMWare ESXi). Responsible for influencing and enabling end-user customers, ISVs, OEMs, and partners to adopt Intel products and technologies, Steve played a pivotal role in accelerating Intel adoption by Financial Services Industry (FSI) end users by collaborating with customers and the ecosystem to architect and validate solutions that solve complex business problems.

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.

Harvey Stein

Senior VP, Research Analytics and Methodology, Two Sigma

Harvey J. Stein is the Senior VP, Research Analytics and Methodology, Two Sigma Labs. He was previously head of the Quantitative Risk Analytics Group at Bloomberg LP, where he did research on credit risk modeling, market risk analysis, and portfolio optimization, and worked on the 'Risk-Neutral Risk Measures' project. Harvey's experience and expertise include Numerical Analysis, Risk Management, Stochastic Processes, High-Performance Computing, Asset Pricing, Quantitative Finance, Parallel Programming, Numerical Methods, and Stochastic Analysis.

Dr. Heather West

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.

Steve Yatko

CEO, Oktay Technology

Steve Yatko has 20+ years of experience in advanced enterprise IT systems and technology as a result of leading his own boutique consulting firm, Oktay Technology, as well as his software development firm, Renascence Development, and Credit Suisse. Over the past eight years, he's led Oktay Technology into multi-year strategic technology advisory partnerships with some of the most innovative and successful firms on Wall Street, including Bank of America, Wells Fargo, Bloomberg, Deutsche Bank, Morgan Stanley, Credit Suisse, Liquidnet and Bridgewater Associates. Over the past year, he's also led Renascence Development into strategic partnerships with the development of modern informatics, AI, and Machine Learning systems for firms on Wall Street.

Alexei Zhukov

VP, Technology Solutions, EPAM

Alexei Zhukov is an industry veteran with a diverse range of experience, always maintaining focus on the thriving edge of technology innovation. His career has spanned many roles, including delivery head and solution/enterprise architect, always emphasizing data and analytical solutions in the Financial and Business Information industries. Today Alexei leads the Data Science and AI practice at EPAM, helping to define AI strategy for EPAM and numerous EPAM clients across business verticals and locations. Alexei's group has delivered 100+ AI-enabled solutions to production, bringing to life Advanced Analytics, Optimization and Automation, Computer Vision, NLP, privacy, security, compliance, and MLOps solutions.
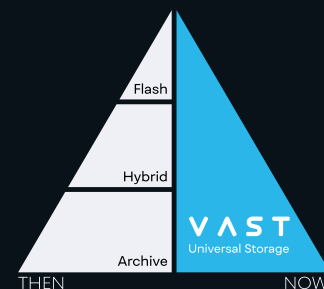
SPONSORS

PLATINUM

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.

Dell Technologies

Dell Technologies helps organizations and individuals build their digital future and transform how they work, live, and play. The company provides customers with the industry's broadest and most innovative technology and services portfolio for the data era.

HPE

HPE is the global edge-to-cloud company built to transform your business. How? By helping you connect, protect, analyze, and act on all of your data and applications wherever they live, from edge to cloud, so you can turn insights into outcomes at the speed required to thrive in today's complex world.

Intel

The Intel® comprehensive HPC technology portfolio helps customers achieve outstanding results for demanding workloads and the complex problems they solve. Intel® HPC hardware is designed to scale from stand-alone workstations to supercomputers with thousands of nodes. Deep learning acceleration is built right into the chip, so Intel® processors are a great choice for innovative solutions that combine HPC and AI. Intel also provides Intel® oneAPI cross-architecture software development tools for HPC. Intel's open software approach and collaboration with a wide range of partners worldwide support a broad ecosystem that helps developers realize more value from technology investments. Learn more at: intel.com/hpc.

GOLD

Cambridge Computer

Cambridge Computer specializes in storage, networking, data protection, and digital archiving technologies. For 30 years, we have helped our clients identify their needs, devise long-term storage management strategies, and compare technologies from different hardware and software manufacturers. Cambridge Computer Services' HPC team provides services such as full design and system architecture, procurement, installation, post-implementation support, application optimization, and benchmarking. Our team of specialists has deep HPC industry knowledge and subject matter expertise in Machine Learning, AI, Computer Science, Computational Chemistry, Bioinformatics, Life Sciences, Physics, and Mechanical Engineering, among others.

DDN

DDN is the world's largest private data storage company and the leading provider of intelligent technology and infrastructure solutions for Enterprise At Scale, AI and analytics, HPC, government, and academia customers. Through its DDN and Tintri divisions, the company delivers AI, Data Management software and hardware solutions, and unified

analytics frameworks to solve complex business challenges for data-intensive, global organizations. DDN provides its enterprise customers with the most flexible, efficient and reliable data storage solutions for on-premises and multi-cloud environments at any scale. Over the last two decades, DDN has established itself as the data management provider of choice for over 11,000 enterprises, government, and public-sector customers, including many of the world's leading financial services firms, life science organizations, manufacturing and energy companies, research facilities, and web and cloud service providers.

Lenovo

Lenovo is a global leader in HPC and AI, with the most entrants on the TOP500 list of supercomputers, industry-leading liquid cooling, and energy-aware software solutions. Lenovo servers have been #1 in reliability for seven years running. We operate in 180 markets with 63,000 employees around the globe.

MemVerge

MemVerge is pioneering Big Memory Computing for a multi-cloud world. Major gaps exist in today's cloud infrastructure for data-intensive high-performance applications. MemVerge® Memory Machine™ delivers software-defined, composable memory and intelligent memory service to bridge these gaps. As a software leader in the CXL ecosystem, MemVerge composable memory technology provisions, tiers, disaggregates, and pools heterogeneous memory to scale memory capacity and decrease memory cost. MemVerge ZeroIO™ in-memory snapshot services transparently checkpoint, clone, replicate, and restore running applications anytime, anywhere in a multi-cloud computing environment. To learn more about MemVerge, visit www.memverge.com.

Penguin Solutions

Penguin Solutions accelerates customers' digital transformation by harnessing the power of emerging technologies in HPC, AI, and IoT with solutions and services that span the continuum of edge, core, and cloud.

Run:ai

Run:ai helps organizations accelerate their AI journey — from building initial models to scaling AI in production. The Atlas Platform brings cloud-like simplicity to resource management, providing researchers with on-demand access to pooled resources for any AI workload. An innovative cloud-native operating-system — which includes a workload-aware scheduler and abstraction layer — helps IT simplify AI implementation, increase team productivity, and gain full utilization of expensive GPUs. Using Run:ai's Atlas AI Compute Platform, companies streamline development, management, and scaling of AI applications across any infrastructure (on-premises, edge, cloud).

Starfish Storage

Starfish is a unique software application for managing unstructured data at a very large scale. Starfish combines a file system metadata catalog with a parallelized data mover and batch processor. You make discoveries and reports using the catalog. You move data fast and furiously and take other actions using the batch processor. The software is agnostic to storage vendors and it works great with HPC file systems, enterprise NAS, object stores, and archives. Starfish allows your users to participate in storage management. Use cases include ROT cleanup, duplicate detection, backup, archiving, reporting, chargeback, content classification, and much more.

VAST Data

Banks and hedge funds have struggled to keep pace with the needs of their quant researchers because of the inefficiencies of mechanical media and the classic tiered-storage data management model. When insights are dependent on rotational media, portfolio managers lose the ability to maximize alpha. VAST breaks the decades-old storage

performance and capacity tradeoff to enable at-scale processing of market data and train AI data sets at any scale. With VAST Universal Storage, all of your market data archive can be stored on one tier of scalable, affordable flash — making it possible to backtest and train all trading models in real time.

SILVER

CDW Corporation

CDW Corporation, a Fortune 500 company and member of the S&P 500 Index, provides integrated Information Technology (IT) solutions and services to 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.

Cornelis Networks

Cornelis Networks™ is the leading provider of scalable, highly-performant interconnect solutions that accelerate High Performance Computing (HPC), High Performance Data Analytics (HPDA), and AI applications in the Cloud and in the Data Center. As the home of Omni-Path Express™, Cornelis supports scientific, academic, governmental, and commercial segments. You can access Cornelis Networks' end-to-end interconnect solutions worldwide through an established set of server OEM and channel partners.

BRONZE

Agonstiq

Agnostiq is developing Covalent, an open source workflow orchestration platform designed to help users manage and execute tasks on heterogeneous compute resources.

Data Vortex Technologies

Data Vortex Technologies is a privately held proprietary network company in Austin, Texas. The novel and unique Data Vortex topology provides a memory semantic, load-store solution for complex large-scale data movement, management, and analysis. Problems that require fine-grained, high radix, low latency, congestion-free communication between nodes are elegantly addressed within the self-routing Data Vortex switch, allowing for true scalability without performance degradation. Application areas include consensus for blockchain, large-scale graph analytics, machine learning, and AI. The parameterizable DV Network-on-Chip, designed around a densely connected high performance topology, is implementation independent, enabling customizable solutions across today's converging fabric-driven IT environment.

GigaIO

GigaIO accelerates time to insight for HPC + AI workflows by seamlessly composing an entire rack of resources into an agile compute platform, bringing cloud flexibility to your data center. No longer is the server the basic unit of compute. With GigaIO's dynamic memory fabric, FabreX, your entire rack is your server. Build impossible servers — from storage to accelerators to memory — at a fraction of cloud TCO by optimizing the utilization and efficiency of your existing hardware. Let your workload define your compute instance by transforming static infrastructure into software-defined hardware where resources communicate over a native PCIe memory fabric for the lowest possible latency, just as they would if they were plugged directly into the server motherboard.

Liqid

Liqid's composable infrastructure software platform, Liqid Matrix™, unlocks cloud-like speed and flexibility plus higher efficiency from data center infrastructure. Now IT can configure, deploy, and scale physical, bare-metal servers in seconds, then reallocate valuable accelerator and storage resources via software as needs evolve. Dynamically provision previously impossible systems or scale existing investments, and then redeploy resources where needed in real-time. Unlock cloud-like datacenter agility at any scale and experience new levels of resource and operational efficiency with Liqid.

NetApp

NetApp Cloud Services enable HPC and AI applications to run optimally in Microsoft Azure. The NetApp portfolio includes intelligent cloud services, data services, and storage infrastructure that helps organizations manage applications and data everywhere across hybrid environments.

Panasas

Panasas builds a portfolio of data solutions that deliver exceptional performance, unlimited scalability, and unparalleled reliability — all at the best total cost of ownership and lowest administrative overhead. The Panasas data engine accelerates AI/ML and high performance applications in manufacturing, life sciences, energy, media, financial services, and government. The company's flagship PanFS® data engine and ActiveStor® storage solutions uniquely combine extreme performance, scalability, and security with the reliability and simplicity of a self-managed, self-healing architecture. For more information, visit www.panasas.com.

Quantum

Quantum technology, software, and services provide the solutions that today's organizations need to make video and other unstructured data smarter — so their data works for them and not the other way around. With over 40 years of innovation, Quantum's end-to-end platform is uniquely equipped to orchestrate, protect, and enrich data across its lifecycle, providing enhanced intelligence and actionable insights. Leading organizations in cloud services, entertainment, government, research, education, transportation, and enterprise IT trust Quantum to bring their data to life, because data makes life better, safer, and smarter. Quantum is listed on Nasdaq (QMCO) and the Russell 2000® Index. For more information, visit www.quantum.com.

Red Barn HPC

Founded in 1996, Red Barn Technology Group (RBTG) has become a leader in HPC solutions for research institutions, industry, and academia, and is proud to have earned the distinction of being an Intel Platinum Technology Provider and an NVIDIA Preferred Solution Provider. Our expertise gives us the unique ability to offer complete solutions encompassing all aspects of HPC — including a wide array of custom Linux and open architecture server and storage solutions. We believe in open source. We believe in open source applications, and that the use of industry-standard non-proprietary hardware yields at least two main benefits — cost savings and increased performance.
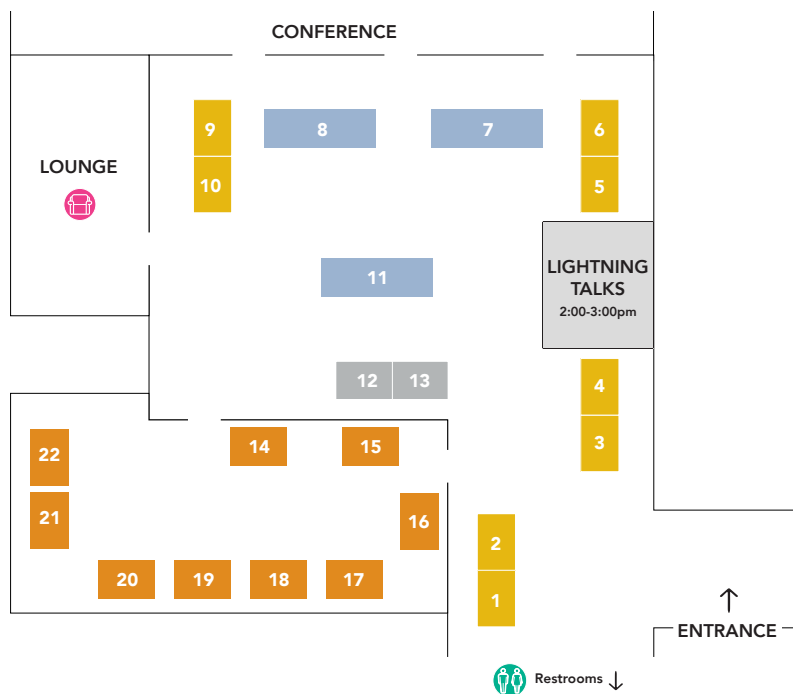
Tabor Communications

Tabor Communications Inc. (TCI) is a media and services company dedicated to high-end performance computing. As publisher of a complete advanced computing portfolio that includes *HPCwire*, *Datanami*, *EnterpriseAI*, *QCwire*, and *HPCwire Japan*, TCI is the market leader in online journalism covering emerging technologies within the high-tech industry, and provides events, audience insights, and other services for those engaged in performance computing in enterprise, government, and research. www.taborcommuications.com.

EXHIBITOR BOOTH GUIDE

CONFERENCE

LOUNGE

| 9 | 8 | 7 | 6 |
| 10 | | | 5 |

| 11 |

LIGHTNING TALKS
2:00-3:00pm

| 12 | 13 |

| 4 |
| 3 |

22		14	15	
21				16
	20	19	18	17

| 2 |
| 1 |

↑ ENTRANCE

Restrooms ↓

PLATINUM SPONSORS
- 7   AMD
- 11 HPE
- 8   Intel
    Dell Technologies

GOLD SPONSORS
- 3   Cambridge Computer
- 2   DDN
- 1   Lenovo
- 10 MemVerge
- 9   Penguin Solutions
- 4   Run:ai
- 5   Starfish Storage
- 6   VAST Data

SILVER SPONSORS
- 12 CDW Corporation
- 13 Cornelis Networks

BRONZE SPONSORS
- 19 Agnostiq
- 14 Data Vortex Technologies
- 20 GigaIO
- 16 Liqid
- 22 NetApp
- 18 Panasas
- 15 Quantum
- 21 Red Barn HPC
- 17 Tabor Communications