STAC BENCHMARK COUNCIL™

STAC® Summit

November 13, 2014
STAC Exchange (exhibits) opens: 8:30am
Conference starts: 9:00am

New York Marriott Downtown
85 West Street at Albany Street
New York, NY

Platinum Sponsor:

Gold Sponsors:

STAC Exchange

New to STAC Summits, the STAC Exchange is an area of exhibit tables near the conference hall where attendees can have 1:1 discussions with vendors before the conference and throughout the day.

Vendors lining up to participate in the STAC Exchange include:

AGENDA

STAC update on STAC-A2 (risk computation) [slides/video]

- *Peter Lankford, Founder & Director, STAC*

The STAC-A2 Benchmark suite is the industry standard for testing technology stacks used for compute-intensive analytic workloads involved in pricing and risk management. Peter will provide a brief update on the latest learnings and the latest activities of the STAC-A2 Working Group.

Key decisions when going parallel [slides/video]

- *David O'Shea, Financial Services ISV Segment Manager, Intel.*
- *Evgeny Fiksman, Staff Engineer, Intel.*

By now, most financial firms know that the way to harness Moore's Law (which continues, by the way) is to parallelize their applications. But that's easier said than done. Huge amounts of existing code are still serial or not parallelized well. What are some of the strategies a firm can take to deciding which applications and libraries to parallelize first? Once the decision is made, how can a developer decide what approach to take to parallelize a given piece of code, such as the programming model, compiler, and tools? Evgeny will propose answers to these questions.

Better performance via better language technologies [slides/video]

- *Eric Sedlar, Vice President & Technical Director, Oracle Labs*

Compilers have made great advances in the last decade in their ability to optimize programs. But as generic tools, standard compilers still frequently make decisions that are not optimal for a given domain. In this talk, Eric will provide insights from Oracle's research into how to endow compilers with enough knowledge to choose the right algorithm or data structure for a given task. Eric contends that compilers need dynamic feedback from the execution of program code to improve choices for a particular use case. He will argue that this generally requires a higher-level, domain-specific language (DSL) and will illustrate this via two use cases: 1) low-latency pre-trade risk calculations in an FPGA-based network card, and 2) a graph-analytics engine. Noting that programmers usually need to use general purpose languages (e.g., R, SQL, Java, JavaScript) in conjunction with DSLs but that cross-language linking often hurts performance, Eric will also unveil research into a high-performance, multi-lingual runtime and explain how compilers can re-use domain-specific and algorithm-specific optimization knowledge more widely for user-written code. This talk is aimed primarily at developers who stand to benefit from these technologies, not compiler writers.

STAC update on low-latency research [slides/video]

- *Peter Lankford, Founder & Director, STAC*

Peter will review the latest STAC activities related to low-latency workloads.

Innovation Roundup – Round 1

"Reducing latency by continuously predicting session performance in real time" [slides/video]	*Tony Pettipiece, Global Head of Sales and Marketing, Cape City Command*
"Introducing the ExaLINK Fusion: a new ultra-low-latency switch and application platform" [slides/video]	*Dr Matthew Chapman, CTO, Exablaze*
"Four Trends Transforming the Market Data Landscape" [slides/video]	*Lee Fisher, VP, Marketing, Redline Trading Solutions*

BREAK

Innovation Roundup – Round 2

"Beyond the first layer" [slides/video]	*Dave Snowdon, Founder, co-CTO, Metamako*
"Introducing NovaTick, NovaSparks Third Generation Pure FPGA Feedhandler" [slides/video]	*Olivier Baetz, COO, NovaSparks*

How PayPal uses DSPs to process event streams: Finding the signal in the noise (literally) [slides/video]

- *Ryan Quick, Principal Architect , PayPal*

PayPal faces a computing problem familiar to many financial firms: they need identify patterns of interest in large volumes of event streams and take action in real time. The events range from well-structured data to natural language/unstructured content and everything in between. Some events are ordered but many are not, and relationships between events must be inferred. While the problem is common, the architecture PayPal chose is not. Rather than distributing data across many CPUs for parallel analysis, the PayPal team transforms event streams into signals that can be handled by conventional digital signal processors (DSPs). They claim that this approach yields microsecond latencies, a radically lower power footprint, and much lower acquisition costs. As the primary architect behind this solution, Ryan will take us through the key challenges PayPal needed to overcome, such as the event ontology, conversion methodology, and mapping to DSP conventions such as frequency, amplitude, modulation, filtering, FFT, etc. He will also discuss the cross-platform programming model that makes this approach practical, and he'll outline the performance and cost benefits of the approach.

Rethinking network capture [video]

- *Eric Powers, Applied Engineering Lead - Global Markets & Equities, Deutsche Bank*
- *Matthew Knight, Marketing Director Financial Services, Solarflare* **[slides/video]**
- *Glenn Wright, Systems Architect, DataDirect Networks* **[slides/video]**

Financial firms today rely on network captures for a range of crucial analytics. But capturing data remains a fairly expensive and rigid process, handled by specialty cards and appliances. How well does this model fit today's customer needs? Would a more open-systems approach be superior--or just more complicated? If a firm takes over the responsibility to engineer its own capture solutions, how should it handle some of the major considerations (e.g., constructing a storage architecture with sufficient capacity and the ability to capture during traffic bursts, what form to capture data in, how to accommodate realtime analytic needs, how to satisfy compliance requirements, etc.)?

NETWORKING LUNCHEON

Analytic opportunities and challenges from big network data [video]

- *Eric Powers, VP: Applied Engineering Lead – Global Markets & Equities, Deutsche Bank*
- *Ken Jinks, Director of Product Marketing, Corvil*
- *Michel Debiche, Financial Services Practice Lead, ThinkBig Analytics*

Many financial firms understand the usefulness of network packet inspection for monitoring latency and network health. But today, network-captured data has potential value in a broader range of analytic needs that depend on integrating information across disparate systems. After all, network protocols are sometimes the only thing that these siloed systems have in common. What are the compelling new analytic use cases for captured network data? What's both technically possible and economically feasible today? What does an analytics architecture look like that can handle these big, fast data sources? What role, if any, should "big data" technologies play? Our panelists will weigh in.

STAC update on big data benchmarking (including tick analytics and backtesting) [slides/video]

- *Peter Lankford, Founder & Director, STAC*

Peter will summarize the latest activities in areas relating to data-bound workloads such as enterprise credit risk, strategy backtesting, and tick analytics.

Innovation Roundup – Round 3

"Can SQL handle tick data with speed?" [slides/video]	*Chris Mureen, COO, McObject*
"Simplifying Big Data Analysis" [slides/video]	*Rod Cope, CTO, Rogue Wave Software, Inc*

Sustainable extensibility and durable data in a complex, volatile world [slides/video]

- *Rupert Brown, CTO, Financial Services, MarkLogic*

Financial institutions face a perfect storm of politically driven scrutiny of their activities coupled with incessant shareholder demands for greater operational efficiency. Adding to this is the exponential growth in data fueled by electronic trading and the ever-increasing range of digital channels to market. The Basel Committee's BCBS 239 paper imposes a new aspect to this problem by focusing on the sustainability of the process rather than the outcomes (which are the focus of FINREP, COREP, and IFRS). In this talk, Rupert will review the key information management and technology challenges that are required to:

- Dynamically adapt to arbitrary changes in internal organization and client corporate entities and keep a consistent auditable record of them

- Correctly capture the complex nuances of derivative instruments and their overarching deal structures

- Determine the true impact of the "force majeure" clauses in client and provider contracts that surface as a result of M&A or outsourcing activities.

BREAK

Best practices in multi-tenant big data infrastructures [slides/video]

- *Gordon Sissons, Senior Manager, Product Marketing – IBM InfoSphere® BigInsights, IBM*

As STAC documented in our white paper on big data use cases in financial services, many banks want to implement multi-tenant Hadoop clusters but face several challenges in doing so. This leaves them with multiple siloed clusters. In this session, Gord will discuss recent client engagements aimed at deploying multi-tenant big data platforms. This includes an insurer that combined over two dozen big data applications--with diverse workloads ranging from click-stream analytics and fraud detection to underwriting and vehicle telematics--in a single, shared cluster. He will discuss what IBM has learned from such implementations and what it believes are emerging best practices.

You Snooze you Lose: Lessons from Real-Time Bidding [slides/video]

- *Brian Bulkowski, Founder & CTO, Aerospike*

Much like financial trading, the online ad market is now dominated by automated decision making. Every time a user opens a web page with ad space, computers from multiple firms engage in an enormous amount of communication, analysis, and bidding behind the scene to decide which company's ad will make it onto the page—all within milliseconds. Similar to many processes in a bank, realtime bidding platforms require a data architecture that can handle terabytes of data and hundreds of thousands of transactions per second. How do the leading platforms achieve this and turn insights into action? Brian will describe his experiences building and operating modern ad-tech platforms, combining an in-memory NoSQL database for predictable high performance with Hadoop based analytics platforms in the back, an architecture that he says is rapidly becoming the gold standard in the Age of Context.

BREAK

The current and future Kafka [slides/video]

- *Joe Stein, Founder, Principal Consultant, Big Data Open Source Security LLC*

Only a handful of industries used to be concerned with streaming data, such as defense, sensor-driven manufacturing, and of course capital markets. Today, stream processing is a topic in many more industries, from retailing to utilities to social media. And as with so many data-intensive problems today, web companies are creating and open sourcing a large amount of code to handle them. Technologists in investment banking, retail banking, and wealth management are considering these open source tools to deal with new classes of problems as well as old ones. One of the tools getting uptake is Apache Kafka. Open sourced by LinkedIn, Kafka is a publish-subscribe message bus designed for high throughput and reliability. In this talk, Joe will introduce the motivation and architecture of Kafka, how it compares to other messaging systems, and what lies ahead in its roadmap.

Innovation Roundup – Round 4

"1 Billion transactions/second on a $25k cluster " [slides/video]	*Nikita Ivanov, Founder and CTO, GridGain Systems*
"Exploiting open source for real time: It's all about the metadata" [slides/video]	*Herman Fick, CEO, fraXses*
"Active Pivot - A Unique In-Memory Big Data Realtime Aggregation engine" [slides/video]	*Jean Safar, co-founder and CTO, QuartetFS*

New approaches to realtime analytics [video]

- *Paul Cao, Director, Data Services, Wells Fargo*
- *Herman Fick, CEO, fraXses*
- *Nikita Ivanov, Founder and CTO, GridGain Systems*
- *Jean Safar, co-founder and CTO, QuartetFS*

Silicon Valley is hopping with activity around "realtime analytics". At the same time, many financial services firms are looking for new approaches to problems like realtime risk, point-of-trade analytics, and fraud detection. Our panel will explore some of the questions that arise, like: what is possible today that was too difficult or too expensive in the past? what role should open source products play in this equation? how have the economics of solutions changed?

NETWORKING RECEPTION

Speaker Biographies – Feature Sessions

Rupert Brown, CTO, Financial Services, MarkLogic. Rupert has spent more than 25 years in the Investment Banking sector working for a number of Tier 1 US and European Investment Banks designing and delivering Global Application, Market Data and Infrastructure Solutions. He has been instrumental in the evolution of the derivatives market platforms in Fixed Income, Equities and Commodities, the transition from host based to distributed and mobile technologies and deciphering the core data requirements of Regulators and Operational Risk managers in the post Crunch era. He graduated from Imperial College London and is a Senior Member of the Association for Computing Machinery (ACM).

Brian Bulkowski, Founder & CTO, Aerospike. Brian is a founder of Aerospike, CTO & Product, networking whiz, innovator and high performance expert. 'My family has a long and varied history in science and tech, so I wound up shipping code in high school. One of the great things about software is you can build something with it. You don't need plywood or welding. Computers are an easy way to start creating stuff. My first taste of networking was in 1989. I knew there was a whole world out there waiting. A computer that's not connected to a network is kind of dull.' Brian became a Lead Engineer at Novell, and then Chief Architect of Cable Solutions at Liberate – where he built a high-performance, embedded networking stack, as well as the high scale broadcast server infrastructure. As Director of Performance at Aggregate Knowledge, Brian had direct experience with the scaling limitations of sharded MySQL systems. 'It wasn't hard to see that there was a huge need for a new distributed database, because they all sucked. Everyone was struggling with what was available. That led to the idea for Citrusleaf – which then became Aerospike.' When he's not busy creating stuff without plywood or a welding torch, Brian plays cello in a band called Rosincoven. He also writes about cuisine for the San Jose Metro."

Paul Cao, Director, Data Services, Wells Fargo . Paul is an Information Technology leader in the financial services industry, focusing on large scale, enterprise data initiatives across asset classes. As head of Data Services at Wells Fargo Securities, he's leveraging various big data technologies to enable various business units to improve efficiency and profitability while meeting regulatory, compliance and surveillance obligations.

Michel Debiche, Financial Services Practice Lead, ThinkBig Analytics. Michel earned M.S. and Ph.D. degrees in Geophysics from Stanford and Princeton Universities. He has been involved in all aspects of quantitative trading since 1991. He worked in proprietary trading groups at Credit Suisse and Daiwa Securities America before creating the Global Equity Statistical Arbitrage desk at CIBC World Markets in New York. In 2002, Michel formed Quantia Capital, an investment advisor that has been involved in building several systematic trading operations. He has also consulted to financial services firms as well as their vendors in the areas of quantitative analysis, algorithmic trading, high-performance systems development, Complex Event Processing systems and enterprise risk management platforms. After a stint as Head of Quantitative Technology at First New York, a proprietary trading shop. Michel joined Think Big Analytics in 2013. Michel sees the rise of the Big Data ecosystem as a golden opportunity to apply the lessons learned in quantitative trading, a quintessential data driven business, to help large enterprises move towards driving greater business value from their growing data streams.

Herman Fick, CEO, fraXses. Herman has a strong background in IT Management (ex Group CIO), practical business solutions, general IT, and Big Analytics. As a founding member of Quantum System Integrators in 1993 and then of the Intenda Group of Companies in 2001, Herman is not afraid to put his energy behind his beliefs in the IT space. He combines his entrepreneurial approach to life with his love for solutions designing. The new fraXses technology and methodology for Big Analytics is a prime example. Herman has built the global 'go-to-market' strategy and the product roadmap while acquiring and integrating a new Machine Learning and Neural Technology company for predictive analytics. In the process, Herman has taken fraXses to the next level--from technology stack to application stack.

Evgeny Fiksman, Staff Engineer, Intel. After joining Intel in 2006 during and working on optimization of video enhancement algorithms for x86 platforms, Evgeny acquired expertise in multi-threading and low level programing. For the last 5 years Evgeny was leading engineer and architected for the implementation of OpenCL runtime for Intel CPUs (Core, Xeon & Atom) and Xeon PHI co-processors. Recently, Evgeny've joined a software enabling team, which is focused on financial applications. Prior joining Intel Evgeny lead development of a naval team training simulator. Evgeny's holding B.Sc and M.Sc in Electrical Engineering from the Technion – Israel institute of Technology, Haifa, Isreal.

Nikita Ivanov, Founder and CTO, GridGain Systems. Nikita is founder and CTO of GridGain Systems, started in 2007 and funded by RTP Ventures and Almaz Capital. Nikita has led GridGain to develop advanced and distributed in-memory data processing technologies – the top Java in-memory data fabric starting every 10 seconds around the world today. Nikita has over 20 years of experience in software application development, building HPC and middleware platforms, contributing to the efforts of other startups and notable companies including Adaptec, Visa and BEA Systems. Nikita was one of the pioneers in using Java technology for server side middleware development while working for one of Europe's largest system integrators in 1996. He is an active member of Java middleware community, contributor to the Java specification, and holds a Master's degree in Electro Mechanics from Baltic State Technical University, Saint Petersburg, Russia.

Ken Jinks, Director of Product Marketing, Corvil. Ken has been at the forefront of the Corvil product since he joined the company in 2003, with extensive experience in designing and creating products for the enterprise and financial technology markets. Ken works closely with parties in the trading loop to enable effective analysis of their high performance trading systems and products. Ken holds a degree in Electronic Engineering.

Matthew Knight, Marketing Director Financial Services, Solarflare. Before joining Solarflare in January 2014, Matthew was the Company President of Accensus, a company building an ultra-low-latency hybrid software/FPGA trading platform. Prior to that he worked at DRW Trading in Chicago in the role of Head of Labs focused on ultra-low-latency technology and before that he worked at STAC in its early days, following almost a decade at Reuters.

Peter Lankford, Founder & Director, Securities Technology Analysis Center. Peter leads STAC®, which provides hands-on technology research and testing tools to the finance industry and facilitates the STAC Benchmark Council™, a group of leading financial institutions and vendors that engages in technical dialog and specifies standard ways to assess technologies used in finance. Prior to STAC, Peter was SVP of the $240M market data technology business at Reuters. Prior to Reuters, Peter held management positions at Citibank, First Chicago Corp., and operating-system maker IGC. Peter has an MBA, Masters in International Relations, and Bachelors in Chemistry from the University of Chicago.

David O'Shea, Financial Services ISV Segment Manager, Intel. David has been with Intel for more than 15 years, the last 10 of which he has served as Intel's ISV lead within financial services. During that time, Intel platforms matured from file and print applications to the mission critical work horse of the industry. David has pioneered Intel's efforts to meet industry requirements for risk management, market data, and high frequency trading. Prior to Intel, David worked with leading software providers on RISC and PC platforms. David is married, has three children, and is an avid golfer with a high handicap.

Eric Powers, VP: Applied Engineering Lead – Global Markets & Equities, Deutsche Bank. Eric has over 19 years of product development and systems integration experience in the financial services industry. At DB, Eric leads the applied engineering team performing systems engineering functions for the Global Markets and Equities business. Subject Matter Expert in precision timing, network monitoring and instrumentation. Eric has a history of advising venture capital firms, leading industry publications and working with numerous startups and established companies on enhancements to existing technologies and new product creation.

Ryan Quick, Principal Architect, eBay/Paypal. Ryan is an expert at scale-out systems, unix kernel design and profiling, and has been recognized for innovation in hardware, application, and object-marshaling. His current work brings machine learning, real-time eventing, set-selection, and digital signal processing technologies to predictive analytics in self-healing and command and control systems. He has focused on distributed systems for the last 20 years, with special attention to the interaction between applications, operating systems, and the hardware and networks underlying them and has experience in government, health-care, financial, educational, manufacturing, and Internet industries. Ryan holds patents for messaging middleware systems, and is a pioneer in bridging High-Performance Computing technologies with enterprise best-practice infrastructure. Ryan received degrees in English and Philosophy from Vanderbilt University, and went on to study American Christian Ethics at Yale Divinity School. Ryan enjoys fishing and birding, and is an active Mormon, involved in religious education and lectures.

Jean Safar, co-founder and CTO, QuartetFS. Jean has a wide expertise in building, designing and managing innovative Object Oriented Distributed trading systems in software houses as well as in Tier I banks. Jean started his career in 1989 at Renaissance Software, a Startup in Californias, as the first developer of Opus, the first Object Oriented Enterprise OTC Trading system. Jean also held the role of VP Head of Architecture for Dresdner Kleinwort Wassertein in Tokyo. Prior to joining Quartet FS, Jean was Director of Professional Services for Grid Computing at DataSynapse

Eric Sedlar, Vice President & Technical Director, Oracle Labs.

Gordon Sissons, Senior Manager, Product Marketing – IBM InfoSphere® BigInsights, IBM. Gord is the Product Marketing Manager for InfoSphere BigInsights at IBM Corporation focused on Hadoop and other analytic workloads. Gord has more than 20 years of experience in product management, high-performance distributed computing and internet technologies. Formerly Director of Technology at Sun Microsystems in Canada, and founder of NeatWorx Web Solutions Inc. Gord has held several roles in IT services and technology consulting throughout his career. Gord is a graduate of Carleton University in Ottawa, Ontario with a degree in Systems and Computer Engineering.

Joe Stein, Founder, Principal Consultant, Big Data Open Source Security LLC. Joe is an Apache Kafka committer and PMC member. A frequent speaker on both Hadoop and Cassandra, Joe is the Founder and Principal Architect of Big Data Open Source Security LLC a professional services and product solutions company. Joe has been a distributed systems developer and architect for over 12 years now having built backend systems that supported over one hundred million unique devices a day processing trillions of events. He blogs and hosts a podcast about Hadoop and related systems at All Things Hadoop.

Glenn Wright, Systems Architect, DataDirect Networks (DDN). Glenn has 20+ years of technical experience in the High Performance Computing industry, working on both server, networking and storage elements of extreme performance solutions. Glenn has been with DataDirect NetworksTM (DDN) for the last 3 years, during which time he has helped clients across HPC industries solve the new "big data" high performance/low latency problems. Glenn has recently focused on technical infrastructure solutions in high capacity analytics for environments with the need for extreme levels of I/O between the storage and server components of the solution. Prior to DDN, Glenn was a senior architect at QLogic (Infiniband group) and also held many varied/international roles at Sun Microsystems.