



STAC Benchmark Council: Recap of 2012

22 January 2013

2012 was another productive year for the STAC Benchmark™ Council. STAC® has highlighted some of the main achievements in this document, with links to further information for members.

Note that some of the links may require you to be logged in (individuals without a login can create one for free). In addition, access to some content may depend on the membership level of your firm. To have information about your firm's membership level sent to your email address, along with contact details for the individuals who manage the Council relationship for your firm, [click here](#) (your email address must be identifiable as a domain belonging to a [Council member firm](#)). If you have confirmed your firm's membership level and are denied access to materials to which you believe your firm is entitled, please [submit a permissions request here](#).

If you would like to discuss your firm's membership level or have any other questions, please contact council@STACresearch.com.

EXECUTIVE SUMMARY

In 2012, membership in the STAC Benchmark Council grew to 249 organizations. 80% of these were trading organizations, with the balance being the vendors who serve them. [Click here to see the list.](#)

The Council held six STAC Summits across New York, London, and Chicago, with the New York and London Summits being expanded, full-day events.

Research in 2012 covered both Fast Workloads (low-latency trading, real-time messaging) and Big Workloads (big data and big compute, such as risk management, tick databases, and unstructured data tasks).

With a theme of "Deeper and Broader", 2012 saw several firsts:

- The first tick-to-trade benchmark
- The first results from the industry standard market-risk technology benchmark
- The first STAC Report of a Hadoop benchmark
- The first STAC Application Performance Boot Camp, a course on how to write high-speed code for low-latency trading and number crunching.

A total of 25 reports were issued during the year (17 of which were available only to members with access to the STAC Vault™, along with other information such as detailed configuration disclosures from benchmark projects). Along the way, we made it easier to access content from the STAC Vault. Content that doesn't require the consent of a third party (such as a disclosing vendor) is automatically available to premium members at www.STACresearch.com/myVault.


The theme for 2013 is: "Bigger. Faster." This represents our continued focus on both big and fast workloads that matter most to members. It also signifies our plan to expand the Council and to produce more information for more members more quickly. We look forward to your input to that process.

Thanks,

Peter Lankford
Founder & Director
STAC

FAST WORKLOADS

Market-interface stacks



- STAC-T1™ debuted in 2012. STAC-T1 is a tick-to-trade test suite that presents a SUT with market data and simulated execution facilities while measuring wire-to-wire latency. STAC developed STAC-T1.EMINI™ using futures data, with the help of the CME and a proprietary trading firm active on the CME. Redline Trading Solutions, HP, and Intel collaborated on [the first published STAC-T1 results](#), marking Redline's transition to a pure-software solution. Three STAC Reports were released (2 public, 1 in STAC Vault). Corvil provided their CorvilNet latency monitor for the tests. STAC-T1 is now up for review by the Council. We plan to [convene a working group](#) in 1Q13.
- The STAC-M1™ working group established a new [NASDAQ TotalView ITCH addendum](#) for STAC-M1 v2, the standard test suite for market data feed handlers. The new addendum defines parameters for using STAC-M1 with a 10 Gbps TVITCH feed and feed handlers that perform market-by-price aggregation to deliver top-N-of-book updates to applications.

Messaging middleware stacks

- Juniper released [STAC-M2™ results of their QFabric system with IBM Low-Latency Messaging software](#). STAC-M2 is a test suite for messaging middleware stacks that is highly sensitive to network latency. This report was the first published STAC-M2 project to utilize a datacenter-scale network fabric as the interconnect, rather than a single switch.
- STAC released v1.2 of the STAC-M2 Advanced Test Harness. The enhancements in this release made it more cost effective to simulate hundreds or thousands of consuming applications. The STAC-M2 Advanced Test Harness is used primarily to perform [lights-out regression testing of Thomson Reuters Enterprise Platform for Real Time and RMDS systems](#).
- STAC released v1.2 of the STAC Pack for RFA. This release enables many mounts per consuming application and supports RFA 7.0.x and 7.2.x.

STAC
SECURITIES TECHNOLOGY ANALYSIS CENTER

Applying STAC-M2 to RMDS

STAC Benchmark Council members can now run industry standard benchmarks as well as customized test cases on their RMDS configurations using two products:

- STAC-M2 Pack for RFA
- Advanced STAC-M2 Test Harness

REDUCE RISK:
Make rigorous performance tests part of your change-control process. Use community-developed testing standards.

IMPROVE PERFORMANCE:
With testing made easy, you can test many more configurations. Change the stack, push a button, get a report. Repeat.

REDUCE COST:
Free your engineers to focus on more valuable tasks than debugging test tools, executing tests, and collating data.

Your RMDS configurations
↓
STAC-M2 Test Harness

The STAC-M2 Pack for RFA enables you to run standard STAC-M2 Benchmarks or customized variants on virtually any RMDS configuration using Marketfeed or OMM. The harness provides complete statistical and graphical analyses of throughput and latency in a variety of scenarios, including normal flow, disconnects, image-request blasts, and three kinds of slow consumer. Results are provided in the familiar STAC Report™ PDF format, as well as formats for analysis in Excel or other tools. The harness is software only. No time sync or capture hardware is required.

The Advanced STAC-M2 Test Harness provides "push-button" automation of STAC Benchmark execution so that (for example) the tests can run as a nightly cron job that leaves a PDF report and Excel files waiting in your inbox.

No human intervention is required to step through scenarios, find max message rates, control producers and consumers, or generate reports.

A custom mode enables you to create your own test cases with an essentially unlimited number of producers and consumers.

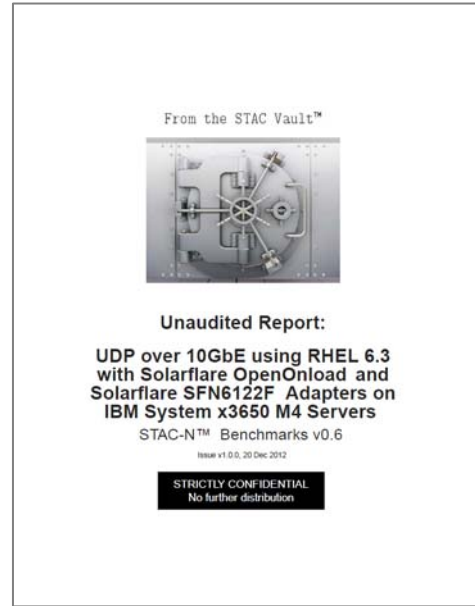
Upgrades included

The STAC-M2 Pack for RFA and Advanced STAC-M2 Test Harness are optional add-on subscriptions to your STAC Benchmark Council Agreement. These subscriptions enable you to upgrade that support new versions of RFA, new versions of supported operating systems, and enhancements to STAC-M2.

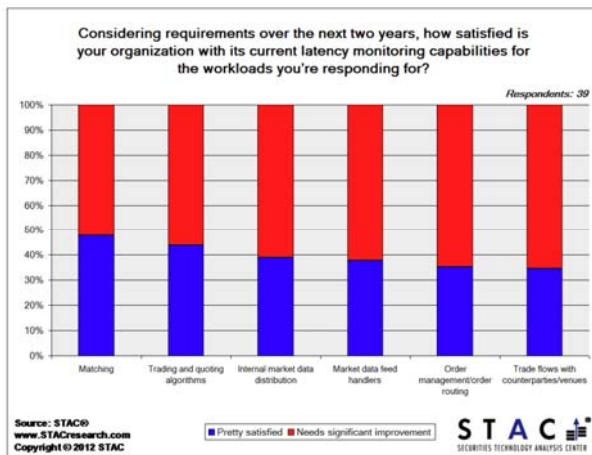
For more information, contact us at:
sales@STACresearch.com

Network stacks

- End-user members of the [STAC Network I/O Special Interest Group](#) specified a list of network stacks that it wished to have tested using the STAC-N™. STAC-N directly tests a network stack without requiring intervening commercial software such as middleware or feed handlers. The SIG's test agenda focused on analyzing the differences in latency at varying message rates for different versions of Linux, in the context of both kernel and kernel-bypass stacks using different vendor hardware. Red Hat played the lead role in providing the data, in cooperation with Solarflare and Mellanox. The results—including [summary slide decks](#), [detailed STAC Reports](#), and [a data file to enable further analysis](#)—are in the STAC Vault. We are working through the final configurations in the list and soliciting the next round of tests of interest to end users. [Click here to get in the loop and give us your wish list](#).
- STAC released v0.6 of the STAC-N test harness. The main enhancement was support for nanosecond precision. The STAC-N harness consists of vendor- and technology-neutral core libraries, along with bindings for UDP, TCP, RDMA, and RDS (including source code). [Click here to request it](#).



Latency monitoring / Time sync survey



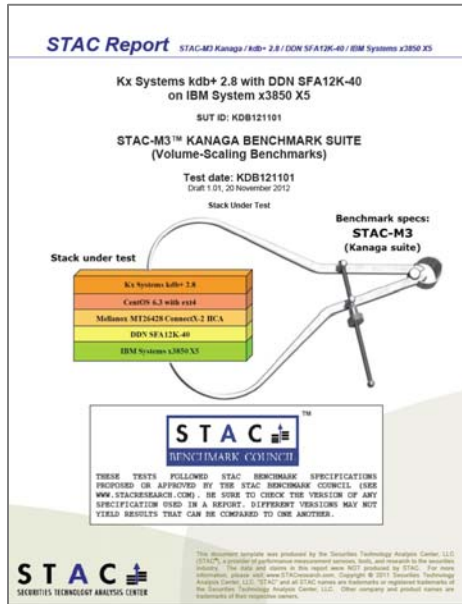
In 2012 STAC issued its [second survey of end-user firms regarding their latency monitoring and time synchronization strategies](#) (available in the STAC Vault). 60 respondents from prop shops, exchanges, hedge funds, and brokers answered business and technical questions regarding their requirements, plans, and viewpoints. As with the first survey two years earlier, the results proved to be of keen interest to many members. While the lack of satisfaction that firms have with their current solutions hasn't changed, what they plan to do about it has. Budgets and build-vs-buy perspectives have evolved. So have views on required accuracy (hence latency tolerances), as well as on key technical strategies.

Training

STAC began offering [training in how to code applications for low latency](#). See the STAC Application Performance Bootcamp, below.

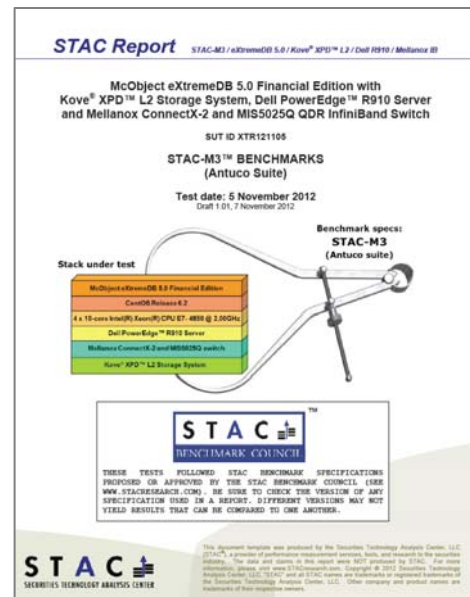
BIG WORKLOADS

Tick database stacks



- STAC-M3™ results continued to roll in (see below). [STAC-M3](#) is the industry standard benchmark suite for tick database stacks. It consists of a baseline suite (Antuco) and an optional scaling suite (Kanaga).
- STAC delivered proposed enhancements to the [STAC-M3 Kanaga suite](#). Kanaga now tests how well a solution scales with the number of users, as well as with the volume of data. As before, Data Direct Networks (DDN) was the first to demonstrate performance at scale with STAC-M3 Kanaga.
- There were several tests of new storage systems using the kdb+ database from Kx Systems, including products from Violin Memory, Texas Memory Systems, ION Computer Systems, Veritas, and DDN. Some of these reports were [public](#), while some are only in the [STAC Vault](#). Note that in 2012 we removed the system configuration details from public STAC Reports and put them into separate Configuration Disclosures available in the STAC Vault.

- In November of 2012, McObject became the [first database vendor other than Kx Systems to publish STAC-M3 results](#). Until then, Kx was the only software vendor to publish results. McObject released results for their eXtremeDB embedded database product deployed on Dell servers using DRAM-based storage from Kove.
- Both the [STAC-M3 Pack for kdb+](#) and the [STAC-M3 Pack for eXtremeDB](#) are available to premium members of the Council. These STAC Packs (together with the STAC-M3 Test Harness materials) include the software necessary to test and analyze a stack using kdb+ or eXtremeDB, respectively.



Bi-temporal data stacks

- The [STAC Bi-Temporal Data \(BTD\) SIG](#) looks at issues related to datasets where changes are recorded over two independent dimensions of time ("valid time" and "transaction time"), which can be a challenge with traditional databases. The SIG proceeded along two tracks: an information stream and a benchmark stream.
- The information stream enjoyed discussions by industry experts such as Hugh Darwen, a pioneer in bitemporal RDBMS. [Click here to view replays of those discussions.](#)
- [Phase 1 performance and functional benchmarks](#) have been outlined, but implementation has proceeded slowly due to job changes among SIG members.

Risk management stacks



- The [STAC-A2™ benchmark specifications](#) are nearly complete. STAC-A2 is the standard benchmark suite for technology stacks that can be used to perform market risk calculations (specifically, Monte Carlo-based derivation of options Greeks using a Heston model). In 2012, the working group added multiple scaling tests, enabling the suite to test how well a solution scales as the workload increases (number of assets, number of paths) and as the solution scales (number of cores, boards, servers, etc.).
- Intel was the first to implement STAC-A2 and [presented preliminary results](#) at several STAC Summits.
- STAC, Intel, and Catskills Research were invited to present a [paper on STAC-A2 at the 5th annual Workshop in High-Performance Computing for Finance at SuperComputing 2012 \(SC12\)](#).

Hadoop


- STAC released its [first benchmark of Hadoop MapReduce stacks](#). In this project, we compared Apache Hadoop to Apache Hadoop accelerated by IBM Platform Symphony Advanced Edition in the same hardware environment. We chose Berkeley's SWIM tools for this, along with a set of MapReduce jobs based on production Hadoop traces from Facebook in 2010.

Big Data SIG

- In 2012, we determined that there was sufficient demand to start up a Special Interest Group to examine issues related to "big data" (data management tasks made challenging by virtue of some combination of volume, variety, velocity, and value). The Big Data SIG is expected to discuss these challenges as well as new technologies with potential to help (be they NoSQL, NewSQL, Not-Only-SQL, or whatever). The SIG will also define benchmark specifications where appropriate and discuss benchmark results. It is expected to launch in the first quarter of 2013. If you're interested in participating, [please let us know](#).
- At the Fall 2012 STAC Summits, Calxeda outlined [proposals for several Big Data benchmarks](#). These included analytic workloads whose inputs range from highly structured data to unstructured (schema-less/schema-last) content common in trading organizations.

Training

STAC began offering [training in how to code compute-intensive applications for high-throughput](#). See the STAC Application Performance Bootcamp, below.



OUTREACH

STAC Summits

In addition to numerous working group and SIG meetings, the Council held six conferences in 2012, spread across New York, London, and Chicago. New York and London were full-day affairs covering both Big Workloads and Fast Workloads. [Click here to view slides and videos](#) (where available) from these conferences, or navigate via the Appendix, which contains a list of topics discussed at these meetings.



STAC Application Performance Boot Camp

- In 2012, STAC introduced [a training course on modern computing architectures and how to get maximum application performance from them in a trading environment](#). The STAC Application Performance Boot Camp debuted in Chicago, then went on to New York and London. The two-day Boot Camp was split into two parts. The first provided an overview of modern hardware systems, aimed at developers, quants, system engineers, and managers. The second day was an intensive, code-driven course for application developers and quants who want to improve the performance engineering of their apps.
- 130 individuals attended this training. We were gratified by feedback like:
 - "I learned some new things and have been extending on that myself, and applying certain things in our production systems...the payoff is very visible to me..."
 - "Extremely knowledgeable presenter, able to field any question."
 - "Deep technical detail - War stories and real examples"
 - "The content, tips and tricks taught in class were very helpful."
 - "No brainer to recommend the course to colleagues and friends in this area"
- We plan to repeat the STAC Application Performance Boot Camp in 2013 and may add more courses of interest. [Let us know if you're interested, or give us your thoughts on additional training that would be useful.](#)



2013: Bigger, Faster

2013 promises to be another exciting year for capital markets technology. Trading organizations continue to compete through technology, even as they face budgetary and regulatory headwinds—a combination that creates even more demand for innovation.

Vendors seem ready to respond. From the roadmaps we've seen, 2013 will be interesting indeed. CPUs are advancing in terms of both speed and power efficiency. At the same time, servers and application software are getting better at exploiting those processors. Non-traditional processors such as GPUs and FPGAs are also advancing, while an increasing number of players make those processors accessible to a broader audience through abstractions, partnerships, and pre-packaged solutions. New distributed data management software is challenging traditional IT economics. Storage and network providers are continuing to push the innovation envelope, with many of them not only improving component-level performance but also taking a more holistic look across the solution stack at specific workloads. And as in other industries, more software and infrastructure is being offered as a service.

Most of these new propositions can be tested using end-user specified standards that the STAC Benchmark Council has already developed or will bring out in 2013. So if you're an end user or vendor interested in running STAC Benchmarks yourselves, we're happy to help. Or if you're an end user who would like a vendor to use these tests to back up their claims, please send them our way or let us know to get in touch with them.

STAC's mission for the year is "Bigger, Faster". Expect to see more research on both Big Workloads and Fast Workloads as new products emerge. In addition to research on cutting edge technologies with potential to shift paradigms, expect to see more research from us on the incremental—but nevertheless crucial—updates and upgrades to well-established technology stacks that are in common use at trading firms today. Many of you have told us that you appreciate having to do less of this routine testing yourselves and also like having another set of results you can compare your own results to.

Meanwhile, stay tuned for announcements on new STAC Reports, the next round of STAC Summits, more training courses, and SIG and working group meetings. And if there's anything we don't seem to be focusing on that you think is important, please let us know. The Council's agenda is driven by its members.



APPENDIX: Topics at 2012 STAC Summits, with Links

(Note: If you are not permitted to access the links below, please [submit a request](#) to have your login associated with your firm's membership.)

Business context

- Panel: "Technology and Career Trends." NYC: [Steven J. Sadoff, Knight Capital Group](#); [Lyle Hayhurst, DRW Trading](#). London: Tim Lipscomb, Bank of America Merrill Lynch; Dr. Tony Chau, UBS; Dominic Connor, P&D Quant Recruitment.
- Panel: "Quality Control of Trading Algorithms." NYC: ["Post Card from the Future – The Very Visible Hand." Thomas Chippas, Barclays](#); ["Introducing AT 9000", Greg Wood, Deutsche Bank](#). Chicago: Peter Nabicht, Allston Trading; [Ben Van Vliet, IIT](#). London: Paul Marks, Citi Futures and OTC Clearing; Sam Tyfield, Katten Muchin, Rosenman UK LLP; Blake Stephenson, Futures and Options Association; Paul Willis, ABN AMRO Clearing.
- "Post Card from the Future – The Very Visible Hand." London: [David Litner, Barclays](#).
- ["Postcard from the Future – How Real-time Big Data Analytics Changed Financial Markets."](#) NYC: Marc Andrews, IBM.
- Panel: "Infrastructure Outsourcing." NYC: [Jerry Nelligan, Integral Derivatives](#); [Nigel Kneafsey, OptionsIT](#); [Bill Ruvo, Thomson Reuters](#); [Michael Dunne, ACTIV Financial Systems](#)

Application development

- Panel: "Gloves Off: Top developer/technologists go toe-to-toe," NYC: [Jeffrey M. Birnbaum, 60East Technologies](#); [Niall Dalton, Calxeda](#).
- Panel: "Java In Low-Latency Trading." NYC: [Ryan Eavy, CME Group](#); [Gil Tene, Azul Systems](#); [Cameron Purdy, Oracle](#). London: Martin Thompson, LMAX; Peter Lawrey; Guillermo Lopez-Taboada, University of A Corunna, Spain; Cameron Purdy, Oracle. Chicago: Jesse Fugitt, Informatica; Gil Tene, Azul Systems.
- "Parallel Programming Case Study," NYC: [Robert Geva, Intel](#). London: [Stephen Blair-Chappell, Intel](#).
- ["Advances In Number Crunching,"](#) NYC: Rick Carlin, Intel. London: Ian Lloyd, Intel.
- ["Detecting Memory Errors, Leaks and Concurrency Violations in Kernel Modules and C/C++ Applications with Deterministic Precision and Zero False Positives,"](#) Chicago: Himanshu Shukla, Parallocity.
- ["Tools and Libraries for the Intel® Xeon Phi™,"](#) NYC & London: Scott Lasica, Rogue Wave. Chicago: Stephane Reynaud, Rogue Wave.
- ["Improving Single-Thread Performance with Cilk\(TM\) Plus."](#) Chicago: Robert Geva, Intel.

Low-latency infrastructure

- Panel: "Not Your Father's Network." NYC: [Pramod Srivatsa, Cisco](#); [Asaf Wachtel, Mellanox](#); [Glenn Dasmalchi, Juniper Networks](#); [Nick Ciarleglio, Arista](#). London: Colin Constable, Juniper Networks; Gordon Hirst, Cisco; Doron Arad, Mellanox; Paul Goodridge, Arista. Chicago: Glenn Dasmalchi, Juniper Networks; Will Ochandarena, Cisco; Doron Arad, Mellanox; Nick Ciarleglio, Arista.
- Panel: "Nearer to 'c' than Thee? A discussion on low-latency wireless." NYC: [S. Jay Lawrence, NeXXCom](#); [Gene Callahan, CCSi](#); [Ed Kopko, ULL Networks](#); [Mike Schonberg, Quincy Data](#). Chicago: S. Jay Lawrence, NeXXCom; Gene Callahan, CCSi; Mike Schonberg, Quincy Data. London: S. Jay Lawrence, NeXXCom; Ed Kopko, ULL Networks; Stephane Tyc, Quincy Data.

- Technical Briefing: "[The Sandy Bridge Difference for Tick-to-Trade Latency](#)," NYC: Daryan Dehghanpisheh, Intel. London: David Barrand, Intel. Chicago: [David O'Shea, Intel](#).
- Point of View: "[Sockets Acceleration in 2013](#)," NYC: Steve Pope, Solarflare Communications. London: Martin Porter, Solarflare Communications.
- "[Optimizing ProLiant Gen8 Systems for Ultra Low Latency](#)," NYC, London & Chicago: Lee Fisher, HP.
- "[Cisco Innovations for High-Frequency Trading Workloads](#)," NYC: Pramod Srivatsa, Cisco. London: Gordon Hirst, Cisco. Chicago: Will Ochandarena, Cisco.
- "[A Faster Trading Infrastructure](#)," NYC & Chicago: Jeffrey Margolis, Mellanox Technologies. London: Asaf Wachtel, Mellanox.
- "[Low latency, congestion avoidance enterprise networks](#)," NYC: Bob Fernander, Gnodal. Chicago: Jim Preasmyer, Gnodal. London: Fred Homewood, Gnodal.
- "[Layer 3 -- How not to do it](#)," Chicago: David Snowdon, Zeptonics.
- "[Arista Innovations in Low Latency platforms](#)," NYC: Ashwin Kohli, Arista. Chicago: Darrin Machay, Arista. London: Alex Nichol, Arista.
- "[Gaining a Competitive Advantage with the Cisco Nexus 3548](#)," NYC: Lucien Avramov, Cisco. Chicago: Nimish Desai, Cisco. London: Will Ochandarena, Cisco.
- "[PHY latency. Breakdown and why it matters](#)," Chicago: Nikolaj Hermann, Fiberblaze.
- "[Optimizing Network Performance and Enabling Algorithm Refinements with FastStack™ DBL™ coupled with Sniffer10G™](#)," Chicago: Brian Grant, Emulex. London: Chris Hill, Emulex.

Hardware acceleration

- Panel: "Accelerator Boards: Making Hardware Softer or Software Harder?" NYC: [Bruce Tolley, Solarflare](#); [Nikolaj Hermann, Fiberblaze](#); [John Lockwood, Algo-Logic](#); [Matthew Knight, Accensus](#); [Bob Doud, Titera](#). London: [David Riddoch, Solarflare](#); [Nikolaj Hermann, Fiberblaze](#); [Mohammad Darwish, AdvancedIO](#). Chicago: [Bruce Tolley, Solarflare](#); [Matthew Knight, Accensus](#); [Mohammad Darwish, AdvancedIO](#); [Kelly Masood, Intilop](#).
- "[Paradigm Shift: Accelerating Financial Applications Via The Programmable Network](#)," NYC & Chicago: Nick Ciarleglio, Arista Networks.
- "[New Techniques and Product Architectures for Resilient Microwave Links: A Use Case for the Solarflare ApplicationOnload\(™\) Engine](#)," Chicago: Dave Parry, Solarflare Communications.
- "[Less is More: Hardware Accelerated Filtering for Optimal Software Performance](#)," NYC & Chicago: Mohammad Darwish, AdvancedIO Systems.

Low-latency software

- "[Redline & Intel: Optimizing the Tick-to-Trade Path with Mainstream Technology](#)," NYC: AI Maillet, Redline Trading Solutions. Chicago: Chris Karpinsky, Redline Trading Solutions.
- "[Achieving true determinism for market data normalization, book building and distribution](#)," NYC & Chicago: Olivier Baetz, NovaSparks. London: Matthew Gadd; NovaSparks.
- "[The Power of Real Time Monitoring With MemSQL](#)," NYC & Chicago: Eric Frenkiel, MemSQL.
- "[Introducing AMPS and 60East Technologies](#)," Chicago: Brand Hunt, 60East Technologies.
- "[Innovations in High Performance Messaging](#)," NYC: Roderick Clemente, Informatica. London: James Andrews, Informatica. Chicago: Bob van Valzah, Informatica.
- "[Performance Like No Other](#)," NYC & Chicago: Bill McLane, TIBCO. London: Basil Kajcovski, TIBCO.

- "[Accelerating Fast Workloads throughout the Front, Middle and Back Office](#)," Chicago: Bill Romano, Solace Systems.
- "[Advancements in High-Performance Enterprise Messaging](#)," NYC & Chicago: Bill McLane, TIBCO. London: Vasil Kajcovski, TIBCO.
- "[Scalable Low Latency Infrastructure](#)," Chicago: Amir Halfon, Oracle.

Time sync and latency monitoring

- "[Delivering and verifying Sub microsecond time at the Linux application layer](#)," NYC & Chicago: Paul Skoog, Symmetricom.
- "[Delivering Sub-Microsecond Accurate Time to Linux Applications Around the World](#)," London: Paul Skoog, Symmetricom.
- "[WAN PTP – the good, the bad and the ugly](#)," NYC & London: Henry Young, TS-Associates.
- "[Indoor GPS Timing Sync Solution](#)," NYC: Jeremy Onyan, Spectracom.
- "[Low Latency to High IOPS: Wire to Storage Solutions](#)," NYC: Mike Spencer, Emulex. London: Steve Perkins, Emulex. Chicago: Brian Grant, Emulex.
- "[Accessible Analytics - Building on CorvillNet](#)," Chicago: Raymond Russell, Corvil.
- "[Network microbursts and switch queues: the hidden truth](#)," NYC & London: Rony Kay, cPacket.

Big Data approaches

- Panel: "Breaking 'Big'". NYC: [Vinod Kutty, CME Group](#); [Dino Vitale, Morgan Stanley](#); [Lokesh Khosla, ParAccel](#). London: [Lee Pollington, Marklogic](#).
- Panel: "Getting Smarter Faster: Building better platforms for research and back-testing." NYC: [Niall Dalton](#); [Bob Gaines, DataDirect Networks](#); [Asaf Wachtel, Mellanox](#); [Gerald A. Hanweck, Jr., Hanweck Associates](#). London: Philip Beasley-Harling, Bank of America Merrill Lynch; James Coomer, DataDirect Networks; Doron Arad, Mellanox.
- "Platforms for research and backtesting." Chicago: [Niall Dalton](#).
- Data Management Frontiers: "[Scaling Graph Analytics](#)," NYC: Venkat Krishnamurthy, YarcData.
- Data Management Frontiers: "[Fast Data Load and Analysis with a Distributed SQL System](#)," NYC: Nikita Shamgunov, MemSQL.
- Data Management Frontiers: "[Integration and Transformation of Front Office, High-Speed Applications and Market Data](#)," NYC: Ronen Schwartz, Informatica. London: Charles Hughes, Informatica.
- "[The MemSQL Launch Announcement](#)," NYC: Eric Frenkiel, MemSQL.
- "[The Shortest Path to the Fastest Data Management](#)," NYC: Christopher Mureen, McObject.
- "[Big Data meets Fast Data](#)," NYC: Mike Stolz, VMware Inc.
- "[Taming Big Data with SAP](#)," NYC: Mark Mumy, Sybase.
- "[Everything You Hear About Big Data Is Wrong](#)," NYC: Tom Deutsch, IBM.
- "[Modeling the Big Data Beast](#)," NYC: Niall Dalton, Calxeda.
- "[Benchmark: Comparison of Two MapReduce Implementations](#)." London: Simon Waterer, IBM; Peter Lankford, STAC.

Big Data Infrastructure

- Technical Briefing: "[Opening up the I/O bottleneck in backtesting](#)." Chicago: Bob Gaines, DataDirect Networks.
- "[Real world solutions for high-scale market data analytics](#)," NYC: Keith Miller, DataDirect Networks. Chicago: Glenn Wright, DataDirect Networks. London: Glenn Wright, DataDirect Networks.
- "[Advances in scale-out computing and clustered file systems](#)," Chicago: Joshua Blumert, IBM.
- "[Accelerating Financial Analytics and Simulation with Parallel I/O](#)," London: Derek Burke, Panasas.
- "[Managing Distributed Big-Workloads across LAN, WAN and Web](#)," NYC: Shawn McAllister, Solace Systems. London: Benjamin Taieb, Solace Systems.
- "[Arista Innovations in Big Data](#)," NYC: Ashwin Kohli, Arista. London: Alex Nichol, Arista.
- "[Accelerating Performance by Transforming the way compute-intensive applications use network data](#)," NYC: Bruce Tolley, Solarflare. London: David Riddoch, Solarflare.
- "[5 Minutes on Big Data Network Designs and Experiences](#)," NYC: Ashwin Kohli, Arista.
- "[The Challenges that Big Data Presents for your Network](#)," London: Alex Nichol, Arista.
- "[Accelerating Big Workloads with WAN-Optimized Data Grids](#)," NYC: Ken Overton, Solace Systems. London: Benjamin Taieb, Solace Systems.