



## **STAC Benchmark<sup>TM</sup> Council: Recap of 2011**

25 January 2012

2011 was an active year for the STAC Benchmark Council. To help you keep track, STAC<sup>®</sup> has captured key points 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 ensure your login has your firm's permissions, please go to [www.STACresearch.com/permissions](http://www.STACresearch.com/permissions).

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

## EXECUTIVE SUMMARY

In 2011, membership reached 222 organizations. Three quarters of these were trading organizations, with the rest being the vendors who serve them. A total of 41 reports were issued during the year, with 35 reports available only to members with access to the STAC Vault™. The Council held 8 conferences across New York, London, and Chicago.

2011 was marked by intense activity in Fast Workloads (low-latency trading, real-time messaging) as well as the ramp-up of a new stream on Big Workloads (big data and big compute, such as risk management, tick databases, and various unstructured data tasks). We use these terms throughout this recap.

Plans for 2012 include ongoing tool development and research in Fast Workloads, particularly oriented around the new generation of processors, network interfaces, and switches, as well as additional tests incorporating trade execution. In Big Workloads, we plan to pay attention to large back-testing tasks, to put into practice new benchmarks based on market- and credit-risk management, and to catalyze research and discussion in areas of keen interest to customers such as bi-temporal data and NoSQL.

# STAC Benchmark Council members

### Trading organizations

- ABN AMRO Clearing
- AHL
- Algo Engineering
- AlgoPartners LLC
- Alston Trading
- Archelon Group
- Aristas Securities
- Aspect Capital
- Athena Capital Research
- Automat Limited
- Bank of America/Merrill Lynch
- Barclays Capital
- BATS Exchange
- BAXTER-FX
- Belvedere Trading, LLC
- BGC Partners
- Bds Quant
- Blue Capital Group
- Bluefin Trading
- Bluefire Capital
- BMO Capital Markets
- BNP Paribas
- BNY Coverex
- Boston Options Exchange
- BTG Pactual
- Buttonwood Group Trading
- Cantor Fitzgerald
- Capital Fund Management
- Capstone Holdings Group
- Caxton
- Celeritas Markets LLC
- Chicago Board Options Exchange
- Chicago Trading Company
- Citi
- Citadel Investment Group
- CityIndex Ltd
- CMC Markets
- CME Group
- CMT Asia
- CNSX Markets Inc.
- Cognitive Capital LLC
- Commerzbank AG
- Crab Capital Management LLC
- Credit Suisse
- Daiwa Institute of Research
- DE Shaw
- Deutsche Bank
- Deutsche Boerse
- DRW Trading
- EchoTrade
- Eclipse Options (HQ) Ltd
- EdgeCI
- E-Trade Financial
- Eurex Frankfurt AG
- FF&F&F
- Fidelity Investments
- Final
- Flow Traders US
- Forex Capital Markets (FXCM)
- Fox River Execution
- FTSE Group
- FXAll
- Gain Capital Group, LLC
- Gelber Group
- Getco LLC
- GFI Group
- Gloucester Research Ltd
- Goldman Sachs
- GSA Capital
- Hard Eight Futures LLC
- High Speed Analytics
- HSBC
- ICAP Services North America
- IG Index
- IMC Financial Markets
- ING
- Integral Derivatives
- Intercontinental Exchange (ICE)
- International Algorithmic Trading GmbH
- ITG
- ITG/IDE Trading Corporation
- JP Morgan Chase
- Jump Trading
- Kapital Trading
- Kinneipf Capital Management
- Knight Capital Europe Limited
- Kohl Trading
- Komaco, LLC
- Kyte Group
- Lime Brokerage, LLC
- Liquid Capital Group
- Liquidnet
- Lloyds Banking Group
- LMAX
- London Metals Exchange Ltd
- Macquarie Group Limited
- Marquette Partners
- Marshall Stone LFP
- Matlock Capital, LLC
- Maven Securities Limited
- MF Global
- Millennium Partners
- Morgan Stanley
- MWD Trading
- Myra Trading LLC
- NASDAQ OMX
- National Stock Exchange (NYSE)
- National Commodity and Derivatives Exchange Ltd
- Wells Fargo
- Nomura
- Nobilis Capital
- NYSE Euronext
- Optiver B.V.
- Ortus Capital Management
- Panther Energy Trading LLC
- Peaks Investments LP
- Peregrine Financial Group
- Pershing LLC
- Pineau Investment Group, LLC
- Platform Capital LLC
- Puppetmaster Trading
- Quantlab Financial LLC
- Questrade
- Quiet Light Securities LLC
- RBC Capital Markets
- Realitiq LLC
- RGM Advisors LLC
- RiverCross
- ROFIX
- Ronin Capital
- Rooftop Trading
- Rotella Capital Management
- Royal Bank of Scotland
- SAC Capital Advisors
- Schneider Trading Associates Ltd
- Societe Generale
- Source Capital AG
- Speed Trading Group
- Spot Trading, LLC
- Standard Bank
- Standard Chartered
- Storm Loan Marketplace
- Sunningdale Capital LLP
- Sun Trading LLC
- SurfaExchange
- Swedbank
- Susquehanna International Group
- TD Ameritrade
- Tenth Sense Group Pty Ltd
- Teza Technologies LLC
- Tibra Capital
- TIK Group
- TradeForecaster LLC
- TradeLink LLC
- TradeStation Technologies
- Tradeweb LLC
- Trading Machines
- Tradition North America
- Transmarket Group LLC
- Tullett Prebon
- Two Sigma Investments
- UBS
- Volant Trading
- Warsaw Stock Exchange
- Wellington Management LLP
- Wells Fargo
- WH Trading
- Wolverine Trading
- XHedge LLC
- Xpert Financial Group

Council members do not necessarily endorse or approve of any of STAC's products, specifications or statements.  
 All company names and logos are trademarks or registered trademarks of their respective owners.

## Fast Workloads

### *Feed handler stacks*

- In 2011, a newly reconvened STAC-M1™ working group reviewed proposed enhancements to the STAC-M1 Benchmark specifications. The result was STAC-M1 v2, which provides a framework for specifying specifies benchmark based on additional datafeeds and defines interfaces for integration with the latency-monitoring products that many trading firms have already deployed. TS-Associates was the first monitoring vendor to integrate their products (TipOff and Application Tap) with STAC-M1 v2.
- SR Labs was the first to undergo STAC-M1 v2 benchmarking, using NASDAQ TotalView ITCH. The SUT included their MIPS software, an IBM server, and Myricom NICs.

### *Messaging middleware stacks*

- TIBCO, RTI, and IBM all [published STAC-M2™ results](#) in 2011. Most of these were on 10 GbE, using Juniper, IBM (BNT), and Arista switches, along with Solarflare and Mellanox NICs. Two of the benchmarks used RoCE. IBM was the most active server provider in these benchmarks, but Dell also made their debut in the low-latency club. All of these benchmarks were on Red Hat.
- There were no revisions to the [STAC-M2 Benchmark specifications](#) for messaging middleware in 2011, as the focus was on maintaining stability to preserve historical comparisons. In 2012, we will call for interest in revisions to the STAC-M2 standard.
- STAC delivered significant tools enhancements in 2011. For example:
  - Version 1.1 of the STAC-M2 Advanced Test Harness added significant customization features and the ability to run unattended tests on systems of arbitrarily large scale. These enhancements were driven by firms using this harness with the STAC-M2 Pack for RFA to perform [lights-out regression testing of their RMDS systems](#).
  - STAC-M2 was privately ported to new platforms to support members' internal research. Some of these ports may be made available as product in 2012.

### *Network stacks*

- In late 2011, STAC produced "[STAC Network API Study A](#)," a first step in responding to demand from Council members for additional research and tools that focus directly on the network stack without requiring intervening commercial software such as middleware or feed handlers. Collaborating with Mellanox, Solarflare, and the University of New Hampshire, we developed new test-harness software and then used it to study the differences between RDMA, Sockets, and RDS and their implementations on QDR InfiniBand and 10 Gigabit Ethernet. This study explored the tradeoffs in performance, programming complexity, and resource utilization of these APIs. It yielded a summary report, 23 individual STAC Reports, and a data file to enable further analysis.
- We dubbed the test methodology used above STAC-N™, with a v0.5 designation." The test harness itself is vendor neutral and technology neutral but comes with bindings for UDP, TCP, RDMA, and RDS (including source code). To request it, see [www.STACresearch.com/network](http://www.STACresearch.com/network).
- We put out the call for a new special interest group called the STAC Network I/O SIG. In addition to providing feedback on the STAC Network API Study A and determining priorities for further research and tool development, this SIG will discuss issues related to optimizing applications and network stacks for low latency and high throughput. 25 people from 15 firms have signed up so far. To express your interest, go to [www.STACresearch.com/network](http://www.STACresearch.com/network).

## Big Workloads

### *Tick database stacks*

- The STAC-M3™ Working Group delivered a [complete baseline benchmark suite \(Antuco\)](#), which presents a variety of I/O-intensive and compute-intensive workloads based on pre- and post-trade analysis of US equities data. IBM were the first to publish results against STAC-M3 Antuco. HP, Kove, DDN, and Ion Computer also [released baseline STAC-M3 results](#) of systems using kdb+. Over the course of the year, [STAC-M3 results showed dramatic improvement](#) as new technologies were brought to bear and as the industry learned more about optimizing for these tick database workloads.
- The working group also created an optional set of volume-scaling tests called the STAC-M3 Kanaga suite, which essentially extend to an arbitrary size the STAC-M3 Antuco dataset and a few of its queries. DDN was the first vendor to publish [results of the STAC-M3 Kanaga suite](#), demonstrating performance from just ~3 TB to ~26 TB.
- The working group even managed to create one more suite: [STAC-M3 Vesuvius](#). Like STAC-M3 Kanaga, this suite is aimed at illuminating how solutions handle large datasets, anywhere from ~10 TB to ~160 TB. STAC-M3 Vesuvius is motivated by algorithm back-testing and draws from client experience in the foreign exchange markets, where data volumes are exploding. The first projects involving STAC-M3 Vesuvius are in swing now. Once we have some experience with this suite under our belt, the working group will need to decide whether two volume-scaling workloads are really necessary.
- Kx Systems provided updates to the [STAC-M3 Pack for kdb+](#), covering the Antuco, Kanaga, and Vesuvius suites. This well-scripted collection of q code enables a user to easily run the STAC-M3 Benchmarks on in-house kdb+ configurations, for exact comparison to STAC Reports available publicly or in the STAC Vault.

### *Risk management stacks*

- The [STAC-A2™ Working Group](#) started up in 2011 for the purpose of defining technology benchmarks based on risk-management workloads. The group scoped out high-level approaches to both market- and credit-risk workloads, then detailed the market risk benchmarks first. The suite deconstructs typical risk management functions to analyze performance by component calculations, then looks at them end-to-end in an arbitrarily scalable way. As with all STAC Benchmark specs, the STAC-A2 Benchmark suite is technology agnostic, with three vendors of processors currently writing implementations. As STAC-A2 has come to be identified with market risk, we will likely spin off the credit-risk related benchmarks as STAC-A3™ in 2012.

### *Bi-temporal data stacks*

- Toward the end of the year, we launched the [STAC Bi-Temporal Data \(BTD\) SIG](#). A dataset is bi-temporal if changes are recorded over two independent dimensions of time: "valid time" (the period during which a fact is true) and "transaction time" (the period during which a fact is stored in a database). Trading firms use BTD in areas such as risk management, back-testing, P&L Explain, and regulatory reporting. Common challenges with today's BTD solutions include high costs, poor performance, difficult maintenance, and complexity in areas such as derived data and schema migration. But vendors are introducing promising new BTD solutions. The STAC BTD SIG is a group of end-user organizations and vendors that meets periodically to understand and influence the future of BTD. 20 people from 12 firms participated in the inaugural (end-user-only) in-person meetings in New York and London.

## Events

In addition to numerous working group and SIG meetings, the Council had 8 [conferences](#) in 2011, spread across New York, London, and Chicago.

The June round included STAC Performance Summits™ (focused on fast workloads) in New York, London, and Chicago, plus the first-ever STAC Analytics Technology Conference (focused on big workloads) in New York.

The October/November round featured STAC Performance Summits in London and Chicago, the first-ever STAC Analytics Technology Conference in London, and a combined STAC Technology Conference in New York that covered both fast workloads and big workloads.

Appendix A contains a list of topics discussed at these meetings. For those topics you missed, links to slides (and video, where available) are included.

## 2012: Deeper and Broader

We have carried the 2011 momentum into the new year, which is off to a brisk start. Vendor roadmaps include a lot of innovative processors, storage, networking, servers, and software, and projects are active in nearly every benchmark domain. We will continue to help the Council define new benchmark specifications to respond to the latest customer needs. And the newly created Special Interest Groups look like they have interesting agendas for the year.

Meanwhile, new initiatives are also in progress. For example, in the Fast Workloads area, we are completing work on the first STAC-T1 installation. [STAC-T1™ is a tick-to-trade test harness](#) that presents a SUT with market data and simulated execution facilities and measures wire-to-wire performance (think of it as the front of STAC-M1 married to the back of STAC-E1™). As with STAC-N, we have taken a tools-first approach in this case. That is, we are enabling Council members to get experience with the customizable test software before convening them to formulate standard benchmarks that use it. We have also just issued our second [survey of end-user firms regarding their latency monitoring and time synchronization strategies and requirements](#). The results of the first such survey, performed two years ago, have proven to be of keen interest to many members.

In the Big Workloads area, we expect to release the first STAC-A2 (market risk computation) results while building out the credit risk benchmarks in parallel. We are also picking up significant interest in a number of new “big data” workloads that are potentially amenable to new software such as analytic databases, Hadoop, and other No SQL approaches, as well as new server, storage, and networking designs. If you're interested in participating in a potential Big Data SIG, [please let us know](#).

Meanwhile, stay tuned for announcements on the next round of conferences. And if there's anything we don't seem to be focusing on that you think is important, please let us know.

Sincerely,

Peter Lankford  
Founder & Director  
STAC

## Appendix: Topics at 2011 STAC Conferences

(Note: This content is reserved for member firms. To ensure your login has your firm's permissions, please go to [www.STACresearch.com/permissions](http://www.STACresearch.com/permissions).)

### Fast workloads context

- ["Global Market Structure Evolution and its Impact on Technology,"](#) by Thomas Chippas, Barclays Capital.
- ["European Market Structure Evolution and its Impact on Technology,"](#) by Rishi Karir, Barclays Capital.
- Discussion: "What's On the Minds of Chicago Trading Technologists," with Doug Puenner, Blue Capital Group, and Craig Mohan, CME.

### Low-latency app dev

- Panel: "Optimizing x86 Apps." NY: [Zach Zimmerer, Barclays Capital](#); [Niall Dalton, Cantor Fitzgerald](#); [Michael McCool, Intel](#); [Peter Godman, Corensic](#); [Stephane Raynaud, Rogue Wave](#). London: David Litner, Barclays Capital; Stephen Blair-Chappell, Intel; Peter Godman, Corensic; Scott Lasica, Rogue Wave Software. Chicago: Michael D'Mello, Intel; Stephane Raynaud, RogueWave Software, Peter Godman, Corensic.
- ["Intel's Next-Generation of Instructions - AVX, AVX2, and FMA,"](#) NY: Mark Davis, Intel. London: Evgueny Khartchenko, Intel. Chicago: Arch Robison, Intel.
- ["Identifying Kernel-Mode Data Races,"](#) by Peter Godman, CEO of Corensic.
- ["Squashing bugs: Innovative techniques for parallel code,"](#) NY: Stephane Raynaud, Rogue Wave. London: Dean Stewart, Rogue Wave. Chicago: Scott Lasica, Rogue Wave
- ["How to detect and fix hard-to-find concurrency errors in multi-threaded systems,"](#) by Eric Scollard, Corensic.
- ["The Performance Benefits of Optimizing Processor Cache,"](#) by Stephane Raynaud, Rogue Wave
- ["Hybrid Instrumentation,"](#) by Henry Young, TS-Associates.
- [Technical Brief: "Using OFED APIs,"](#) by Richard Croucher, Informatix Solutions.
- ["STAC Network Research with RDMA, Sockets, and RDS,"](#) by Peter Lankford, STAC.

### Low-latency infrastructure

- Panel: "Navigating Low-Latency Network Choices." NY: [Asaf Somekh, Mellanox Technologies](#); [John Moore, Juniper Networks](#); [Vinit Jain, IBM Systems Networking](#); [Nick Ciarleglio, Arista Networks](#). Chicago: Tony Verga, HFT Technologies (*moderator*); Jeff Margolis, Mellanox; John Moore, Juniper Networks; Dean Nebrig, Arista. London: Rob Howes, Juniper Networks; Paul Goodridge, Arista.
- ["Mellanox: Delivering the World's Fastest Trading Solutions,"](#) by Jeffrey Margolis, Mellanox Technologies.
- ["Cisco Innovations for High Frequency Trading,"](#) NY: Will Ochandarena, Cisco. London: Yang Yang, Cisco.
- ["Advances in low latency market data messaging,"](#) by Ásgeir Eiriksson, Chelsio.
- ["Low Latency 10Gb Ethernet With Chelsio WireDirect™,"](#) by Troy Leedberg, Chelsio.
- ["FSI-HPC: 5 HP updates in 5 minutes,"](#) by Chris Ferrari, HP.
- ["The Extreme Performance Platform for Trading and Risk,"](#) by Ian Pearl, Oracle Financial Services.
- ["Building a scalable 100% FPGA based trading infrastructure,"](#) Matthew Gadd, NovaSparks.
- Panel: "Whither Hardware Acceleration?" NY: [Scott Parsons, Exegy](#); [Mark Skalabrin, Redline Trading Solutions](#); [Yves Charles, NovaSparks](#). Chicago: Mike Mayhew, Trading Technologies, Inc.; David Taylor, Exegy; Mark Skalabrin, Redline Trading Solutions, Olivier Baetz, NovaSparks.
- ["Need For Speed: Taming Latency,"](#) NY: Raja Daita, SR Labs; Chicago: Sujay Subhash Lele, SR Labs.
- Panel: ["Managing Low-Latency Trading Infrastructure,"](#) with Scott Caudell, Interactive Data 7ticks; Raymond Russell, Corvil; Wil Tirado, CFN Services.

## Messaging middleware

- Panel Discussion: "The Future of Messaging Middleware." NY: [Alessandro Petroni, Deutsche Bank](#); [Shawn McAllister, Solace Systems](#); [Todd L. Montgomery, 29West \(part of Informatica\)](#). London: Mike Barker, LMAX; Ben Newton, Citihub; Shawn McAllister, Solace Systems; Todd L. Montgomery, 29West (part of Informatica).
- "[Informatica Ultra Messaging: Low Latency / Broad Functionality Across the Enterprise](#)," by Todd Montgomery, 29West (part of Informatica).
- "[Avoiding Post Trade Pitfalls – Increasing Simplicity, Capacity and Transparency](#)", by Bill Romano, Solace Systems.
- "[Every nanosecond counts: aiming at nanosecond inter-machine messaging](#)" Chris Kime, Solutions Architect, IBM.
- "[TIBCO FTL® Goes Global](#)." NY: John Page, TIBCO; Chicago: William McLane, TIBCO.
- "[Automated testing of RMDS](#)," by Peter Lankford, STAC.

## Time sync

- Panel: "Time Sync for Our Time." NY: [Paul Skoog, Symmetricom](#); [Will Ocharandena, Cisco](#); [Bruce Tolley, Solarflare](#). London: Victor Yodaiken, FSMLabs; Paul Skoog, Symmetricom; Nikolaus Kero, Oregano Systems, David Riddoch, Solarflare, Yang Yang, Cisco.
- "[Latency: not just a number](#)," by Raymond Russell, Corvil.
- "[Implementation experience and performance results with the Solarflare 10GbE precision time network adapter](#)." NY: Bruce Tolley, Solarflare. London: David Riddoch, Solarflare.
- "[Timekeeping Accuracy on Servers Under Load](#)," by Paul Skoog, Symmetricom.
- "[Precision and Accuracy ... Faking Exposed ...](#)," by Rony Kay, cPacket Networks.
- "[How to Achieve Sub-Microsecond Synchronization Accuracy on Application Level](#)," by Nikolaus Kerö, Oregano Systems.
- "[Why your traders are driving the demand for precision timing](#)," by Victor Yodaiken, FSMLabs.
- "[Solarflare: The platform for precision timing](#)," by Steve Pope, Solarflare.

## Latency monitoring

- Panel: "Latency Monitoring: Adult, Adolescent, or Infant?" NY: [Raymond Russell, Corvil](#); [Henry Young, TS-Associates](#); [Nick Ciarleglio, Arista Networks](#). London: Henry Young, TS-Associates; Rony Kay, cPacket Networks; Ken Jinks, Corvil; John Peach, Arista Networks.
- "[New ways to correlate latency measurements with trading success](#)," by Fergal Toomey, Corvil.
- "[Enhancing Network Instrumentation](#)," by John Peach, Arista Networks.
- "[Real time network monitoring with EOS LANZ Streaming](#)," by Dean Nebrig, Arista Networks.
- "[When ultra-low latency is not fast enough: Monitoring at the speed of trading](#)," by Sharon Besser, NetOptics.

## Big Compute

- ["Standard Technology Benchmarks for Market Risk Platforms,"](#) by Lars Ericson, Citi.
- ["Issues in Computing Counterparty Risk,"](#) by Harvey Stein, Bloomberg.
- ["Breaking the Logjam in Credit and Counterparty Risk Management,"](#) by Adam Vajda, Platform Computing.
- ["Head in the clouds? From Naivety to Clarity in 512 days: Strategies for Grid and HPC in and out of the Cloud for Finance,"](#) by Adam Vile, Excelian.
- ["Impossible on a Grid—But Very Possible,"](#) by Claudio Albanese, Global Valuation.
- ["Intel Many Integrated Core \(MIC\) Architecture,"](#) by David O'Shea, Intel.

## Big Data

- ["The Near Future of HPC and Big Data on Wall Street,"](#) Andy Brown, UBS.
- ["Watson: The Big Picture,"](#) Dr. Carl E. Abrams, IBM Research.
- Panel: ["Big Data in the Front and Middle Office,"](#) with Rick Farnell, Think Big Analytics; Amir Halfon, Oracle Financial Services; Scott Howser, Vertica, an HP Company; Adam Vajda, Platform Computing.
- ["Big Data, Big Insights,"](#) by James Lynch, Netezza, an IBM Company.
- ["Hadoop for Trading Firms: Beyond the Science Project,"](#) by Jack Norris, MapR.
- ["Moving \(BIG\) Data Faster,"](#) by Asaf Somekh, Mellanox.
- ["Network Implications of Big Data Workloads,"](#) by Will Ochandarena, Cisco.
- ["Arista - Big Compute / Big Data,"](#) by Ashwin Kohli, Arista Networks.
- ["Bitemp 2.0 – Avoiding a Bi-Temporal Bottleneck,"](#) by Paul Smeaton-Russell, Bank of America Merrill Lynch and Craig Baumunk, BitemporalData.com.
- ["Big Data Solutions in Capital Markets - A Reality Check,"](#) by Andrew Delaney, A-Team Group.

## Storage

- ["The STAC-M3 Effect: How standard tick-database benchmarks are precipitating engineering breakthroughs,"](#) by Peter Lankford, STAC.
- Panel: "The New Frontier in Storage Performance." NY: [Larry Jones, DataDirect Networks; John Overton, Kove; Keith Josephson, Ion Computer, ION Computer.](#) London: [Andy McDonald, Isilon and Larry Jones, DataDirect Networks.](#) Chicago: Christoph Lameter, Jump Trading (moderator); Larry Jones, DataDirect Networks; John Overton, Kove; Jeff Margolis, Mellanox.
- ["RAMCloud: Scalable High-Performance Storage Entirely in DRAM,"](#) John Ousterhout, Stanford University.
- ["Accelerating market data processing and analysis with parallel storage,"](#) Larry Jones, DataDirect Networks.
- ["Fast methods for using DRAM clusters as storage,"](#) by John Overton, Kove.
- ["Scalability and Speed: Proof That You Can Have Both,"](#) by Larry Jones, DataDirect Networks.
- ["Spinning Rust is Dead."](#) NY: Cameron Campbell, ION Computer;
- ["Kove DRAM Storage: The \(STAC\) Results Are In,"](#) by John Overton, Kove.