



STAC[®] Summit

June 16, 2014

STAC Exchange (exhibits) opens: 9:00am

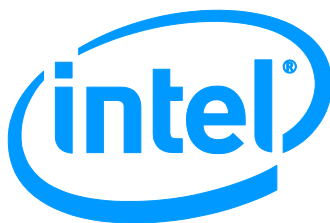
Conference starts: 10:00am

New York Marriott Downtown

85 West Street at Albany Street

New York, NY

Platinum Sponsor:



Gold Sponsors:



DataDirect[™]
NETWORKS



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 inaugural STAC Exchange so far include:



CONFERENCE AGENDA

Big data use cases in investment and retail banking [\[slides/video\]](#)

- *Introduction by Gordon Hughes, Global Sales Director, Financial Services Alliances, Intel*
- *Jennifer Costley, big data study lead, STAC*
- *Peter Lankford, Founder & Director, STAC*

Investment and retail banking often appear near the top of the list of industries investing in "big data" technology. Yet information about how banks are using that technology is sparse. With the help of seed funding from Intel, STAC recently worked with several of the global banks in the STAC Big Data Special Interest Group to identify concrete use cases that pose big data challenges. By interviewing staff with direct knowledge of these cases, we were able to characterize the workloads involved and understand the business problems that arise with traditional technologies. We were also able to learn about the advantages and challenges of the new approaches that banks were taking to these workloads. Along the way, we published [a white paper on the themes that had emerged](#). After an introduction from Gordon, Peter will review the use cases and those themes, as well as discuss how these are influencing potential STAC Benchmarks in big data domains.

Democratizing big data analytics: realities, visions, and hurdles [\[video\]](#)

- *Gary Bhattacharjee, Head of Enterprise Information Solutions, Wealth Management, Morgan Stanley*
- *Luigi Mercone, Director, Investment Banking Engineering, Credit Suisse*
- *Carl Reed, Managing Director, Goldman Sachs*
- *M.C. Srivas, CTO and Co-Founder, MapR*

As we documented in the white paper mentioned above, many banks and securities firms are turning to new "big data" technologies to enhance their powers of analysis, increase their agility, or reduce their costs. The use cases are real, important, and growing. Some banks have implemented or are implementing centralized analytics platforms using these technologies to serve multiple use cases. Their ultimate goal is to "democratize" analytics--that is, not only serving elite data scientists but also a much broader swath of employees who have ideas worth exploring. However, doing this has many infrastructural implications. For example, how can the many security and entitlements requirements be handled? How can firms deal with "ontological sprawl" so that users don't have to be experts in the eccentricities of specific datasets? And as more users crowd onto the system, how does the platform prevent one of them from spoiling the service levels to others? Our panelists will share their views on what has been achieved, what they would like to achieve next, and what it will take to get there.

Changing What "In-Memory" Means for Business. [\[video\]](#)

- *Brian Bulkowski, Founder & CTO, Aerospike* [\[slides/video\]](#)
- *Miko Matsumura, VP Marketing and Developer Relations, Hazelcast* [\[slides/video\]](#)
- *Tom Macdonald, VP and Director, Non-Volatile Solutions Group, Intel* [\[slides/video\]](#)

Financial firms have long exploited memory to store data when storage systems were too slow for the workload. But today's landscape of in-memory computing is different from what existed just a few years ago. New use cases are going in-memory. New software has emerged, some proprietary and some open source. And the very definition of memory itself is changing, as new non-volatile RAM products (NAND flash and its successors) blur distinctions with main memory and storage. Our panelists will explain this new landscape and discuss how in-memory architectures are transforming the capabilities and economics of financial solutions. What capabilities can in-memory software products now offer? What can finance learn from other industries? What should solution architects consider today that wasn't possible or economical before? How can performance-obsessed application developers best take advantage of flash today (memory channel? kernel bypass? NoFTL?). Where is this all heading? Our panel of experts will weigh in.

NETWORKING LUNCHEON

STAC Update - Tick data and backtesting [\[slides/video\]](#)

- *Peter Lankford, Founder & Director, STAC*

Peter will summarize the latest STAC Benchmark Council activities in areas relating to backtesting and tick data, including STAC-M3, STAC-A3, and research on flash storage.

Innovation Roundup – Round 1

"Getting the most out of multi-year and multi-source trading history" [slides/video]	<i>Glenn Wright, Systems Architect, DataDirect Networks</i>
"NetApp Enterprise Flash Solutions" [slides/video]	<i>Rob Ryan, Manager, Business Solutions Architects, NetApp</i>
"EMC ScaleIO for Financial Data Management Applications" [slides/video]	<i>Ken Taylor, Director Software Engineering, EMC</i>
"Enterprise Solutions from HGST, division of Western Digital" [slides/video]	<i>Doug Snyder, Regional Manager, HGST</i>

Point of View: A faster way to get strategies to market [\[slides/video\]](#)

- *Bryan Lewis, Chief Data Scientist, Paradigm4*

Trading firms today require faster cycle times for development and backtesting of trading strategies. Bryan believes that a key to this is finding more agile and less code-intensive approaches using high-productivity languages like R and Python. He believes the other key is to use shared-nothing architectures on the back-end that exploit hundreds or thousands of commodity hardware nodes in a cloud or local cluster. The crucial design imperative for large data is to minimize data movement by moving computation to the data. In Bryan's view, many SQL and Hadoop solutions fall short for strategy development and backtesting directly in the database. Even good distributed SQL databases often end up largely requiring ETL to an analytics environment like R or Python because of lack of in-database analytics. And Hadoop solutions typically sacrifice powerful data management concepts like join. As a consequence, developers are compelled to learn and use multiple toolsets, losing considerable time moving data between them. Taking the open source SciDB database as an example, Bryan will argue that it's possible to use an array language like R with inexpensive, scale-out hardware to create a single environment for tick storage, strategy development and backtesting.

Point of View: Putting the big data ecosystem to work for tick data [\[slides/video\]](#)

- *Michel Debiche, Financial Services Practice Lead, Think Big Analytics*

As a quantitative trader for over 20 years, Michel built one backtesting platform after another for a series of trading firms. Each time, the requirements were essentially the same. And each time, the platform had to be built largely from scratch. With the advent of "Big Data" technologies like Hadoop, he suddenly felt liberated. In his view, combing through huge volumes of historical data to search for signals and backtest trading strategies is a natural match for these new platforms. Not only because the map-and-reduce paradigm is a close fit to requirements, but also because these tools make it easy to integrate a huge range of content, from market data and news to fundamental data, research reports, and more. He believes that such scalable, open source tools also open up analytic opportunities for groups outside the front office. Risk, compliance, technology, and security groups are among those who must test assumptions and scenarios, and then use the intelligence gleaned from historical research to effectively monitor developments in real time across many dimensions. In this talk, Michel will start from first principles to explain how the "Big Data" ecosystem maps to requirements. He will then address some of the key questions for quants and engineers: what's it like for the developer? how does it scale for different sorts of backtests? when does it make sense to combine "Big Data" with "Cloud Computing" in financial services?

Point of View: The new meaning of TCA – and how to deal with it [\[slides/video\]](#)

- *Louis Lovas, Director of Solutions, OneMarketData*

Dispersal of liquidity across lit and dark pools has pushed firms to expand their hunt for alpha across brokers and borders. But the disparities among markets create barriers to efficient execution. As a consequence, a revolution in trade-performance measurement has been occurring in parallel with the race for high-speed execution. Demanding asset managers and institutional investors have redefined best execution beyond just price, pushing brokers to understand and manage complete trade performance and opportunity costs. This means that both brokers and the buy-side need tools that provide customized, in-depth, cross-market analysis of market participation and implementation shortfall. Louie will argue that such tools require several things: 1) advanced software to capture, clean, and store tick data, 2) powerful array-based tools such as the open source R language, 3) a framework that brings compute to the data, 4) the same state-of-the-art, high-speed platforms used for high-frequency trading, 5) intuitive visual dashboards.

BREAK

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.

Practical Parallelization in C++ [\[slides/video\]](#)

- *Evgeny Fikshan, Staff Engineer, Intel*

Parallelization continues to be a focal point for C++ programmers in finance. Some have not yet achieved parallelism in their code, while others debate the best ways to do so. As a former technical lead for Intel's OpenCL implementation and now as a specialist helping financial customers improve their application performance, Evgeny understands the challenges, benefits, and tradeoffs of parallelization. In this talk, Evgeny will provide new data on the benefits of parallelism in financial applications, as well as recent insights on best practices in parallelizing specific financial computations, including how Intel achieved its recent record-setting STAC-A2 Benchmark results.

Accelerating Stochastic Volatility Model Calibration in R Using GPUs [\[slides/video\]](#)

- *Matthew Dixon, Term Assistant Professor, University of San Francisco*

The R statistical software environment is popular with quants partly because it facilitates application design space exploration. However, despite offering a rich set of native numerical and statistical functionality that is useful for options-related analytics, a typical R based implementation of a stochastic volatility model calibration on a CPU does not meet the performance requirements for sub-minute level trading, i.e. mid- to high-frequency trading. Calculations rely on native optimization packages that are difficult to precisely replicate outside of R but are bottlenecked on computationally intensive kernels that can easily be replicated. In this talk, Matthew will present a new R library that dramatically accelerates options-related calculations by offloading the most computationally intensive part of the volatility model calibration to a GPU via an easy-to-use map-reduce interface. He will provide evidence that the new R-based implementation performs comparably to GPU-based calibration code written in C/C++.

Innovation Roundup – Round 2

“Static code analysis: Finding flaws as they’re typed” [slides/video]	<i>Scott Lasica, VP, Field Technical Services, Rogue Wave Software</i>
“Start Fast & Stay Fast: Priming Java for Market Open with ReadyNow!” [slides/video]	<i>Gil Tene, CTO, Azul Systems</i>
“Big Data meet Big Network” [slides/video]	<i>Kai Backman, CEO, Airstone Labs</i>

BREAK

STAC Update: Low-latency research [\[slides/video\]](#)

- *Peter Lankford, Founder & Director, STAC*

Peter will review the latest STAC activities in areas such as low-latency networking.

OpenSPL - A new way of thinking about FPGA [\[slides/video\]](#)

- *Ryan Eavy, Executive Director, Architecture, CME*

The benefit of FPGAs for certain tasks in financial markets is now widely accepted. But as we've discussed at many STAC Summits, one of the challenges slowing FPGA adoption is programmer productivity. A consortium of firms led by the CME, Chevron, and Juniper have set out to change that with a programming framework called the Open Spatial Programming Language, or OpenSPL. OpenSPL is an effort to standardize a "spatial programming" approach for FPGAs and other computing technologies. Whereas traditional programming techniques focus on execution in just one dimension (time), a spatial programmer thinks of execution in two dimensions: time and space. In this talk, Ryan will provide a brief overview of OpenSPL, illustrating how it eases programming of devices like FPGA.

Not Just Server Consolidation: Virtualization for Big and Fast Workloads [\[slides/video\]](#)

- *Matthew Knight, Technical Marketing Director, Financial Services, Solarflare*

While virtualization dominates the enterprise data center, it has been unrealistic for Tier 1 financial applications due to performance limitations. That may be about to change. In this talk, Matthew will examine some emerging cases in which applications achieve near-native performance by leveraging the latest virtualization technologies such as SR-IOV, multiple PFs and VFs, PCI passthrough, and a preview of Linux containerization that clears a path toward more elegant integration of 10/40GbE, Linux and KVM into OpenStack Clouds. As part of this, he will show how running IEEE 1588 (PTP) in a VM and leveraging ingress and egress time stamping enables real-time latency measurements without external equipment.

Innovation Roundup – Round 3

“MetaConnect: Proven performance.” [slides/video]	<i>Dave Snowdon, Founder, co-CTO, Metamako</i>
“Fiberblaze Update” [slides/video]	<i>Nikolaj Hermann, CTO, Fiberblaze</i>
“New Developments in Exchange Latency Management” [slides/video]	<i>Tony Pettipiece, Global Head of Sales and Marketing, Cape City Command</i>
“Achieving ultra low latency with the Exablaze ExaNIC network card family” [slides/video]	<i>David Mann, Firmware Engineer, Exablaze</i>

NETWORKING RECEPTION

Speaker Biographies– Feature Sessions



Gary Bhattacharjee, Head of Enterprise Information Solutions for the Wealth Management business, Morgan Stanley. Gary brings a wealth of experience in the field of information management, banking and capital markets, business analytics and technology strategy. As a thought leader, Gary holds a patent in the “Management of data via cooperative method and system”, a wiki-based approach for managing structured data. Gary started his career at IBM Australia, where he was involved in building key technology capabilities in a product, that later came to be known as CoreBank®. He was also a significant contributor to the industry standard model called FSDM® (Financial Services Data Model). In the late nineties, he played a crucial role in CitiGroup to lead the consolidation of information across the Global Corporate Banking portfolio.

As the CTO of the Great Atlantic & Pacific Tea Company, a grocery retail company, Gary was responsible for defining the strategic architecture and managing business-IT alignment. As the Practice Director of information management consulting in the Financial Services division at Hewlett-Packard, he established a reputation for developing innovative and strategic solutions for the industry. Gary earned his Bachelors in Electronics and Telecommunications from the Indian Institute of Technology.



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.”



Jennifer Costley. Jennifer is a scientifically-trained technologist with broad multi-disciplinary experience in enterprise architecture, software development, line management, and infrastructure operations. Following 31 years of technology leadership in organizations like Credit Suisse, Bankers Trust, and Double Click, Jennifer now consults to companies, non-profit organizations, and individuals in areas related to data, governance, and sustainability. This includes roles with the IEEE and the STAC Benchmark Council. Jennifer has a PhD in Chemical Physics from Columbia University and a Bachelors in Physics and Chemistry from Brandeis University.



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.



Matthew Dixon, Ph.D., FRM, Assistant Professor of Analytics, University of San Francisco.

Matthew is an Assistant Professor in the School of Management at the University of San Francisco and specializes in financial modeling, machine learning and high performance computing. Matthew began his career as a quantitative developer at Lehman Brothers in London before pursuing academics and consulting for financial institutions in quantitative risk modeling. Matthew is a chartered financial risk manager and currently serves as consulting director of risk at HedgeFacts. He holds a Ph.D. in Applied Mathematics from Imperial College (2007), a Master of Science in Parallel and Scientific Computation with distinction from the University of Reading (2002) and has held postdoctoral and visiting professor appointments at Stanford University and

UC Davis respectively. He has published several academic papers at the intersection of financial modeling and high performance computing, chairs the workshop on high performance computational finance at SC and is co-founder of the Thalesians.

**Ryan Eavy, Executive Director, Enterprise Architecture, CME Group.**

Ryan leads the Emerging Technology function as part of the Enterprise Architecture team at the CME Group. He is responsible for ensuring the company's technology is aligned with business strategy through application research and development, technical solution assessment and systems roadmap definition. Among Ryan's many other accomplishments, he has recently led the design and implementation of the Exchange's ultra-low-latency middleware strategy and common abstract middleware layer. Prior to this, he most recently served as a Senior Application and Security Architect for the Chicago Board of Trade. Before joining the CBOT in 2003, Ryan worked as a Senior Applications Developer at Quotes Plus where he designed and developed trading analysis

software and as a Senior Applications Engineer at Dewpoint. Ryan earned a bachelor's in computer science from the University of Michigan and is currently pursuing a master's in Computational Finance from DePaul University.

**Evgeny Fikshan, 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 programming. 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, Israel.

**Gordon Hughes, Global Sales Director, Financial Services Alliances, Intel.**

Gordon is responsible for Intel's worldwide strategy for eco-system engagements in the company's focus areas in banking and capital markets. Gordon has been with Intel for over 8 years and has a pedigree in Alliance Management with both software and hardware vendors, including Tandem Computers and Informatica. He holds a B.Sc (Hons) in Applied Mathematics from the University College of North Wales.

**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.



Bryan Lewis, Chief Data Scientist, Paradigm4. Bryan is Paradigm4's Chief Data Scientist and creator of SciDB-R. Bryan pioneered the concept of "Personal Supercomputing," introducing the first integrated desktop cluster system as co-founder of Rocketcalc LLC. He is the author of a number of R packages and was also the founding Director of Engineering at Revolution Computing (now Revolution Analytics), working on many performance-related aspects of R, especially related to parallel computing. His research areas are in iterative methods for inverse problems and numerical analysis.



Louis Lovas, Director of Solutions, OneMarketData. Louis' responsibilities include strategic business development and delivering targeted solutions for quantitative research and trading systems. Louis brings over 20 years of experience in developing cutting edge solutions for financial markets and is a leading voice on industry trends, trading technology and Complex Event Processing (CEP) for Capital Markets. Prior to joining OneMarketData, Louis was the Chief Architect and Fellow of Apama products at Progress Software. Earlier in his career, Louis held numerous positions in the Financial Systems and Corporate Program Management divisions at Unisys Corporation (formerly Burroughs). Louis has a B.S. in Business Administration/Management.



Thomas R. Macdonald, Vice President, Technology & Manufacturing Group, Director, Non-Volatile Solutions Group, Intel. Tom leads a corporate-wide team responsible for the strategy, execution, and ramp of next generation non-volatile memory products and technology at Intel Corporation. Previously he was general manager of the Platform Solutions Group in the Datacenter and Connected Systems Group in the Intel Architecture Group. In that role he was responsible for the datacenter chipset business as well as the development of Intel® Atom™ processor-based system-on-chip products for servers, including the Intel Atom processor S1200 family and Avoton processors. He also had responsibility for manageability firmware and storage software products. Since joining Intel in 1988 as a product manager for the Intel386™ microprocessor, Tom has held a variety of general management and marketing management

positions. He was general manager of the Fabric Components Division and was co-chair of the InfiniBand* trade association. Previously, he was marketing director for Intel's enterprise server group. In this role, he was responsible for all platform and product marketing and planning, including launching Intel's first server-optimized processor, the Intel® Xeon® processor. Tom has also served as marketing director for the Intel486™ processor family and the Intel® Itanium® processor family. Tom received his bachelor's degree in mechanical engineering from Stanford University, and his MBA from the Kellogg Graduate School of Management, Northwestern University.



Miko Matsumura, VP Marketing and Developer Relations, Hazelcast. Miko is a Vice President at Hazelcast, the leading open source in-memory data grid company. Hazelcast is a venture backed startup out of silicon valley that seeks to disrupt the database world through providing compact and easy to use in-memory software that re-imagines Enterprise Software Architecture to be fully distributed and elastic. Previously he was Senior Vice President of Platform Marketing & Developer Relations at Kii Corporation. Kii provides a Mobile Backend-as-a-Service (MBaaS) that serves the largest phone carriers and leading app makers on Android and iOS platforms. Kii Cloud serves over 25 million users and drives over \$250M in industry revenues. Miko's career spans over 15 years of entrepreneurial technology leadership, having built and led technology platform marketing and product organizations across a wide range of software organizations from startup to billion-dollar levels. Miko is based in silicon valley and has created shareholder value through

buying and selling many companies. He is a Limited Partner with Focus Ventures, a VC firm with over a half billion dollars under management. Prior to his current role, Miko served as Vice President and Chief Strategist at Software AG, where he played a key role in the M&A team that drove Software AG from under a half billion to over one billion dollars in annual revenue through acquisitions. Miko holds an MBA from San Francisco State University and a master's degree in neuroscience from Yale University. Miko was named at Cloud Expo 2009 West as one of the [World's 30 Most Influential Virtualization Bloggers](#).



Luigi Mercone, Director - Investment Banking Engineering, Credit Suisse. Responsible for architectural thought leadership, transparency and quality governance of strategic book of work with critical work streams that spans industry leading ultra-low latency trading, extreme transaction volume and applying agile/cloud computing to metascale infrastructure. Success measured in the product management balance of accelerating uptake of strategic platforms (compute, virtualization, storage, database) and in leading a principle-driven evolution of those platforms towards emerging business needs. Driving increasing efficiency in total cost of ownership forms one pillar; pragmatic, applied and rapidly executable innovation the other--particularly in areas of harnessing predictive analytics and leveraging enterprise "big data", extreme performance compute densification and elastic data and compute fabrics in support of delivering next

generation platforms. Prior roles at Credit Suisse have included Equities Front office business development in European Electronic Markets, Global Head of Equities Engineering and Asia Pacific Head of Infrastructure. Additionally experience includes tenures in the Hedge Fund sector (Bridgewater) and Enterprise Infrastructure Software (Veritas.)



Carl Reed, Managing Director, Goldman Sachs.



M.C. Srivas, CTO and Co-Founder, MapR. Srivas ran one of the major search infrastructure teams at Google where GFS, BigTable and MapReduce were used extensively. He wanted to provide that powerful capability to everyone, and started MapR on his vision to build the next-generation platform for semi-structured big data. His strategy was to evolve Hadoop and bring simplicity of use, extreme speed and complete reliability to Hadoop users everywhere, and make it seamlessly easy for enterprises to use this powerful new way to get deep insights. That vision is shared by all at MapR. Srivas brings to MapR his experiences at Google, Spinnaker Networks, Transarc in building game-changing products that advance the state of the art.

Srivas was Chief Architect at Spinnaker Networks (now NTAP) which built the industry's fastest single-box NAS filer, as well as the industry's most scalable clustered filer. Previously, he managed the Andrew File System (AFS) engineering team at Transarc (now IBM). AFS is now standard classroom material in operating systems courses. While not writing code, Srivas enjoys playing tennis, badminton and volleyball. M.C. has an MS in Computer Science from University of Delaware, and a B.Tech. in electrical engineering from IIT Delhi.