

***PPASCA Department Meeting  
October 8, 2010***

***XLDB4***

***<http://www-conf.slac.stanford.edu/xldb10/>***

***Anders W. Borgland***

# XLDB

## Attendance - Rough Breakdown

Invitational workshop			conf		
xldb1	xldb2	xldb3	xldb4		
43%	41%	45%	37%	24%	Data-intensive scientific users
21%	19%	15%	26%	42%	Data-intensive industrial users
30%	19%	23%	18%	9%	Vendors, incl. startups
6%	21%	15%	15%	24%	Academia, db research & programmers
0%	0%	2%	4%	1%	DOE, NSF, EC
<b>53</b>	<b>62</b>	<b>52</b>	<b>57</b>	<b>150</b>	← Total count

*XLDB4 attendees:  
North & South Americas, Europe, Asia, Australia*

# Workshop Program

## Workshop Program

 [Printer-friendly version \[PDF\]](#)

TUESDAY, OCTOBER 5, 2010		
08:00 AM	Continental Breakfast	
9:00 AM	<b>Welcome</b>	Jacek Becla
9:10 AM	<b>Oil/Gas</b>	moderator: Oliver Ratzesberger
10:10 AM	Coffee Break	
10:30 AM	<b>Finance</b>	moderator: Raghu Ramakrishnan
11:30 AM	<b>Medical / Bioinformatics</b> Including medical bioinformatics, genomics, perspectives from science and industry.	moderator: Kian-Tat Lim
12:30 PM	Lunch, discussion in small groups about expanding and strengthening XLDB community	
1:30 PM	<b>Science Benchmark</b> Current status, next steps	Michael Stonebraker
2:00 PM	<b>Approaches to Big Data Statistical Analytics</b>	moderator: Curt Monash
3:00 PM	Coffee Break	
3:20 PM	<b>Emerging Challenges</b> Solid state disks, GPUs, very-many-cores, fast networks	moderator: Alex Szalay
4:20 PM	<b>Writable Extreme Scale Databases</b> Extremely large non-read-only databases (beyond bulk load)	moderator: Magdalena Balazinska
5:00 PM	<b>Plan Collaborative Projects and Next Steps</b> Planning collaborative projects (within USA and international), future funding, next steps	moderator: Jacek Becla
5:30 PM	Adjourn	

# Conference Program

## Conference Program

WEDNESDAY, OCTOBER 6, 2010		
08:00 AM	Continental Breakfast	
9:00 AM	<b>Welcome -- Science and Computing at SLAC</b> Official "welcome" and introduction to new exciting science and computing challenges at SLAC.	Donald Lemma (SLAC CIO and Computing Division Director)
9:15 AM	<b>Conference Introduction and Logistics</b> Main objectives, logistics and agenda.	Jacek Bocia (SLAC, XLDB4 chair)
9:30 AM	<b>Complex Scientific Analytics at Extreme Scale</b> A comprehensive overview of how big science approaches complex analytics at extreme scale.	Gregory Dubois-Felsmann (HEP), Andrew Connolly (astronomy), John Caron (atmospheric research), Bill Howe (ocean sciences, bio), Jacek Bocia (summary)
10:50 AM	Coffee Break	
11:10 AM	<b>Complex Industrial and Government Analytics at Extreme Scale</b> A comprehensive overview of how data intensive industries approach complex analytics at extreme scale. It will highlight similarities and differences comparing to approaches taken by big science.	Irina Vayndinar (MITRE), Steve Hirsch (NYSE Euronext), Mike McIntire (Yahoo!), Damian Reeves (Quantcast) moderator: Klan-Tat Lim (SLAC)
12:30 PM	Lunch	
1:30 PM	<b>Operational Issues with Managing Large Database Clusters</b> Practical, operational issues with managing large database clusters, based on experiences from at least two large-scale industrial setups.	Oliver Ratzesberger (eBay), Jeffrey Rothschild (Facebook)
2:10 PM	<b>Behind the Scenes of Big Science Projects</b> A talk explaining how big and long-term scientific projects get started, what the decision processes are, how requirements and data volumes are decided, how vendors are evaluated.	Amber Boehnlein (DOE)
2:55 PM	<b>Existing Scientific Tools/Formats - netCDF, HDF5, fits, etc</b> A set of short talks to make non-scientific communities aware of most commonly used scientific formats and related tools custom-built by scientists.	Daniel L. Wang (SLAC)
3:15 PM	Ice Cream Social - Poster Session for Gold Sponsors	
3:55 PM	<b>Existing Scientific Tools</b> A continuation of short talks to make non-scientific communities aware of most commonly used scientific tools custom-built by scientists. Considering covering: root, cstor, xroot.	Richard Dubois (SLAC)
4:15 PM	<b>Lightning Talks (8 x 5 min)</b> 1. <i>A (Hypothetical) Data to Discovery Engine</i> 2. <i>Big data analytics - use of a highly scalable MPP database on a cloud-based multi-instance platform to reliably deliver near real-time analytics</i> 3. <i>Exabyte plus Initiative</i> 4. <i>Damasc</i> 5. <i>Horizontal virtualization on commodity hardware without requirements for database optimization techniques</i> 6. <i>Array Versioning System</i> 7. <i>Smorgasbord of Real World Extreme Scale Database Analytics</i> 8. <i>Introduction to MongoDB</i>	1. Mark Stalzer, Caltech 2. DJ Das, Third Eye Cloud 3. Leon Guzenda, Objectivity 4. Nooklis Polyzotis, UCSC 5. Stefan Groschupf, DataMover 6. Philippe Cudre-Matoux, MIT 7. Andrew Lamb, Vertica 8. Shulman, William, MongoDB
5:00 PM	Adjourn	
6:30 PM	Reception and dinner	

THURSDAY, OCTOBER 7, 2010		
08:00 AM	Continental Breakfast	
8:40 AM	<b>Announcements and Logistics</b>	Jacek Bocia
8:45 AM	<b>Emerging Technologies for Complex Extreme Scale Analytics</b> Unifying emerging technologies such as map/reduce, streaming databases, and workflow management into a coherent tool set for extreme scale data analytics.	Jeff Hammerbacher (Cloudera)
9:25 AM	<b>Emerging Scientific Tools - SciDB</b> Lessons learned from studying scientific use cases and interacting with scientific communities by the SciDB team. SciDB response to these needs.	Mike Stonebraker (MIT)
9:45 AM	<b>Science Benchmark</b> Introducing the idea behind science benchmark, current status and plans.	Mike Stonebraker (MIT)
10:00 AM	Coffee Break	
10:40 AM	<b>Extreme Scale Architectures and New Hardware Trends</b> Impact of new hardware trends such as solid state disks, GPUs, servers with very-many-cores.	Alex Szalay (JHU)
11:25 AM	<b>Data Preservation and Integration</b> Challenges related to integrating data from multiple sources and preserving petabytes of data.	Jane Mandelbaum (Library of Congress)
12:05 PM	<b>Automated Information Extraction, Content Curation and Machine Learning</b> Needs in the area of automating data processing and analytics. Perspectives from industry and science.	Raghu Ramakrishnan (Yahoo!), Kirk Borne (GMU)
12:45 PM	<b>Closeout</b> Next conference planning, final conclusions and closeout.	Jacek Bocia
1:00 PM	Lunch	
2:00 PM	Adjourn	