

The Official User Conference of The Apache Software Foundation

APACHECON US 2009



Conference Program

[updated 20091103]

November 2 - 6

Oakland Marriott City Center

Oakland, California

<http://us.apachecon.com>

Presented by: The Apache Software Foundation

Produced by: Stone Circle Productions, LLC.

Sponsored by :

Microsoft

t thawte



sourceforge

COLLABNET



lucid
IMAGINATION

spring
source

10 YEARS

2009: Ten Years of The Apache Software Foundation
Celebrating a Decade of Open Source Leadership

Thank You to Our Sponsors

Platinum Sponsors



Gold Sponsors



Silver and Bronze Sponsors



Community Sponsors and Exhibitors



Media and Community Partners



ApacheCon US 2009



Happy Birthday Apache! We're delighted to be back in the Bay Area to celebrate **ten years** of great code and great community, and we're thrilled that you're here to celebrate with us.

We've got loads of activities planned for the week, including two days of BarCamp, three days of great content, and a whole week of community-building and celebrations! Keep an eye on the registration desk and the wiki for last-minute updates, and don't forget to log on to CrowdVine and build your very own ApacheCon program.

This year, for the first time, our Project Management Committees have planned their programs themselves. This means you're getting the very latest and greatest content on the topics that matter most to you. Whether you want hands-on answers to the practical problems, or a sneak-peek at what our projects are planning for the next ten years, you'll find it all at ApacheCon.

If you weren't around in 1999, make sure to check out our Pioneer's Panel keynote, featuring stories from the people who started it all. If you want to know what's going on right now, check out the evening MeetUps for community-focused, late-breaking content. And whatever you do, make sure you join us for our unmissable Big Feather Birthday Bash on Wednesday night!

We're delighted that you're here - so please don't hesitate to join in and have some fun!

Noirin Shirley and the ApacheCon US 2009 Planning Team

Table of Contents

• Sponsors & Exhibitors	4
• Community & Media Partners	14
• Trainings	18
• Keynotes	23
• Schedule at a Glance	24
• Special Events & MeetUps	27
• Sessions	29
• Speakers	55
• Thanks	70

Microsoft®

Microsoft - Platinum Sponsor + Exhibitor

At Microsoft, we're motivated and inspired every day by how our customers use our software to find creative solutions to business problems, develop breakthrough ideas, and stay connected to what's most important to them. We are committed long term to the mission of helping our customers realize their full potential. Microsoft's open source strategy is grounded in recognition of the value of openness to working with others-including commercial open source vendors, developers and open source communities-to help customers and partners succeed in today's heterogeneous "world of choice."

For more information, visit

» <http://www.microsoft.com/opensource/>



thawte - Platinum Sponsor + Exhibitor

thawte is a leading global Certificate Authority (CA). Our SSL and code signing digital certificates are used globally to secure servers, provide data encryption, authenticate users, protect privacy and ensure online identities go through stringent authentication and verification processes. Our SSL offerings include Extended Validation (EV), Server Gated Cryptography (SGC), SANS, and wildcard certificates. The focus of *thawte* is to allow companies to establish a trusted relationship with anyone, anywhere on the Internet. *thawte* is committed to the egalitarian ethos of the Internet and is dedicated to ensuring world-class multi-lingual customer service, and providing outstanding value for money.

For more information, go to

» <https://www.thawte.com>

Gold Sponsors

**HotWax Media - Gold Sponsor + Exhibitor**

As the global leader in Apache Open For Business (OFBiz), HotWax Media provides web-based enterprise resource planning (ERP), e-commerce, and business automation consulting services. From enterprise-level e-commerce to warehouse, manufacturing, and financial management, HotWax Media and OFBiz help optimize business processes. Open source is the smart choice today for businesses ranging from small-and medium-sized organizations to global Fortune 100 companies. Using Apache OFBiz, HotWax Media delivers top quality enterprise software for a fraction of the price of outdated, proprietary vendors.

For more information visit
» <https://www.hotwaxmedia.com>.

**SourceForge - Gold Sponsor**

"The Hub of Software Innovation." SourceForge is the world's largest open, collaborative platform enabling millions of technology innovators around the globe to develop and distribute cutting edge, emerging technologies amongst a massive universal community of casual consumers and leading industry professionals seeking and evaluating software solutions. With over 230,000 open source software projects, delivering over 3 million downloads every day, expert help, and a global reach penetrating every area on planet Earth, SourceForge has become the world's largest reliable collection of open source software tools and applications on the net today.

For more information visit
» <http://www.sourceforge.net>

Silver Sponsors

**CollabNet - Silver Sponsor + Exhibitor**

CollabNet is the leader in application lifecycle management (ALM) platforms for distributed software development teams. CollabNet TeamForge is the industry's most open ALM platform, supporting every environment, methodology, and technology. With an integrated suite of easy-to-use tools that share a centralized repository, it is the only ALM platform that enables a culture of collaboration, improving productivity 10-50% and reducing the cost of software development by up to 80%. As the corporate sponsor of the open source Subversion project, the best version control and software configuration management (SCM) solution

for distributed teams, collaborative development is in CollabNet's DNA. Millions of users at more than 800 organizations, including Applied Biosystems, Capgemini, Deutsche Bank, Reuters, and the U.S. Department of Defense, have transformed the way they develop software with CollabNet.

For more information, visit

» <http://www.collab.net>.



i n v e n t

Hewlett-Packard - Silver Sponsor + Exhibitor

HP Open Source and Linux solutions are designed to solve your mission-critical IT problems today! We sell over a million Linux servers a year, delivering high performance, scalability, security, storage, and backed by our commitment to the open source community. Learn more about our market leading Linux platforms and Open Source Middleware Stacks at www.hp.com/go/linux. Also be sure to check out our new FossBazaar and Fossology initiatives which are aimed at further encouraging the adoption of Open Source in the enterprise. Our solutions are delivered and supported by over 6500 Open Source and Linux professionals across 160 countries. Come visit us at our booth to hear how HP and the Open Source community are working together.

For more information, visit

» <http://www.hp.com/go/linux>



Lucid Imagination - Silver Sponsor + Exhibitor

Lucid Imagination is the first commercial entity exclusively dedicated to Apache Lucene/Solr open source technology. As an active participant in the enormous community using Lucene/Solr, Lucid Imagination offers certified distributions of Lucene and Solr, commercial-grade support, training, high-level consulting and value-added software extensions. The company's web site will serve as a knowledge portal for the Lucene community, with information and resources to help developers build and deploy Lucene-based solutions in a more efficient and cost-effective manner. The Lucid Imagination founding team consists of several key contributors and committers to the Lucene project, as well as experts in enterprise search application development. A privately held company, Lucid Imagination is based in San Mateo, CA, USA.

For more information, visit

» <http://www.lucidimagination.com/>




SpringSource - Silver Sponsor + Exhibitor

SpringSource, a division of VMware, Inc., (NYSE: VMW) and the leader in Java application infrastructure and management, provides a complete suite of software products that accelerate the entire build, run, manage enterprise Java application lifecycle. SpringSource employs the open source leaders who created and drive innovation for Spring, the de facto standard programming model for enterprise Java applications. SpringSource also employs the Java and Web thought leaders within the Apache Tomcat, Apache HTTP Server, Hyperic, Groovy and Grails open source communities. Nearly half of the Global 2000, including many world's leading retail, financial services, manufacturing, healthcare, technology and public sector clients are SpringSource customers.

For more information, visit
» www.springsource.com

Bronze Sponsors



HIPPO

Hippo - Bronze Sponsor + Exhibitor

Hippo develops the Apache licensed Hippo CMS and Hippo Portal application. Hippo is an avid believer of the Apache way of developing, and participates pro-actively in ASF Projects. Hippo's developers are stimulated to get involved in Apache in order to become committed in relevant projects, sharing knowledge with others in true Apache spirit. You can find code initiated by Hippo people in many projects, amongst others: Apache Cocoon, Apache Jetspeed and Apache Slide. Hippo is a member of the Java Community Process and its developers help define new Java specifications.

Hippo CMS is already widely used in the Netherlands and across Europe. With the acquisition of the assets of Blue Sunrise, another great company active in the Apache community, Hippo obtained a strong threshold into the USA. Hippo is now well on the way to become a household name in the US market space, and is rolling out vastly to set the global standard in open source content management and portal applications.

For more information, visit the company website at
» <http://www.onehippo.com>.

or get involved in the Hippo community at
» <http://www.hippocms.org>.



Progress Software - Bronze Sponsor + Exhibitor

The FUSE Community (www.fusesource.com) was created to help users adopt and deploy Apache ServiceMix, Apache ActiveMQ, Apache Camel, and Apache CXF in enterprise environments. The FUSE distributions are tested, certified versions of these popular Apache projects and are backed by world-class support and services. Progress Software employs many of the key committers to these projects.

For more information, visit

» <http://fusesource.com/>

Community Sponsors



Apache Software Foundation

Community Sponsor + Exhibitor

Established in 1999, the all-volunteer Foundation oversees more than sixty-five leading Open Source projects, including Apache HTTP Server -- the world's most popular Web server software. Through The ASF's meritocratic process known as "The Apache Way", nearly 300 individual Members and 2,000 Committers successfully collaborate to develop freely available enterprise-grade software, benefiting millions of users worldwide: thousands of software solutions are distributed under the Apache License; and the community actively participates in ASF mailing lists, mentoring initiatives, and ApacheCon, the Foundation's official user conference, trainings, and expo. The ASF is funded by individual donations and corporate sponsors including Google, HP, Microsoft, Progress Software, SpringSource, and Yahoo! » <http://www.apache.org/>



Aster Data Systems - Community Sponsor + Exhibitor

Aster Data Systems is a proven leader dedicated to providing the best data analytics and management platform for frontline data warehousing - the first DBMS to tightly integrate SQL with MapReduce - providing rich insights on data managed on clusters of inexpensive commodity hardware. The Aster *n*Cluster database cost-effectively powers rich analytic applications for companies such as Coremetrics, MySpace, aCerno (an Akamai company), and ShareThis. Running on low-cost off-the-shelf hardware, and providing 'hands-free' administration, Aster enables enterprises to meet their data warehousing and analytics needs within their budget.

» <http://www.asterdata.com/>

How secure are you feeling these days?



See us at our booth
www.thawte.com

Are you experiencing
security and trust issues?
We can help.



cloudera - Community Sponsor + Exhibitor

Cloudera (www.cloudera.com), the commercial Hadoop company, develops and distributes Hadoop, the open source software that powers the data processing engines of the world's largest and most popular web sites. Founded by leading experts on big data from Facebook, Google, Oracle and Yahoo, Cloudera's mission is to bring the power of Hadoop, MapReduce, and distributed storage to companies of all sizes in the enterprise, Internet and government sectors. Headquartered in Silicon Valley, Cloudera has financial backing from Accel Partners, Greylock Partners and angel investors who include Diane Greene (former CEO of VMware), Marten Mickos (former CEO of MySQL), and Jeff Weiner (CEO of LinkedIn).

For more information, visit
» <http://www.cloudera.com/>

facebook

Facebook - Community Sponsor + Exhibitor

Founded in February 2004, Facebook is a social utility that helps over 300 million people communicate more efficiently with their friends, family and coworkers. Over the past few years, the company has developed a number of open source infrastructure technologies to support the site's growth (<http://developers.facebook.com/opensource.php>) including Cassandra, Hive, Scribe, and Thrift as well contributed to numerous open source projects including Hadoop. Facebook engineers actively participate in various Apache projects, developing technologies that facilitate the sharing of information through the social graph.

For more information, visit
» <http://www.facebook.com/>



HALO Worldwide - Community Sponsor

HALO Worldwide is a consulting firm specializing in promotion, publicity, and brand integration programs that span strategic communications and positioning, brand planning and translation, public relations and outreach, and image management. HALO principals have gained a reputation for being among the most competent in the business, leveraging more than 30 years' combined out-of-the-box thinking, intuitive marketing, and media relations expertise to develop highly effective campaigns across a range of media resources. With a clientele that ranges from startups to blue chip organizations, we are passionate about

managing reputations, driving brand awareness and recognition, and developing valued relationships with credibility and gravitas.

For more information, visit
» <http://www.haloworldwide.com/>



LinkedIn - Community Sponsor

Your professional network of trusted contacts gives you an advantage in your career, and is one of your most valuable assets. LinkedIn exists to help you make better use of your professional network and help the people you trust in return. Our mission is to connect the world's professionals to make them more productive and successful. We believe that in a global connected economy, your success as a professional and your competitiveness as a company depends upon faster access to insight and resources you can trust.

To enable companies to develop professional applications for LinkedIn's over 45 million members, LinkedIn has developed the Intelligent Application Platform (InApps) which is based on OpenSocial and leverages the Apache Shindig project.

For more information, visit
» <http://www.linkedin.com/>



Mark Logic - Community Sponsor

Mark Logic Corporation is a leading provider of infrastructure software for information-centric applications, serving industries including media, government, software, transportation, and healthcare. Mark Logic was recently named the fourth-fastest growing information technology company in Silicon Valley. The company's award-winning product, MarkLogic Server, is the industry's leading XML server and includes technology protected by multiple US patents. Mark Logic is privately held with investors Sequoia Capital and Tenaya Capital.

For more information, to download a trial or community version, or to read the award-winning Mark Logic CEO Blog, go to
» <http://www.marklogic.com>



MuleSoft - Community Sponsor

MuleSoft, formerly MuleSource, provides enterprise-class software, support and services for the world's most popular open source application infrastructure products, Mule ESB and Apache Tomcat. With the lightweight and powerful Mule ESB and MuleSoft Tcat Server, MuleSoft is bringing simple yet powerful infrastructure to today's dynamic Web applications. Today, MuleSoft products boast more than 1.5 million downloads and over 2,000 production deployments by leading organizations such as Walmart.com, Nestlé, Honeywell and DHL, as well as 5 of the world's top 10 banks. MuleSoft is headquartered in San Francisco with offices worldwide.

For more information, please visit
» <http://www.mulesoft.com>.



Ning - Exhibitor

Ning is the social platform for the world's interests and passions online. Based in Palo Alto, Calif., Ning offers an easy-to-use service that allows people to join and create Ning Networks. With more than 1.5 million Ning Networks created and 35 million registered users, millions of people every day are coming together across Ning to explore and express their interests, discover new passions, and meet new people around shared pursuits. Ning continues to experience strong growth with over 4,500 new Ning Networks being created daily and 1 million new registered users being added every 15 days. The privately held company was founded in October 2004 by Gina Bianchini and Marc Andreessen.

For more information, please visit
» www.ning.com.



OSS Watch - Community Sponsor

OSS Watch is an open source software advisory service for the Higher and Further Education sector. We are funded by the Joint Information System Committee (JISC), and our services are free to the HE and FE sector. We help institutions and education related projects in the use and development of free and open source software. In addition we create links between research software projects and open source projects like those found within the Apache Software Foundation.

For more information, please visit
» <http://www.oss-watch.ac.uk/>



Sun Microsystems - Community Sponsor

Sun Microsystems has held open systems to be the cornerstone of its business philosophy since the beginning. From desktops to supercomputers, and from development tools to productivity suites, Sun is dedicated to delivering hardware and software based on open industry standards and open source software. Today, through shared technology innovation, Sun's continued commitment to open source is reflected in its leadership and key contributions to the many projects including Apache Tomcat, Apache HTTPD, Apache Derby, OpenSolaris, OpenOffice.org, OpenJDK, GNOME, Glassfish, NetBeans, Mozilla and myriad others.

For more information, please visit

» <http://www.sun.com/software/opensource/>



ThinkGeek - Community Sponsor

ThinkGeek started as an idea: to create and sell stuff that would appeal to the thousands of people out there who were on the front line and in the trenches as the Internet was forged. Since 1999, we've stayed true to that idea by serving up awesomesauce every day to programmers, engineers, students, lovers of open source, and the geeky masses. If you're passionate about technology, science fiction, caffeine, video games, or well... zombies, chances are there's a more than a little something waiting for you at ThinkGeek.

Members of the Zombie Monkey Army are standing by » <http://www.thinkgeek.com/>



WSO2 - Exhibitor

WSO2's developer portal, the WSO2 Oxygen Tank, is a vast resource pool for open source Web services and SOA projects. Encouraging community build up and facilitating collaboration among users, developers and other enthusiasts, The Oxygen Tank is the ideal place to get your open source SOA deployments up and running. On the Project Home pages you will find technical information and artifacts related to the specific WSO2 projects released under the Apache Software Foundation License, while the Library section comprises articles, tutorials, interviews, webinars, presentations, podcasts and white papers for everyone from beginner to advanced developers and users.

Check it all out here

» <http://wso2.org>



Feather Cast - Media Partner

There are a number of search-based projects at Apache, and they all center around Lucene. Grant will be giving a training class on Lucene at ApacheCon in just four weeks. Then, there's the largest collection of search technology talks we've ever had at an ApacheCon. Grant talks about his training class, about Lucene, and about the related projects.



Linux Journal - Media Partner

Linux Journal, first published in 1994, is the original publication of the global Linux community, delivering readers the advice and inspiration they need to get the most out of their Linux systems. With its award-winning Linux how-tos, tutorials, reviews, in-depth reporting, tips and tricks, and market analysis, Linux Journal has become well-known as the industry standard.



Linux Pro Magazine - Media Partner

In Linux Pro Magazine you'll find the tools, tutorials, reviews, and concrete technical discussions you'll need to unlock the secrets of Linux - and not just from the server side. Linux Pro delivers solutions for real users with Linux on real desktops.

Methods & Tools

Methods and Tools - Media Partner

Methods and Tools is a free magazine providing practical knowledge for the software developer, tester and project manager



NLUUG - Community Partner

NLUUG is the association of (professional) Open Source and Open Standards users in the Netherlands. Since the late seventies, the NLUUG has brought together the community of systems administrators, programmers, researchers and IP network professionals. The primary goal of the NLUUG is to extend the application of, and knowledge about, open systems and UNIX.

For more information, please visit
» <http://www.nluug.nl/events/nj09/index.html>



openSUSE - Community Partner

The openSUSE Conference is for anyone and everyone making a contribution to openSUSE. Whether you're a developer, packager, translator, artist, or ambassador, you belong at the openSUSE Conference if you want to help promote, improve, and influence the openSUSE Project. Conference

activities include keynotes, tutorials, Birds of a Feather (BoF) sessions, roundtable discussions, and hackfests. The conference will consist of equal parts planned sessions and “unconference” activities for attendees to pick their own topics and set the agenda.

opensource3

Opensourc3 - Media Partner

Opensourc3 is the premier Open Source and Unified Computing magazine for Information Technology Professionals. Published on a monthly basis, opensourc3 is available for FREE download in PDF format, or can be read on-line. Opensourc3 focuses on emerging technologies, their relationship within the Unified Computing model and technologies that free organizations from vendor lock-in. Regular columns focus on Virtualization, Networking, Storage, Server / Computing along with a Feature column.

Subscribe or download the latest issue at
» <http://www.opensourc3.org>



OStatic - Media Partner

OStatic aims to help users find the right Open Source software for their personal and business needs, evaluate Open Source alternatives to Proprietary Source applications and services, and collaborate with a network of trusted peers. As a member of the award-winning GigaOM Network, the OStatic blog delivers originally created content from a team of experienced editors and technologists. OStatic also provides a set of tools and services to help increase the development and adoption of Open Source Software and in 2008 OStatic was selected as PC Magazines Top 100 Sites on the Web.

Slashdot

Slashdot - Media Partner

News for Nerds. Stuff that Matters.” Slashdot is the single hottest source of peer-tested news content and discussion, the only site of its kind, providing a ritualistic daily dose of technology news, culture and humor for an insatiably loyal, near zealous, audience dialed into every aspect of the technology world. For over ten years, the content has remained peer driven, straight from the source, and relatively unfiltered, giving a heightened sense of overall trust and quality to Slashdot, defining the standard by which similar sites are judged.

» <http://slashdot.org>



SourceSense - Community Partner

SourceSense is the leading Open Source Systems Integrator in Europe. Our international team contributes to major Open Source projects, and our customers leverage our ability to provide dependable solutions based on Open Source software.



The BitSource - Media Partner

The BitSource, Technology's Premier Source of Information, is an online publication and research lab providing in-depth enterprise software and service product reviews, interviews and tutorials. The BitSource brings together best-in-class research, reviews, and most importantly: the people of the technology world.



WSO2 Oxygen Tank - Community Partner


WSO2's developer portal, the WSO2 Oxygen Tank, is a vast resource pool for open source Web services and SOA projects. Encouraging community build up and facilitating collaboration among users, developers and other enthusiasts, The Oxygen Tank is the ideal place to get your open source SOA deployments up and running. On the Project Home pages you will find technical information and artifacts related to the specific WSO2 projects released under the Apache Software Foundation License, while the Library section comprises articles, tutorials, interviews, webinars, presentations, podcasts and white papers for everyone from beginner to advanced developers and users.

Check it all out here

» <http://wso2.org>

HP wishes to extend our congratulations to the ASF for
10 years of leadership in the open source community.
And to the hundreds of Apache developers
for the great software you create...

Thanks!



Absolutely committed
to open source.

For over a decade more companies have chosen
HP servers to run Linux than any other vendor.



See what's possible at hp.com/go/linux

Apache Maven, End-to-end

Trainer : Brett Porter

Day : 11/2 (One Day Training)

Time : 9:00 - 17:00

This training session will walk through the lifecycle of developing a typical Java application from creation to deployment, and show how to use Apache Maven most effectively to manage the build and development process. In addition to the fundamental building blocks of the project, the session will cover testing, day-to-day development in the IDE, application of Maven best practices, effective dependency management, establishing a release process, using profiles effectively, setting up documentation, tracking development reports and practices. Effective use of continuous integration (illustrated with Apache Continuum) and repository management (using Apache Archiva) as a part of development infrastructure for team and enterprise environments will be demonstrated. This course will be suitable both for those that are looking to get the most out of their existing Maven projects, and those that are looking to use Maven for the first time. Time is reserved for addressing specific situations that attendees have encountered in exist projects.

Basic Hadoop Training

Trainer : Aaron Kimball

Day : 11/2 (One Day Training)

Time : 9:00 - 17:00

This basic training session, designed for programmers and data analysts who have never worked with Hadoop before, will allow you to understand the answers to the following questions:

- What must our organization do differently to effectively use large-scale data?
- What tools help us analyze large-scale data and extract meaningful

results, and how do we use them?

- How can I become more proficient with these tools?

Solr Day

Trainer : Erik Hatcher

Day : 11/2 (One Day Training)

Time : 9:00 - 17:00

Solr has emerged as THE way to add search services to applications written in any language. Lucene powers Solr's lightweight HTTP web services. This all-day tutorial will begin with a Solr overview, and move through installation and configuration. The bulk of session will be spent diving deeply into integrating Solr with various technologies. Solr features caching, replication support for load balancing, faceting, highlighting, more-like-this, simple HTTP protocol supporting XML, JSON, and other formats. Erik will discuss his Solr experiences on the applications he has built with it.

Media & Analyst Training

Trainer : Sally Khudairi

Day : 11/2 (Half Day Training)

Time : 9:00 - 12:15

Prepare yourself for a variety of media opportunities with communication skills that resonate with editors and analysts and help generate positive coverage. Learn how to craft, control, and bridge messages and gain confidence during the interviewing process. Present your position clearly and become more persuasive to effectively target editors and influence the analyst community. Hone your presentation techniques -- from public speaking to analyst briefings -- through group coaching, individual interviews, and breakout practice sessions. Takeaways include identifying common mistakes and tips on how to avoid them, press do's and don'ts,

analyst relationship management, developing thought campaigns, positioning previews and launch tactics, and more. A must-attend for those interested in promoting their Open Source project and becoming a resource for the media community at large.

.....

Intermediate Media & Analyst Training

Trainer : Sally Khudairi

Day : 11/2 (Half Day Training)

Time : 13:45 - 17:00

Perception is reality! Step up your basic media skills to become a more effective spokesperson and take control of the interviewing process. Handle sensitive topics, aggressive interviews, and misquotes with gravitas and confidence. Present yourself credibly in a variety of scenarios - press conferences, analyst briefings, communications crises, and more. Class format includes group coaching, individual interviews, and breakout practice sessions. Hands-on, tactical tips on improving your written media skills to mitigate the reality of overworked reporters and make your efforts more credible to analysts and the community at large. NEW! guidelines on working with the ASF Public Relations Committee. NOTE: Restricted class size; limited to participants who have already completed basic media/analyst training.

.....

Apache HTTP Server - Nuts to Bolts

Trainers : Jim Jagielski and Rich Bowen

Day: 11/2 & 11/3 (Two Day Training)

Time : 9:00 - 12:15

This training will provide a Nuts to Bolts Guide to Apache HTTP Server 2.2. This includes the basic design and architecture of the server, an

overview of the API, deep dives into various useful modules. It also looks at some of the deficiencies of the current design and what is on the horizon for 2.4/3.0. Join Jim Jagielski and Rich Bowen for this in-depth two day training.

.....

Everything Tomcat -Administering, Tuning, Troubleshooting & Developing

Trainer : Filip Hanik

Day: 11/2 & 11/3 (Two Day Training)

Time : 9:00 - 12:15

The fastest way for you to learn everything you need to know about Tomcat. If you are an administrator, developer, release engineer, or perhaps a bit of all, then this training is for you. Whether you are an expert or a beginner, you are guaranteed to learn something new in this class. The curriculum spans from basic Tomcat configuration to advanced tuning of both Tomcat and the JVM. We'll go over several different deployment methodologies, with and without web server connectors, load balanced or not, and how to maximize the performance of Tomcat in your infrastructure. You'll learn how to troubleshoot Tomcat, how to classify different types of error messages and how to correctly reconfigure your server based on the actual error. Towards the end we will take a deep dive into the core of Tomcat, learning both architecture and code, and how to write your own Tomcat components.

.....

Lucene Boot Camp

Trainer : Grant Ingersoll

Day: 11/2 & 11/3 (Two Day Training)

Time : 9:00 - 12:15

Apache Lucene is a high-performance, cross-platform search API in production use in a large number of applications. This tutorial

will be an in-depth, hands-on class targeted at developers who wish to integrate Lucene Java into their applications. The two day tutorial will cover the key concepts behind Lucene, along with code examples, documentation and resources. Concepts covered are: Lucene setup; indexing content, indexing customization; searching, including how to use the many Lucene Query APIs to achieve better search results. Additionally, we will cover advanced topics like Lucene performance and threading issues, Span Queries, filters, term vectors, and sorting, all of which can help solve common problems in search applications. The class will gain hands-on experience with Lucene and leave with the know-how to build a Lucene-based application. No prior Lucene knowledge is required, but participants must be able to code in the Java programming language and come prepared to code. Concepts learned can also be applied to Solr. Students are encouraged to bring their own content to index/search. Please refer to the Lucene Boot Camp website (<http://www.lucenebootcamp.com>) for more information on the class..

ActiveMQ & Camel Tutorial

Trainer : Rich Newcomb
Day: 11/3 (One Day Training)
Time : 9:00 - 17:00

This one day training will start with an overview of ActiveMQ and Camel. It will walk through the architecture of these projects and then delve into hands-on exercises. This course will provide you with the knowledge you need to configure and run ActiveMQ and Camel for your messaging and integration needs.

Advanced Hadoop Training

Trainer : Aaron Kimball
Day : 11/3 (One Day Training)
Time : 9:00 - 17:00

Designed for users already familiar with Hadoop, and building on the basic training session offered on Monday, this advanced training session will answer the following questions:

- How can we reorient our data generation and collection processes to enable more powerful analysis later?
- What's different about high-scalability algorithms for large data analysis, and how are they employed?
- How do I write custom logic to extend Hadoop to deal with data from diverse sources?

Wicket Core Principles - Full Day Training

Trainer : Jeremy Thomerson
Day : 11/3 (One Day Training)
Time : 9:00 - 17:00

You go to classes because you want to learn something, and leave with a new skill. Even in a one-day class, you will have coded a working application and been introduced to the fundamentals of developing applications with Wicket. The class will consist of fast-moving explanations of core design principles, Wicket components, and "The Wicket Way", and each section will be followed by a coding practice where you can put into use what you just learned. We will focus on laying a foundation - how to use Wicket, create pages, organize your application, and create a Wicket application. We will cover the following: - The fundamentals of Wicket - Handling data / working with objects and models - Standard components provided by framework - Containers / Application / Session / Page - Effective code reuse strategies

Web Application Security Boot Camp

Trainer : Christian Wenz

Day:11/3 (Half Day Training)

Time : 9:00 - 12:15


According to a recent study, 9 out of 10 web sites have security vulnerabilities. Try to be website #10 and come to this half-day tutorial which covers common and not-so-common (but still dangerous) security vulnerabilities, how they are exploited today, and which countermeasures are feasible. Topics covered include XSS (Cross-site scripting), CSRF (Cross-site request forgery), SQL injection, attacks on Ajax-y websites, protection against automated HTTP requests, and more. You will also hear war stories from real-world security audits, and get an overview of suitable development processes and testing tools.



The Industry's First MPP Data Warehouse with Applications-Within™

Aster Data announces the first 'Massively Parallel Data-Application Server' that for the first time allows applications to be run inside the database itself. The result? Dramatically faster, deeper insights on massive data sets.

HOW? BY PROVIDING:

-  Always-Parallel Performance
-  Always-On Availability
-  Embedded Applications & SQL-MapReduce
-  Dynamic Workload Management

Write custom applications in Java, Perl, Python, Ruby, .NET, C++, C#, and other popular languages. Use Aster Data to push-down and parallelize the applications for ultra-fast performance. The power of Aster Data's Massively Parallel Data-Application Server is recognized by partners like SAS, Informatica, and MicroStrategy and customers such as Coremetrics, MySpace, LinkedIn, ShareThis, Full Tilt Poker and Specific Media.



Visit the Aster Data booth

to hear how you can get faster, richer insights from your 'Big Data' using Aster Data's patent-pending SQL-MapReduce framework inside our MPP data warehouse.

Open Source
Enterprise

Custom **E-commerce** &
ERP Platform Consulting



Built on Apache Open For Business

Bring open source front-and-center in your enterprise by partnering with the global leader in Apache Open For Business. From business analysis through design and implementation, HotWax Media is with you every step of the way.

HotWax Media

(877) 736-4080 www.hotwaxmedia.com

Since 1997, delivering leading edge e-commerce, creative services and enterprise automation.

Wednesday : 10 years of The Apache Software Foundation Apache Pioneers' Panel

2009 marks the 10th Anniversary of The Apache Software Foundation. What started as a formalization of methods of collaboration for the world's most popular Web server has grown into a very popular watering hole on the savanna of Web-aware software development. Join several pioneering ASF members in this very special panel where we'll hear stories from the early days of the Apache Way.



Thursday : Standing Out in the Crowd Kirrily Robert

How far have we come in ten years? Kirrily Robert has been working as a web developer and Internet technologist for more than a decade. In that time, the world wide web -- driven by Apache -- has become ubiquitous, and with that ubiquity has come a broadening of the market: from geeks and specialists to anyone and everyone.

But while half of all web users are women, women make up only a tiny fraction of the open source developers whose work underpins most websites. A 2006 survey showed 1.5% of developers are female, compared to 24% in IT roles across the US. Many women in open source have had the experience of being the only one on their project or at an event. Kirrily will talk about the experiences of women in open source, and present two new open source projects that have 75% or more women contributors. What can we learn from these majority-female projects, and how can it improve open source overall



Friday : How Open Source Developers Can (Still!) Save The World Brian Behlendorf

Open Source communities like Apache have done far more than just create awesome, no-cost, flexible software - they have demonstrated a whole new way for communities of interest to band together and solve big problems. Beyond operating systems and web servers, Open Source software is now being used to organize disaster relief efforts, enable microfinance banks in the developing world, integrate electronic healthcare systems, and change the way citizens collaborate with their government. How well are these projects going? What should they emulate or attempt to inherit from the more "infrastructural" kinds of projects and organizations - such as Apache - and what new kinds of challenges emerge? And where can you help?



COLLABNET®-stewarded

SUBVERSION®



...a perfect fit!

Open source to the power of 10

- 10-year anniversary of the Apache Software Foundation
- 10-year anniversary of CollabNet, founder of Subversion
- 10th year of development for Subversion

November 4, 2009

*CollabNet-Founded Subversion Submitted to
Become a Project at The Apache Software Foundation*

www.collab.net • info@collab.net • +1-650-228-2500

Thursday

**The Conference Schedule
has been updated.**

Please refer to the
Conference-at-a-Glance
signage on-site for the
most current schedule.

Friday

**The Conference Schedule
has been updated.**

Please refer to the
Conference-at-a-Glance
signage on-site for the
most current schedule.

Special Events

ApacheCon's Special Events bring the conference to life, and Oakland will be no exception, with many free events open to the public. The fun continues with stimulating receptions, lively presentations, quick-witted sessions such as the Lightning Talks, and our amazing Birthday Bash.

Hackathon | 2-3 November

If you have a desire to contribute to Open Source software, whether you're a developer, writer, designer, tester, or anyone else, please join us for this free Hackathon event. Meet Apache project contributors and work on projects together all day long.

BarCamp | 2-3 November

BarCampApache is our free Apache-focused unconference. It will gather people interested in The Apache Software Foundation and its various projects to share and learn in a completely open environment. BarCampApache will be intense with discussions, demos and interaction from participants packed into an ad-hoc, participant-driven event. BarCampApache is free, but due to a limit in size, pre-registration is required.

Big Feather Birthday Bash | 4 November

The Apache Software Foundation is 10 years old, and we're celebrating in style Wednesday evening with a Big Feather Birthday Bash! Be sure to join us for this action packed reception and party.

Join us Monday, Tuesday, and Thursday evenings for FREE MeetUps!

Sponsored by:      

MeetUps give projects the chance to organize their own events, and talk about what they're interested in. MeetUps cover everything from developers attacking a bug list, to user stories and sample applications, to introductory nights to introduce new users to a project.

MeetUps start at 8pm on Monday, Tuesday and Thursday nights. Check the ApacheCon site for updates. Currently planned MeetUps include:

Monday

- » NoSQL
- » Tomcat

Tuesday

- » Lucene
- » OSGi/Felix
- » Traffic Server Podling

Thursday

- » Content Technologies
- » Web Services
- » Shindig/SocialSite/Widget
- » Hadoop
- » Subversion

PROUD SPONSOR OF APACHECON



SUPPORT
FOR APACHE
ACTIVEMQ
SERVICEMIX
CAMEL & CXF

PROGRESS
SOFTWARE

BUSINESS MAKING PROGRESS™

www.fusesource.com

© 2009 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

Welcome to ApacheCon

William A. Rowe

9:00-9:10 | West Hall

Welcome to ApacheCon US 2009, and the Oakland Convention Center. We are delighted to have you here, to celebrate ten years of innovation and to kick off the next ten years of progress.

State of the Feather

Jim Jagielski

9:10 -9:30 | West Hall

A review of the last 6-12 months at the ASF.

10 years of The Apache Software Foundation

Apache Pioneer's Panel

9:30 - 10:30 | West Hall

2009 marks the 10th Anniversary of The Apache Software Foundation. What started as a formalization of methods of collaboration for the world's most popular Web server has grown into a very popular watering hole on the savanna of Web-aware software development. Join several pioneering ASF members in this very special panel where we'll hear stories from the early days of the Apache Way.

Introduction to Apache Hadoop

Owen O'Malley

11:00 - 11:25 | East Hall

Apache Hadoop provides a framework for running applications that process large amounts of data (hundreds of terabytes), on large clusters (thousands of computers), of commodity hardware. The Hadoop framework transparently provides both reliability and data motion to applications. Hadoop implements MapReduce, and provides a distributed file system similar to GFS. This presentation looks at the motivation and approach for Hadoop, an overview of the components and architecture, and an overview of the

tools built on top of Hadoop, such as Apache Pig and Apache Hive.

State of the Elephant

Christophe Bisciglia

11:25 - 11:50 | East Hall

Ever wonder how many clusters there are out there? How big are they on average? How much data are they processing? Where does your cluster fit in the mix, and what can you learn from others?

In this talk, we will look at how the Apache Hadoop community has grown and changed over the past year. We'll share some aggregate data we have gathered over the year, and walk through a few interesting case studies based on companies that have made a big difference in their business/operations with Hadoop. We'll also provide a summary of the last year's development, and outline what users (and developers) can expect to see in 2010.

Tomcat community overview

11:00-11:50 | Jewett A - C

Mladen Turk

The presentation gives an up-to-date overview of the Apache Tomcat Project and its direction into the future. The session explains how we work as a community of developers and how we interact with our community of users. It gives brief overview of the technologies developed and maintained within the project, the different versions of Tomcat, connectors, and much more. It will also cover our development methods, how we use Bugzilla, and even a bit about security in Tomcat. In short, everything you need to know to get the most out of your interactions with Tomcat and its community!

Tapestry 5: Java Power, Scripting Ease!

Howard M. Lewis Ship

11:00-11:50 | Jewett D & E

The Apache Tapestry web framework has been making a name for itself in terms of innovative features and ease of use. In this session, we'll introduce the goals of the Tapestry framework: Simplicity, Consistency, Efficiency and Feedback. We'll discuss the challenges of building modern, Web 2.0 applications, and the two main approaches: action-oriented and component-oriented. We'll show the advantages of the component-oriented approach, especially as implemented in Tapestry.

Tapestry includes many developer productivity features, including live class reloading and convention over configuration: we'll show how these require minute amounts of code to accomplish big goals, and how Tapestry brings scripting language productivity within reach of Java developers without sacrificing any of Java's inherent speed and power.

SCA, Java EE, Spring and Web 2.0 Come Together - Service assembly with Apache Tuscany SCA

Luciano Resende

11:00-11:50 | Jewett F

Apache Tuscany provides an easy-to-use, open source, services infrastructure for building, assembling, deploying and running SOA solutions based on the Service Component Architecture (SCA) specifications from OASIS. The Apache Tuscany project goes beyond the SCA specification and is an environment for innovative ideas around SOA. For example, it extends SCA to work with Web 2.0 and OSGI. Apache Tuscany is integrated with other Apache technologies such as Apache Tomcat, Apache Axis2, Apache Geronimo, Apache BSF, Apache

Ode, Apache XMLBeans and Apache Abdera.

This talk will provide a short overview of SCA and the Apache Tuscany project, and will mainly focus on an enterprise integration example to demonstrate how to take advantage of SCA and Tuscany to describe, assemble and deploy an end-to-end SOA solution. The talk is directed at those who are building distributed solutions from connected services, and who want to understand how SCA can help.

The presenters will use their experience of working with the Apache Tuscany project and its users to illustrate partitioning of the application into components and services, exploitation of a variety technologies to implement components including Java EE, Spring and Web2.0, and use of different communication technologies. They will also look at configuration of policy to control consistent quality of service across the distributed application, as well as deployment to distributed runtimes with varying capabilities.

Panel: The Business of Open Source - Power, Prestige, and Propulsion

Sally Khudairi

11:00-11:50 | 208

The ApacheCon Business Track has expanded over the past eight conferences, addressing an array of key business, marketing and legal/licensing issues in Open Source. Our panel of influencers --Hewlett-Packard's Scott Lamons, Progress Software's Debbie Moynihan, Microsoft's Sandy Gupta, and RedMonk's Michael Coté-- will answer your questions on The Business of Open Source, that includes customer requirements, application opportunities, deployment challenges, best practices, product development,

standards compliance, business model disruptors, and more. Moderator Sally Khudairi will create a lively, interactive dialogue by inviting comments from the audience throughout the 50-minute session.

Becoming a Pig Developer

Alan Gates

13:30 - 14:20 | East Hall

Apache Pig is a sub-project of Apache Hadoop. It provides a high level language, Pig Latin, that enables the programmer to describe parallel data flows which the Pig engine can execute. It provides the standard relational operators (group, join, etc.) and allows users to plug in their own code at any point in the data flow.

Developing a high level language that supports arbitrary data flow graphs while maintaining performance close to the level of Java MapReduce code is challenging! This talk will discuss the architecture of Pig, review the design of some of the performance improvements that have been added recently, and discuss future work and directions the Pig team is considering. It will be a great starting point for those interested in contributing to the Pig project, as well as informative for advanced Pig users who want to understand more about the system and its possible future directions.

Introduction to Apache Tomcat 7

Mark Thomas

13:30 - 14:20 | Jewett A

In addition to implementing version 3.0 of the Servlet 3.0 specification, Apache Tomcat 7 has a number of new features, and improvements on old favourites. This presentation will provide both system administrators and web application developers with the information they need to get the most out of Tomcat 7.

The topics covered will include a brief overview of Servlet 3.0 features with an emphasis on configuration, a look at improvements to management and monitoring via JMX, discussion of security enhancements to the Manager application, and a review of the logging enhancements made in Tomcat 7.

An Introduction to Apache Velocity 1.6

Henning Schmiedehausen

13:30 - 14:20 | Jewett D & E

Apache Velocity is mostly considered an exciting alternative to the ubiquitous Java Server pages for web apps. But Velocity is much more than that! As a general purpose templating engine written in Java, it is at the heart of code generators, mass mail senders, and documentation frameworks. This talk gives an overview of the language features, and looks at how to integrate Velocity into your own applications.

Distributed OSGi with SCA using Apache Tuscany

Raymond Feng

13:30 - 14:20 | Jewett F

OSGi is no longer confined to a single JVM in its service invocations, since the introduction of RFC 119 - Distributed OSGi. This RFC enables an OSGi bundle deployed in a JVM to invoke a service (either OSGi or non-OSGi) in another JVM or process, potentially on a remote computer accessed via a network protocol. Meanwhile, an OSGi service deployed in another JVM or a non-OSGi program such as a Web Service client, can find and access a service running in the "local" OSGi JVM. The distributed computing functionality is added to the OSGi programming model without additional APIs or concepts, and the distribution layer is mostly transparent to OSGi developers by

configuration.

Service Component Architecture (SCA) provides a technology-neutral approach to abstracting business logic into components, and assembling them into composite applications. It greatly simplifies the component communications using declarative bindings. QoS requirements can be uniformly declared as SCA intents, which can be mapped and realized using different stacks. SCA is a great fit as a distribution provider for distributed OSGi.

Apache Tuscany is an open source project that implements the SCA specifications. It provides integrations of implementation types (such as Java, Scripting, BPEL, Spring, OSGi and JEE) and binding types (such as Web Service, JMS, EJB, CORBA, RMI, JSONRPC, and ATOM) as the infrastructure for SCA programming. We recently added support for OSGi RFC 119.

In this session, we will teach you how to develop a distributed OSGi application to leverage the SCA capabilities using Apache Tuscany. A calculator scenario will be used to demonstrate the distributed OSGi service invocations using RMI and Web Service protocols. We will also explain how to model an OSGi bundle as an SCA component and configure the SCA composite to provide communications between services. You'll leave this session with an understanding of the basic ideas behind distributed OSGi and the power of Tuscany SCA for service composition.

Apache License as a Business Model - Challenges and Opportunities

Paul Fremantle

13:30 - 14:20 | 208

Conventional wisdom in the open source business community is

that the formula for successful open source businesses is GPL + dual license. The Apache license in particular is viewed as "too challenging" to monetize due to its "overly friendly" nature. In this talk we examine the challenges and opportunities of running a business based on the Apache license.

This talk is a must-see for anyone involved in open source businesses around Apache.

Apache Hadoop in the Cloud

Tom White

14:30-15:20 | East Hall

Apache Hadoop is particularly well suited to take advantage of the cloud, but not all "clouds" are equally well suited to take advantage of Hadoop. In this talk, we will compare various options for running Hadoop in the Cloud.

We will compare cost, performance, elasticity, etc for various cloud platforms. We'll also go over tools that make Hadoop deployment in the cloud easier - things like Amazon EC2 images for Hadoop. In this talk, you'll learn how the cost structures and performance characteristics of various clouds align with their work loads, and gain the knowledge needed to make deployment decisions that are right for you!

Asynchronous servlet processing in Apache Tomcat 7.0

Filip Hanik

14:30-15:20 | Jewett A

Apache Tomcat 7.0 is a Servlet container for the Servlet 3.0 specification. One of the more exciting features in this new specification is the standard support for asynchronous servlet processing. In this session we will take a look at the asynchronous servlet API, the reasoning behind it, its benefits and shortcomings.

We will also take a look at the extensions and different methods Apache Tomcat 7 provides to make up for some of these shortcomings. Filip is a Senior Software Engineer for SpringSource and a key participant in SpringSource's Apache Tomcat initiatives. Filip brings 12 years of extensive experience in architecture, design and development of distributed application frameworks and containers and is recognized for his top-quality system development skills and continuous participation in Open Source development projects. Filip is a committer to the Apache Tomcat project where he is a leading authority on Tomcat clustering and a key contributor to the core of the platform. He also authored the online book Hitch-Hiker's Guide to Tomcat. Prior to SpringSource,

Introduction to Wicket

Jeremy Thomerson

14:30-15:20 | Jewett D

Apache Wicket is a component-oriented Java web application framework that brings object oriented programming back to the Web. With just HTML and Java, Wicket gives designers and programmers a tool they can work in together. In this session you'll be introduced to Wicket's programming model, and we'll quickly go from simple components, form processing and providing feedback to integrating with Spring and Guice, as well as securing, testing and deploying your application. A lot of ground to cover but the best way to learn everything you wanted to know about Wicket.

Event Driven Architectures with Apache Synapse

Paul Fremantle

14:30-15:20 | Jewett F

Event Driven Architecture (EDA)

is a highly scalable approach for building large scale distributed applications and integrating across large organizations. In an EDA, many systems can manage their own destiny in connecting to the rest of the applications, and this approach means that new applications can be much more self-reliant. EDA has been widely adopted in the financial sector, but is gaining ground in many other areas. This presentation will introduce EDA and the benefits, challenges and best practices for building EDA, and also show how to architect real-life EDA systems. In addition, we will explore Complex Event Processing, which adds a whole new set of capabilities to find patterns, react in real time to data, and correlate across high volume data feeds.

We will also look in detail at how the Apache Synapse ESB can be used to build an EDA, both by itself and in conjunction with other projects such as Apache Qpid and Apache ActiveMQ. We will examine how to build topic spaces, create routing and subscription patterns, and how to add Complex Event Processing to an EDA. This talk is aimed at coders and system architects who want to understand how event architectures can improve their systems and how to use Apache technology to build them.

Building a Company on Distributed Computing with Apache Hadoop and Friends

Bradford Stephens

14:30-15:20 | 208

Visible Technologies is a Social Media Analysis and Monitoring company that collects hundreds of millions of blog entries and bits of social media daily, and performs data mining and heavy semantic analysis on them. We then provide advanced business intelligence on that data. We have quite a few

Fortune 100 companies as clients.

We used to be MS-exclusive, but now our entire core infrastructure is built on Apache. Apache Hadoop, Apache Nutch, Apache Lucene, Apache Mahout, Apache Solr, Apache Zookeeper, Apache HBase, and more. It's provided our company with amazing agility and leverage for what we invest in R&D and development. I am the Lead Engineer on the Data Platform team, in charge of our distributed computing core.

The presentation will cover leveraging Hadoop, Lucene, Mahout, Nutch, Zookeeper, HBase, and more. I'll give a business and technology overview of moving from Microsoft to Open Source, as well as discussing how to use Apache software to increase the effectiveness of small development teams. The presentation will also look at how distributed computing affects business processes and development lifecycles, and the importance of "scale-out" over "scale-up".

Practical HBase

Michael Stack

16:00-16:50 | East Hall

Apache HBase is the Apache Hadoop database. It is an open-source, distributed, column-oriented store, modeled after Google's Bigtable. This talk will offer practical tips on getting the most out of your HBase install. Topics covered will include performance tuning, migrating your data from an RDBMS, and description of production HBase deployments. Michael Stack is on the Hadoop Project Management Committee, and is sponsored by Powerset (part of Microsoft) to work full-time on Apache HBase. HBase, which is an important component of Powerset's development, is developed as part of the Apache Software

Foundation's Hadoop project, and runs on top of the Hadoop Distributed File System, providing BigTable-type capabilities.

Securing your Tomcat installation

Tim Funk

16:00-16:50 | Jewett A

Even though we do our best to ensure that Apache Tomcat is secure "out of the box", every user makes configuration changes that can potentially impact security.. In this talk, we'll give you hints to make sure your changes don't expose you to security issues. Whether you are new to Tomcat or have used it for a long time, we'll walk through ways to ensure that your installation is secure. We'll also cover items developers should and should not do in their code to keep the installation secure.

What's new in Roller 5.0

Dave Johnson

16:00-16:50 | Jewett D & E

In this talk, we'll cover the features of Apache Roller 5.0, and how it compares to other blog server options. We'll also look at typical deployment architectures, and examples of Roller sites in production today. The talk will discuss in detail some of the ongoing, student-driven work on OpenID and onbetter Media Blogging support. You'll learn how blog software is adapting to incorporate Social Networking features such as profiles, friend relationships and social gadgets, as well as how you can use Apache Shindig and Project SocialSite to add these features to Roller.

Integration made simple with the Synapse ESB

16:00-16:50 | Jewett F

Asankha C. Perera

The Apache Synapse Enterprise Service Bus (ESB) is a lightweight

and high performance ESB, which is designed to be simple to use and manage. The ESB is configured via a simple human readable configuration file, and can be managed via JMX for management and monitoring. It provides support for fully non-blocking http/s based on the Java NIO extensions of the Apache HttpComponents project, and thus can manage thousands of concurrent connections with standard hardware. It also support transactional JMS, File systems based on Apache VFS, Mail, FIX etc, and supports multiple message formats, with support for extension. This presentation introduces the Synapse ESB, and how one can immediately make use of it to fulfill integration requirements of an enterprise. We will look at some sample scenarios and how they are made possible with ease, using the Synapse ESB.

Open Source Business for Hackers

Ross Turk

16:00-16:50 | 208

The current market is ready for aggressive open source promotion and adoption, but some of us are more comfortable coding projects than commercializing them. How do you make money for your efforts (and enable members of your community to do the same)? How can you market your software to the people who care, and harmonize all of that with the unique ideological needs of your communities? And, even more importantly, why should you care?

In this talk, Ross Turk, Director of Community at SourceForge, will talk about these issues and more, as he discusses what every coder needs to know about the business behind open source.

Apache Hive: SQL and Data warehousing on Apache Hadoop

Ashish Thusoo

17:00-17:50 | East

Apache Hive is a data warehouse built on top of Apache Hadoop, that provides tools to enable easy data summarization, adhoc querying and analysis of large datasets data stored in Hadoop files. It provides a mechanism to put structure on this data and it also provides a simple query language called QL which is based on SQL and which enables users familiar with SQL to query this data. QL also allows traditional MapReduce programmers to be able to plug in their custom mappers and reducers to do more sophisticated analysis which may not be supported by the built in capabilities of the language.

Hive is used extensively in Facebook for all of its data processing needs. Currently the Hive installation at Facebook contains over 1 PB of data, and processes thousands of jobs daily on a 3200 core cluster. In this presentation we will talk about Hive, the new capabilities and optimizations that it has incorporated over the last year, how its developer community is growing, the various applications that are being served by it, and our plans for the future.

mod_jk / mod_proxy and others

Jean-Frederic Clere

17:00-17:50 | Jewett A

Apache Tomcat is often used in a cluster, or as the back-end to a front-end reverse proxy. Several front-ends can be used, including mod_jk, mod_proxy, mod_serf and mod_cluster. This presentation will give a quick overview of each of the available modules, including the features they offer in terms of load balancing, failover, QoS, and more.

Apache Stonehenge - Interoperability through Community

Kent Brown & Paul Fremantle

17:00-17:50 | Jewett F

Apache Stonehenge is a community-driven solution to the challenges of web services interoperability. There is now strong support for the WS-* standards in several open-source and proprietary web services stacks. However, it sometimes requires fiddling with the bits or playing around with config files to achieve interop for specific scenarios. The goal of Stonehenge is make web services interoperability easier, by serving as a forum for resolving interop questions and by sharing the solutions as documents and samples.

Stonehenge started with a sample application contributed by WSO2 and Microsoft. It has grown to add implementations on three more web services stacks (including Sun


Metro). And the future direction promises to attract even more web services vendors and users. This talk is directed at those who are interested in building distributed cross-vendor solutions. The presenters will review Stonehenge goals and activities to date, and announce the latest vendors to participate in Stonehenge. We will also demo the latest version of the StockTrader sample application, using Claims-based security (WS-Federation), and discuss the future direction of Stonehenge as a forum for interoperability testing.

Apache in the Enterprise

Debbie Moynihan

17:00-17:50 | 208

The session will discuss a real world example of how and why an enterprise selected open source software for its infrastructure and the importance of the Apache license and community as part of the selection process.



Microsoft



it's everybody's business to
PLAY WELL with others.

communication from the open source community at Microsoft

port25.technet.com

Munging URLs with mod_rewrite

Rich Bowen

9:00-9:50 | East Hall

It's the bane of every website administrator, and most of us have done it wrong more often than we've done it right. mod_rewrite mangles and munges URLs in order to make them easy to type, easy to remember, and attractive to put on your brochures. In this session, Rich covers the basics of mod_rewrite syntax, a little bit of the regular expression lexicon you need to know, and a few recipes for common scenarios.

Rich is the author of 'The Definitive Guide to Apache mod_rewrite'

Content Driven Portals with Jetspeed and Jackrabbit

David S Taylor

9:00-9:50 | Jewett A - C

If you are looking for a solution that provides end to end content management and portal based delivery, all under a business friendly, open-source Apache license, this presentation is for you. Target audience is software developers, system integrators and managers interested in Apache's offerings of Open Source Portal and Enterprise Content technologies.

This presentation will be short on theory and high on live demonstrations of the Apache Jetspeed and Apache Jackrabbit projects. (Jetspeed is the portal, Jackrabbit is the Java Content Repository (JCR)). Features discussed and demonstrated include everything from content creation, delivery and management, through

portal site management, to the tools available and how to use a JCR Site Toolkit to rapidly develop content-based applications that run transparently in a portal and in plain websites. Using Jetspeed + Jackrabbit + Site Toolkit, a demo will be given showing how to rapidly develop a collaborative enterprise application.

Introduction to the Apache Lucene Ecosystem

Grant Ingersoll

9:00-9:50 | Jewett D & E

Apache Lucene has grown immensely from its early days as a pure search library within Apache Jakarta, to an Apache top level project. The project encompasses a variety of tools aimed at solving problems related to structured and unstructured data. In this overview talk, we'll introduce each of the tools and provide use cases and information on getting started using the Lucene ecosystem.

Apache Web Services Flyover (A Bird's Eye View)

Glen Daniels

9:00-9:50 | Jewett F - H

In this session we'll introduce the Apache Web Services project, and cover the many subprojects and how they all fit together. You'll learn not only the purpose and uses of each piece, but also who's involved, the current status, and what plans are in store for the future. We'll also talk about some of the broader goals for the project itself. You'll come away from this talk fully up-to-speed on all the activity in this busy project.

Making Sense of Open Source Licenses

J Aaron Farr

9:00-9:50 | 208

Are all open source licenses equal? Why are there so many? How are they different? Which should I choose? If you've ever asked these questions, don't worry. You're in good company. "License proliferation," or the increasing number of free and open source licenses in use, has many managers and developers alike confused. In this session, we'll review the most common open source licenses in a simple way that will help you make sense of them. So the next time someone asks you about the difference between the GPL, LGPL, and MPL, you'll know just what to say.

Hardening Enterprise Apache Installations Against Attacks

Sander Temme

10:00-10:50 | East Hall

Enterprise installations of Apache are particularly attractive targets for malicious attacks, including Denial of Service, defacement, theft of data or service and installation of zombies or viruses. Hardening your deployment against such attacks calls for some special techniques and tactics.

Come to this session to learn about attack detection techniques, server protection, secure deployment of multiple servers, configuration of firewall "demilitarized zones", and judicious use of SSL encryption. How do you deploy an off-the-shelf application that insists on writing to the file system? And what steps do you take to securely deploy

Apache on Windows or UNIX? This presentation will explore solutions to these very real situations.

JCR in Action - Content-based Applications with Jackrabbit

Carsten Ziegeler

10:00-10:50 | Jewett A - C

The Java Content Repository API (JCR) is the ideal solution to store hierarchical structured content, and to develop content-oriented applications. This session provides a practical introduction to help you get started using JCR in your own application. To demonstrate the basic architecture of such applications, a sample content-based application will be developed during the session. Basic techniques will be explained, including navigation, searching, and observations, using the Apache Jackrabbit project.

Apache Lucene Basics and New Features

Michael Busch

10:00-10:50 | Jewett D & E

In the first half of this presentation we will talk about the fundamental concepts of Apache Lucene Java. We'll provide a technical overview, as well as tips about how to get started with Lucene. In the second half, we'll introduce the major new features coming with the 2.9/3.0 releases of Lucene, and briefly look at the enhancements that the Lucene developers are currently working on.

The Axis2 Landscape

Eran Chinthaka

10:00-10:50 | Jewett F - H

This session focuses on Axis2, the project's core message engine. We'll

cover the architecture and history of Axis2, and describe how all the pieces fit together. This includes both core components used by Axis2 such as Axiom, and plugins that fit on top, like our modules for reliability and security. After this talk you'll have a good grasp of how Axis2 works under the covers, and how everything works together to provide a flexible, fast, and interoperable Web Services platform.

Life in Open Source Communities

Bertrand Delacretaz

10:00-10:50 | 208

Open Source communities often seem to have their own unwritten rules of operation and communication, their own jargon and their own etiquette, which sometimes make them appear obscure and closed to outsiders. In this talk, we'll provide recommendations on how to get touch with, and how to join, Open Source communities. Based on ten years of experience in various Open Source projects, we will provide practical information on how to communicate effectively on mailing lists, how to formulate questions in an effective way, how to contribute in ways that add value to the project, and generally how to interact with Open Source communities in ways that are mutually beneficial. This talk will help Open Source beginners get closer to the communities that matter to them, and help more experienced community members understand how to welcome and guide newcomers.

Standing Out in the Crowd

Kirrily Robert

11:30 -12:20 | West Hall

How far have we come in ten years? Kirrily Robert has been working as a web developer and Internet technologist for more than a decade. In that time, the world wide web -- driven by Apache -- has become ubiquitous, and with that ubiquity has come a broadening of the market: from geeks and specialists to anyone and everyone.

But while half of all web users are women, women make up only a tiny fraction of the open source developers whose work underpins most websites. A 2006 survey showed 1.5% of developers are female, compared to 24% in IT roles across the US. Many women in open source have had the experience of being the only one on their project or at an event. Kirrily will talk about the experiences of women in open source, and present two new open source projects that have 75% or more women contributors. What can we learn from these majority-female projects, and how can it improve open source overall?

Scalable Internet Architectures

Theo Schlossnagle

14:00-14:50 | East Hall

We will dive deep into traditional web architectures and discuss what makes them go. Looking at real-life architectures, we'll analyze weak points and discuss what will happen to them under the intense pressure of prolonged exposure on Digg, NYTimes or even MSN (when traffic can jump from 15Mbits/second to over 1Gbit/second inside 30 seconds). In a friendly

atmosphere, you'll get to enjoy a mix of anecdotes, war stories, good and bad design principles and more tips and tricks than you could use in a year.

Topics will include the importance of visualizing data, web protocols (HTTP/HTTPS) and browser behavior, firewalls and load balancers, HTTP requests, sizes, pipelining and what that means in packets, remedial queue theory and why apps just "stop working", solid caching techniques for correctness and performance, scary caching techniques for emergencies, techniques for service isolation (quality of service protection), databases, caching and their torrid love affair, and service queues, how they work and why they usually don't!

Rapid JCR applications development with Sling

Felix Meschberger

14:00-14:50 | Jewett A - C

Apache Sling is an OSGi-based, scriptable applications layer, using REST principles, that runs on top of a JCR content repository. In this talk, we'll see how Sling enables rapid development of JCR-based content applications, by leveraging the JSR 223 scripting framework. We'll also look at the rich set of OSGi components provided by Sling. We will create a simple application from scratch in a few minutes, and explain a more complex multimedia application that does a lot with just a few lines of code. This talk will help you get started with Sling and understand how the different components fit together.

Apache Solr: Out of the Box

Chris Hostetter

14:00-14:50 | Jewett D & E

Apache Solr is an HTTP-based enterprise search server, built on top of the Lucene Java search library. In this session, we will see how quick and easy it can be to install Solr, and configure it to provide full-text searching of structured data. There's no need to write any custom code for this! We will cover various built-in features including data loading from CSV files, tolerant parsing of user input, faceted searching, highlighting matched text in results, and retrieving search results in a variety of formats (XML, JSON, etc....).

We will also look at Solr's Administrative interface during the session. This interface can help us understand how different text analysis configuration options affect our results, and why various results score the way they do against different searches. No previous Solr experience is expected.

Web Services in the Real World

Afkham Azeez

14:00-14:50 | Jewett F - H

In this session, we'll tell some user stories about utilizing Apache Web Services technologies in production environments. We'll cover both the good and the bad in a series of interesting case studies, and demonstrate how the work we do at Apache translates directly into customer successes.

How 10 years of Apache has changed my life

Ted Leung

14:00-14:50 | 208

This talk will be a tour through my involvement with the ASF, starting back in 1999 with the Apache XML project. Along the way, we'll visit major milestones in the history of the foundation, reflect on the ways that the organization and community have changed, and try to look ahead at some of the challenges facing us in the next 10 years.

Scalable Internet Architectures, continued

Theo Schlossnagle

15:00-15:50 | East Hall

CouchDB from 10,000 feet

J. Chris Anderson

15:00-15:50 | Jewett A - C

This presentation takes a look at CouchDB from 10,000 ft. We try not to lose you in technical details, but rather to paint the big picture to help you understand CouchDB's strengths and weaknesses. CouchDB is a document oriented database. It does not adhere to the relational principles of traditional databases. In this session, you will learn what that means for your application design. CouchDB's replication support solves the problems of high-traffic web sites, distributed peer-to-peer, and offline applications all at the same time. We show you what kind of applications you are able to build with that.

However, CouchDB is no Silver Bullet and there are cases where it is a poor fit. You will learn to decide when to use it for your project

and when you are better off with a different solution.

Introduction to Apache Nutch

Andrzej Bialecki

15:00-15:50 | Jewett D & E

Introduction to Apache Nutch - the future of Open Source web search. Nutch is a framework for building web-scale crawlers and search applications. This talk will present the challenges related to Web search, solutions to these challenges implemented in Nutch, and how Nutch is used in vertical search. The talk will also outline future directions of the project.

Axis2 Quickstart

Deepal Jayasinghe

15:00-15:50 | Jewett F - H

Apache Axis2 is one of the most successful open source Web service frameworks. Axis2 has a number of different features which can be used to build a full SOA stack. It is also fully inter-operable with commercial Web service frameworks like Microsoft .Net. It supports common Web service specifications, and provides easy integration with existing applications over the protocol of your choice. Above all, Axis2 provides an easy-to-use API to enable you to develop and consume Web Services.

In this session, you'll learn how to use Axis2 to develop and consume Web services. We'll look at Axis2 iPOJO, an easy way to make a Web service from a Java class, and learn how to invoke a simple service-- whether you prefer to use REST, ServiceClient, or RPCServiceClient. We'll look at examples of how to generate and deploy server-side

code from a WSDL definition, and how to generate client stubs and invoke a service. We'll also do a quick introduction to Axis2 pluggable extensions, from Deployers to MessageReceivers, from Observers to Transports, and look at sample code to use the Axis2 module extensions.

Teaching and Learning About Open Development

Ross Gardler

15:00-15:50 | 208

The management structure of open source software projects is unlike the management structures found in most workplaces. Similarly, the process of open source software development is unlike the process taught in almost all universities and colleges, or the processes practiced in many organisations. A newcomer to an open source project is faced with seemingly insurmountable barriers, including the need to engage openly with a wide community. Exposing your weaknesses to an unknown public body, and having those weaknesses recorded in public archives or indexed by search engines, is a scary prospect for even the most confident of developers. Even after this initial fear is conquered, a contributor must face the fear of exposing unfinished, and potentially buggy, code to peers for public review.

In this presentation, Ross Gardler summarises his experiences as a Google Summer of Code administrator for Apache since 2005. He will present his experiences in the context of his role as manager of OSS Watch, an

open source advisory service to the UK education sector. Ross will also look at various approaches to teaching open source. Finally, he will go on to suggest ways in which the mentoring model can be more widely applied, to ensure that open source projects are better equipped to guide newcomers in the art of open development.

Recent Developments in SSL and Browsers

Rick Andrews

16:30-17:20 | East Hall

Certificate Authorities have developed variations on SSL server certificates with degrees of identity verification and authentication, necessitating changes to browser security UIs. This talk will survey and explain the many changes that have occurred in this space, describing the types of attacks that have been successful in the past, and the changes that have been made to mitigate those attacks.

Apache POI recipes

Paolo Mottadelli

16:30-17:20 | Jewett A - C

The Apache POI project provides Open Source Java APIs for the manipulation of Microsoft Office format files. It was developed to provide OLE2 Compound Document format support. POI support for the new format was necessitated by the proliferation of new Office Open XML (OOXML) documents, due to its standardization. As a result, a common challenge emerged for projects that leverage POI to read and write Excel, Word, and PowerPoint documents: supporting the new format while maintaining backward compatibility with the

earlier one.

This session provides an overview of how the new POI architecture makes that challenge easier, using the common interfaces package and their double implementation. Participants will also learn about the main new features provided by POI towards support of the new OOXML format. To demonstrate POI's features, this session will also drive through a collection of practical recipes to solve the tough problems of integrating Office documents in your enterprise applications.

Apache Lucene and Apache Solr Performance Tuning

Mark Miller

16:30-17:20 | Jewett D & E

Apache Lucene and Apache Solr are known to be quite fast out of the box, but some insider knowledge can still lead to big gains. In this talk, Mark Miller will cover many of the factors at play when tuning Lucene/Solr performance and how to configure them for maximum speed. The talk will also cover Lucene's role in Solr, as well as tips on setting up Solr's schema and solrconfig settings.

Securing Web Services with Apache WS

Nandana Mihindukulasooriya

16:30-17:20 | Jewett F - H

Security is one of the challenges that can prevent widespread adoption of Web Services. Despite the popular security model defined

in WS-Sec* specifications, many users still don't fully understand how to ensure Web Services security. In this session, we will cover the basic concepts of security in Web Services--from authentication and authorization, through confidentiality and integrity, to non-repudiation.


Apache Rampart is the security module of Axis2 web services framework. It is a Web Services security implementation build on top of Apache security projects including Apache WSS4J and Apache XML Security. This session will include examples of how to deploy and configure Rampart to secure an Axis2 web service. It will also introduce the WS Security Policy language and the Rampart configuration language, which are used to configure Apache Rampart.

Great Code comes from Great Community: Meritocracy in Action

Chris J. Davis

16:30-17:20 | 208

This presentation discusses how to use the meritocracy model to foster community participation in open source and free software. We will take a look at the dynamic community that can be built around a project or idea, and how this inclusive model can accelerate development, attract quality contributors, and most importantly, keep those contributors engaged.



Use ur brain.



Think Geek

\$10 OFF
\$40 ORDER

Code: APACHECON09
Offer excludes shipping

Smart gift ideas from thinkgeek.com

Implementing an Information Retrieval Framework for an Organizational Repository

Sithu D Sudarsan

9:00-9:50 | East Hall

Successful Information Retrieval (IR) frameworks for large repositories have been reported in recent times. Invariably, all of them have used machine readable repositories, where plain text availability is the norm. However, organizations with legacy archives need to develop a framework which first converts the non-electronic archive to an electronic archive, and then extracts machine readable text with an acceptable error rate.

The Food and Drug Administration (FDA) has electronic images of their document repository. These documents were collected as part of their charter to approve and monitor products related to health care. The documents date back multiple decades, and are in formats ranging from microfiche through early optical character recognition to recent electronic formats.

We believe that a large knowledge base hidden in the FDAs document repository could be mined. To mine this knowledge base, we are developing a semantic-mining framework using open source tools such as Apache Lucene, Apache PDFBox, Apache Solr, Apache Poi, and Java.

Challenges include determining the quality of text being extracted and the ability to handle documents containing formatted text in part. The text itself may contain specific vocabularies from medical, legal,

engineering and scientific domains, and terminology that evolves over time. Careful thought needs to be given to selecting analyzers for indexing and retrieval and implementing a framework for heuristics useful to domain experts as well as novices.

An initial prototype, with a sample size of over 100,000 documents and 70GB of data, is currently being evaluated for different extractors, analyzers and search heuristics, with multiple indices for each document stored in a distributed fashion.

Welcome to the Future!

William A. Rowe

9:00-9:50 | Jewett A - C

This panel will cover all the latest and greatest new features scheduled for Apache httpd 2.4. An unmissable round-up of all the ways httpd 2.4 will make your life easier, presented by the people who wrote the code.

ApacheDS: what's in the box?

Emmanuel Lecharny

9:00-9:50 | Jewett D & E

Apache Directory Server is an embeddable LDAP server, written in Java and certified by OpenGroup. This presentation will cover some of the features of Apache DS, and improvements we have made recently. We will also discuss how Apache DS compares to other LDAP servers.

Some of the features we'll present include support for LDAP and Kerberos, multi-backend support, OpenLDAP replication and multi-master replication, triggers and stored procedures, and the

ChangeLog that allows you to roll back any operations! Finally, we'll take a look at Apache Directory Studio, which is not specifically dedicated to ApacheDS, but is a complete set of GUI tools to manage any LDAP server.

OSGi and SOA

Paul Fremantle

9:00-9:50 | Jewett F - H

OSGi is changing the way we build Java applications and solutions. In this talk, Paul Fremantle will examine how OSGi can help build better SOA applications and integration. The talk will explore how OSGi is changing the fundamental design of middleware, and especially how open source SOA middleware is being re-engineered to avoid the pitfalls of large proprietary platforms.

Looking at how OSGi is being used in Apache projects, this talk will examine the benefits and challenges of using OSGi to build large distributable enterprise systems. Paul will also examine real life experiences building an Open Source SOA platform - Carbon - based on the Eclipse Equinox, Apache Axis2 and Apache Synapse. Looking at the challenges, benefits and lessons learnt, this session will be a practical introduction to developers on how to get the best out of OSGi and the systems that use it.

Getting Started with Apache OFBiz in 5 Easy Steps

Ruth Hoffman

9:00-9:50 | 208

Apache OFBiz brings Enterprise Resource Planning (ERP) to the

masses, with a no cost, Open Source, enterprise-ready suite of business automation applications. Built on best-of-breed technology, the database-agnostic OFBiz ERP platform makes cost-effective, reliable and secure business automation a reality. No longer are ERP systems beyond the reach of the small and medium sized business owner. In this session, we'll discuss getting OFBiz up and running quickly, to maximize return on investment and minimize upfront costs. Using a real life example, we'll demonstrate that it is possible for any organization to build a sound business using the Apache OFBiz Open Source ERP platform.

Apache Mahout - Going from raw data to information

Isabel Drost

10:00-10:50 | East Hall

It has become very easy to create, publish, and collect data in digital form. The volume of structured and unstructured data is increasing at a tremendous pace. This has led to a whole new set of applications that can be build to solve the problem of turning raw data into valuable information. Possible applications include everything from discovering new trends out of a stream of weblog entries, to automatic learning approaches that supplementing market research processes for new products. Machine learning provides tools for building these applications.

A large community of researchers has been working on the topic of learning from data. Although solutions to common problems are publicly available, scaling these

solutions into the range of terabytes and petabytes is an open issue. To scale algorithms to such dimensions it is vital to distribute data as well as computation. The mission of the Apache Mahout project is to build a suite of scalable machine-learning algorithms that can cope with today's quantities of data. Mahout is built on top of Apache Hadoop.

This talk provides a beginner-friendly introduction to the topic of machine learning. It presents a broad set of applications that benefit machine learning, as well as a high-level overview of Mahout. It also covers the types of tasks that can be solved with each algorithm, and the pitfalls to look out for along the way.

Deciphering mod_ssl: Using SSL with the Apache HTTP Server

Joe Orton

10:00-10:50 | Jewett A - C

mod_ssl is one of the most complex modules shipped with the Apache HTTP Server. This presentation will start by explaining how to set up an SSL server using httpd and mod_ssl, then discuss in detail how support for the SNI extension allows use of name-based virtual hosting with SSL. Finally, we'll cover use of advanced features from client certificate authentication and fine-grained access control through to session caching and performance.

The OpenLDAP/Apache Directory Project Alliance

Alex Karasulu

10:00-10:50 | Jewett D & E

The Apache Directory Project and the OpenLDAP Community both have BSD-derived licenses. Both

projects are conducted with the same spirit of collaboration, to produce open source software that benefits our downstream users. Both are dedicated to the increased and proper utilization of directories, while looking towards new advances that improve ease of use and adoption.

Recently the projects have been working together to enable a high degree of interoperability between ApacheDS and OpenLDAP. The award winning Apache Directory Studio has been altered to serve as a tooling platform for OpenLDAP in much the same way as it does for ApacheDS. A powerful alliance has formed between the two organizations that will impact LDAP users, the market, and how we utilize LDAP in this modern age, decades after it's inception. We will discuss this alliance, the changes it's brought about, its impact on our users, and what it means for other LDAP vendors.

Apache Felix iPOJO - OSGi Made Simple

Karl Pauls

10:00-10:50 | Jewett F - H

The OSGi Service Platform provides an excellent foundation for building dynamically extensible Java-based applications. However, such power and flexibility does not come without cost. In addition to new APIs and packaging formats, OSGi developers must shift their thinking to be able to design highly dynamic applications. Early efforts, such as the Service Binder and Service Tracker, attempted to alleviate some of these issues. There have also been more recent efforts,

such as Declarative Services and Spring Dynamic Modules. Another such effort in this area, called Apache Felix iPOJO, is a service-oriented component model created specifically for the OSGi Service Platform.

The purpose behind iPOJO is two-fold. The first goal is to simplify OSGi development to the bare necessities. This means that providing services, using services, and dealing with dynamism should require as little effort as possible from the developer. The second goal is to push beyond support for basic OSGi mechanisms and try to more seamlessly incorporate advanced features, including component configuration, synchronization, and composition.

This presentation will briefly discuss how iPOJO achieves these goals, with the main focus being on how iPOJO looks from the perspective of the developer. Particular emphasis will be placed on features that are unique to, or more easily achieved, in iPOJO.

Collaboration on Requirements and Designs for Apache OFBiz

David E Jones

10:00-10:50 | 208

To date, most of the collaboration in the Apache OFBiz project has been on implementation. What to implement and how is frequently discussed, but there have been very few efforts to produce artifacts that represent requirements that can drive designs, and designs that can drive implementation. In this presentation David will offer an overview of UBPL, the HEMP

artifacts it is based on, designs based on parts of UBPL and applications being developed based on those, and overlap analysis between UBPL and what exists in OFBiz. With these tools, business people can directly participate in the project and developers can more effectively communicate and understand the context in which they are working. The result of these efforts will be software that more directly meets the needs of a variety of end-users.

MIME Magic with Apache Tika Jukka Zitting

11:15-12:05 | East Hall

Apache Tika aims to make it easier to extract metadata and structured text content from all kinds of files. Tika is a subproject of Apache Lucene, and leverages libraries like Apache POI and Apache PDFBox to provide a powerful yet simple interface for parsing dozens of document formats. This makes Tika an ideal companion for Apache Lucene, or for any search engine that needs to be able to index metadata and content from many different types of files.

This presentation introduces Apache Tika and shows how it's being used in projects like Apache Solr and Apache Jackrabbit. You will learn how to integrate Tika with your application and how to configure and extend Tika to best suit your needs. The presentation also summarizes the key characteristics of the more widely used file formats and metadata standards, and shows how Tika can help deal with that complexity. The audience is expected to have basic

understanding of Java programming and MIME media types.

mod_lua in Apache 2.4

Brian McCallister

11:15-12:05 | Jewett A - C

Mod_lua lets you write efficient and fast modules in Lua for the multi-threaded Worker and Event MPMs. Lua is a very fast scripting language specifically designed for embedding and extension -- making it a perfect match for the Apache HTTP Server. Learn how to setup, write a module, optimize it, and easily move from Lua to C and back again with mod_lua!

Apache Geronimo 2.2

David Jencks

11:15-12:05 | Jewett D & E

This talk will provide an overview of the new features found in Apache Geronimo 2.2, and discuss Geronimo Custom Assemblies and Application Security. Geronimo makes it easy to assemble a custom server around your applications, including only those features needed to run your apps. We'll see how to use Apache Maven to integrate this into your build process. This session will also cover the clustering options available with Geronimo, and how to configure load-balancing and failover of web applications using mod_jk & Apache HTTP Server.

After completing this session, the audience will be comfortable with the clustering and farming features provided by Geronimo, running multiple instances of Geronimo server from the same installation, preparing their applications for deployment on the cluster and

configuring Geronimo clusters to suit their needs.

Tuscany: Applying OSGi modularity after the fact

Luciano Resende

11:15-12:05 | Jewett F - H

Apache Tuscany is an open source project that simplifies the development, deployment and management of distributed applications built as compositions of service components. It is based on the Service Component Architecture specifications being defined by the OASIS Open SCA Collaboration. Tuscany was built with a modular architecture, using a different approach from that of OSGi. In 2008 an effort was started to integrate Tuscany with OSGi. As part of this we undertook an investigation into how to apply OSGi modularity to the Tuscany runtime, so that clean boundaries between modules are enforced and different versions of the same library can coexist. We started by analyzing the existing Tuscany modularity, which turned up lots of interesting information about the linkages between the various sub-components! This understanding was then used to determine how to map Tuscany into a suitable form for use as OSGi bundles, including prototyping various levels of decomposition granularity.

This presentation will share the experiences of analyzing and modularizing an existing project using OSGi. It will discuss what to expect when approaching modularization for existing projects, the tools (and caveats) for aiding

analysis, as well as best practices for applying OSGi modularity. Beyond enabling OSGi for the runtime, we are also going to cover how to integrate OSGi and SCA at the application level.

Selling Open Source E-commerce and ERP

Mike Bates

11:15-12:05 | 208

As open source becomes more mainstream, the pool of potential open source e-commerce and ERP clients continues to grow. In this presentation Mike Bates, CEO of HotWax Media, discusses the act of including open source in the sales process with different target audiences; what works, what doesn't, the anatomy of a success story, and perhaps a catastrophic failure or two.

How Open Source Developers Can (Still) Save the World

Brian Behlendorf

12:15 - 13:05 | West Hall

Open Source communities like Apache have done far more than just create awesome, no-cost, flexible software - they have demonstrated a whole new way for communities of interest to band together and solve big problems. Beyond operating systems and web servers, Open Source software is now being used to organize disaster relief efforts, enable microfinance banks in the developing world, integrate electronic healthcare systems, and change the way citizens collaborate with their government. How well are these projects going? What should they emulate or attempt to inherit from

the more "infrastructural" kinds of projects and organizations - such as Apache - and what new kinds of challenges emerge? And where can you help?

Building Intelligent Search Applications with the Lucene Ecosystem

Ted Dunning

14:00-14:50 | East Hall

Apache Lucene has evolved in recent years beyond a core search library, into a top level project that contains a whole suite of tools for working with content. Starting with Solr, which builds on the core Lucene search library, we can add in open source libraries like Tika, Mahout, and Droids to build intelligent search applications. This talk will focus on how to leverage the various components of the Lucene Stack, to build out intelligent search applications. These applications can better enable users to find what they are looking for, in today's sea of content.

Writing modules for Apache httpd

Jim Jagielski

14:00-14:50 | Jewett A - C

This session will provide an overview of the Apache HTTP server API and its interaction with the Apache Portable Runtime (APR) library. The concepts of hooks, providers and optional functions will be discussed as they relates to the Apache startup and request-response cycle. In addition, filters and bucket-brigades will be introduced, all to provide the attendees with a basic understanding of how Apache and modules interact. With this

information, people will be able to understand how current modules work and will have the necessary information to write their own Apache httpd modules.

OSGi Blueprint Container Specification and Geronimo

Jarek Gawor

14:00-14:50 | Jewett D & E

This talk will provide a general overview of the Blueprint Container Specification and demonstrate some of its core features using Apache Geronimo's Blueprint Container implementation. The talk will also describe some more advanced features of the Geronimo implementation such as Configuration Admin Service support and custom namespace handlers. Future plans will be discussed, including how the Blueprint Container might be used and integrated within the existing Geronimo Kernel.

Dynamic Deployment with Apache Felix

Marcel Offermans

14:00-14:50 | Jewett F - H

The OSGi framework allows you to build your applications out of dynamically updatable components, giving you a lot of possibilities to automatically install, update and remove these components without restarting your applications. In this talk we will show how standard OSGi services can be used to implement a system to install and update software components, as well as updating configuration and other data. Apache Felix, which implements the OSGi specification, includes deployment and resource processing components that will

be used as part of our application. This session will also demonstrate the installation of both OSGi and non-OSGi applications onto various target devices.

Functional Testing Apache OFBiz Applications with Selenium XML

Brett Palmer

14:00-14:50 | 208

Functional testing for ERP applications, like those in Apache OFBiz, can be very laborious and time consuming. In this presentation, Brett explains how OFBiz users can create effective functional tests for their applications using a framework called Selenium XML. Selenium XML is an extension of Selenium that allows a user to create and run functional tests in a simple XML format. The framework supports multiple browsers and provides functional and visual verification of a real user experience. This presentation will cover recording tests and converting them to Selenium XML, parameterizing those tests, combining functional and system level tests, and automating test runs. Brett will also present possible future directions, including converting test cases to functional documents, auto-generating test cases based on OFBiz artifacts, and adding scripting support and AJAX libraries to Selenium XML.

Realtime Search

Jason Rutherglen

15:00-15:50 | East Hall

What is near-realtime search in Apache Lucene? How can it be used to enhance your existing search applications? What are the limitations? We'll go over how near-

realtime search is implemented in Lucene, include use cases for when to use it and when not to, and discuss how to replicate index changes with near-realtime performance in environments with a high volume of queries.

Testing Apache Modules with Python and Ctypes

Markus Litz

15:00-15:50 | Jewett A - C

Writing tests for your Apache module is often a developer's least favourite task, especially if you don't like using the Perl Apache Testing Framework! One of the main reasons for this is that it's difficult to test your C code on-the-fly without a running Apache server.

Using ctypes, you can test your modules without a running httpd, just by writing and using simple Python scripts! In this talk we will look at how to compile a specific version of the Apache webserver for testing, and how to use a common Python unit testing framework to write and set up tests. We will also cover examples of writing simple test cases for the Catacomb WebDAV Server module.

Apache Geronimo 3.0: OSGi and Java EE6

David Jencks

15:00-15:50 | Jewett D & E

This talk will provide a preview of the new features in Apache Geronimo 3.0. This will include the new OSGi support with the Geronimo kernel, and new Java

EE 6 features: Profiles, JPA 2.0, Servlet 3.0, EJB 3.1, etc. We'll also include demonstration of current functionality and discuss the plans for further development.

Embrace OSGi - A Developer's Quickstart

Carsten Ziegeler

15:00-15:50 | Jewett F - H

In theory, the first choice for highly modular, dynamic, and extensible applications is OSGi technology. The theory sounds very tempting, but what about the real world? Starting with the basics of OSGi, this session is focused on practical examples, tools, and procedures for a rapid adoption of OSGi in your own projects. Learn how to avoid the typical traps and how to get the most out of OSGi.

Roundtable Discussion: Usability, Content Management, Current Areas of Apache OFBiz Activity

David E Jones

15:00-15:50 | 208

This roundtable discussion will focus on a variety of aspects of the Apache OFBiz project. The discussion will provide an opportunity for people to share related ideas, become familiar with some of the ideas that have been discussed on the mailing lists, and find out which ideas are popular! The goal of the session is to share experiences, opinions, and ideas about what opportunities exist, and what sorts of tools and practices we can use to go after them.



The open source SOA company

Enterprise Service Bus

Governance Registry

Identity Server

Web Services Application Server

Data Services Server

Mashup Server


Web Services Framework

Gadget Server

Business Process Server

Business Activity Monitor

Business Rules Server



**Lean open source middleware
for your enterprise**

Powered by Apache Projects

<http://wso2.com>

Developer portal <http://wso2.org>

Aaron Kimball

Aaron Kimball has been working with Hadoop since early 2007. Aaron has worked with the NSF and several other universities nationally and internationally to advance education in the field of large-scale data-intensive computing. He helped create and deliver academic course materials first used at the University of Washington, which were later adopted by many other academic institutions, as well as Hadoop training materials used by several industry partners. Aaron has also worked as an independent consultant focusing on Hadoop and Amazon EC2-based systems. Aaron holds a B.S. in Computer Science from Cornell University, and an M.S. in Computer Science and Engineering from the University of Washington.

Afkham Azeez

Afkham is a committer on the Apache WS and Apache Synapse projects, a member of the Project Management Committee for Apache WS, and a Member of the Apache Software Foundation. He currently works as a software architect at WSO2, and is the Product Manager of WSO2 Web Services Application Server.

Andrzej Bialecki

Andrzej Bialecki is a member of the Apache Lucene Project Management Committee, and a committer to the Apache Nutch subproject. He is also the author of Luke - the Lucene Index Toolbox. He participates in many commercial projects that use Lucene, Apache Solr and Apache Hadoop to implement enterprise and vertical search.

Asankha C. Perera

Asankha is a member of the Project Management Committee, and a committer, for several Apache projects, including Apache Synapse ESB, Apache Web Services and Apache HttpComponents. He is also an independent consultant providing

services around Apache Synapse ESB.

Ashish Thusoo

Ashish is a member of the Apache Hadoop Project Management Committee, and a committer on the Apache Hive project. He works on the Data Infrastructure team at Facebook. His interests are in the areas of Cloud Computing, Distributed Systems and Data Processing, and Querying frameworks. In the past he worked at Oracle, building components in the Oracle RDBMS, and made contributions and holds patents in the areas of Parallel Query Execution, SQL execution and XML Databases.

Bertrand Delacretaz

Bertrand Delacretaz works as a senior developer in Day Software's R&D group, using open source tools to create world-class content management systems. Bertrand is a member of the Apache Software Foundation and has been or is involved in the Apache Cocoon, FOP, Solr, Sling, Tika and ESME projects.

Bradford Stephens

Bradford is the Lead Engineer on the Data Platform team at Visible Technologies. He works with massively distributed systems to analyze Social Media. Technologies at the core of his business include Apache projects such as Hadoop, Nutch, Lucene, Solr, Mahout, and more!

Brett Palmer

Brett has been developing software professionally for the past 15 years. He currently works as a software consultant helping companies and governments implement large ERP systems. He has been developing custom business applications using Open for Business (OFBiz) since 2002. For the 5 years, Brett has been working on several offshore projects combining software developers in the US with

think

“...before you send that flame.”

“...about Sponsoring Apache.”

“...on how community should work.”

“...about the other person’s shoes.”

“Those who would give up essential Liberty, to purchase a little temporary Safety, deserve neither Liberty nor Safety.” ⇔ Benjamin Franklin, 1755



offshore teams working in India. Brett holds an Electrical Engineering degree and a Masters of Business Administration from Brigham Young University. He also holds a Masters in Computer Science from Utah State University.

Brett Porter

Brett Porter is a Member of the Apache Software Foundation. He serves on various Project Management Committees, and has been working with Apache Maven since 2003. He lives in Sydney, Australia, and his dayjob is VP, Product Development at G2iX. He is the co-author of the books "Apache Maven 2: Effective Implementation" and "Better Builds with Maven".

Brian McCallister

Brian McCallister is a programmery kind of person who loves that folks are willing to pay him to write

code. He works as a Software Architect at Ning, a platform for building your own social websites. Brian has worked variously as a programmer, technical writer, and systems administrator for over a decade, on projects ranging from telecommunications GIS systems to loom (weaving) control software, and as an English teacher and canoeing guide to boot. He is proud to be a Member of the Apache Software Foundation.

Carsten Ziegeler

Carsten Ziegeler is senior developer for JEE and portal applications at Day Software. He is a member of the Apache Software Foundation and has participated for more than fifteen years in several open source projects. Carsten is a member of several Apache communities and Project Management Committees including Cocoon, Felix, Sling, Excalibur and Portals.

Chris Hostetter

Chris Hostetter is a Apache Member who serves on the Apache Lucene Project Management Committee. As a Principal Software Engineer at CBSi he has spent the last 10 years of his life thinking about searching "structured data" that is never as structured as it should be. If he is not sending typo-riddled emails to the Lucene-Java and Solr mailing lists, he is probably welding something shiny out of stainless steel and silicon bronze.

Chris J. Davis

Chris J. Davis is a Developer living in North Texas. Currently Chris is a Senior Developer for Blastoff Communications, where he is building a thing that will change the way you think about this other thing. When not changing games, Chris is a core developer of the Habari Project and a published author. You can find his ramblings at Sillyness, werd.

Christian Wenz

Christian Wenz is an author, trainer and consultant with a focus on web technologies and web application security. He has written or co-written more than 100 books, and frequently speaks at developer conferences around the globe. His "real" job is as one of the founders of Arrabiata Solutions GmbH, where he leads web projects and conducts security and usability audits.

Christophe Bisciglia

Christophe Bisciglia joins Cloudera from Google, where he created and managed their Academic Cloud Computing Initiative. Starting in 2007, he began working with the University of Washington to teach students about Google's core data management and processing technologies - MapReduce and GFS. This quickly brought Hadoop into the curriculum, and has since resulted in an extensive partnership with the National Science Foundation (NSF) which makes Google-hosted Hadoop clusters available for research and education worldwide. Beyond his work with Hadoop, he holds patents related to search quality and personalization, and spent a year working in Shanghai. Christophe earned his degree, and remains a visiting scientist, at the University of Washington.

Dave Johnson

Dave is a software developer, blogger and web geek from Raleigh, North Carolina. He's a member of the Apache Software Foundation, working on the Apache Roller open source blog server project. And, until recently, he was a social software architect at Sun Microsystems, working on Roller, SocialFish and Project SocialSite.

David E Jones

Collaboration on Requirements and Designs for Apache OFBiz, Roundtable Discussion: Usability, Content

Management, Current Areas of OFBiz Activity.

David Jencks

After studying mathematics in college and graduate school, David Jencks spent 14 years building baroque musical instruments. Returning to an early interest in programming, he has been involved with open source projects since around 2000. A founding member of Apache Geronimo, he is also a committer on ActiveMQ, Directory, Portals, and OpenEJB at Apache and Jaybird, Jencks, Mojo, and Tranql elsewhere. He has been an IBM employee since 2005.

Debbie Moynihan

As director of Progress Software's FUSE Community, Debbie has extensive experience with the joys and challenges of the business aspects of open source. The FUSE Community was created to help users adopt and deploy Apache ServiceMix, Apache ActiveMQ, Apache Camel, and Apache CXF in enterprise environments. The FUSE distributions are tested, certified versions of these popular Apache projects and are backed by world-class support and services. Progress Software employs many of the key committers to these projects. Prior to joining Progress Software, Debbie was responsible for Product Line Management for the WebSphere Process Integration portfolio. She has a B.S. in Mechanical Engineering and an M.S. in Management from MIT.

Deepal Jayasinghe

Deepal is a computer science graduate student at Georgie Institute of Technology. Before joining Georgia Tech he worked at WSO2 for about three years, where he actively contributed to Apache Software Foundation and a number of WSO2 projects. Deepal has been involved in the design, development and implementation of Apache Axis2 since its inception. Deepal is also an Apache Member, and a PMC member for both Apache Web Services and



COLLABNET®

Supporting Open Source for 10 Years and Counting...

- 25% of all code written by CollabNet engineers donated to open source
- Helping more than 800 enterprises run their own open source-style collaborative development communities
- Hosting over 2 million users across thousands of open source projects

Tigris.org

250,000 users | 1400 projects

Projects: TortoiseSVN, Subclipse, ViewVC, ArgoUML, SCons, Scarab, SCPlugin, WebSVN, tsWebEditor, more...

openCollabNet

375,000 users | 250 projects

Projects: AnkhSVN, desktop-vs, desktop-eclipse, truMerge, SVNbinaries, SharpSVN, ClearSVN, geek-track, more...

- > Open Source in the enterprise: www.collab.net/oss_ent
- > Start your project on  **OpenCOLLABNET**: www.open.collab.net/cmmtty

www.collab.net • info@collab.net • +1-650-228-2500

Apache Synapse. He is also a frequent speaker at ApacheCon and author of "Quickstart Apache Axis2". His current research work involves Cloud Performance, Cloud Scalability, and Web services.

.....

Emmanuel Lecharyn

Emmanuel is currently the chair of the Project Management Committee for the Apache Directory Server project, and an ASF member. With more than 20 years in the field, he still likes to develop, which is one of the reasons he participates in the Directory project! System architecture and optimizing existing applications are what keep him fed on a daily basis.

.....

Eran Chinthaka

From Ambalangoda, Sri Lanka, Eran Chinthaka is a member of ASF and also a PMC member of Web services and Synapse projects. He is also a PhD candidate in Indiana University, Bloomington, Indiana, researching on improving scientific job executions in large scale systems. In addition to his involvements in open-source and research he also loves photography, playing cricket and table tennis.

.....

Erik Hatcher

Erik Hatcher is an Apache Software Foundation member, and an active committer on the Apache Lucene and Apache Solr projects. Erik co-authored the award-winning book Java Development with Ant, and the well-reviewed Lucene in Action. Erik has spoken frequently at industry conferences, including JavaOne, ApacheCon, OSCON, and the No Fluff, Just Stuff symposium circuit. Erik co-founded and is a member of the technical staff of Lucid Imagination, a full service Lucene/Solr shop.

.....

Felix Meschberger

Felix Meschberger works as a senior developer for Day Management AG, creating content management systems

using open source and internal tools. Felix is a committer to the Apache Jackrabbit and Apache Felix projects and contributed Apache Sling to the Incubator in September 2007.

.....

Filip Hanik

Filip is a Senior Software Engineer for SpringSource and a key participant in SpringSource's Apache Tomcat initiatives. Filip brings 12 years of extensive experience in architecture, design and development of distributed application frameworks and containers and is recognized for his top-quality system development skills and continuous participation in Open Source development projects. Filip is a committer to the Apache Tomcat project where he is a leading authority on Tomcat clustering and a key contributor to the core of the platform. He also authored the online book Hitch-Hiker's Guide to Tomcat. Prior to SpringSource, Filip was a Senior Software Engineer for Covalent Technologies, which was acquired by SpringSource in early 2008. Previously, Filip worked as a Senior Software Engineer/Architect for La Quinta Corporation. Filip has also made contributions to software initiatives for Walmart.com, Sony Music, France Telecom and has held a variety of senior software engineering positions with technology companies in both the United States and Sweden. He received his education at Chalmers University of Technology in Gothenburg, Sweden where he majored in Computer Science and Computer Engineering.

.....

Glen Daniels

Based in Arlington, MA, Glen is an independent consultant working primarily in the SOA and Open Source spaces. He previously worked for both large and small companies (including Macromedia, WSO2, and Progress Software) in various leadership roles, and participated in many of the Web Services standards groups. He also co-authored 2 editions of the award-winning "Building Web Services with

Java". When not writing code, advising clients or talking about technology, he enjoys creating and experiencing both fine food and groove-oriented music. Glen is a Member of the ASF, and currently chairs the Project Management Committee for Apache Web Services.

.....

Grant Ingersoll

Grant Ingersoll is a committer on the Apache Lucene Java, Apache Solr and Apache Mahout projects and a member of the Lucene PMC. Grant has used Lucene and Solr for implementing many different search and text-based applications, including cross-language information retrieval systems, question answering, sentiment analysis, collection analysis and customization applications. Grant is also a contributor to the Apache Mahout machine learning project.

.....

Henning Schriedehausen

Henning Schriedehausen is a contributor to a number of Apache Java projects. He works as a freelance consultant, architect and software developer using the J2EE platform and admits under torture that he can program in PHP and perl. When not sitting in front of a computer, Henning enjoys traveling with his wife around the world, sports (both active and passive) and moonlighting as a 11th level barbarian at his local D&D group.

.....

Howard M. Lewis Ship

Tapestry 5: Java Power, Scripting Ease!
Howard Lewis Ship is the creator of the Apache Tapestry and Apache HiveMind projects, and a frequent speaker at JavaOne, NoFluffJustStuff, ApacheCon and other conferences. He expertise covers web application development, enterprise architecture, dependency injection and improving developer productivity in general.

Isabel Drost

Isabel Drost is a contributor to Apache Mahout. She is employed at neofonie GmbH, a company building enterprise search engines. Isabel interned at Google for six months in 2005/06. She has also worked as research assistant in Berlin. Isabel holds a master's degree in computer science from the University of Applied Sciences, Mittweida.

.....

J Aaron Farr

J Aaron Farr has contributed to the Apache Avalon, Excalibur, Incubator and Labs projects. After developing J2EE and Eclipse RCP software for Sony Electronics and Siemens, he founded The JadeTower Corporation. Aaron currently resides in Hong Kong.

.....

J. Chris Anderson

J. Chris Anderson is an Apache CouchDB committer and co-author of the forthcoming O'Reilly book "CouchDB: The Definitive Guide". He is a director of couch.io, offering commercial hosting, support, consulting, and custom development. He enjoys working on JavaScript CouchApps which can be peer-replicated just like any other data. Chris is obsessed with bending the physics of the web to give control back to users.

.....

Jarek Gawor

Jarek Gawor is a software engineer at IBM. He is an Apache Member, and a committer and PMC member for the Apache Geronimo, Apache Axis2 and Apache CXF projects. He has been working on Apache Geronimo since 2007.

.....

Jean-Frederic Clere

Jean-Frederic has spent more than 20 years writing client/server software. He is a committer to the Apache APR, Jakarta, Httpd and Tomcat projects, and he likes complex projects where different languages and machines are involved. Born in France, Jean-

Frederic lived in Barcelona (Spain) for 14 years. Since May 2006 he has lived in Neuchatel (Switzerland) where he works for RedHat in the JBoss division.

Jeremy Thomerson

The first web application that I published commercially was purchased by eBay. I stayed on with eBay for several years, assisting and mentoring teams that were migrating several applications over to a scalable Java-based architecture. After those were complete, I was off to new adventures - which led me to Wicket. I have been using Wicket since version 1.1, and have deployed over two dozen successful web applications of various sizes with it.

Jim Jagielski

Jim's been active on the 'Net since the early 80's, starting as editor of the A/UX FAQ. He worked on the NCSA server and joined the Apache Group (as it was called back then) at a very early stage. He actively contributes on HTTPD, APR and Tomcat, but also hacks on other projects (ASF and others) as well in addition to mentoring many ASF incubator podlings. In addition to being a charter and core member of the ASF Jim serves as Director and Chairman for the foundation. His real job is as Chief Architect and Principal Software Engineer for the Covalent Products Division of SpringSource. Previously, Jim served as CTO for Covalent Technologies, which was acquired by SpringSource in early 2008.

Joe Orton

Joe Orton has been writing free software for ten years and is the author of a number of open source WebDAV tools such as "cadaver". He works in the UK for Red Hat as maintainer of web-server packaging, from Apache httpd to PHP and beyond. Joe is a member of the Apache Software Foundation, and is a committer on the httpd and APR projects.

Jukka Zitting

Jukka Zitting is an active open source developer who is currently involved in too many Apache projects. He sees software development as modern day craftsmanship and believes in open source as the best way to practice the craft. Jukka spends his free time exploring Switzerland or visiting his fiancée and their two cats back in Finland.

Kent Brown

Kent spent 20 years in software development and consulting, mostly on middle-tier and integration technologies. He has published numerous technical articles, spoken at technical conferences, and founded developer user groups. Kent joined Microsoft in 2008, is the Technical Product Manager for Windows Communication Foundation, and oversees cross-vendor web service interoperability efforts.

Karl Pauls

Karl Pauls is the head of Mobile Applications and OSGi Development at akquinet AG. He leads projects in the OSGi and mobile application area and is a member of the Apache Software Foundation. With more than six years experience, he is a long time OSGi enthusiast and committer and member of the Apache Felix PMC. He is also co-authoring the "OSGi in Action" book.

Luciano Resende

Luciano Resende is a Senior Software Engineer at IBM Silicon Valley and a member of The Apache Software Foundation. Currently Luciano is a PMC Member of the Apache Tuscany project contributing to the Service Component Architecture area, he also contributes to other Apache projects such as Apache PhotArk. Luciano has interest in SOA, Web 2.0 and Cloud Computing areas, and some of his thoughts can be found at <http://lresende.blogspot.com>.

Marcel Offermans

Marcel is a software architect at luminis with broad experience in different operating systems, languages and applications, ranging from data-oriented enterprise applications to embedded and distributed systems. At luminis iQ products, he is responsible for the architecture and development of an OSGi based software provisioning product. Marcel is a PMC member of the Apache Felix open source project and interested in all Java and OSGi related development.

Mark Miller

Mark Miller has spent the last two years building a large newspaper archive system with Apache Lucene, which is used by many large newspapers in the US and Canada. He now works at an open source Lucene/Solr startup. He is an active member of the Lucene/Solr community.

Mark Thomas

Mark has been using and developing Tomcat for over five years. He first got involved in the development of Tomcat when he needed better control over the SSL configuration than was available at the time. After fixing that first bug, he started working his way through the remaining Tomcat bugs and is still going. Along the way Mark has become a Tomcat committer and PMC member, volunteered to be the Tomcat 4 release manager, created the Tomcat security pages, become a member of the ASF and joined the Apache Security Committee. He also helps maintain the ASF's Bugzilla instances. Mark has a MEng in Electronic and Electrical Engineering from the University of Birmingham, UK and is currently a Senior Software Engineer with the Covalent Products Division of SpringSource.

Michael Busch

Michael Busch is a Software Architect at the IBM Silicon Valley Lab in San Jose, California. He has been working for 3+ years on Enterprise Search

applications, such as IBM's OmniFind and eDiscovery products. He has a Bachelors Degree in Computer Science from the University of Paderborn, Germany. Michael is a committer on the Apache Lucene Java project.

Michael Stack

Michael Stack is on the Hadoop Project Management Committee, and is sponsored by Powerset (part of Microsoft) to work full-time on Apache HBase. HBase, which is an important component of Powerset's development, is developed as part of the Apache Software Foundation's Hadoop project, and runs on top of the Hadoop Distributed File System, providing BigTable-type capabilities.

Mike Bates

Mike founded HotWax Media in 1997, and his career in sales, marketing, design and web development spans a dozen years. Mike is focused on business development, acquiring and growing key accounts, and developing a profitable business model around open source software. He has led high-profile projects for numerous national and global brands and has taught web development courses at the graduate level. After graduating Summa Cum Laude from the Monterey Institute of International Studies with a B.A. in International Studies, he earned his MBA in International Management from the Fisher School.

Mladen Turk

Mladen Turk is a Principal Software Engineer at Red Hat Europe, where he is responsible for native integration. He is a long time ASF member and committer for the Apache Jakarta, Apache Tomcat, Apache Httpd and Apache Portable Runtime projects. He is currently acting as chair of the Apache Tomcat PMC.



GET HIPPO TIZED

The Hippo stack gives you all the tools you need to develop large scale content driven web sites, portals and intranets for enterprise collaboration. Based on the Apache Java stack* it gives you all the freedom you need. And if you or your

customer needs any help, call Hippo. We'll support you.

Want to know more? Visit the Hippo booth in the Exhibition Hall for a live Demo; attend one of our speaker sessions or go to www.onehippo.com



HIPPO

* Jackrabbit • Wicket • Jetspeed • Cocoon • Lucene • Tomcat • Maven

SALES@ONEHIPPO.COM WWW.ONEHIPPO.COM

USA • SAN FRANCISCO HIPPO USA INC. • 185 H STREET, SUITE B • PETALUMA CA 94952 • +1 707 773 4646
EUROPE • AMSTERDAM HIPPO B.V. • OOSTEINDE 11 • 1017 WT AMSTERDAM • +31 20 522 4466

Nandana Mihindikulasooriya

Nandana Mihindikulasooriya is a committer and a PMC member of Apache Web Services project, and mainly works on Apache Rampart and Apache Axis2. Currently he is studying for his Masters degree on Software Engineering in Blekinge Institute of Technology, Sweden. Prior to starting his masters programme, Nandana worked as a Technical Lead in WSO2 Inc.

enterprise workflow design. Paolo is an active committer and Project Management Committee member for the Apache POI project, involved in the development of its support for Office Open XML. As Sourcesense's presales manager, he is committed to educating customers on Enterprise Open Source adoption. The four-time Italian national decathlon champion is currently in training, hanging out at tracks and fields worldwide.

Owen O'Malley

Apache Hadoop

Owen O'Malley is a software architect on Apache Hadoop working for Yahoo's Grid team, part of Yahoo's Cloud Computing & Data Infrastructure group. He has been contributing patches to Hadoop since before it was separated from Nutch, and is the chair of the Apache Hadoop Project Management Committee. Before working on Hadoop, he worked on Yahoo Search's WebMap project, which builds a graph of the known web and applies many heuristics to the entire graph that control search. Prior to Yahoo, he wandered between testing (UCI), static analysis (Reasoning), configuration management (Sun), and software model checking (NASA). He received his PhD in Software Engineering from University of California, Irvine.

Paul Fremantle

Paul Fremantle is CTO of WSO2, where he works on Open Source projects in Apache, including the Apache Synapse and Incubator projects. He has contributed to Apache since the first Apache SOAP project. While at IBM, he was instrumental in starting up the Apache WSIF and Apache Woden projects, as well as being heavily involved in the AxisC/C++ initiative, where he led IBM's involvement. Paul was a Senior Technical Staff Member in IBM, where he was the lead architect and co-creator of IBM's Web Services Gateway. Paul is the co-chair of the OASIS WS-RX technical committee and led the JSR110 committee (JWSDL). Before joining IBM, Paul worked as a consultant in the pharmaceutical industry. Publications include co-authoring "Building Web Services in Java, 2nd Edition", articles on Web Services and SOA, and a redbook - "The XML Files: Using XML and XSL in WebSphere". Paul has presented many times at ApacheCon, Colorado Software Summit, XML Europe, Software Architecture and other industry conferences. Paul has an M.A. in Mathematics and Philosophy and an M.Sc in Computation from Oxford University. He also plays the tin-whistle.

Paolo Mottadelli

Paolo Mottadelli is a senior developer at Sourcesense, an Open Source systems integrator. He has been focusing on Open Source content management for five years, gaining a deep knowledge of world-class CMS solutions such as Magnolia, Hippo, and Alfresco ECM, as well as the JCR content repositories standard. Over the last two years, he has been designing complex architectures for document management, including documents capture, processing and indexing, and

Raymond Feng

Raymond Feng is a senior software engineer with IBM. He is now a PMC member and committer of Apache Tuscany, working on the Service Component Architecture (SCA) runtime. Raymond has been developing



MuleSoft

the web middleware company

Creators of:

- ▶ **Mule ESB** — the #1 open source enterprise service bus
- ▶ **iBeans** — dead-simple integration for web applications
- ▶ **Tcat Server** — enterprise Tomcat with provisioning and diagnostics

Congratulations to the ASF

10 years of open source leadership

www.mulesoft.org

www.mulesoft.com



The Leading Provider of Apache Tomcat Support and Expertise Celebrates the 10th Anniversary of the Apache Software Foundation



The Enterprise Java Leader

Full Commercial Support Available

Apache HTTP • Apache Tomcat • Apache ActiveMQ

www.springsource.com

A division of **vmware**

SOA platforms for more than 7 years and is one of the pioneers at implementing SCA. Prior to this role he was a key developer and team lead for WebSphere

Rich Bowen

Rich Bowen is the IT Director for ClearMyRecord.com in Lexington, Kentucky. Rich is the author of Apache Cookbook and The Definitive Guide to Apache mod_rewrite. He is a member of the Apache documentation project and of the Apache Software Foundation, and is one of the voices on FeatherCast.

Rick Andrews

Rick Andrews is the Product Line Architect for SSL products at Thawte (a VeriSign company). His primary responsibility is to design and maintain the software architecture of Thawte SSL and Code Signing products and

services. In addition, he chairs an internal Software Architecture Board and manages the team that performs security reviews of Thawte products. He's been with Thawte (VeriSign) for 12 years, the last 5 in the SSL group. A long time ago, he received a BSEE from The Cooper Union in NYC, and a MSCS from UCLA.

Ross Gardler

Ross has filled many roles during a varied working life ranging from the music industry (managing a popular Dub Reggae act in the late '90's), research and teaching (computer science and management and business administration) in both the UK and West Indies, and a freelance software engineer working with open source software in order to deliver solutions well beyond the normal capabilities of a micro-business. Today he leads a team providing open source advice and guidance to the UK Higher and

Further Education sector, assisting with both institutional procurement and open source software development in research projects. He is an active contributor to a number of open source projects and is a Member of the Apache Software Foundation.

.....

Sally Khudairi

Active in the Web since 1993, Sally Khudairi is the publicist behind some of the industry's most prominent standards and organizations. The former deputy to Sir Tim Berners-Lee and long-time champion of collaborative innovation, she became involved with The Apache Software Foundation in 1999, and was elected its first female and non-technical member. Sally is chief executive of luxury brand communications consultancy HALO Worldwide.

.....

Sander Temme

Sander Temme is an Enterprise Solutions Engineer for a security company whose clients include Fortune 500 companies, financial services companies and government agencies. He is a member of the Apache Software Foundation and is active in the httpd, Infrastructure and Gump projects. Sander is owned by Murphy, the wonder cat.

.....

Sandy Gupta

Sandy Gupta is the Director of the Platform Strategy Team in Microsoft's Server and Tools organization, where he leads strategic initiatives for the Server platform. Sandy's primary focus is to drive Mission Critical workloads on Windows Server and also enable key Open Source software like Java, AMQP, Linux/Unix virtualized guests, and others on the Windows Server and Virtualization platform by working together with Microsoft development teams, strategic partners and open source communities. Prior to joining Microsoft in 2008, Sandy spent more than 16 years in the Unix industry,

starting his career as developer for Unix kernel and device drivers at ICL. Subsequently Sandy worked as a senior executive at other Unix companies, where he drove their technology architectures and product roadmaps. Sandy holds a Bachelor Degree of Computer Engineering from the University of Delhi and is an alumni of Kellogg School of Management.

.....

Suresh Marru

Suresh Marru is a senior researcher in the Pervasive Technology Institute at Indiana University and leads the scientific workflow efforts. Mr. Marru is a Co-Principal Investigator on the Linked Environments for Atmospheric Discovery (LEAD) and Open Grid Computing Environments (OGCE) projects. Mr. Marru also co-ordinates the TeraGrid science gateways program at IU and has been working very closely with GridChem and UltraScan Gateway projects. Suresh received a B.E. in Electrical and Electronics Engineering from Osmania University, India and an M.S in Electrical and Computer Engineering from University of Oklahoma. Marru's research interests include distributed, grid and cloud computing and building complexity hiding interfaces and workflows for various science domains.

.....

Ted Dunning

Ted Dunning is CTO at DeepDyve, a startup offering specialized search capabilities to knowledge workers. Previously, he was Chief Scientist at Veoh Networks (an on-line video provider), ID Analytics (the leading identity fraud modeling company), Musicmatch (named product of the year 7 years in a row by PC magazine) and Aptex (early developer of internet based intelligent applications). Before working at startups, he was an academic working on a broad range of advanced computing problems. The common theme throughout has been pattern discovery and the analysis of symbolic sequences in language, genetic sequences, web-browsing behavior, musical preferences,

purchasing behavior and financial transactions.

Ted Leung

Ted Leung is a principal engineer in the Cloud Computing division of Sun Microsystems. He is an ASF member and has been involved with Apache since 1999.

Theo Schlossnagle

Theo Schlossnagle is a Principal Consultant at OmniTI Computer Consulting, where he designs and implements scalable solutions for highly trafficked sites and other clients in need of sound, scalable architectural engineering. He is author of Scalable Internet Architecture, published by Sams. Theo is also the author and maintainer of the mod_backhand load-balancing module for Apache, an author and maintainer of the Wackamole IP failover tool, the initiator of the Backhand Project and an active participant in a plethora of open source projects.

Tim Funk

Tim Funk has been a Tomcat user for over 7 years and a committer since 2003. Tim currently resides in Pennsylvania and works for Armstrong World Industries in a variety of roles, enjoying the role of developer the most. Tim has a Masters of Software Engineering from Penn State.

Tom White

Tom White has been an Apache Hadoop committer since February 2007, and is a member of the Apache Software Foundation. He is writing "Hadoop: The Definitive Guide" for O'Reilly. He works for Cloudera, a startup providing Hadoop support. Previously he worked as an independent consultant specializing in Hadoop. Tom has a Bachelor's degree in Mathematics from the University of Cambridge, and a Master's degree

in History and Philosophy of Science from the Universities of Leeds, UK, and Florence, Italy.

William A. Rowe


William's contributions to Apache include numerous enhancements to the Win32 port of the HTTP Server project, such as CGI, security, service control, file system support, and APR design targeted at the Win32 native API. He provides Win32 hints to Apache related lists, and has been a speaker at previous ApacheCon events. As a member of the ASF and the Apache httpd and APR projects, and a Senior Software Engineer with the Covalent Division of SpringSource, his work on Apache continues in areas such as integration of Apache 2 with the Win32 security model and add-in modules such as mod_aspdotnet. William started his career developing an array of customized and revenue document imaging systems. Prior to joining Covalent, he provided consulting services in revenue document generation and management, data transformation, application integration and Web interface services.

Thank You to Our Track Sponsors

Apache Tomcat:



Httpd.conf:



Visit us at the
Lucene Meetup!

Tuesday
Nov. 3rd
8:00pm
rooms 1&2



Lucid Imagination is the commercial entity exclusively dedicated to Apache Lucene/Solr open source search technology. We offer the Lucene/Solr community a wide range of online resources, including:

- in-depth technical articles
- case studies of successful implementations
- search.lucidimagination.com
complete online Lucene community knowledge resources

We also offer products and services that can help you and your solution partners develop and deploy search solutions with confidence, including professional training, SLA-based support subscriptions, best practices consulting, free certified distributions of Lucene and Solr and value-add software.

Thinking Lucene ▼ Think Lucid.



Check out our Lucene & Solr
Technical Resources at
<http://bit.ly/lucid-lucene>

lucid
IMAGINATION

Proud Silver Sponsor of ApacheCon US 2009
Apache Lucene and Apache Solr are trademarks
of the Apache Software Foundation.

www.lucidimagination.com
tel 650.353.4057



The **Apache**
Software Foundation
<http://www.apache.org/>

would like to thank...

The Apache Software Foundation could not exist without the continued generous support from the community.

We would like to take this opportunity to acknowledge and thank our Sponsors.

Platinum Sponsors

YAHOO![®]

Google[™]

Microsoft[®]

Gold Sponsor



Silver Sponsors



Bronze Sponsors

AirPlus International
BlueNog
Intuit

Joost
Matt Mullenweg
Two Sigma Investments

Infrastructure Support

Oregon State University Open Source Lab
SURFnet

APACHECON

US 2009



Leading the wave of open source

10 YEARS

Thank you for joining us this year to celebrate 10 years with The Apache Software Foundation and for making ApacheCon US 2009 a great success!

We look forward to welcoming you to our future

Apache Roadshow Asia 2009

Beijing, China November 28th & 29th 2009
Colombo, Sri Lanka December 3rd through 5th 2009

Apache Retreat

Wicklow, Ireland - April 9th-12th, 2010
<http://apache.eventbrite.com/>

ApacheCon Europe 2010

Late Spring or Early Summer 2010
<http://eu.apachecon.com>

2010 brings a change to ApacheCon US as it officially becomes ApacheCon North America

ApacheCon North America 2010

Atlanta, Georgia • November 1 through 5, 2010
<http://na.apachecon.com>

ApacheCon North America 2011

Vancouver, Canada • November 6 through 11, 2011
<http://na.apachecon.com>