

# UK IUA Conference 2011

## Afternoon Break-Out Sessions

Note that this agenda is subject to change. Presentation times and speakers may change.

### OpenROAD Track

**14:00 Sean Thrower, Ingres Europe**

#### **OpenROAD Displays with no-code Gradients, Roundels, Highlighting<sup>†</sup>, Multi-Image Bitmaps, and More—it's Open Season!**

Display styles change, enduser tastes evolve, eventually each product gets trapped in its own display cabinet...

But with customer collaboration driving change, and some lateral thinking, the future can change too—in the coming OpenROAD release you can map your display scheme (defined as a set of bitmap images) chosen without lock-in to any development product, directly into the frame display: fields can display the corners you want, gradients can be of any orientation, and include special effects; a single bitmap can support fields of any size; bitmaps can contain multiple images, for smart highlighting, three-way selection, embedded dropdown—and all of this out of the box.

<sup>†</sup> (ok, maybe a line of code to switch the highlighting ...)

And you can capture, image-search, manipulate and apply your bitmaps at runtime.

And more. And more to come...

**14:45 Kim Ginnerup, Bording Data A/S & Roy Deal Simon (Consolidate Systems)**  
**Don't Miss The Train! Actionable advice for migrating Ingres codebases from ABF to n-tier OpenROAD on to HTML5**

Do you have mission critical ABF and wish it was in OpenROAD and/or on the web? Do you have mission critical 2-tier OpenROAD and wish to exploit the OpenROAD Application Server? Do you have mission critical 3-tier OpenROAD and wish to exploit HTML5 & CSS3 front-ends? Well, listen in and see where you can hop on the train!

**15:45 Bodo Bergmann, Ingres Europe**

#### **OpenRoad XML and Hash Table Classes unleashed**

With the event of OpenROAD 5.1 there are a bunch of new system classes provided, which offer a wealth of additional possibilities.

This session will explain their usage and provide several examples and ideas of how to exploit the new features, e.g. for:

- Implementing fast lookups
- Including performance comparisons with commonly used techniques as find methods of ArrayObject and ChoicList.
- Serializing objects into/from XML
- Use for platform-independent parameter passing
- Exchange with other XML based technologies

**16:30 Joe Kronk, Ingres Corporation**

#### **OpenROAD Status and Roadmap**

In this session we will look at exciting new features that have just been released in OpenROAD 5.1.

Many of these features were started at code sprints held around the world. The foundational feature of OpenROAD 5.1 is XML support which will open up many development opportunities. I will examine how we take advantage of XML in Workbench to enable you to prepare your application for Unicode. Once we are finished with this tour of OpenROAD 5.1 features, we will peer into the short term roadmap for OpenROAD. Here I will discuss the planned OpenROAD 6.0 release in which we cross the chasm from Ingres 2.6 to Ingres 9.2 code base that OpenROAD 6.0 is built on. Finally we will look at some future features that are currently under development. Hopefully we will have some time at the end of this session for a few questions.

## Ingres DBMS & DBA Track

### 14:00 James K. Lowden, AllianceBernstein Ingres as the Next Big Thing

Ingres is uniquely situated to rescue the Relational Model from a market that would drown it if it could. To do so, it must embrace a new query language and do more to facilitate open-source contributions.

Relational database technology stagnated long ago, and the promise of the Relational Model has nearly disappeared from the minds of today's programmers, IT management, and even DBAs. In its place the ignorati propose or are adopting pre-relational "successors" that lack even a rudimentary theoretical foundation.

The greatest challenge to Relational technology isn't technological, but educational. The number of people in the industry who lack both a computer science education and knowledge of database management history is large enough to constitute a market. Following Santayana's famous dictum, they are repeating history, reinventing hierarchical databases from the ground up. It is a sad, sad story.

The second greatest challenge is SQL. SQL is inadequate and inappropriate to today's Relational databases. Designed in the 1970s as an end-user query tool by people who didn't understand the Relational Model, it is today used -- when not ignored entirely -- only by programming professionals. SQL's disregard for Relational precepts and implicit disdain for its users is keenly felt even by programmers who can't quite articulate the problem. It is hardly surprising that a language invented by IBM in the Age of COBOL finds little favor among adherents to, say, Ruby-on-Rails. SQL has become a barrier to adoption of Relational database technology!

Changing this will not be easy. What role can Ingres play?

Unique to Ingres among database servers is the combination of its language-independent relational algebra layer and its open-source licensing. Ingres supports 1.5 non-SQL languages: QUEL and Tutorial D. QUEL could use some updating, and the Tutorial D implementation is acknowledged by the author to be incomplete.

QUEL and Tutorial D are both more faithful to Relational Algebra than is SQL. To the expert programmer, educated in Relational theory, these languages are more powerful than SQL because they implement relational constructs directly. By the same token, they present to the server a smaller query optimization challenge. The result is more powerful queries, written more quickly, executed more predictably.

As the \*best\* open-source database server technology, not least because of its multi-language architecture, Ingres can, with very little direct investment by Ingres Corporation, become a non-SQL alternative demonstration platform. Ingres can permit and promote full implementation of Tutorial D and the modernization of QUEL, thereby providing a basis for comparison to SQL believers and doubters alike.

### 14:45 Richard David, Hewlett Packard Partial Ingres Upgrade/Migration

AAH Pharmaceuticals used to run both Ingres applications and databases on one tier (Ingres 2.6 on

Sun Solaris 32bit). Today the applications remain on the 'old' hardware, but the databases have been migrated to modern 64bit Linux servers i.e. Client/Server.

This solution isn't ideal by any means, and there are plans to migrate the applications to 64bit Linux in time, but it has enabled AAH to enjoy cheaper running costs and significantly better performance with little or no changes to their applications.

There are 'gotchas' from running Ingres 2.6 Clients against Ingres 9.2 and Ingres 10 and most of these are understood and accepted.

### 15:45 **Dr. Wojtek Rappak, Rational Commerce Ltd.** **Life of Brian—The Story of a VectorWise Appliance**

About a year ago Rational Commerce launched a range of Vectorwise appliances. Each of these is a Linux machine fully configured from the hardware up through Linux, with the latest version of Vectorwise running on top. A mid-range, moderately priced prototype was built and installed on a company LAN. For reasons which have nothing to do with Monty Python, its node name is Brian. This presentation will be about how we built Brian, what we learned building him and on the impressive performance we achieved from Vectorwise running on Brian.

When he's not out doing a demo, Brian's favourite place is a 6U rack on a server farm. He is also set up to run the official TPC-H benchmark up to scaling factor 300 (300GB). Rational Commerce is in the software business, but we do love hardware. So we wanted to design a machine which would use the best components money could buy to get top performance out of Vectorwise at the lowest possible price.

Brian is mostly built from Intel components: chassis, motherboard and processors are from Intel. He has two Xeon 5500 processors providing 8 cores; RAM stands at 32GB. A key feature is the use of SSDs (Solid State Drives) which takes Vectorwise performance up to an even higher level. We tried a number of SSD configurations before settling on what we think is the best one.

We learned a lot from Brian.

### 16:30 **Doug Inkster & Karl Schendel, Ingres Corporation** **Comparing and Contrasting the Architectures of Ingres and VectorWise**

[details pending]

## VectorWise Track

### 14:00 **Marty Bowes, Oxford University CTSU** **VectorWise Speed Thrills**

This presentation is simply an introduction to vectorwise concepts. There will be a brief discussion on installation, configuration and getting started with tables. Current limitations will be discussed. The new SQL command `call vectorwise (combine ...)` will also be mentioned. The paper will then show the results of some performance tests on some reports generated from large tables. These reports are essentially table pivots where a user may select any number of fields to pivot the tables around.

### 14:45 **Dan Koren, John Galloway & Rilson Oscar do Nascimento, Ingres Corporation** **VectorWise Tuning (Part 1)**

[details pending]

### 15:45 **Dan Koren, John Galloway & Rilson Oscar do Nascimento, Ingres Corporation** **VectorWise Tuning (Part 2)**

Continuation of previous session.

**16:30 Mark Whalley, Ingres Europe**

## **Analysing High-Speed Streams of Data**

More than just imagining business analytics at the speed of thought, with VectorWise we now have a revolutionary database engine that leverages vector-based processing and on-chip memory that does just that.

In real-life business scenarios, as well as querying large volumes of static data, there is the ever increasing need to capture and record continuous data feeds often referred to as “streaming” data.

Such examples of high volume streaming data will include stock market tick data, EPOS, network traffic, telecommunications, ATM transactions, sensor data et al.

This presentation demonstrates how, using Vectorwise tools, multiple streams of data can be loaded into VectorWise tables at rates in excess of 100,000 rows per second whilst at the same time servicing high performance sub-second queries.

## **Blue Skies Track**

**14:00 Alex Hanshaw, Ingres**

### **Getting Started with Ingres in the Amazon Cloud**

The presentation explains what the cloud is, goes into some specifics for the Amazon Cloud and runs a demo starting at login to the Amazon Cloud and running through to connecting to a running Ingres AMI via the Ingres Database Workbench (IDW).

**14:45 TBA**

**15:45 TBA**

**16:30 TBA**

## **Tools Track**

**14:00 Chris Clark, Ingres Corporation**

### **Unicode Migrations**

Do you have an existing application that needs to support Unicode data? Migrating an application and your data isn't quite as straight forward as you might imagine. Discover areas you'll need to understand before going down the Unicode path in moving your application and your data into a Unicode world. Learn strategies for using Unicode enabled applications with old non-Unicode enabled applications and data sources.

**14:45 Roger Whitcomb, Ingres Corporation**

### **The New Ingres Management Studio**

Last year the design of the new Ingres Management Tools was presented along with a working prototype. Many changes to the design have been made in the last year, and much progress has been made on the real implementation. So, the current state of the Ingres Management Studio (IMS as it is now being called) will be presented, along with an overview of the design and features. A working demonstration of the IMS being used to browse and control multiple Ingres installations will be made.

This presentation will present the overall design and feature set of the Ingres Management Studio as currently specified; lay out the current state of the development of the IMS including a glimpse of the roadmap for the product, and finally, demonstrate the current working version

**15:45 Steve Shiel & Niels Bouwman, HVR Software by  
Replication into VectorWise using HVR**

[details pending]

**16:30 TBA**