Dear Tridex Db2 LUW Members,

We are pleased to announce the Tridex Db2 LUW October 11th meeting agenda. As a reminder, you must register for the meeting at the new Tridex website www.tridex.org. Email registration will no longer be accepted.

TRIDEX DB2 for LUW

Featured Speakers:

David Kalmuk, Senior Technical Staff Member & Master Inventor, IBM
Dave Williamson, Business Development, Emerging Analytics Technology, IBM
George Baklarz, Worldwide DB2 Technical Support, IBM
Dale McInnis, Senior Technical Staff Member, IBM

Third Quarter User Group Meeting DON'T JUST KEEP UP - STAY AHEAD!
All Meetings are FREE. You must Pre-register to Attend. Walk-in registration will not be accepted. A photo ID is required for building access.

Third Quarter Meeting October 11th, 2018

Location: IBM Corporation 590 Madison Ave (between 56th & 57th St.)
Room 1219, 12th Floor

Meeting Agenda

8:30 – 9:00	Coffee (provided by IBM)
9:00 – 10:15	David Kalmuk, IBM ***Advanced Performance Diagnostics for SQL***
10:15 – 10:30	Break
10:30 – 11:45	David Kalmuk, IBM ***Adaptive Workload Management in Db2 Warehouse***
11:45 – 12:30	Lunch (provided by IBM)
12:30 – 1:30	Dave Williamson, IBM ***Data Virtualization and Project Queryplex***

1:30 – 2:45 George Baklarz, IBM
RESTful Jupyter notebooks in a Cloudy Environment

2:45 – 3:00 Break

3:00 – 4:15 Dale McInnis, IBM
DB2 HADR Multiple Standby Deep Dive

Please distribute the fliers to others in your organization who may wish to attend. Pre-registration is required for security reasons.

Sincerely Tridex BOD,
Celia Gahagan, Angela Arettines, Joe Geller, Ulf Heinrich, Umesh Mehta, Chris Tsounis, Steve Tsounis, Nate Murphy

Abstracts & Biography

Session title: Advanced Performance Diagnostics for SQL

Session abstract: Db2 provides a number of important features for diagnosing and solving your SQL problems. Through a series of practical examples, this session will introduce you to the latest Db2 capabilities available for SQL diagnostics and show you how you can leverage them to solve query problems in your own environment. You will learn how to identify your most expensive SQL statements using the package cache table functions, and how to use DB2's "time spent" metrics to analyze where your problem queries are spending their time. You will also learn how to use the latest sort memory metrics to analyze query memory consumption in order to tune your analytic workloads. Finally you will learn how to use the runtime explain capabilities in conjunction with the activity event monitor to capture actual plan cardinalities as well per object metrics for individual plans.

Session title: Adaptive Workload Management in Db2 Warehouse

Session abstract: IBM's Db2 Warehouse (and Db2 Warehouse on Cloud) has recently rolled out a new Adaptive Workload Management technology that allows you to easily manage your database resources across workloads and ensure they meet their desired performance objectives. This exciting new technology automatically manages the scheduling and execution of work in your system based on resource targets you assign, ensuring stable and reliable performance even in the face of the most highly concurrent workloads. This session will introduce you to these new capabilities and show you how you can leverage them in your own environment.

Speaker biography: David Kalmuk is a Senior Technical Staff Member & Master Inventor currently responsible for the MPP scale-out, Workload Management, and Monitoring areas of the Db2 family products. David has contributed to the development of numerous technologies in Db2 over the years including BLU Acceleration, Workload Management, Monitoring, as well as much of Db2's Processing and Communications architecture. He is currently leading development efforts focused on Db2's Hybrid Warehousing Architecture. David has been a member of the Database team at IBM since 2000.

Session title: Data Virtualization and Project Queryplex

Session abstract: Regardless of the industry, data has become a fundamental driver in today's business decisions. Businesses pivot and remain competitive based on the analytics performed on that data. Increasingly that data comes from more and more sources, its volume is increasing at an exponential rate, and need for insight is increasingly real-time. Data is dispersed in silos across the enterprise, derived from external sources, and increasingly coming from remote locations or devices that are internet connected. And as the number of data sources increase, that data becomes increasingly heterogeneous in nature. To date the physics have required complex data architectures where data is copied, moved and replicated to satisfy the analytics needs of business users. This has resulted in increased cost and complexity, loss of data integrity and accuracy, and governance headaches. Data Virtualization is a means of avoiding needless copying and moving of data to perform analytics. Instead analytics are performed real-time on the data where it resides, and both physical structure and location are made irrelevant to users through virtual views. IBM's Project Queryplex has been in-progress now for a couple of years, and it is a game changing breakthrough in Data Virtualization. The technology has been in a beta trial period since early 2018, and the first GA of the technology will form the Data Virtualization capability within the IBM Cloud Private for Data platform in November. Subsequently this technology will replace existing Federation technology across IBM's Hybrid Data Management product line. This session will overview the Data Virtualization problem domain, Project Queryplex, the release roadmap, a live demo, followed by Q&A.

Speaker biography: Dave Williamson leads Business Development for Emerging Analytics Technology (including Data Virtualization). Dave joined IBM in 2011 and has over 35 years of experience in the Software Industry. He has held a number of Sales leadership roles in IBM Canada and at the IBM North America level, and has been in the Analytics organization since 2014. He holds a degree in Computer Science and Statistics, and started his career in software development of mission critical systems for air-traffic control, aerospace and telecommunications. Prior to IBM he was the senior Sales and Services executive for several Software startups (leading to both acquisitions and IPOs).

Session title: RESTful Jupyter notebooks in a Cloudy Environment

Session abstract: Did you know that Db2 on Cloud has a RESTful interface that let's you build apps without any database drivers? This session will show you how to build apps using RESTful calls, explore the use of Jupyter notebooks and finally give you details of how you can sign up for a free proof of technology to try this all by yourself!

Speaker biography: George Baklarz, B. Math, M. Sc, Ph.D. Eng., has spent many years in IBM working on various aspects of database technology. He has worked in every department in the Toronto DB2 Development lab and generally can't keep a job for more than 18 months. George currently works in the Worldwide DB2 Technical Sales group and creates powerpoint slides for a living.

Session title: DB2 HADR Multiple Standby Deep Dive

Session abstract: In this presentation we will examine the HA and DR options available for DB2. This will be followed by a deep drive in the workings of DB2's HADR technology with a specific focus on how running HADR with Multiple Standbys changes operations. We will walk to several failure scenarios describing how the HADR technology will react.

Speaker biography: Dale McInnis is a Senior Technical Staff Member (STSM) at the IBM Toronto Canada lab. He has a B.Sc.(CS) from the University of New Brunswick and a Masters of Engineering(M.Eng) from the University of Toronto. Dale joined IBM in 1988, and worked in the DB2 development team from 1992 - 2014.

Dale's area of expertise includes DB2 for Linux, UNIX and Windows Kernel development, where he led teams that designed the current backup and recovery architecture and other key high availability and disaster recovery technologies. He has obtained eleven (11) patents in the area of database resiliency. Dale is a popular speaker at the International DB2 Users Groups (IDUG) conferences worldwide, as well as DB2 Regional users groups and IBM's Information On Demand (IOD) / IBM World of Watson / IBM Think conferences. His expertise in the area DB2 availability area is well known in the information technology industry. Dale spent 10 months as the IBM Cloud Data Services Chief Availability Architect, and now Dale is now on the North American Hybrid Data Management Technical Sales team.

Registration: Db2 for LUW Users Group October 11th 2018 meeting

www.tridex.org

Enrollment Instructions:

New members must first register on the site:

- **Go to https://www.tridex.org/join-tridex/ to register. (This is a one-time process. You will be emailed a starting password, along with a link to modify that password.)**

Existing members:

- **Go to https://www.tridex.org/login/ to enter your email address and password**
  - **Click Login**
  - **Click Meetings**
  - **Click Upcoming Meetings**
  - **Click Db2 LUW**
  - **Click October LUW meeting**
  - **In the middle of the page – Click Confirm RSVP**
  - **You will receive an email confirmation of your registration**
- **If you need to cancel your registration at a later time, you can also do it from here.**

If you have any problems with the registration process, please contact **celia.gahagan@tridex.org**

Special thanks to IBM for sponsoring our October meeting.