

Central Canada DB2 User Group – June 2018 Abstracts

Keynote

Db2 and the Future of Data

In a world of increasing data scale, an expanding number of data sources and the increasing applicability of data science and machine learning, Db2 is evolving in dramatic ways. At the same time, while the industry's needs evolve many companies face the practical reality that they are unable to move all of their systems to a new paradigm at once. With so much changing the need for trusted incremental transformation, rather than high risk corporate IT revolution, becomes a critical consideration organizations in their adoption strategy. In this talk IBM Fellow and Db2 technology executive will preview a vision of how our industry is changing, and some of the remarkable technology evolving within Db2. Sam will also discuss how IBM is investing to enable users to move to new technologies incrementally when the timing is right for them.

Sam Lightstone, IBM Fellow

Sam Lightstone is an IBM Fellow and Master Inventor in the IBM Analytics group. He leads a number of technical teams in product development for relational databases, data warehousing & big data, cloud computing, analytics for IoT, data virtualization, ground to cloud data movement, and machine learning. He co-founded the IEEE Data Engineering Workgroup on Self-Managing Database Systems. Sam has more than 60 patents issued and pending and has authored four books and over 30 papers. Sam's books have been translated into Chinese, Japanese and Spanish. You can follow Sam on social media. His Twitter handle is "samlightstone".

DB2 for z/OS

Maryela Weihrauch, IBM

Microservices Application and Data Architecture Overview

Almost all companies have ongoing IT projects to support the transformation to a Digital Enterprise. The application architecture discussion is centered around implementing micro services and exposing them as APIs. This emerging workload is frequently called System of Engagement to distinguish it from traditional OLTP workloads which are called System of Records, Microservices assume that the service has complete control of the data that it is providing the services for and it is being decentralized.

Maryela shares her observation on functional and non-functional requirements that drive data architecture decision in support of System of Engagement and consequences from chosen data architecture. She discusses application patterns and data architecture options that would keep the data centralized but would give the customers the flexibility they need in this re-architecture journey.

Db2 RESTful Service Support and z/OS Connect

The Application Programming Interface (or short API) Economy has changed how developers build application and how organization deploy software as a service. Mobile devices have accelerated the shift of providing services via API. Maryela will show how SQL and stored procedures in Db2 on z/OS can be made available in the era of application assembly and integrate with z/OS Connect EE. She will introduce the new DB2 11 for z/OS functionality to define SQL and Stored Procedures as services that can be call from anywhere through RESTful service calls with the availability, scalability and performance that Db2 user expect from distributed Db2 applications.

Speaker biography

Maryela Weihrauch is an IBM Distinguish Engineer and z Systems Analytics world-wide technical sales and client champion. She has extensive experience with DB2 regarding systems, application, and database design. Maryela is very engaged in enterprises world-wide, driving the adoption of new analytical technologies. Her most recent role in DB2 for z/OS Development has involved determining DB2 z/OS strategy for HTAP (Hybrid Transaction and Analytics Processing). This includes IDAA strategy and implementation, as well as DB2's application enablement strategy. Maryela is a member of the IBM Academy of Technology and frequently shares her experience at conferences.

Sheryl Larsen

Revival of the SQL Tuner

The pool of DBAs with SME knowledge of SQL tuning has shrunk dramatically while the number of SQL statements needing attention has grown and will continue to grow. DB2 performance tooling can help identify what needs to be tuned. But what do you do then? This presentation will help fill the skills gap by reviving the role of an SQL Tuner. Specifically, you will learn what access paths are good for what types of queries, what indexes are best, and how to change the DB2 optimizers mind.

Objective 1: Witness the access path choices with graphic animations and understand which paths are good/bad for certain queries

Objective 2: Learn how to read a visual EXPLAIN

Objective 3: Learn how to use traditional query tuning techniques

Objective 4: Learn how to experiment with more extreme tuning techniques

Objective 5: identify SQL tuning opportunities and fix them with confidence

Speaker Biography:

Sheryl M. Larsen is a Sr. DB2 Product Specialist at BMC. Previously she worked for IBM, Sheryl M. Larsen, Inc., and Platinum Technology. She is known for her extensive expertise in SQL Tuning and has performed detailed DB2 Performance Health Checks for many fortune 500 clients. Sheryl has over 25 years' experience in DB2, has published many articles, white papers and co-authored a book, DB2 Answers, Osborne-McGraw-Hill, 1999. Currently she is the President of the Midwest DB2 User Group -www.mwdug.org, and a member of the Northern Illinois University Computer Science Alumni Council.

Saghi Amirsoleymani & Timm Zimmerman, Rocket Software

Db2: A Single-Entry Point for Data Virtualization

Today, organizations face a proliferation of data types – different formats, structures, multiple platforms – growing in volumes and needing to be integrated faster to deliver the real-time insight required by business leaders. It's a daunting challenge complicated by aging tools and outdated data architecture. But what if all your enterprise data could be accessed through a single API, use a single metadata model or catalogue, and utilize a single security model? In this session, we will explore IBM's Data Virtualization Manager for z/OS and its ability to transform Db2 into a complete data virtualization facility. Learn how any application that connects to Db2 can have virtualized access across a wide range of mainframe and non-mainframe data sources. Discover how to turn Db2 into your single-entry point for data virtualization.

Presentation Bullet Points:

1. A Single API, catalog and security model
2. Brings data virtualization layer to Db2
3. Makes Db2 a single point of entry for mainframe/non-mainframe data

Speaker biographies

Saghi Amirsoleymani is a Director, WW Analytics at Rocket Software and IBM Champion, where she evangelizes solutions for customers and partners on IBM technologies. Prior to joining this team, Saghi was part of the IBM DB2 z/OS technical support team and worked onsite premier support for almost 20 years. She began her career at IBM as an intern while earning her degree at the University of Southern California, and therefore is a diehard Trojans fan. Although she's a SF Bay Area native, raised in Mountain View, Saghi has earned her nickname from friends and colleagues as Ms. Worldwide. She travels throughout the world presenting on various topics related to Db2 at conferences, users groups, webcasts and events.

Timm Zimmerman is a Director for Data and Analytics R&D at Rocket Software. He has worked in the enterprise data management for almost two decades - service in a number of roles including Db2 for z/OS development and consulting global mainframe customers worldwide. Timm began his career as part of Db2 for z/OS Development, He also has been part of the Db2 for z/OS SWAT team supporting large customer to enhance and improve the enterprise installations. As of 2014, Timm is part of Rocket Software, where he focused on R&D for several enterprise data management features.

Tony Andrews, Themis

Advanced SQL and the Power of SQL Rewrites

Often times there can 4,5,6 different ways to write an SQL query and get the same results back. What makes one better than any of the others, and is there any ones that are always better, or always worse? Sometimes rewriting 1 predicate in an SQL statement can cause optimization to change. This presentation breaks down many examples of query rewrites and how they can affect performance and optimization. Being strong in SQL is a great skill when it comes to performance and tuning of queries. This presentation is great for anyone who works with Db2 on z/OS or LUW, and especially great for any Db2 developer.

Speaker biography

Tony Andrews has more than 25 years' experience in the development of relational database applications. Most of this time, he has provided development and consulting services to Fortune 500 companies and government agencies. Tony has written literally thousands of queries and programs during his development years, and has also served as a database analyst. For the last 10 years, Tony has been splitting his time between performance and tuning consulting engagements along with training. His main focus is to teach today's developers the ways of RDBMS application design, development and SQL programming -- always with a special emphasis on improving performance. He is a current IBM Champion, and regular speaker at many regional user groups, IDUG NA, and IDUG EMEA. He is also the author of a book for developers titled 'DB2 SQL Tuning Tips for z/OS Developers'.

Frank Rhodes, BMC

An XML Document's Life: Dr. Node!

XML continues to be an emerging technology and more features are being added to DB2 to support it. So what is this new XML data type?

Follow the adventures of 'Dr. Node', an XML document as he gets inserted into a DB2 table. Not only will you see how he is related to his parent row but you will learn about his hierarchical tree structure and the catalog tables involved in supporting him. We will then put the parsed XML document under a microscope, and examine the actual data that is stored. You will also see how to use the new DB2 10 XMLMODIFY command to manipulate individual nodes within Dr. Node, and the effects of the command. We will also look at how a Reorg impacts the DOCTOR. We will end with a discussion on XML indexes and demonstrate how to create indexes with the new DB2 10 Date and Timestamp data types.

Speaker biography

Frank Rhodes started with BMC Software in 1995. In his current role he works to promote cross product coordination between the Database Administration products and the other BMC DB2 products. The products that he has primary responsibility for are ALTER for DB2, Change Manager for DB2, Catalog Manager for DB2, DASD Manager Plus for DB2, NGT STATS, NGT Utility Manager as well as Workbench for DB2..

Frank was a Developer on the Change Manager for DB2 product for over 10 years. His main areas of responsibilities being the development and maintenance for the Analysis Engine as well as the product's user interface. Prior to joining the Change Manager Product team, he worked on the TIS Install for the Administrative products. He was the lead developer in charge of the initial implementation of a common install for all of the DB2 products.

Prior to BMC, Frank worked as a Systems Programmer for IBM in support of NASA JSC's administrative processors. He supported installation and maintenance of MVS products and was in charge of implementing Automated Operation for IBM's internal systems as well as NASA's systems.

Frank holds a Bachelor's Degree in Computer Science from the Texas A & M University.

Michael Cotignola, BMC

What's New in Db2 12....or How Do I keep up with all this stuff

With Db2 12 being GA for well over a year now...time for a refresh. We'll take a cursory look at some of the major new features and enhancements to IBM's Db2 12, including the new branding ! Technically, this release delivers enhancements in all the usual categories, and introduces us to the Continuous Delivery Model. We'll conclude by taking a closer look at LOB compression, a much anticipated feature with lots of promise.

Speaker Biography:

Michael Cotignola is a Principal Software Consultant supporting BMC's DB2 Mainframe Solutions. In this role, he helps his clients maximize the value they receive from BMC's DB2 products. Michael has worked in IT since 1975 in numerous IT disciplines, including operations, application development, and database administration. He has been working with DB2 since version 1.1. Prior to joining BMC Software in 1996, Michael was a Managing Consultant for a large telecommunications and technology firm in Canada. Michael was president of IDUG and continues to be involved with IDUG in several capacities, and nominated as an IBM Analytics Champion for 2018.

Sam Poon, IBM Canada

Transforming IT with Machine Learning

The world economy experienced an enormous demand and growth in big data and analytics in the past decade. This trend will continue transforming how business is conducted. In recent years, machine learning breakthroughs already bring new ways of analyzing and usage of data. Companies, such as Uber and Netflix, leverage machine learning and data analytics to predict what customer needs and wants. Financial institutions enhance the fraud detection techniques to fight against fraud with machine learning technology. This session will focus on machine learning applying to IT operations, which includes introduction of the machine learning, reviewing machine learning essential capabilities in enterprise computing, and recent use cases specific for IT operations.

Speaker Biography:

Sam Poon is an IBM Executive IT Architect, specializing in DB2, data virtualization, and machine learning. He is certified as a Open Group Distinguished Certified IT Specialist, and has more than 20 years of IT consulting experience, including database experiences on DB2 for z/OS, DB2 distributed databases. Sam participated important roles in worldwide DB2 projects, including application enablement, benchmarking, and performance tuning. His recent focus on technologies, such as data virtualization and machine learning.

Jonathan Sloan, IBM

Maximize ROI of your Enterprise Data and Machine Learning Investment

Data scientists' skills are valuable and the models they develop are vital to the business. Maximizing the impact of your data scientist skills is dependent upon the currency of their data source, latency of execution and the way models are deployed and monitored over time. The lowest latency and freshest data exists within engagement applications on transactional systems. Integrating machine learning models with transactional applications gives you the ability to optimize every interaction and decision.

Often, more sensitive, personal data represents a unique combination of information about your clients and business that only your organization possesses. Leveraging this data for machine learning gives you the insight you need to outperform competitors

Keep your enterprise data securely in place while also using it for machine learning to deliver a unique customer experience.

Please join Jonathan Sloan, Senior Analytics Architect, IBM Analytics will discuss how to:

- Transform your engagement applications with machine learning
- Enhance the ROI of your machine learning efforts
- Better manage machine learning models and continually monitor model quality
- Maximize the impact of your data scientists' time

Speaker Biography:

Jonathan Sloan is a Senior Certified Executive Software IT Specialist, with more than two decades experience successfully leading, selling and developing business analytic and data warehouse applications as a vendor and implementer. He joined IBM in 2000 and specializes in helping customers understand the application of IBM analytics technology to business problems.

Jonathan has experience in the health care, financial services and consumer packaged goods industries. He is currently a subject matter expert and team lead for data analytics on the IBM z Systems.

He is passionate about helping organizations drive greater insight and value from data. Excels working directly with customers and in providing leadership within team environments.

Roy Boxwell, SEGUS

Db2 12 – 12 months agile production experience with more focus on the Data and less on the base

On October 21st 2016 Db2 12 went GA and several customers have quickly started to take a look at this new version. About a year ago some migrated their production systems and started to gain real life experience.

In July IBM also decided to do a massive rebranding throughout the z – ups, sorry – the Z portfolio. Db2 since then comes with a lowercase b, to place all emphasis on "Data"—your data. With all of this in mind we've created this presentation, bundling a year of experience to share best practices and tips for the Data-focused, agile database.

Speaker Biography:

Roy Boxwell has more than 32 years of experience in MVS, OS/390, and z/OS environments – 30 of those in Db2. He specializes in installation, migration, and performance monitoring and tuning. Roy leads the SEG development team responsible for the real time database maintenance solutions. He is also an active participant, speaker and contributor on the IDUG Db2 Listserv and sends out a monthly Db2 z/OS Newsletter.

DB2 for LUW

George Baklarz

Magic with JSON: How to integrate JSON into your Db2 Database

In 2017 we announced the availability of JSON functions within Db2 that were previously hidden. This session will take you through the functions, how to use them in Jupyter notebooks, and demonstrate some of the ways you can use our Db2 Jupyter notebook extensions to build your JSON SQL. We will also cover some of the directions that we are taking with JSON and how this will make it even easier to use JSON in your Db2 applications.

Speaker Biography:

George Baklarz, B. Math, M. Sc, Ph.D. Eng., has spent 30 years in IBM working on various aspects of database technology (it seems longer, trust me). 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 has a Doctorate in Computer Engineering, which took him only 13 years to get. The Engineering Faculty didn't like the fact that he worked at IBM - they felt that part-time students would never finish their degree if they had to work for a living! IBM on the other hand wondered how he could work while trying to get a degree! George currently pretends to work as part of the Core Database Technical Sales group (aka DB2) and creates PowerPoint slides for a living.

Jim Bean, CIGNA

How to Make Up Time

The main purpose of this presentation is two-fold. First, show our colleagues that significant updates can be made to a very active application with a minimal amount of down time, and while maintaining failback capability. And secondly show that some relatively straight-forward activities can make a huge difference to overall performance when combined.

The presentation would cover project activities completed this past summer and fall of 2017, specifically focused on OS and DBMS upgrades, application changes, database changes, etc. resulting in large performance improvements:

Various charts will be shown reflecting the huge improvements we realized from these changes. For this 6,000+ concurrent user CRM application, we eliminated 3¼ to 5¾ days of processing and wait time per day, or 24 to 36¾ days per week, a reduction in overall database processing time of 90.4% to 94.2%. Also will show impacts to the Db2 buffer pools, physical IO, and other savings such as Db2 licensing. And I will probably find some other charts to present as well.

Speaker Biography:

Jim Bean is an IT Principal in the Performance Management area of Cigna, a U.S.-based global health services company. Jim has been with Cigna for 25 years as a developer / team lead for a major application, a DBA, and a database product support technician, and most recently has specialized in database and application performance. Jim has worked in IT 34 years, including prior positions in IT operations and development at a machine tool distributor, and application development and support at a large aircraft engine manufacturer. He has presented both internally at these companies and externally at other conferences. He has a master's degree in Computer Science from Rensselaer Polytechnic Institute (RPI), and a bachelor's degree in Computer Engineering from the University of Hartford School of Engineering.

Keri Romanufa

DB2 LUW : The Latest from the Lab

This session will cover the latest new advances on DB2 LUW, including the V11.1.1.1 release. This includes a number of improvements in BLU acceleration, security, SQL compatibility, and manageability. Keri will also cover a sneak peak at some of the labs current work that might show up in future releases.

Speaker Biography

Keri has worked on Db2 for 25 years, and is currently an STSM and Db2 Chief Architect. She spent her career in a variety of teams from SVT to Performance to Development. Most of her development career has been spent in the Db2 kernel (buffer pool, logging, data management) and on availability such as PureScale and HADR..

Thomas Chu, IBM

IBM Analytics' Hybrid Data Management Strategy

This presentation gives an overview about IBM Analytics' Hybrid Data Management Strategy, with a focus on the future of Netezza and how it relates to Db2 and the Common SQL Engine.

Speaker Biography

Thomas Chu is the Director of Offering Management in IBM Analytics. He oversees the Hybrid Data Management portfolio that span cloud, private cloud, and on-premises integrated solutions. Thomas has joined IBM since 2001 and throughout his career, he has served in a variety of development management roles for IBM database and data warehousing software, as well as integrated data management systems.

John Hornibrook

Improved HTAP Use Cases with BLU Acceleration Indexes

Db2 with BLU Acceleration provides industry leading performance for analytical applications running complex queries. Db2 now further extends its capabilities into the hybrid transactional and analytics processing (HTAP) space with full support for indexes on column-organized tables. We will begin this session by covering some key use cases for the enhanced technology. Then we will cover the technical aspects. This includes the ability for queries to use both indexes and BLU Acceleration to enable more applications with transactional characteristics to run alongside analytic workloads on the same column-organized tables. BLU Acceleration indexes can also significantly improve the extract, transform and load (ETL) processing time for analytic applications, allowing analytics to be performed on data much closer to real-time. The addition of index support to column-organized tables is a significant milestone on Db2's roadmap for the HTAP solution. Attend this session to understand the types of use cases enabled by BLU Acceleration indexes and learn how this new technology works.

Speaker Biography

John is a Senior Technical Staff Member responsible for relational database query optimization on IBM's distributed platforms. This technology is part of DB2 for Linux, UNIX and Windows, dashDB and the Big SQL feature of Big Insights. John also works closely with customers to help them maximize their benefits from IBM's relational DB technology products.

Paul Bird

GWLM: A helpful free tool to unlock the complexities of Db2 WLM

Tired of trying to figure out your workload management configuration from dense db2look output? Having trouble making even simple changes to your system without mistakes? Unable to remember all the DDL options available? Come learn about the free IBM Graphical Workload Manager (GWLM) tool that helps to unfold and make clear the mysteries of Db2 workload management in a simple, visual manner!

Speaker Biography

Paul Bird is a senior technical staff member (STSM) in the Db2 development organization. For the last 25+ years, he has worked on the inside of the DB2 for Linux, Unix, and Windows product as a lead developer and architect with a focus on such diverse areas as workload management, monitoring, security, upgrade, and general SQL processing. You can reach him at pbird@ca.ibm.com.

Dale McInnis

DB2 Recovery Use Cases

This Presentation examines the best practices for recovery DB2 database objects. We will walk through several scenarios evaluating the different options with the pros and cons of each being discussed. It will cover what is new in the field of recovery within DB2.

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. 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) conference. 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 Core DB Technical Sales team.

Michael Roecken

Performing DB2 HADR Updates and Upgrades Made Easy

Moving to a new version or fix pack of DB2 should not be a scary event for databases using the high availability disaster recovery (HADR) feature. Fear of an outage or re-initialization of your standby is no longer a concern. This presentation will introduce to you and detail the procedures to perform a rolling update and a major release upgrade of your HADR single standby, HADR multiple standby and HADR pureScale databases. A detailed step by step analysis, with examples, from start to end so that you can get your database to the latest versions of DB2 with the least amount of concern.

Discuss the process to perform an HADR single standby and multiple standby fix pack rolling update.

Discuss the process to perform an HADR pureScale online member rolling updates.

Discuss the process to perform an HADR single standby and multiple standby major release upgrade.

Discuss the process to perform an HADR pureScale major release upgrade.

Speaker Biography

Michael Roecken is a senior software developer with Db2 for Linux, UNIX, and Windows platforms at the IBM Toronto Lab. Michael has worked since 2000 designing, implementing and supporting various features and capabilities in the areas of: backup/restore, crash/rollforward recovery, high availability/disaster recovery, logging/transaction management and upgrade/migration.

Pavel Sustr

Db2 LUW Forensic Investigation: Guide to Diagnosing Performance/Hang Problems

Learn how to efficiently analyze instance/database hangs and performance problems. Get a deep understanding of diagnostic logs, techniques, and advanced diagnostic tools shipped with the product. Explore the product's FODC capabilities. Learn how and when to use these troubleshooting tools for a quick problem resolution. The presentation contains a live demo during which the audience will be given an opportunity to perform an interactive investigation of real problems reported in the field.

Objectives

Identify and collect the must-gather information for hangs or performance problems

Understand and analyze Db2 diagnostic data

Learn about documented and undocumented problem determination tools

Master shortcuts leading to problem resolution

Get an insider view of problem determination techniques used by Support teams

Speaker Biography

Pavel Sustr is a Senior Manager and Senior Software Engineer with IBM Db2 LUW development, responsible for multiple core Db2 kernel components. Always thrilled to work on hard-to-crack puzzles. Expertise in Db2 LUW kernel architecture, configuration and administration, advanced problem determination, memory architecture, memory leak troubleshooting, and assembly language. Hands-on development experience with buffer pool management, storage, prefetching, page cleaning, transaction logging, recovery, monitoring, and problem determination. As a member of the Db2 team, Pavel spent years in Db2 L2/L3 advanced support (over 1,500 resolved cases), then transitioned to Db2 LUW kernel development. In his past life Pavel was an application developer mostly using C++, SQL, .NET, Oracle, MS-SQL, and Informix on Windows, Linux, Solaris, and HP-UX.

Calisto Zuzarte

SQL Query Writing Tips To Improve Performance in Db2 Warehouse

Data managed by databases has been growing exponentially and the trend has been to push much of the application processing close to the data. Applications often can no longer pull the data to do some processing as part of the application logic. Much of this logic is now expressed as one big SQL statement and executed in the database engine. Complex and sophisticated SQL queries, particularly analytical queries, can be a challenge to database optimizers.

There are many ways to write SQL variations that get the same result. The optimizer has the goal of relieving the user from having to write the perfect version of the SQL statement. The optimizer usually does an outstanding job but there are times when it may not have some knowledge about the data and can do with some help. This session looks at some SQL query writing tips and tricks to improve performance in the dashDB and DB2 database systems

Speaker Biography:

Calisto Zuzarte is a Senior Technical Staff Member (STSM) in the Db2 development organization. His focus area is database query performance. He also represents IBM on the ANSI and ISO SQL Standards committee.

IMS for z/OS

Sandy Sherrill & Rick Engel, IBM

The View from Here – the latest from the IMS and IMS Tools Labs

IMS OM Trends & Directions

In this opening segment, Sandy Sherrill, IMS Market Enablement, will share what is happening in IMS from the perspective of Offering Management: What new programs and workshops are available to customers; what all of the transformation means to how we design, develop, and deliver IMS to you; and how you can influence how we do all of those to make you more successful with IMS.

IMS Tools Trends & Directions

Many of the world's mission critical applications run on IMS, the most secure, highest performing, and lowest cost database management system for online transaction processing (OLTP).

How do customers keep this key subsystem available, running efficiently and delivering peak performance to their business operations?

The session focuses on actual customer case studies and associated IBM IMS-based software solutions used to address issues with managing critical IMS environments.

Speaker Biographies:

In her 30 years at IBM, Sandy Sherrill has worn many hats. Currently she is responsible for IMS Market Enablement - a title that also requires her to wear many hats. She writes a lot (blogposts, white papers, executive presentations), ghost writes a bit less, plans and executes IMS conference participation, coordinates IMS User Group activities for the GTA and Quebec, contributes to the IMS Developer Center, manages the IMS web site, and is currently trying to uncover the easiest, fastest, and least disruptive channel to keep all IMS customers up-to-date on what is happening in IMS. If you have suggestions, she is interested in hearing them.

Rick Engel is a 25 year IBM veteran, working the past 15 years as an IBM CTP (Client Technical Professional) and Team Leader for IBM's IMS Tools business. He has spent his entire 35 years IT career experience supporting IBM mainframe systems, his first position being an MVS systems programmer in his home area of St. Louis, Mo.

He has attained several levels of IBM certifications and is a regular speaker at IBM and customer technical events throughout the US and Canada.

Deepak Kohli, IBM

IMS Masterclass: An Afternoon with Deepak

(An in-depth IMS Masterclass)

Join Deepak Kohli for an in-depth exploration of all things IMS. Extending across all three afternoon sessions, Deepak will start by leading you through the key aspects of IMS Architecture; with that foundation, he will then highlight how IMS modernization is changing how we manage and exploit our IMS systems and how the product is positioned today to meet the technical and business challenges of its second half-century.

In keeping with the Masterclass format, if there are areas of IMS capability that are of particular interest to you, make your voice heard – email registration@ccdb2.ca and your interests and questions will be passed on to Deepak. So, whether your interest lies in Transaction Management, esoteric database queries, language support (ancient and modern), how 50 years of history has been exploited to enhance today's capabilities, or any other topic, this is your opportunity to get your questions answered.

Speaker Biography:

Deepak Kohli is a Senior Software Engineer at IBM's Silicon Valley lab. Currently, he is part of the IMS Product management team helping customers modernize and SOA-enable their IMS systems. Deepak also provides IMS Level 2 technical support specializing in IMS On Demand and TM areas. He has a Masters in Computer Science from NYU. In the area of IMS, Deepak's experience ranges across applications, DBA, Systems, performance tuning and problem determination. He has taught various IMS and DB2 courses. His presentations have won him rave reviews at conferences worldwide.