StrateSQL's Presentation Abstracts

Jason Strate
e: jasonstrate@gmail.com
t: @stratesql
c: 651-285-8885

Big Data

Name	SessionState	Abstract	Prerequisites	Goals
All Your Data's Shuffling, Understanding Data Movement in Parallel Data Warehouse	*Abstract*	To fully leverage the power of the Parallel Data Warehouse (PDW), data is distributed across all of the nodes within the appliance. Ideally, the data that each node requires will be located on that node. Unfortunately, there are a number of situations where this isn't the case. To solve this, PDW sometimes moves data between nodes within the applicance. At times this movement is desired and at other times, it is not, resulting in possible performance issues. In this session, we'll review the different types of data movement within the PDW, explaining the need and benefit of each type of movement. At the end of the session, you'll understand each type of data movement and understand how to identify when the data movement isn't desirable with options on how it can be mitigated.		Describe the need for data movement in PDW Describe the type of data movements available in PDW Demonstrate data movement operations in PDW
Exploring Implementation and Performance of Parallel Data Warehouse (Pre-Con)	*Abstract*	If you need to design a data warehouse that scales for today and into the future, then you need to know about Microsoft's Parallel Data Warehouse. In this full-day session, we'll review the architecture of PDW with an emphasis on what makes the PDW the solution of choice for data warehouses. We'll discuss what it takes to migrate a data warehouse to PDW and what bottlenecks you could expect to see in a migration. Through live demonstrations, you'll witness the power and performance of PDW, demonstrating the performance impact of PDW's data distribution. The presenters will discuss, from personal experience, the successes that clients already on PDW are seeing and why it is the key to future data warehouses. By the end of the session, you'll have a clear understanding of the value that PDW provides when it comes to performance and scale.		Describe architecture of Parallel Data Warehouse Provide understanding on how to deploy and query data warehouses on Parallel Data Warehouse Describe the advantages of MPP data warehouses over SMP designs
Introduction to Parallel Data Warehouse Distribution Theory	*Abstract*	Microsoft's Parallel Data Warehouse utilizes a distributed SQL Server environment. In order to properly leverage this environment, developers and query writers must have an understanding of the underlying distribution theory. This distribution theory governs how data is distributed and plays a key role on returning results. In this sessions, we'll review the rules that comprise PDW's distribution theory with demonstrations on how these rules function on PDW. Attendees can expect to learn the rules that comprise distribution with a clear understanding of their impact on data distribution and query execution.		Outline each of the rules for PDW's distribution theory Provide an understanding of the rules and how they effect query execution Demonstrate how the rules effect data distribution

Big Data

Name	SessionState	Abstract	Prerequisites	Goals
Recognizing the Parallel Data Warehouse (PDW) Workload	*Ready*	If you are like many, you've been hearing a lot about Parallel Data Warehouse (PDW) lately. PDW is Microsoft's appliance solution with SQL Server 2012 that allows massively parallel processing (MPP) of database workloads. But how do you know that you have a workload that requires or could benefit from MPP over the standard symmetric multiprocessing (SMP) architecture? In this session, we'll review that question and more with a run through of how to identify MPP workloads? As part of the session, we'll also review some of the key components and differentiators of the PDW platform. You'll leave this session with a clear understanding of the value that PDW can bring to an organization and the types of workloads that it is best suited for.		Describe architecture of Parallel Data Warehouse Define data warehouse workloads Provide understanding of workload best suited for Parallel Data Warehouse

Business Intelligence

Name	SessionState	Abstract	Prerequisites	Goals
Building Business Analytics Solutions Without the IT Department	*Abstract*	Have you ever had an idea for a Business Analytics (BA) solution that you just wanted to try out? Usually these proof-of-concepts (POC) need hardware resources to test out and that becomes the project bottleneck. You're stuck waiting for IT to make an under powered, ancient server available so that you can just try something out. What if you could skip over the IT department all together and start building the BA solution today and start working on figuring out if the solution works tomorrow? In this session, we'll show you how you can do this by leveraging the Microsoft Business Intelligence stack and the Azure platform. By the end of the session, you'll have the tools and resources you need to get that new BA POC going and finding out if the investment in resources is worth the value of the project.	Understanding of general data warehouse architectures	Describe value in leveraging Azure environment Define components available for developing business analytics solutions in Azure Demonstrate how to deploy data warehouse environment through Azure

Business Intelligence

Name	SessionState	Abstract	Prerequisites	Goals
Bypassing the IT Department with Business Intelligence Solutions	*Abstract*	Business Intelligence (BI) solutions need to move at the speed of business. Unfortunately, in most cases, they are stuck working at the speed in which IT resources can be allocated. What would you do, if I told you that you could deploy an entire BI infrastructure in a couple hours and start loading data into it by the end of the day. All this without having to ask the IT department for any assistance. In this session, we'll demonstrate how to leverage Microsoft tools and the Azure cloud environment to build out the BI solution that you need. In the time of the session, you'll gain an understanding of the capabilities of Azure and how you can start building a new BI proof-of-concept for your managers next week. All this without having to ask the IT department for anything.		Describe value in leveraging Azure environment Define components available for data warehousing in Azure Demonstrate how to deploy data warehouse environment through Azure

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
5 Amazing Reasons DBAs Need to Love Extended Events	*Ready*	Extended events provide DBAs with a powerful tool that can be used to troubleshoot and investigate SQL Server. Throughout this session, you'll walk through five great reasons, with demos. By the end of the webcast, you'll be itching to grab the scripts from the demos to start building your own extended event sessions today.		
Are You Following Your Own Best Practices?	*Ready*	Everyone has their own best practices that they try to follow. But often times we aren't sure how well we are following our own or industry best practices. In this session we'll go into monitoring and managing best practices throughout your environment. The session will look at how to create a solution using policy based management and other tools to report on your compliance to best practices.	Understanding of SQL Server Best Practices	Learn the concepts behind policy based management Identify how to translate best practices into policies Demonstrate how to deploy and monitor policies based on best practices in a SQL Server environment
Backups! Don't Get Caught With Your Pants Down!	*Ready*	Did you hear about the blog site that disappeared when their system failed and they didn't have a backup to restore from? Or how about the pet supplier that went out of business after a DELETE statement removed everything from their database and there wasn't a backup on hand. These aren't urban myths but rather typical stories that make the news every couple months. Is your company one database failure from collapse? In this session, we'll go over the basis of backups. And also go over the basics and best practices in creating a backup and recovery strategy.	Understanding of basic T-SQL coding skills Understand basics of SQL Server backups	Define the reasons for database recovery plans Understand how to fit database recovery plans to database backup strategies Demonstrate methods to ensure a recoverable database

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
Benchmarking Your Database System	*Ready*	Your databases are typically not static entities. Data is often being added, updated, and deleted from your databases. New applications are being written and existing applications are being modified. Throughout all of this the database is expected to perform as well as it did when it was first unwrapped and places into production. But do you know how if performed that day? Or last month? In this session, we'll review common areas that should monitored for SQL Server databases. Also, we'll review key metrics that should be analyzed over time.	Understanding of performance benchmarks Understanding of Database Administration concepts	Identify key performance metrics that need to be included in performance benchmark Demonstrate process for collecting benchmark data Discuss process for analyzing benchmark data
Error Logs and Deadlocks, Oh My!	*Draft*	Tired of turning to the SQL Server error log to find out what's happened? Would you rather just know about a deadlock when it's occurred rather than later in the day when the manager as you why HRs bonus distribution application failed? Find out how to get this information and more delivered to your inbox so that you can head off issues instead of drowning in them. In this session we'll look at Event Notifications and find out how to stop having to monitor SQL Server error logs today.		
Extended Events, Work Smarter Not Harder	*Ready*	There are many ways to performance monitor your SQL Server environment. In this session we'll review Extended Events, which is one of the newer SQL Server monitoring platforms. Learn the ins and outs of how to get detailed information on the errors and events that occur within SQL Server and how to dig into the information. With a few T-SQL statements, issues that could take weeks to research can be investigated in minutes.	3. Understanding of performance issues that can occur in SQL Server 4. Mid-level skills writing T-SQL code	1. Define the architecture for Extended Events 2. Identify scenarios where Extended Event can be used to troubleshoot performance 3. Demonstrate solutions to real-world performance scenarios
Extending Data Platform Availability Into Windows Azure	*Abstract*	Over the lsat few years, there's been a sea change in the options available for providing high availability and disaster recovery for SQL Server. From the development of AlwaysOn Availability Groups to the ability to backup directly to Windows Azure Storage. In this session, we will look at the available options and review the key criteria for choosing between these options. As part of the session, we'll how you can extend AlwaysOn Availability Groups into Azure and demonstrate the ease in configuring this availability option.		Describe options for extending on-premise to azure architectures for availability

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
Getting To First Base	*Abstract*	It's going to happen someday. The phone will ring and someone will ask about the performance of some database application or another. Will you have the information on hand to know that this is either an isolated event or the point where hardware resource constraints have peaked? Along with that, to properly implement consolidation and virtualization projects you need to know where resource surpluses exist. In this session, we'll look at building a performance baseline for your SQL Server environment. Discover the tools that are already available and what information you need monitor to manage your SQL Server environment.	Understanding of SQL Server resource utilization Understanding of database capacity planning	Understand the need for base lining the SQL Server environment Learn statistics that are useful in base lining SQL Servers Demonstrate how to base line the SQL Server environment
Getting To Know Your Indexes	*Ready*	Without proper indexing SQL Server can be hard pressed to create efficient and performant execution plans. Dynamic Management Views (DMV) and system views provide a slew of information about indexes that can be used to analyze indexes within SQL Server. In this session we'll go under the hood of SQL Server to look at DMVs and system views to know what indexes you have, should have, and how they feel about the way applications are treating them.		Identify methods to analyze current and potential indexes Learn how to alleviate stress found on indexes Demonstrate methods for tuning indexes
Leveraging the Value in Dynamic Management Views	*Ready*	Dynamic Management Views (DMVs) and Dynamic Management Functions (DMF) have been around since the release of SQL Server 2005. But for many people DMVs and DMFs are still a mystery. Step under the hood and review some of the main DMVs and DMFs in SQL Server. From sessions to caches, this session will provide a walk-through of some of the more useful DMVs and a practical look at how they can be used?	Understanding of database administration functions	Provide a definition of key DMVs within SQL Server Learn how to use DMVs to discover statistics on the state and performance of the SQL Server Demonstrate how to use DMVs to collect and utilize information to manage and performance tune SQL Server
Performance Tuning With Extended Events	*Ready*	SQL Server 2008 saw the launch of a new performance tuning tool was made available. That tool was Extended Events. Where some previous tools were clumsy and intrusive, Extended Events is flexible and light. It's able to extract exactly the information that you need when you need it. In this session, we'll explore some common performance issues related to SQL Server and demonstrate method to resolve those issues.	Understanding of performance tuning	Understand the use cases for Extended Events Demonstrate methods for performance tuning Describe some best practices for using Extended Events

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
Performance Tuning with SQL Server's Plan Cache	*Ready*	Execution plans are stored after execution in the plan cache. This metadata about how queries are executed can provide insight into how your SQL Server environment is functioning. By using XQuery to browse and search the plan cache you can find potential performance issues and opportunities to tune your queries. This information can be used to help reduce issues related to parallelism, shift queries from using scans to using seek operations, or discover exactly which queries are using what indexes. All of this and more is readily available through the plan cache. In this session we will explore the plan cache and start you on the road to discovery.	Understanding of basic T-SQL coding Understanding of basic XQuery statements	Discuss information available in an execution plan Demonstrate use of XQuery to query plan cache Demonstrate methods to use the plan cache for performance tuning
Preventative Maintenance for SQL Server	*Abstract*	Murphy's Law states that whatever can go wrong will go wrong. What are the best practices for maintaining and backing up your database? What should you do to keep your database healthy? This session will look at some of the basic monitoring and maintenance that you should be doing to make sure Murphy doesn't interrupt your weekend.	Basics understanding of the SQL Server architecture Understand of basic T-SQL coding	Understand the need for proper database maintenance Learn the proper tasks necessary for database maintenance Demonstrate methods for basic database maintenance
Snapshots Every Day To Keep The Users At Bay	*Ready*	SQL Server maintains a large amount of information on SQL Server instances in Dynamic Management Objects (DMOs). Over time this information can be used to isolate potential performance issues, identify patterns in usage of your databases, and build a performance benchmark to identify when problems may be around the corner. This session will investigate and define some of the common DMOs that can be used to monitor SQL Server instances to address performance issues before users are addressing you with them.	Understanding of basic T-SQL Understanding of Database Administration concepts	Discuss scenarios where snapshots can be used to monitor SQL Server Identify DMOs that provide information useful for snap-shotting Demonstrate methods to acquire and access snapshots of system information
Strategies for SQL Server Index Analysis	*Abstract*	Properly managing indexes for a database is a common usse in many environments. When reviewing an index, should it have a single or multiple columns? Should the table be a heap, a clustered index or clustered columnstore index? These considerations are often at the forefront when analyzing your indexes. In this session, we'll look at easy methods for identifying new indexes, we'll review patterns for index consolidation, and discuss how and when to remove indexes. At the end, you'll have a strategy that helps design indexes to improve performance.		

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
Strategies for SQL Server Index Design	*Abstract*	Properly designing indexes for a database is a common desire in many environments. When designing an index, should it have a single or multiple columns? Should the table be a heap, a clustered index or clustered columnstore index? These considerations are often at the forefront when designing indexes. In this session, we'll look at easy methods for identifying new indexes, we'll review patterns for index consolidation, and discuss how and when to remove indexes. At the end, you'll have a strategy that helps design indexes to improve performance.		Key take away will be strategies for identifying new, justifying existing, and removing un-needed indexes.
The Creepy DBA, How to Stalk Your Users	*Ready*	Do you know when your users are using the production login accounts from their workstations? Are you aware of changes being made to the development servers that may impact future deployments? Has the new Junior DBA start giving out access like candy? As DBAs, we are responsible for knowing and acting on all of this and much more. Fortunately, SQL Server provides a number of features that you can use to monitor and track user activity. In this session, we'll look at these features and demonstrate how you can use them to the extent that your users find you creepy!	1. Understanding of database administration	1. Define the responsibilities required of DBAs 2. Discuss methods to monitor activity on SQL Servers 3. Demonstrate tools available to "stalk your users"
Upgrade to SQL Server 2008 – Before This Presentation Finishes	*Ready*	In this session we'll review lessons learned and best practices when upgrading from previous versions of SQL Server to SQL Server 2008. In this session we'll go through the changes to the installation and upgrade to SQL Server 2008. Also, how do you prepare for the upgrade? Do you install now or wait until later? Can your production database be upgraded in less time that it takes to complete this session?	Basics understanding of the SQL Server architecture Understand of basic T-SQL coding	Learn the tasks necessary to prepare for a SQL Server migration Learn the tasks necessary to execute a SQL Server migration Understand the need to prepare for a SQL Server migration
Using XML to Query Execution Plans	*Ready*	SQL Server stores its execution plans as XML in dynamic management views. The execution plans are a gold mine of information. From the whether or not the execution plan will rely on parallelism to what columns are requiring a key lookup after a non-clustered index seek. Through the use of XQuery this information can be available at your fingertips to help determine the value and impact of an index and guide you in improving the performance of your SQL Server databases.	Understanding of basic T-SQL coding Understanding of basic XQuery statements	Discuss information available in an execution plan Demonstrate use of XQuery to query execution cache Demonstrate methods to performance tun

Enterprise Data Management

Name	SessionState	Abstract	Prerequisites	Goals
What Are You Waiting For?	*Ready*	You've spent the afternoon sweating over your T-SQL query. You've wrung out all of the wrinkles. It's deployed, it's running… now its waiting. In SQL Server the most performant query can become a problem query if there aren't any resources in the SQL Server for it to use. In this session, we'll look at wait statistics and what they are. We'll look into how they are accumulated and how they can be monitored. By the end of the session, you'll be equipped with the tools needed to determine if there are resource issues in your environment and methods to start mitigating them.	Understanding of performance tuning	Understanding of wait statistics Demonstrate methods for resolving issues related to wait statistics Demonstrate process for monitoring wait statistics
What SQL Server Are You Talking About?	*Abstract*	In many environments, all of the SQL Server instances are housed in a data center or server room where the DBAs manage and maintain the servers and databases. But also in many of these same environments there are SQL Servers mysteriously hiding under people's desks or in server rooms that don't fall under the prevue of the DBA or that team. In this session, we'll look into SQL Server sprawl and how to discover all of the SQL Server instances in your environment.	Understanding of database capacity planning Understand the existing SQL Server Environment	Understand the need for documenting SQL Server environments Learn how to document and take control of the SQL Server environment Demonstrate methods to discover SQL Server instances

ETL

Name	SessionState	Abstract	Prerequisites	Goals
Solving Business Pains with SQL Server Integration Services	*Ready*	SQL Server Integration Services (SSIS) offers a wide range of features and functionality that can be used to solve business pains within Extract, Transform, and Load (ETL) processes. These pains can be caused by the performance of current ETL process, the maintainability of ETL processes or represent a new pain that SSIS features can resolve. This presentation will provide a demonstration of real world scenarios where features of SSIS were used to solve business pains through modification of existing SSIS packages or through the creation of new SSIS packages. The session will also discuss current best practices in SSIS package design.	Understanding of basic ETL concepts Understand of basic T-SQL coding	Understand flexibility that SSIS features provide Demonstrate application of SSIS functionality to business needs Discuss current best practices in SSIS package design

Professional Development

Name	SessionState	Abstract	Prerequisites	Goals

Professional Development

Name	SessionState	Abstract	Prerequisites	Goals
Do SQL Server Certifications Really Matter?	*Ready*	Microsoft has a number of certifications for SQL Server. They cover the DBA, developer and business intelligence skills. But, does anyone take them seriously? In this session, see the tracks for SQL Server certifications and discover what you need to know about the exams. Which track will be right for you? What steps should you take? How do you get going? And what value will they bring to your career?	Interest in learning about how certifications can affect a career	Define the SQL Server certification tracks Understand the benefits gained from gaining certifications Identify how to obtain certifications and what to do to prepare

SQL Development

Name	SessionState	Abstract	Prerequisites	Goals
5 Ways to Improve Performance through Indexing	*Ready*	Sometimes the hardest performance issues in your environment have simple solutions. One of these simple solutions is through properly indexing the database. In this session, we'll look at five ways that performance of a database can be improved through indexing. When we finish, you'll be armed with scripts and methodologies that you can put into place today to improve database performance.		
ABCs of CTEs	*Ready*	Common Table Expressions (CTE) aren't as common as their name implies. CTEs are often seen as a secret part of the dark art of recursion. This session will explore CTEs to show how they can be extremely useful in improving performance and legibility of T-SQL code. And, of course, we will look at their use in returning recursive data.	Understand of basic T-SQL coding	Provide an understanding of the purpose and use of CTEs for recursion and derived tables Define the pro's and con's to using CTEs in recursion and for derived tables Demonstrate real-world examples detailing methods for using CTEs to solve coding challenges
Become Your Own Picasso: Writing Better T-SQL	*Ready*	While not everyone can be a Picasso, there is an art form to writing great SQL statements. In this session we'll go over tips and tricks in writing SQL. We'll also review some popular bad practices and provide demonstrations to better change those velvet Elvis's to Rembrandts. As every paint brush can't produce a masterpiece, the same goes for SQL statements. Writing good SQL code requires the right brush. After this session, you'll have a few more paintbrushes and probably throw out a few that are no longer useful.	Understanding of basic T-SQL coding skills	Learn methods to review and identify issues with T-SQL code Demonstrate common coding practices that are poor performing Demonstrate new coding practices that provide good performance

SQL Development

Name	SessionState	Abstract	Prerequisites	Goals
Common T-SQL Patterns Affecting SQL Server Performance	*Abstract*	For any requirement, there are multiple ways that T-SQL can be used to provide a solution. In most of these cases, the difference between one solution and the next is unseen and may result in SQL Server processing the data identically in either case. But sometimes, there are bad solutions. These are cases where the choices will dramatically affect performance and lead to serious performance issues. In this session, we'll talk about these T-SQL patterns and discuss how they affect performance and options for solving the same issues with other patterns.		Identify common T-SQL patterns that cause performance issues. Describe how these issues affect performance Demonstrate alternative T-SQL patterns
Necessarily Evils, Building Optimized CRUD Procedures	*Ready*	Every developer loves them and a lot of DBAs hate them. But there are many and valid reasons for creating generic SELECT, INSERT, UPDATE, and DELETE procedures. In this session, we'll go through designing CRUD procedures that utilize new and existing SQL features to create CRUD procedures that are optimized for performance.	Understand application data access requirements Understanding of basic T-SQL coding skills	Identify the uses for CRUD procedures in databases Identify the common problems associated with CRUD procedures Demonstrate methods for writing performant CRUD procedures
Performance Impact of User Defined Functions	*Ready*	User defined functions provide a means to encapsulate business logic in the database tier. Often the purpose of the encapsulation is to provide standard method access segments of data within the database. Unfortunately, not all methods of creating user defined functions are equal. In this session we'll review the types of user defined functions and investigate the performance impact in selecting the different types.		Identify purposes for creating user defined functions Discuss the types of user-defined functions Demonstrate performance impact in selecting different types of functions
Profiling SQL Server for Developers	*Draft*	When performance in SQL Server databases suffer, developers are often tasked with discovering the problems and providing resolution. In many cases, the tools used by DBAs aren't a good fit for developers when they need to profile the applications to discover issue. The issues can be as simple as investigating queries with simple performance issues to more complex issues, such as identifying the causes for connection timeouts. In this session, we'll explore the tools available within SQL Server that allow developers, along with DBAs and sysadmins, to profile activity and discover the root cause of current issues.		Explanation common issues facing developers Introduction to the SQL Server Extended Events feature Walk-through to scripting and using GUI to create extended event sessions Demonstration of how using Extended Events can help resolve issues

SQL Development

Name	SessionState	Abstract	Prerequisites	Goals
Secrets to Improving Report Performance Through Indexing	*Abstract*	Pssst, what the DBAs aren't telling you about indexing is probably hurting your reports! Are you tired of slow reports and having to wait hours, or weeks, for DBAs and the IT Team to "fix" the performance problems with the data warehouse? And when it is fixed, you're told that everything got better because the DBA added just one index. Now you can get around your performance problems and waiting on others by getting a thorough understanding of the basics for clustered, non-clustered, and columnstore indexes. We'll look at how to identify when indexes are needed and demonstrate impressive performance improvements with minimal effort. By the end of the session, you'll be armed with the knowledge needed to request that indexes be built so that you can see improved report performance today.		Provide basic around clustered, non-clustered, and columnstore indexes. Identification of queries that could improve with indexing Understanding of how to build indexes to improve report performance
Secrets to Improving Report Performance Through T-SQL Pattern	*Abstract*	Pssst, just because you can write it like that in T-SQL doesn't mean you should. T-SQL tries to make getting to data easy. It's a declarative language that just wants you to define the output and it'll worry about how to get to the data. Unfortunately, as we have all seen from time to time, it doesn't work out all that well. Come join others for some tricks and techniques for improving the queries you write. Learn simple techniques that will improve the performance of your reports. When you leave, you'll have have example techniques and demo code that you can take back to start imrproving report performance right away.		Describe useful T-SQL patterns for report writing Describe useful T-SQL anti-patterns for report writing Demonstrate how to identify problems with T-SQL queries
Taking a Crack at CLR	*Ready*	A few years back the talk of CLR was all the rage. As time has gone on, the roars of the marketing crows have waned and DBAs fearful of the CLR apocalypse have had to put away theirs signs. With those days behind us, it's time to make sure that you've take a look at CLR and what it can do for your database and your environment. In this session, we'll look at how to build and implement CLR objects. We'll also look at some specific use cases that can help improve the performance of your SQL Server environment and leave that environment more secure.	Basic understanding of C# NET	Reintroduce CLR to the audience Describe appropriate use cases for CLR Demonstrate how to implement CLR

SQL Development

Name	SessionState	Abstract	Prerequisites	Goals
XQuery Basics for the DBA	*Ready*	For the past few years, XML has been making its way more and more into our SQL Server instances. For a moment let's forget that developers often find the need to store XML data in their databases. Outside of those databases, XML has already invaded our SQL Servers. It's made its way into execution plans with the SHOWPLAN XML. It's telling us about deadlock through deadlock graphs. It's also presented to us through service broke, event notifications, and extended events. XML is there and we need to learn how to query it. In this session, we'll explore some of the basic methods for querying XML through XQuery. By the end of the session you'll be equipped with the tools and understanding needed to ease yourself into XQuery.	Understanding of basic T-SQL coding	Discuss the importance of learning XQuery Discuss the XQuery methods Demonstrate methods of using XQuery