Data Platform Summit 2019 - Session List (Final Release)

Final Release. Release Date: July 31, 2019. Session Count: 108

Important Notes. Please read carefully.

1. This is the final release.

2. Final agenda/schedule published. There will be minor changes to the agenda/schedule without prior notice.

3. With 100+ sessions DPS 2019 is Asia's Largest Data AI & Advanced Analytics Conference on Microsoft Data Platform, Advanced Analytics, Artificial Intelligence and Open Source. DPS 2019 brings you the finest & latest content from world's best trainers.

4. DPS Summit & Pre-Con pass price increases on 1st of every month. Block your seat today.

Data Platform Summit 2019
Empowering Data & AI Transformation

	To view, up-to-date sessions list --->>			http://www.dataplatformgeeks.com/dps2019/break-out-sessions/	

Session_Type	Speaker	Track(s)	Session Title	Abstract	Level
Break-Out	Abhishek Narain	BD/DEV	Modern Data Integration in Cloud using Azure Data Factory	Abstract: Azure Data Factory is a hybrid data integration service that provides you with an intuitive visual drag and drop environment (data flow) for code-free ETL so you can quickly and easily move, prepare, transform, and process any data from any source at scale with the fully managed cloud service. In this session, we will showcase a typical Modern Data Warehouse with Data Lake based architecture which is operationalized using Azure Data Factory. We will use data flow in ADF to transform the Big Data from data lake before aggregating it in the Azure SQL Data Warehouse. We will also talk about best practices, customer learnings, and product roadmap. Key Learning: Modern ETL/ ELT on Azure Operationalize a Modern Data Warehouse in Azure Demos: Drag and drop based ETL (ELT) on big data workload in Azure Data Lake Gen2 and Azure SQL DW. Hybrid connectivity to fetch data from on-premise/ cross-cloud at scale.	Intermediate
Chalk-Talk	Abhishek Narain	BD/CL	Migrating PB scale data lake to Azure	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA

Break-Out	Ajay Jagannathan	ARCH/DBA	Azure SQL Database Hyperscale and Serverless - new destination for your modern SQL workloads	Abstract: Hyperscale offers a data tier that has infinite scale, optimized for highly transactional workloads and offers seamless integration with all the game changing benefits of being a fully managed database while still offering all the familiar SQL functionality. Come learn how to integrate, deploy and modernize your data platform to Azure SQL Database Hyperscale. Serverless is a new compute tier that optimizes price-performance and simplifies performance management for single databases with intermittent, unpredictable usage. Serverless automatically scales compute for databases based on workload demand and bills only for compute used per second. Serverless also automatically pauses databases during inactive periods when only storage is billed and automatically resumes when activity returns. Join us for this session to learn about using serverless in SQL Database and save cost while achieving the desired performance. Key Learning: Learn the latest and greatest offerings in the Azure SQL Database family - Hyperscale and Serverless Understand the scenarios where you can leverage Hyperscale and Serverless Demos: Hyperscale: Provisioning Point in time restore Read scale out Serverless: Provisioning	Basic

Break-Out	Ajay Jagannathan	ARCH/DBA	Migrate or modernize your database applications using Azure SQL Database Managed Instance	Abstract: Azure SQL Database provides a fully managed instance deployment model that enables friction-free migration for SQL Server applications running on-premises. It lets you focus on your business without having to worry about infrastructure management or operations or end of life. In this session, you will learn how to start your cloud modernization, using Azure SQL DB Managed Instance, and how to create a hybrid topology with on-premises and cloud database services. You will also learn various best practices for maximum performance of your database instance. Finally, you will learn how to easily build data ingestion using SSIS and ADF, app modernization using new database engine and security capabilities in the service and finally build near real-time analytics and visualization using Power BI and scaling this out to read replicas of instance. Key Learning: Learn best practices for configuring and using Managed instance for database applications Learn about the rich feature set and compatibility with SQL Server Demos: Managed instance network setup Managed instance service configuration - DR, AAD, Advanced Data Security ETL and Power BI against	Advanced
Break-Out	Alicia Moniz	BIA/DS	Real-Time Analytics with Databricks and Cognitive Services	Abstract: Join us for a session on Real-Time Analytics on Azure, utilizing Azure Databricks and Cognitive Services. We begin our session with an overview of the Microsoft Azure Lambda architecture, which provides an impressive suite of tools for data streaming. We will also step through code for an end-to-end Azure Analytics data pipeline via Event Hubs, Databricks & Azure Cognitive Services. After this session, you will know how Azure accelerates delivery of actionable insights to your business! Key Learning: Learn how to create a Cognitive Services Model Learn about Cognitive Services capabilities Learn how to create an Event Hub Learn about Azure Databricks capabilities Demos: Create a Cognitive Services Model. Create an Event Hub. Create a notebook in Databricks. Send data to the Event Hub via Databricks code. Process data from Event Hub via Databricks code and send to Cognitive Services Model.	Intermediate

Break-Out	Amit Banerjee	DBA/DEV	SQL Server 2019: Making your data more Intelligent	Abstract: Microsoft is investing in it's database engine in areas of Availability, Security and Performance to help organizations build a complete data platform that enables Analytics at a massive scale. For the first time, SQL Server 2019 is introducing Spark, HDFS for the first time and enabling querying over multiple data sources using T-SQL. In this session, we will talk about how developers can use a consistent programming surface to query data at the edge, in the datacenter and the cloud. There will be demos which shows how the new features in the recent CTP release of SQL Server 2019 can help make your data management life easier. Key Learning: 1. Learn about the enhancements in SQL Server 2019 2. Learn about new hybrid capabilities of our data platform 3. Learn about how T-SQL can become the language to query all your data Demos: SQL Server Data Virtualization SQL Server 2019 Enhancements	Intermediate
Break-Out	Amit Bansal	DBA/DEV	Real-World Query Tuning Examples	Abstract: Query Tuning is easier said than done. Here is an opportunity to learn some real-world query tuning examples. In this 100% demo, deep-dive session you will learn how you can re-write T-SQL queries using new constructs, tune indexes & deal with statistics to improve query performance. You will also learn a few advanced concepts about execution plans and iterators. This session will be an eye-opener for you and you will learn things that Google cannot find for you. Assured. Key Learning: You will get solutions to the real-world problems that you can apply right away in your production environment. Demos: Entire session is demo-based. All you see are real-world examples	Intermediate
Break-Out	Anand Raman	CL/ARCH	5 industries that are getting disrupted by AI on Cloud and on Edge	Abstract: Join this interactive session to gain insight into what's new in Azure cognitive services and Bot services. Learn about reference architecture, best practices and recent customer success stories. Finally learn about the product roadmap and upcoming exciting features. Demos: 1. Demonstrate what's new in cognitive services and Bots 2. Learn to build cognitive service 3. Learn how to build a AI service with no machine learning experience	Basic
Key-Note	Anand Raman		Keynote by Anand Raman, Microsoft		
Open-Talk	André Melancia	CL/ARCH	Everybody is flying in the Cloud, why aren't YOU?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

				Abstract	
Break-Out	André Melancia	BIA/DS	BI Powers - The A.I. Who Tagged Me	Abstract: So everyone knows how to make some nice graphs in Power BI from Excel or SQL Server... But what about getting "intelligent" results from unusual data sources? Disclaimer: A.I. will eventually take over the world. Unlike Doctor Evil, it won't ask for a ransom. Key Learning: Accessing "unusual" data sources in Power BI; Using ML for analysis. Demos: 75%	Intermediate
Chalk-Talk	Anupama Natarajan	DS/DEV	AI capabilities in Power Platform	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Anupama Natarajan	ARCH/DBA	Migrate an on-premises Data Warehouse to Azure SQL Data Warehouse	Abstract: Azure SQL Data Warehouse is an enterprise-class, distributed database, capable of processing massive volumes of relational and non-relational data. Azure SQL Data Warehouse uses the Microsoft massive parallel processing (MPP) architecture. Come and learn on how to plan, prepare and to efficiently migrate an on-premises Data Warehouse to Azure SQL Data Warehouse. In this session you will learn • Prepare and expose source data from on-premises • Transfer the data into Azure • Load the data into Azure SQL Data Warehouse • Expose the data from Azure SQL Data Warehouse using Power BI • Tips and tricks to perform the migration Key Learning: Learn how to plan, prepare and perform migration of on-premises DW to Azure SQL DW Understand how to efficiently perform the migration Demos: Showcase migrating a sample on-premises Data Warehouse to Azure SQL Data Warehouse and exposing the data using Power BI	Intermediate
Break-Out	Aviv Ezrachi	BIA/ARCH	Microsoft Power BI: Introduction to Embedded Analytics	Abstract: Learn how Microsoft Power BI Embedded analytics gives users access to world class analytics inside applications, portals, and websites so users can make decisions in context. Join this session to familiarize yourself with the best practices for embedding Power BI for your users and products. Additionally, hear one of our customers, Bentley Systems, describe their experience with Power BI Embedded and demonstrate their application with integrated Power BI embedded analytics.	Basic

Break-Out	Aviv Ezrachi	BIA/ARCH	Microsoft Power BI: Introducing Power BI template apps	Abstract: Template Apps are integrated packages of pre-built Power BI content, configured to connect to specific data sources. With them, Microsoft partners can quickly access analytics for the apps and services they provide. Partners can also manage the Template Apps development lifecycle, from dev to marketplace to updates. Key Learning: By the end of this session you will learn how partners using Power BI can accelerate customer acquisition with a self service App building experience Allowing Power BI users to quickly connect to valuable data sources and immediately begin exploring, learning, and acting on their data Demos: Build, a template App Install a Template App and customize it to your	Basic
Chalk-Talk	Ben Weissman	BIA/DEV	ETL vs. Data virtualization – When to use what	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Ben Weissman	BIA/DEV	Let's Build SSIS Packages with Biml – Live!	Abstract: Have you heard about the Business Intelligence Markup Language (Biml)? Maybe you've even seen a session about it before but you still have doubts about how easily you can make something useful out of it. In this session, we'll use Biml to build and populate a staging area including the corresponding SSIS packages. But there won't be any pre-compiled demos - everything is happening live! Starting with a blank staging database, we will end up building a complete solution over the course of this session to prove that you can start from scratch and still quickly be successful. Let's see, how that goes… :) PS: Even if you have not heard about Biml but are still tired of manually building SSIS packages or ADF pipelines, this is the right session for you! Key Learning: What is Biml? How can I use it to automate my SSIS/ETL/adf Needs without any previous code Demos: Intro to Biml Build full staging Environment incl metadata repository from scratch	Intermediate

Break-Out	Ben Weissman	BD/ARCH	Big Data Clusters for SQL Server 2019 – a game changer?	Abstract: Big Data Clusters are probably the biggest feature in SQL Server 2019. They provide additional Artificial Intelligence, Machine Learning, Scale Out and Data virtualization capabilities. But what does that mean? We will take a deeper look at the extended Polybase options (Data Virtualization) as well as the option to integrate and query Big Data in SQL Server with the use of Kubernetes Clusters, Notebooks and more! Key Learning: We will look into the requirements to actually install BDC, the distinction between the different features within BDC and how/when to use them. Demos: deployment of BDC on AKS polybase configuration with remote servers and csvs spark queries against BDC	Intermediate
Chalk-Talk	Casey Karst	DBA/CL	Azure SQL Data Warehouse: New features and what they mean for you!	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Casey Karst	BIA/ARCH	Demystifying the Azure Modern Data Warehouse	Abstract: When building out your next Data Warehouse solution make sure that you are using the latest architecture design to maximize your business impact. In this session we will walk through the components of the Modern Data Warehouse architecture and gain a deeper understanding about how these components work together to create your solution. Key Learning: 1) Understand the Modern Data Warehouse Architecture and it's applications. 2) Have an understanding of the components of an MDW. 3) Learn about new features in SQL DW which enable BI reporting capabilities. Demos: I'll demo ADF, Databricks and Azure SQL DW. Current plan is to show MDW architecture in it's entirety then deep dive on SQL DW.	Intermediate

Break-Out	Charley Hanania	DBA/DS	Machine Learning Services for SQL Server On Premises - Why and How	Abstract: SQL Server comes with support for languages and constructs used for Machine Learning. Now it's viable to run predictive code segments for things such as forecasting, categorisation and anomaly detection in line with your T-SQL Coding objects. In this session, Charley walks through why this is important to DBA's and Developers, as well as how it's installed, configured and works under the hood. Key Learning: Why Machine Learning is so important for databases now. How to install ML Services correctly. Using it correctly so that it's an asset to the application ecosystem rather than a burden Demos: Demo on main installation points, another on using Machine Learning code within SSMS, installing and using the ML Management Reports	Intermediate
Open-Talk	Charley Hanania	DBA/DEV	MS Data Platform Certifications... What's changing and what should I target?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Open-Talk	Damian Widera	ARCH/DEV	Why database modelling matters	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Damian Widera	DS/BIA	Data Science for everybody. Beware of what you are eating	Abstract: This is a session about practical approach to the data analysis. I do not know much about mushrooms but I know that some of them could be poisonous. But what about other mushrooms? Should I eat them or not? If you have the same doubts than you should attend the session. I would like to show you the machine learning and neural networks approach in such type of problems. And then you will see how you can generalize the approach to other types of problems as well. Key Learning: Learn the approach to Machine Learning Learn the area of interests Learn how to start the process Learn the data cleanup - simple Learn how to pick an algotithm Demos: Demo in Python or R - read the data from the dataset - do a simple cleanup - pick a good algorithm - score the algorithm	Basic

Break-Out	Damian Widera	DS/DEV	Introduction to Deep Learning with Keras	Abstract: The session is an introduction to DeepLearning with Keras. I will show you the model of a neuron and how oyou can build a different types of neural networks . Then you will see how Keras can help you with all the tasks you need to do - create, run, score the neural network. Key Learning: What a neural network is What are the area of interests How to build a neural network with Keras Demos: We will be using CIFAR-10 dataset Keras installation Create neural Network with Keras Score the model	Intermediate
Break-Out	David Raj Daniel / Senthil Sundarajan	DS/ARCH	NLP using tensorflow	Abstract: Natural language processing is a sub field of AI involved in understanding human languages by Computer. It involves Natural language Understanding and Natural language generation . The application of NLP spawns across number of areas like Virtual assistant, Translation, Text mining, chatbot, Text sequence generation and many more. Tensor flow is a top deep learning framework which can be used efficiently for NLP as an alternative to traditional NLP. The session will focus on NLP using tensorflow and various deep learning techniques used for the same.	Advanced
Break-Out	Denny Cherry	DEV/DBA	Where should I be encrypting my data	Abstract: In this session, we'll be looking at all the various places within the application stack that data can be encrypted or hashed so that as an application/database administrator/developer/architect you can make the best decision as to where in the application stack to do this work. Key Learning: Encrypting and Hashing data within Application Tier Encrypting data within the Database tier Demos: Viewing some .NET code to encrypt data	Intermediate
Break-Out	Denny Cherry	CL/ARCH	How to Maintain the Same Level of utilities in Cloud Deployments - Securability, Reliability and Scalability".	Abstract: In this session, we will review the differences between deploying Microsoft SQL Server in Microsoft Azure and on-premises from a Security, Reliability and Scalability perspective. We'll review the common mistakes which people make when deploying SQL Server Virtual Machines to Azure which can lead to security problems including data breaches. Key Learning: We'll review the common performance problems which people encounter, and how to resolve them. We'll review the common scalability misunderstandings of Azure and SQL Server Virtual Machines. Demos: Azure Portal Azure Security Center	Basic

Break-Out	Denzil Ribeiro	DBA/ARCH	Improving Availability in SQL Server with Accelerated Data Recovery and Resumable Indexes	Abstract: Are your database application users frustrated from downtime due to common maintenance tasks or have faced long recovery times ? Are you constantly struggling to manage the log space usage and growth? Then this session is for you. We will deep dive into the internals of the new engine enhancements in SQL Server and Azure SQL Database that will drastically reduce database down time during some of the common activities such as database startup/recovery and indexing Key Learning: Learn about Accelerated Database Recovery and how it improves database availability and up-time Deep dive into Engine internals of both accelerated database recovery. Learn benefits and internals of resumable indexes Demos: Demo on Accelerated database recovery and how it helps availability Demo on improving application up-time with resumable Indexes.	Advanced
Break-Out	Denzil Ribeiro / Sanjay Mishra	DBA/DEV	Optimizing Workloads on Azure SQL Database	Abstract: In this session we will demonstrate several workloads and patterns on Azure SQL Database and Managed instances and discuss optimization techniques for each of them. Key Learning: In this session we will demonstrate several workloads and patterns on the different deployments on Azure SQL Database including Single DB and Managed instances. In each scenario we will discuss optimization techniques to increase throughput or reduce latency. Demos: Demos of optimizing multiple types of workloads and tasks on	Intermediate
Break-Out	Devashish Salgaonkar	DBA/ARCH	Practical Experiences of an Site Reliability Engineer with large scale Azure SQL DB Performance Management	Abstract: Azure SQL Database is the intelligent, fully managed relational cloud database service. In Microsoft we use the service as a backend database for other Online services. In this session, we will walk you through some of our practical experiences and key learnings while working with SQL Azure databases deployed as a backend for a large scale service - particularly, detecting and fixing performance issues proactively. Key Learning: To intelligently harvest the power of Log Analytics data to maintain the reliability and performance of Azure SQL databases. Demos: Querying Log Analytics basics Building a Perf Index for a database Setting up alerts in OMS	Intermediate
Key-Note	Dr. Rohini Srivathsa		Keynote by Dr. Rohini Srivathsa, Microsoft		
Key-Note	Gayle Sheppard		Keynote by Gayle Sheppard, Microsoft		

Break-Out	Ginger Grant	DS/DEV	Introduction to Machine Learning with Azure Databricks	Abstract: Azure Databricks, one of the newer components added to Azure, allows users to connect to data sources such as Azure Data Lake, Azure Blob Storage, Azure SQL DW, Cosmos DB, Azure DB, and SQL Server and stream the data using Apache Spark for processing data to create a machine learning [ML] solution and providing the data to Power BI for visualization. Azure Databricks can provide a very quick way of processing data by adding nodes increase performance for tasks, such as analyzing data for a ML solution from an Azure data store. Azure Databricks also includes a collaborative workspace so that using Azure Active Directory, teams of people can create code in a notebook in R or Python and implement the notebook as an Azure Databrick job. The step-by-step demos will include all you need to know to implement Databricks. Key Learning: The audience will learn the different component parts of Azure Databricks, how to use the environment to create and deploy code within a team environment, and when Azure Databricks is the right tool for your machine learning solution. Demos: The demo will walk through the different components of Azure Databricks to show how it can be used to created machine learning solution, and how the solution can be run immediately and on a scheduled basis.	Basic
Break-Out	Ginger Grant	DS/ARCH	Revealing the Differences in Azure Machine Learning Products	Abstract: When searching for Azure ML [AML], the browser returns a number of different products. Which product do you want to use when? Which is the appropriate tool for integrating with SQL Server or data in Azure? Internet searching Azure ML [AML] returns several different products. Learn what the differences are between them, which product to use when, and see how to use AML to create a complete data science solution using Azure and SQL Server. Learn the differences between the Azure ML products and how to combine functionalities to create, deploy and manage ML models deployed in your organization.	Intermediate
Open-Talk	Gogula Aryalingam	ARCH/BIA	How to get businesses started and going with analytics	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

Break-Out	Gogula Aryalingam	BIA/DEV	Azure Data Factory for Ingesting and Integrating	Abstract: When building a modern data warehouse, the mantra for getting data on board is ELT, and not ETL. Building an ELT solution is not as simple as switching two letters. This session will give you a brief of what a modern data warehouse is, why its data integration needs are different, how you would design a solution for this need, and finally how you would build it with Azure Data Factory. Key Learning: A briefing about what a modern data warehouse is and why it's needed. The concepts of integrating data into a modern data warehouse. What the different technologies are for this, and how it can be performed using Azure Data Factory. Demos: The demo would feature using Azure Data Factory to ingest data to an Azure Data Lake, and then integrate that data using Azure Data Factory (SSIS) onto a Azure SQL Database-based	Intermediate
Break-Out	Hamish Watson	CL/DEV	How to deploy SQL Server on Kubernetes in Azure	Abstract: When Microsoft released SQL Server 2017 it allowed us to run SQL Server in docker containers. This radical change provides a wider and open platform that data professionals can choose how to deploy and run SQL Server. This session will demonstrate how containers are a game changer for deploying and managing SQL Server. It will also showcase how the Azure Kubernetes Service (AKS) is a scalable and highly available platform for SQL Server running in containers. Demonstrations will show how easy it is to create containers, deploy SQL Server in them and how to automate and manage your SQL Server containers using Kubernetes on the Azure platform. This session will show you the exciting future of SQL Server and how you can learn about it today. Key Learning: Goal 1: The audience will learn what docker containers are and how to deploy SQL Server in them. Goal 2: The audience will learn how to build a custom SQL Server image and deploy this to a container registry in Azure Goal 3: The audience will learn how deploy and manage containers running on the Azure Kubernetes Service. There will also be discussion and demonstration of running containers on cabernets locally as well. Demos: The DEMOs will showcase SQL Server running in containers and then pushing those containers to Azure and running them on the Kubernetes platform.	Intermediate

Chalk-Talk	Hamish Watson/Alicia Moniz	DEV/DS	How DevOps Can Help Data Scientists	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Harsh Thakur / Kranthi Medam	BD/ARCH	Building the Enterprise Data Lake with Azure Data bricks: A real life case study on best practices and design of the enterprise data lake for people and organizational analytics	Abstract: Are you interested in understanding how Microsoft builds it's enterprise data lake and also does people/organizational analytics, which then helps in understanding various critical aspects of the enterprise. If so, this is a session you wouldn't want to miss. The session will focus on a real life scenario on how we use Azure Databricks along with Azure Data Lake Gen2 to ingest, store and process a large amount of data and build connected insights using machine learning techniques. You will get to learn on using some key design patterns to load data, advanced features like data bricks delta lake, performance tuning, coding in scala and spark, enabling a lambda architecture. In addition you will also get to see how we have segregated compute (data bricks), storage (adls gen2), orchestration (ADF v2) and presentation (PowerBI/Azure AS) to get the most out of the azure stack. Key Learning: In addition to the technology stack audience can expect to know about design considerations for building a data lake, figuring out how to design the Overall workflow to achieve optimal cost and performance, debugging spark and scala code Enterprise data lake Architecture Azure Data lake Gen 2 Azure Data bricks Azure Data Factory Azure Analysis Service+ Power BI Demos: Demos on 1) Writing optimal spark/scala code to load data that is performant. 2) Loading data from multiple files into single parquet 3) Building facts and dimensions using data bricks delta lake. 4) Build Power BI dashboards by connecting to the lake 5) Leverage the power of Azure Analysis services for advanced analytics. 6) Orchestrate the flow of data using ADF v2.	Advanced
Open-Talk	Harsh Thakur / Kranthi Medam	BD/ARCH	Enterprise Data lake solution patterns with Azure Databricks.	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

Break-Out	Heidi Hasting	DEV/DBA	Intro to Azure Data Studio	Abstract: Azure Data Studio, keeping in line with Microsoft's ethos of tooling that works on many platforms this one runs on Windows, macOS and Linux. This session is to give you an introduction to the product give you an idea of where/how it can be used in the Data Platform space. We will also cover some of the main features (code snippets, widgets, extensions) and time permitting Notebooks. Hopefully by the end of the session you will be aware of Azure Data Studio and be ready to give it a go! Level: Beginner Level 100 Key Learning: Tooling that can help with daily job T-SQL Notebooks to help Demos: Demo'ing features mentioned in abstract	Basic
Break-Out	Hugo Kornelis	DBA/DEV	Now Where Did THAT Estimate Come From?	Abstract: The SQL Server Query Optimizer makes its plan choices based on estimated rowcounts. If those estimates are wrong, the optimizer will very likely produce a poor plan. And there's nothing you can do about it. Or is there? In this session, you will learn exactly where these estimates come from. You will gain intimate knowledge of how statistics are built and maintained and how they are used to estimate row counts. But you will also learn how filters and joins influence those estimates. Though the focus of this session is on understanding the cause of bad estimates, you will also learn ways to fix the problems and get better estimates - and hence, better performing queries. Key Learning: After a brief discussion of common and well-known causes of bad estimates, I move on to a discussion of how statistics are stored in SQL Server. I then cover how these statistics are used to estimate cardinalty when various types of filters are used. Just a single filter at first, and then multiple filters. After that I talk about cardinality estimation of a join, by histogram alignment. Demos: During the session I will use short demos to illustrate key elements of my talk	Advanced
Chalk-Talk	Hugo Kornelis	DEV/DBA	Scalar UDF inlining in SQL 2019: How does it work?	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA

| Break-Out | Hugo Kornelis | DBA/DEV | Execution plans ... where do I start? | Abstract: SQL (the language) is not a third generation language, where the developer tells the computer every step it needs to take. It is a declarative language that specifies the required results. SQL Server itself will figure out what steps it takes to get to those results. Most of the time, that works very well. But sometimes it doesn't. Sometimes a query takes too much time. You need to find out why, so you can fix it. That's where the execution plan comes in. In the execution plan, SQL Server exposes exactly which steps it took for your query, so you can see why it's slow. However, execution plans can be daunting to the uninitiated. Especially for complex queries. Where do you even start? In this session you will learn how to obtain execution plans. and how to start reading and understanding them. Key Learning: Attendees will learn where to find execution plans, and how to start reading and interpreting them. They will also learn about some common misconceptions. Attendees will learn the importance of not just looking at the graphical execution plan, but also looking at the properties. Finally, attendees will also get some basic information about a few of the most important execution plan operators. Demos: During the session I will use short demos to illustrate key elements of my talk | Basic |

Break-Out	Janusz Rokicki	DBA/DS	Anomaly Detection in Predictive Maintenance for Sql Server	Abstract: The amount of data we collect, store and process is rapidly growing and goes hand in hand with growing amount of work related to managing it. Automation is a key if we want to stay in control and AIOps represents its next major level with potential to revolutionise IT operations. This session explains how to apply data science to operations - specifically Sql Server administration, provides common use cases and gives practical examples that help to avoid costly Sql Server disruptions and eliminate firefighting. Presented solutions give great insights and addresses limitations of Query Store and improve security monitoring. Key Learning: Structured process how to apply data science to administration of Sql Server to get actionable insights, automate routine tasks and prevent downtime. Demos: I'm going to show 4 demos (real data) that cover different techniques of detecting anomalies: 1. Finding password spraying attacks based on content of ERRORLOG 2. Analyzing how QueryStore affects performance 3. Finding irregularities in flow control (slow clients) 4. Behavioral analytics/finding abnormal user activity based on index operational stats	Advanced
Chalk-Talk	Janusz Rokicki	DEV/DBA	Maturity Model for SELECT query	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Jen Stirrup	DS/DEV	Deep Learning in Azure with Tensorflow	Abstract: What is Microsoft's approach to Deep Learning, and how does it differ from Open Source alternatives? In this session, we will look at Deep Learning, and how it can be implemented in Microsoft and Azure technologies with Tensorflow in Azure. Join this session in order to understand deep learning better, and how we can use it to provide business and technical benefits in our organizations. Key Learning: AI is accessible, and if you know R, you can do AI. Demos: What is Microsoft's approach to Deep Learning, and how does it differ from Open Source alternatives? In this session, we will look at Deep Learning, and how it can be implemented in Microsoft and Azure technologies with Tensorflow in Azure. Join this session in order to understand deep learning better, and how we can use it to provide business and technical benefits in our organizations.	Advanced

Chalk-Talk	Jen Stirrup	BIA/ARCH	Approaches for good data visualization	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Jen Stirrup	DS/DEV	R in AI Applications	Abstract: Are you a business intelligence professional or a developer who works with data? You should consider adding R to your skillset so you can add Artificial Intelligence to your applications. R is a popular for analysis and data visualization, but it also can be used for concepts in AI such as inference and prediction. In this demo-based session, we will go through in-depth examples of using R for the purpose of Artificial Intelligence work. We will also see how we can embed it in for production. Key Learning: AI is accessible, and if you know R, you can do AI. Demos: Demo will show forecasting sales data using AI inspired techniques such as prediction and inference. We will also use natural language processing to understand the data.	Advanced
Open-Talk	Joey Dántoni	DBA/ARCH	Security in the Cloud—Keeping Things Secure	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Joey Dántoni	DBA/ARCH	Containers, PoData Science, and Databases--Learning About the Future of Infrastructure	Abstract: Containers have quietly been taking over the world of infrastructure, especially amongst developers and CI/CD practitioners. However, in the database space, container adoption has been lower. SQL Server 2017 introduced the concept of deploying databases into Docker containers. In this session, you will learn the fundamentals of creating containers, learning about Kubernetes for management, and how to further your learning in this new and emerging space. Key Learning: Kubernetes overview SQL on Kubernetes Availablity Groups on Kubernetes Demos: Upgrade SQL Server Patch SQL Server Failover SQL	Intermediate
Break-Out	Joey Dántoni	DBA/CL	Azure for the DBA	Abstract: Are you getting started with the cloud? In this session you will learn about booth Infrastructure and Platform as a Service offerings in Azure. You will learn about storage and networking and how they apply to your SQL Server and other data environments. You will also learn about the benefits of using platform as a service offerings and how they can improve your productivity. Key Learning: Learn about: Storage in Azure Networking in Azure Choosing the Right VM Demos: Build a VM, Configure storage, configure networks	Basic

Chalk-Talk	John Q Martin	DEV/DBA	Intro on how to design Row Level Security in SQL Server	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	John Q Martin	DEV/ARCH	Building Secure Solutions with SQL Server & Azure SQL DB	Abstract: SQL Server and Azure SQL Database introduced a wealth of new security features over the last couple of years. But which ones should we use in our database development? Together we will look at some of the key security features for developers in SQL Server and Azure SQL Database. understanding when to use them and the threats they mitigate. Ranging from effective use of database roles for securing access to objects and data. All the way through to Dynamic Data Masking, Row Level Security, and Always Encrypted. The SQL Server platform allows us to build more secure solutions that meet the compliance and security needs of our businesses and customers today. Key Learning: Understand when and how to use the different security features of SQL Server. Whether to use one or layer multiple features to build a secure database solution. Demos: 1. Why to avoid the use of Trustworthy databases 2. How to create and use database roles effectively 3. How to use Dynamic Data Masking to restrict sensitive data visibility 4. Some limitations of Dynamic Data Masking 5. How to use Row Level Security in our applications 6. How to use Always Encrypted in our applications 7. How Always Encrypted can impact the size of your databases	Basic

Break-Out	John Q Martin	DBA/DEV	Managing SQL Server with PowerShell DSC	Abstract: Whether you are managing 1 or 1000 servers automation is a vital skill to have today. With PowerShell DSC we can automate the deployment and configuration management of our SQL Server estates. From pushing configurations manually as needed through to having SQL Server systems check in, validate, and remediate their configuration as needed. PowerShell DSC lets us solve the management headache of performing many admin tasks without needing to login to the server. Key Learning: How to manage SQL Server configuraitons via PowerShell DSC. How to work with large numbers of servers without needing large teams. That PowerShell DSC applies to older versions of SQL Server as well as new ones. Demos: Installing SQL Server via PowerShell DSC by pushing a configuration. Configuring SQL Server via PowerShell DSC by pushing a configuration. Adding a SQL Server to a Pull server. Deploying a configuration via Pull Server.	Intermediate

Break-Out	Kanika Gera / Manik Jindal	DEV/ARCH	Adobe I/O Events & Microsoft Flow	Abstract: Introduction: Harness the power and unlock the potential to develop, extend and integrate with Adobe solutions using Adobe I/O - Adobe Developer Platform which is based on servless architecture based on Apache Open whisk.Underlying Technology Stack :A technical deep dive about Adobe I/O Developer Platform and open source technologies it is built on. Solving Technical integration Use-Cases: Some powerful use-cases with a demo (Microsoft Flow + Events) to solve highly technical integrations almost with no code. Demo: Flow into Clouds with Adobe I/O Events & Microsoft FlowStory Line: There's a great story of a programmer who automated his work. His scripts made coffee for him, emailed his wife when he was working late, and sent in sick notes when he had a hangover. (I want to be him when I grow up!)Now, with a little help from Adobe I/O Events and Microsoft Flow. I'm one step closer to that dream even when i'm not a Developer!Let's consider this - A graphic designer uploads an image on creative cloud assets. Image automatically tagged and intelligently described. Once done, email sent to his lead for approval and upon approval the image is automatically posted on twitter and the details of the images are created as a note. (A live "Magical"demo will be shown while this script is being narrated)Key Learning: 1) Introduction to Adobe Developer Platform products based on Open source technologies like OpenWhisk, Kafka, Nodejs, etc. 2) Technology Stack of Adobe Platform 3) Integrating Adobe I/O Events & Microsoft Flow4) Making Event driven integrations without code5) Leverage GDPR Events for the EU-GDPR compliance in your Power Flow into Clouds with Adobe	Intermediate
Break-Out	Kevin Boles	DEV/DBA	Common TSQL Mistakes	Abstract: We are going to examine a variety of mistakes MANY developers fall prey to - some obvious, some pretty subtle and some downright sneaky! Lots of code examples with the bad AND good code presented. I GUARANTEE that you will find things here that will either prevent you from getting bad data, throwing unwanted errors or vastly improving your database application's performance and concurrency. I have given this talk almost 100 times now and it is always very highly rated! Key Learning: Learn as many TSQL coding issues focused on performance and bad data as we can fit into the allotted time! Demos: The session is almost entirely demos	Intermediate

Break-Out	Kevin Boles	DEV/DBA	Know What Your Code is Doing to SQL Server	Abstract: ORMs - oh how I LOVE them! As a SQL Server consultant, when I see them in play at a client I see $$$$$$$$$$!! :-D They can definitely help you code faster. But if you don't Read The Fine Manual, do some things right, and avoid some major gotchas, you will fall prey to one of my favorite Guruisms: "Anything that allows developers to slap code together more quickly is inversely proportional to the performance, concurrency and scalability you will get from that code! Although this talk is based on Entity Framework, most ORMs suffer from a lot of the same flaws. Key Learning: Learn just how HORRIBLE ORMs are for database performance when used without a deep understanding of how to PROPERLY use them Demos: no demos, although there are screenshots of code and code output and messages	Intermediate
Chalk-Talk	Kevin Boles	DEV/DBA	Why are Things SLOW? The Tuning Big Three	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Manohar Punna	DEV/ARCH	PowerApps and Flow Best Practices – Working with SQL Server	Abstract: Microsoft PowerApps and Flow have seen tremendous adoption in recent years. With an ever-growing need for rapid and agile solution development, these two products have the potential to solve business problems with very minimal development efforts. However, as with any product, with great simplification comes greater problems. As a data professional, it is our responsibility to make sure the data layer is accessed with proper security and efficiency. In this session, I will explore different methods to connect to the data from PowerApps and Flow. In this demo filled session, I will cover different scenarios for sending and receiving data in PowerApps using Flows. Key Learning: The attendees will learn design best practices in PowerApps and Flow while working with data layer. Demos: 1. Connecting to SQL Server databases 2. Receiving data into PowerApps using Flow 3. Sending data from PowerApps using Flow and receiving a response	Basic
Chalk-Talk	Martin Cairney	DBA/DEV	How to use dates properly in SQL	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA

				Abstract: If you've dipped your toe into the waters of Azure SQL DB, you'll know that there's no SQL Agent available. Fortunately there are many options available to get you over this hurdle. One of these is Azure Automation, a scheduling engine which runs PowerShell or Python scripts to perform your tasks. As a bonus, these Runbooks can also be triggered from Azure SQL DB alerts. This session will introduce Azure Automation, from a basic, manually executed Runbook to an alert-driven responsive utility that can save you considerable effort. We'll then take Alert responses a step further and use them to trigger an Azure Logic App workflow - a utility that provides a robust workspace with built-in scalability & retry logic along with a multitude of connectors which can expand the response to your alert. Imagine being able to capture an Alert, perform the remedial action, log the issue & response in your ticketing system and switch on your kettle so you can sit back and let Azure do it all for you. Key Learning: Attendees will learn about the new Azure Monitor capabilities for capturing performance metrics from Azure SQL Database and the format of queries used to compare against thresholds and trigger alerts. Attendees will learn the fundamentals of building an Azure Automation Runbook to perform basic tasks against a selected Azure SQL Database. Attendees will see how a comprehensive workflow in response to Azure SQL Database performance or activity related alerts can be built and used to simplify their day to day management tasks. Demos: Demos will include: Creating Azure Automation	
Break-Out	Martin Cairney	CL/DBA		Smart Alerts for Azure SQL Database using Automation & Logic Apps	runbooks Creating Logic App Creating Alerts from SQL DB	Advanced
Open-Talk	Martin Catherall	CL/DEV		Use-cases for Real-World Cosmos DB Implementations	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Chalk-Talk	Martin Catherall	DBA/ARCH	Data Privacy for the Data Driven professional		Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA

| Break-Out | Martin Catherall | DEV/DBA | Get Asynchronous with SQL Server and Service Broker | Abstract: While Service Broker has been in the SQL Server product since 2005 it remains a mystery to a lot of developers and database administrators. This generally stems from the fact that a whole host of issues outside of SQL Server such as firewalls, networks, security and infrastructure can contribute to the management of a Service Broker implementation and this can often to over whelming for the developer or database administrator – especially when troubleshooting an issue. However, it doesn't have to be this way and this session will look at a very simple implementation of Service broker and then scale it to a more advanced concept – highlighting how asynchronous operations can help to scale, optimise and de-couple SQL Server components and ultimately lead to more robust application, enhanced stability and ultimately happier customers. Key Learning: Service Broker can be built simply and then scaled. Service Broker can aid concurrency in the database. Messaging in SQL Server can solve many problems Demos: Build a simply Service Broker implementation - highlighting the basic concepts. Scale this implementation to a more advanced implementation. | Intermediate |

				Abstract: The job of a data professional is evolving rapidly, driving many of us to platforms and technologies that were not on our radar screen until quite recently. I am certainly no exception to that trend. These days we are not just monitoring backups and tuning queries - we are collaborating with teams throughout the organization to provide them data and insights that drive decisions. Cloud providers are democratizing technologies and techniques that were complicated and proprietary not too long ago. This presentation walks you through how a silly idea from a football podcast got me thinking about how Azure Logic Apps, the Cognitive Services API, Azure SQL DB, and Power BI combine to provide potentially powerful insights to any organization with a social media presence. Join me as I walk you through building a solution that can impact your company's bottom line - and potentially yours too! Key Learning: Attendees will learn the following: what sentiment analysis is, what Azure Logic Apps are (and why they're so powerful), and how to monitor sentiment of any written source they ingest into the logic app. They will also learn how to wire this up to Power BI for instantly actionable information. Demos: This talk has two main demos. The first one is a complete walkthrough of creating an Azure Logic App, connecting it to Cognitive Services, connecting it to an Azure SQL Database, and then using that Azure Logic Apps to store and score information about tweets using a specific hashtag. The second demo is using that logic app to stream Twitter and sentiment data to a	
Break-Out	Matt Gordon	BIA/DS	Feelings Quantified: Scoring, Storing, and Exploring Social Media Sentiment	Power BI streaming dashboard in order to monitor live event	Intermediate

Break-Out	Matt Gordon	DBA/CL	Azure-d Availability: SQL Server HA In and To the Cloud	Abstract: Has your manager come to you and said "I expect the SQL Server machines to have zero downtime?" Have you been told to make your environment "Always On" without any guidance (or budget) as to how to do that or what that means? Are you facing pressure to have data in Azure as well? Help is here! This session will walk you through the high availability options in on-premises SQL Server, the high availability options in Azure SQL Database and Managed Instances, and how some or all of those can be combined to enable you to achieve the ambitious goals of your management. Beyond the academic knowledge, we'll discuss frequently seen scenarios from the field covering exactly how your on-premises environments and Azure services can work together to keep your phone quiet at night. Key Learning: 1) Understand all on-premises HA options in SQL Server and which play well with Azure 2) Understand all Azure-only HA options within Azure SQL Database and Azure SQL Database Managed Instances 3) Understand hybrid HA scenarios frequently seen in the field Demos: I will demo setting up replication from an on-premises database to an Azure SQL Database. I will also provide a demo/walkthrough of an Azure SQL Database Managed Instance as many people have not	Intermediate
Open-Talk	Matt Gordon/Alicia Moniz	DS/DEV	Real World Applications for Cognitive Services Text Analytics and Vision APIs	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Open-Talk	Mia Chang	DS/DEV	Data Scientist career path	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Mia Chang	BIA/DS	A data scientist journey in real estate industry	Abstract: As a data scientist, I don't really know lots of knowledge about the real estate industry(also till now). In this session, I would like to share what and how I learned to apply my knowledge in deep learning with a particular domain. And I will also share how to develop a production level service for your AI application. Hope can bring up more people, start your data science path from your interested. Key Learning: - cross-team communication - build data science team work routine from scratch - tips for developing the production level product Demos: None	Basic

Open-Talk	Monica Rathbun	DBA/DEV	Networking 101 for your next Career	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Monica Rathbun	DBA/ARCH	Performance Tuning SQL Server on Crappy Hardware	Abstract: Many of us have to deal with hardware that doesn't meet our standards or contributes to performance problems. This session will cover how to work around hardware issues when it isn't in the budget for newer, faster, stronger, better hardware. It's time to make that existing hardware work for us. Learn tips and tricks on how to reduce IO, relieve memory pressure, and reduce blocking. Let's see how compression, statistics, and indexes bring new life into your existing hardware. Key Learning: How to work with other admins to get SQL Server to run better on their hardware. Check list on what to look for and how to prove hardware is the issue. Demos: None	Basic
Break-Out	Muniraju Pulipalyam	DS/DEV	Advanced AI for Practitioners	Abstract: This session covers recent AI/ML advances and how best to use them in Microsoft Cloud, Edge and AI platforms. The topics covered are AutoMachine Learning, Azure Machine Learning Services, Machine Learning DevOps and model deployment at scale. The session also covers how large scale infrastructure can be setup for highly parallel data processing and data pipeline setup in Azure very easily. Key Learning: Learn about how data scientists and ML engineers can use latest ML advances Demos: AutoML Azure ML Services and Visual ML Designer ML Devops	Advanced
Break-Out	Ola Hallengren	DBA/DEV	Query Store In-Depth	Abstract: Query Store has become one of the most important tools for troubleshooting performance problems. However, there are some quirks that are important to be aware of. In this session I will share my experiences about how Query Store is working in different scenarios. We will also look at some real-world cases. Finally we will look at Query Store in SQL Server 2019. Key Learning: Get a better understanding of Query Store Demos: Several Query Store demos	Intermediate

Break-Out	Peter Myers	BIA/DEV	The Do's and Don'ts of Power BI Relationships	Abstract: This is a presentation about Power BI modeling with a specific focus on relationships. In short, relationships are essential to delivering intuitive, accurate and optimal models. Interestingly, relationship design can be considered part science and part art. You will learn what relationships do and how they can be configured for Import, DirectQuery and Composite models. The introduction to relationships is reasonably straight forward. Theory and demos will progress beyond the basics to explain how to achieve sophisticated model designs that include role playing dimensions, many-to-many relationships (two techniques), ambiguous relationships, and disconnected tables. All DAX functions that leverage or modify model relationships will also be covered. Key Learning: In this session, you will learn: • The purpose for relationships • How to optimally design data models • How to configure relationship properties • Good design practices, including when to use specific relationship patterns Demos: A demonstration of each relationship type.	Intermediate
Break-Out	Peter Myers	BIA/ARCH	Delivering Real-Time Power BI Content	Abstract: This session will provide the know-how to deliver real-time Power BI content. It will cover real-time dashboard tiles, the Power BI REST API, and Azure Stream Analytics (incorporating predictive analytics), and reports based on DirectQuery data models. This session will be relevant to business analysts, developers and IT Pros In this session, you will learn: • How to assemble real-time Power BI dashboards • How to create streaming and push datasets • How to programmatically push data to Power BI • How to integrate prediction results into real-time dashboards • How to deliver low latency results through reports • How to choose the appropriate real-time technique Key Learning: In this session, you will learn: • How to assemble real-time Power BI dashboards • How to create streaming and push datasets • How to programmatically push data to Power BI • How to integrate prediction results into real-time dashboards • How to choose the appropriate real-time technique Demos: A demonstration of each of the three real-time techniques.	Intermediate
Chalk-Talk	Peter Myers	BIA/ARCH	Solid Power BI Model Design Guidance	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA

Open-Talk	Peter Myers	ARCH/BIA	Power BI Reports vs. Dashboards: Which is the Better Fit?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Praveen Srivatsa	DS/DEV	Detecting human pose with AI	Abstract: Understanding human poses has multiple challenges. To start with, we need to interpret the human body structure and identify parts of the body. But most of the time, the human body is not visible as a static pose. The body could be sleeping, sitting, bent over, partly visible etc which adds on to the challenges on identifying it. To add to this, if there are multiple humans (like in a game or a room), identifying each part associated with an individual itself is a challenge.. In this session, we talk about how we can go about detecting human pose and how we can use it for various use cases including security monitoring, skill training and health assessment. Key Learning: Understanding how to approach a problem and map it to an AI/ML solution. Leveraging vision algorithms for detecting human poses Various use cases that can benefit from pose detection of human beings. Demos: Training pose algorithms with TensorFlow Python based algorithms for pose detection at the edge -- for single image -- for multi-image Expost the pose algorithms to other frameworks Real-time pose	Basic
Break-Out	Raj Pochiraju / Mukesh Kumar	DBA/ARCH	Part-1: SQL Modernization best practices and lessons learned (including EOS)	Abstract: Are any of you still running SQL Server 2008/R2 and would like to migrate and upgrade soon? Are you willing to modernize to modern SQL Server platforms to achieve breakthrough performance, maintain security and compliance, and optimize data platform infrastructure? Azure SQL Database and Managed Instances allows you to build globally scalable applications with extremely low latency. Azure SQL Database is the best cloud database offering in market. In this session, we will take a detailed look at the migration life cycle, modernization best practices and share customer stories who already migrated and learnings. We will also cover most commonly seen migration blocking scenarios and demonstrate how our service can unblock your migration to Azure SQL Databases.	Intermediate

Break-Out	Raj Pochiraju / Mukesh Kumar	DBA/ARCH	Part-2: Deep Dive with Data migration tools and Azure Database Migration Service (DMS)	Abstract: In this session, we will take a detailed look at the migration life cycle and show you how we have made it easy to migrate SQL Server instances to Azure with near-zero downtime by using the Azure migrate, Data migration Assistant (DMA), Azure Database Migration Service and related tools. We will take deep dive into our discovery, assessment and migration tools that help you to identify low-hanging fruits to migrate to Azure, determine right Azure data target, size and perform near zero downtime migrations across your data estate, heavy focus on live demos. Key Learning: 1: Understand the phases in Database Migration Journey, DMS service and partner tools your customers leverage to take migration journey to Azure. 2: Deep dive into assessment tools to determine the right Azure target, the right target size and readiness of databases migrating to Azure at scale. Demos: 1. Perform online database migration using DMS 2. Automate Database migrations using PowerShell	Intermediate
Break-Out	Reid Havens	BIA/DEV	Unlocking New Visualizations and Features in Power BI	Abstract: Power BI is a great sandbox environment for report design. However, knowing how to leverage visuals and features in a way that adds additional value can be challenging. New visualizations and features can be created a number of ways including: layering visuals, customizing visual formatting, and/or utilizing DAX measures. The session will include a series of visualization and reporting techniques that you'll be able to leverage in your company's reports to take them to the next level.	Basic

Break-Out	Reid Havens	BIA/DEV	Demystifying Chart Types and Report Design Principles in Power BI	Abstract: Failing to deliver a well-designed Power BI Report can be a common reporting pitfall. What good is quality data if it is not presented in a way that is meaningful or easily understood? Someone without any prior knowledge should be able to quickly understand a report without explanation and be quickly drawn to the key elements you want them to view. This talk will walk through many elements of bad report design. Learn about visual cues and how certain chart types can convey data more accurately than others. Also, learn about the basic dos and don'ts of report design and layout, using easy-to-learn techniques that bring data to life. Key Learning: Goal 1 - Learn about core principles of report design, and practices to avoid. Allowing delivery of more professional looking reports. Goal 2 - Understand how users consume information, and how to apply that knowledge to create more effective reports. Goal 3 - Comprehend the mechanics of visual cues and chart design. Helping to create a more accurate and clearer story for the data being visualized for the report consumers. Demos: A short demo will be presented at the end of the presentation. This demo will showcase many of th design principals and practices implemented into a Power BI report.	Basic
Open-Talk	Reid Havens	BIA/DEV	Power BI Data Model Optimization Practices	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

Break-Out	Rony Chatterjee	BD/DEV	SQL Server 2019 Big Data Clusters: Journey to a Modern Data Platform	Abstract: In this session we will cover scenarios which modern data developers struggle with and how SQL Server 2019 can be the solution to the issues they constantly face. We will discuss in depth the concept of data virtualization and the benefit it provides over conventional ETL. We will also show the architecture on which SQL Server 2019 Big Data Clusters were built and the choice of the platform we used. We will also discuss how we are bringing SQL Server and Spark as the two compute engines over SQL and HDFS and how we built a unified data analytics platform with SQL Server Ml Services and Spark ML and how we deploy models back into the platform. We will also show how this entire product was built thinking about the user experience and how we optimized this experience in Azure Data Studio. Key Learning: Deep Dive On SQL Server and Big Data Clusters. Build intelligent apps and AI with all your data Analyzing all data. Easily and securely manage data big and small Managing all data. SQL Server enables intelligence over all your data. Integrating all data Managing all data Easily deploy and manage a SQL Server + Big Data cluster . Build ML models in Python and Spark. App Deploy. Learn to use Azure Data Studio. Demos: https://www.youtube.com/watch?v=s2G-IuhIaA8&_lrsc=4d56ab7e-87f1-4bb9-9948-7bed018b7fe8 https://www.youtube.com/watch?v=uZRXlkBBRnM&t=51s https://www.youtube.com/watch?v=kCcDnua0Fbk&t=458s	Intermediate

Break-Out	Rony Chatterjee	BD/DS	Introducing Azure Data Studio and Machine Learning in SQL Server 2019 Big Data Clusters	Abstract: Azure Data Studio is a cross-platform data tool for data professionals using on-premises and cloud data platforms on Windows, MacOS, and Linux. We will show the built in experiences for SQL Server, PostgreSQL and other databases and also have introduced Notebook support in Azure Data Studio to not only write your ML models in Python and Spark but have also included the support for your SQL development with Intellisense, code snippets, source control integration and an integrated terminal experience. I will walk through with the experience to virtualize the data using PolyBase and then how we can prepare, visualize, build ML Models using Python and Spark, Machine Learning Model Management, deploying models back into SQL Server. Key Learning: Learn about Azure Data Studio. Introducing PostgreSQL support in Azure Data Studio. Data Virtualization. Introducing SQL Notebooks. Data Preparation using Notebooks. Data Visualization in Azure Data Studio. Building ML Model in Python and Spark. Model Management using MlFlow. Deploying Model back into SQL Server. Demos: https://www.youtube.com/watch?v=uZRXlkBBRnM&t=51s https://www.youtube.com/watch?v=kCcDnua0Fbk&t=458s	Intermediate
Chalk-Talk	Rony Chatterjee	DS/DEV	MLFlow Model Management	Chalk-Talks are 30 minutes' sessions focussing on conceptual & architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	NA
Break-Out	Sandeep Alur	DS/ARCH	Reinforcement Learning – A Guided Tour	Abstract: Artificial Intelligence as a stream of technology is evolving at a rapid pace. While Supervised and Unsupervised learning are helping us solve key scenarios, a type of learning that will have a significant impact on software agents taking real-time action is 'Reinforcement Action' (RL). RL based systems are making their way into mainstream services with an intent to drive consumer experiences. Taking timely decisions leading to positive outcome is key and this is where RL based systems operate to maximize effectiveness of results. Azure Cognitive Services recently added a new pre-built API called 'Personaliser' that is powered by RL based capability. Join this session to understand the core tenets of 'Reinforcement Learning' as a type of Machine Learning, and how 'Personalizer' a cognitive API helps deliver personalised use experience.	Intermediate

				Abstract: This session focuses on the all PaaS solution of Azure SQL Database/Managed Instance and SSIS in Azure Data Factory (ADF) to lift & shift on-premises ETL workloads and modernize ETL workflows in the cloud. We will first show you how to provision Azure-SSIS Integration Runtime (IR) - dedicated ADF servers for running SSIS - with SSIS catalog (SSISDB) hosted by Azure SQL Database server/Managed Instance or local file shares/Azure Files to store your SSIS projects/packages and extend it with custom/Open Source/3rd party components. We will also show you how to configure Azure-SSIS IR to access data on premises using Windows authentication, Virtual Network (VNet) injection, Self-Hosted IR as proxy, etc. Preserving your skillsets, you can then use the familiar SQL Server Data Tools (SSDT)/SQL Server Management Studio (SSMS) to design/deploy/configure/execute/monitor your SSIS packages in the cloud just like you do on premises. Next, we will guide you to modernize your ETL workflows by invoking/scheduling SSIS package executions as first-class activities in ADF pipelines and combining/chaining them with other activities, allowing you to leverage ADF connectivity (+80 connectors), store secrets in Azure Key Vault (AKV), connect using Azure Active Directory (AAD) authentication with ADF managed identity, automatically provision Azure-SSIS IR on demand/just in time, trigger SSIS jobs/runs by file arrivals/deletions, and splice advanced data transformations, such as Power Query Source, Data Quality (Matching/Cleansing) components, etc. Key Learning: - Lift & shift your existing ETL workloads w/ little/no	
Break-Out	Sandy Winarko	BIA/DEV	PaaSport to Paradise: Lifting & Shifting with Azure SQL Database/Managed Instance and SSIS in Azure Data Factory		Intermediate
Open-Talk	Sandy Winarko	BIA/DEV	The Complete SSIS Migration Playbook	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

Break-Out	Shweta Gupta	DS/DEV	AI Engineering - Productionize your AI models in Enterprise Grade Business Application	Abstract: AI Engineering constitutes deploying models built by data scientists at scale, data pipelines that could assist data scientists in initial model development, automated re-training etc., and integrating with other systems or other components within the AI driven systems. These deployments also need to align with core principles of Data Privacy, Security, Governance & Compliance, which are typically well -defined an in enterprise. This session will attempt to be a practical guide to designing, building and maintaining AI driven systems with data privacy, security, governance and compliance integral to the entire process end to end. Key Learning: AI Engineering - the various steps involved in building ML models for productionizing, key technology choices and platforms available to implement. Demos: We will take an existing deep learning model, integrate it with a data pipeline using ML Ops, which will enable automated deployments, re-training and integrating the model with an enterprise application. The demo will deploy this model using ML Ops tools, and also do a code walk through where helpful.	Intermediate
Break-Out	Sivakumar Harinath	BIA/DEV	Designing for high performance in Power BI	Abstract: It's so easy to get started with Microsoft Power BI Desktop and get immediate value. Power BI is a feature-rich platform with many advanced capabilities. However, as solution complexity grows, with this great power comes great responsibility. It's important to carefully consider how your datasets and reports are going to be deployed and consumed so that they perform at scale and save you from future pains. In this session we address common pitfalls and cover top tips to get great performance across the entire platform - from sourcing data to designing reports and dashboards. Bonus material - special considerations for Power BI Premium!Key Learning: Understand what features of Power BI based on business needs and performance implications.Demos: Most of the content will be presentation and guidance.	Basic

Break-Out	Sivakumar Harinath	BIA/ARCH	Microsoft PowerApps and Microsoft Flow: Learn how to create AI Builder models and use them across the Power Platform	Abstract: Learn about the functionalities and best practices to achieve your digital transformation through the Power Platform and AI Builder. AI Builder provides AI templates (Binary Classification, Text Classification, Object Detection, Form Scanning) that you can tailor to your business processes and use in PowerApps or Flows.Key Learning: Understand capabilities and example scenarios of how AI builder can help in digital transformation.Demos: Demo of AI Builder	Basic
Break-Out	Sivakumar Harinath	DS/BIA	AI powered analytics in Power BI	Abstract: Come learn how to use AI in Microsoft Power BI to solve your business problems. There are too many business problems and too few data scientists to solve - so come learn the AI capabilities in Power BI that empower analysts to solve problems that were once reserved for dedicated data scientists. We cover how you can enrich your Power BI reports and dashboards with key driver analysis (what are the things that drive your KPIs), regression and classification analysis, forecasting, as well as text analytics such as sentiment analysis and key phrase extraction. Sound difficult? It's not!Key Learning: Understanding AI capabilities in Power BIDemos: Various demos about AI features in Power BI that analysts can leverage.	Intermediate
Open-Talk	Sunil Agarwal	OS/ARCH	Why PostgreSQL is the OSS DB of choice for migrating workload from Oracle	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	Sunil Agarwal	OS/ARCH	Successful workload patterns on Azure managed PostgreSQL/MySQL/MariaDB and its seamless migration experience	Abstract: Azure OSS databases have been deployed in multiple verticals such as banking/financials, IOTs with stream analytics, multi-tenancy applications. All these applications can then leverage rich Analytics capabilities natively provided in Azure. Many of these Key Learning: After this session, the attendee will learn (1) variety of production workloads that have successfully migrated to Azure OSS databases. These standard patterns can be used as a blueprint to migrate similar workloads (2) Performance of Scale of the production workloads running in Azure (3) How Azure Database Migration Service simplifies migration to azure with minimal downtime These learnings target building confidence of the attendees for Azure Open Source Database platform Demos: We will demo (1) online migration to Orcas Database from on-premise leveraging DMS tool (2) Demo representing customer production scenario	Intermediate

Break-Out	Sunil Agarwal	OS/DBA	Getting Best performance on Azure Open Source databases	Abstract: Are you interested in getting the best performance when running Open Source Database (e.g. MySQL, PostgreSQL or MariaDB) Managed Service on Azure? This session will describe the key performance impacting factors, intelligent insights and the best practices we have learnt by working with many customer workloads. We will show case some of the common performance pitfalls many customers have stumbled upon and how to address them. Key Learning: After this session, the attendee will learn (1) How OSS databases are deployed in the Azure and the common factors that can impact the performance (2) Common pitfalls when migrating to Azure OSS database services (3) Best performance practices Demos: Demos (1) Demo PgBench workload running on a IaaS VM and running the same benchmark on Azure managed service on a comparable configuration to demonstrate resource bottlenecks and the importance of connection pooling an (2) Offloading query workload to read replica (3) Leveraging Querystore to troubleshoot performance problems	Advanced
Break-Out	Sunil Kamath	OS/ARCH	Overview of Azure managed PostgreSQL/MySQL/MariaDB offering and the competitive advantages	Abstract: This session will cover why Azure OSS database platform is a compelling destination to create or migrate your MySQL/PG/MariaDB workload from on-premise or competitive cloud. With built-in HA, workload scale-out with hyperscale (Citus), security and intelligent performance at lower TCO, many enterprises have chosen Azure as the preferred destination for OSS. We will look at the competitive advantage and enterprise readiness features that collectively host thousands of successful customer workloads Key Learning: After this session, the attendee would learn (1) Open Source Database Offerings in Azure (2) Basics of creating OSS Database Service and differenced capabilities (3)Competitive advantage running OSS databases in Azure Demos: Demo (1) Provisioning PostgreSQL Database Service to show case basic capabilities (2) Changing storage/CPU dynamically (3) Read Replica	Intermediate
Open-Talk	Sunil Kamath	OS/ARCH	Why should I choose Microsoft Azure over AWS for Open Source Databases	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA

Break-Out	Sunil Kamath	OS/ARCH	PostgreSQL at any scale using Hyperscale (Citus)	Abstract: Is your application performance slowing down as the size of your database grows? Are you concerned about hitting the scaling limits as workload increases? If these concerns you, Hyperscale (Citus), a new offering for Azure Managed Database Service for PostgreSQL is our answer. Hyperscale allows you to horizontally scale your database across multiple nodes with no application re-write. It is packaged as an extension of PostgreSQL and is always compatible with latest Postgres versions and the broad ecosystem of Postgres tools. This session will provide an architectural overview, targeted scenarios, performance and best practices for running workload on Hyperscale (Citus). Key Learning: After this session, the attendee will learn (1) Hyperscale (citus) architecture and its deployment on Azure (2) How horizontal scaling delivers industry leading performance and scale (3) Workloads that are targeted by Hyperscale (Citus) Demos: Demos (1) Setting up Hyperscale in Azure (2) Side by side performance comparison single PG vs Hyperscale (3) Show case a high scale workload running on Hyperscale (Citus)	Intermediate
Break-Out	Tejas Shah	DBA/ARCH	What's new for SQL Server on Linux and containers in SQL Server 2019	Abstract: Learn about new and exciting features coming for SQL Server on Linux and containers in SQL Server 2019. Key Learning: SQL Server on Linux is a fully compatible and very capable alternative to other database platforms on Linux and containers. Demos: In session. (Using replication, DTC, HA in containers, deployment to Kubernetes/OpenShift etc.)	Intermediate
Break-Out	Tejas Shah	DBA/DEV	Intelligent Query Processor and Security improvements in SQL Server 2019	Abstract: Learn about Intelligent Query Processor and Security features coming soon in SQL Server 2019 (and Azure SQL DB). Key Learning: SQL Server 2019 and Azure SQL DB have world-class intelligent Query processor to improve performance and security improvements to keep your data safe and secure, no matter which platform you are on. Demos: In session. (Using Intelligent QP and security demos).	Intermediate

Break-Out	Tulasi Menon	DS/DEV	Create Scalable, Entreprise-ready AI solutions using Azure Bot and Cognitive services	Abstract: Learn how to create scalable AI solutions using the power of Azure cognitive services and Azure bot service. Understand the spectrum of available Ai services and the technology it is grounded in. See real enterprise ready examples and understand how to create a similar solution for your needs. Walk through the breadth of Azure Cognitive services and bot services. then dive deep into a real use case- by creating an e2e bot for a specific use case, and seeing how to combine all the pieces to create a great enterprise ready solution. Key Learning: 1. Understand the breadth of AI pre-built & custom services available and their readiness for enterprise 2. Identify possible use cases and benefits of using these AI services 3. Learn how to build an e2e solution for a specific use case 4. Understand how to dive deeper and further customize Demos: 1. Creating an e2e bot with bot service and adding cognitive services to it like LUS, QnA Maker, text, Speech..	Basic
Break-Out	Vijay Reddy	DBA/DEV	Partitioned Tables and Indexes	Abstract: Table and Index partitioning is the one of underestimated feature in SQL server. In this session, we will understand Partition concepts and components and the benefits that we achieve by partitioning data and indexes. Key Learning: We will learn how and why to partition table and indexes and also, effective ways to manage partitions. Demos: https://www.youtube.com/watch?v=ZVAZZqFSsgU&t=1462s	Advanced

Break-Out	William Durkin	DBA/DEV	Automation for the DBA: Embrace your inner sloth	Abstract: DBAs and sysadmins never have time for the fun stuff. We are always restoring a DB for a dev or setting up a new instance for that new BI project. What if I told you that you can make all that time consuming busy-work disappear? In this session we will learn to embrace the power of automation to allow us to sit back and relax..... or rather focus on the real work of designing better, faster systems instead of fighting for short time slots when we can do actual work. Along the way we will see that we can benefit from the wide world of automation expertise already available to us and avoid re-inventing the wheel, again! Key Learning: Don't do boring tasks with the GUI, use PowerShell to speed up your life and let you do "real" work Automation is coming, embrace it or lose your professional advantage PowerShell is the future of a DBA's career Demos: Demonstrate automation of "boring" DBA tasks Restore a database Copy a login from one server to another Run a full instance migration with one line of code!	Basic
Open-Talk	William Durkin	DEV/DBA	Data Platform Migrations - Experiences and Thoughts	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	NA
Break-Out	William Durkin	DBA/DEV	It's broken, now what?! (Practical problem solving)	Abstract: We've all experienced weird situations in IT - things break without any real apparent reason. Sometimes, error messages can be helpful, but mostly they are cryptic and lead to no real explanations/solutions. In this session, I will show a variety of problems that I have run into in the past and explain how I approached them. Sometimes finding simple solutions, but sometimes having to be creative and employ methods that may not be so intuitive. You will leave the session with a better understanding on how to approach solving any technical issues you experience at work. Key Learning: Re-consider how to approach solving an issue. Dont give up at the first hurdle, most errors are "easy" to troubleshoot. A clear approach to solving issues will make your life easier Demos: This is a story telling session. I will take real life problem situations and re-tell them and how I approached solving them. Audience interaction is expected too, to get them thinking about how to solve a problem	Basic