Data Platform Summit 2018 - Session List (Final Release)

Final Release. Release Date: July 19, 2018. Session Count: 104

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 2018 is Asia's Largest Data & Analytics Conference on Microsoft Data Platform and Open Source. DPS 2018 brings you the finest content from world's best trainers.
4. DPS Summit & Pre-Con pass price increases on 1st of every month. Block your seat today.

To view, up-to-date sessions list ---------->> http://www.dataplatformgeeks.com/dps2018/break-out-sessions/

Session_Type	Speaker	Track(s)	Session Title	Abstract	Level
Break-Out	Abhishek Narain	BIA/BD	Building your Big Data and Advanced Analytics Pipeline on Azure using Azure Data Factory	orchestration capabilities and how it could meet the ETL needs for Big Data and Advanced Analytics projects. We will use User Interface to create	Intermediate
Open-Talk	Abhishek Narain	BIA/DEV	Azure Data Factory – Customer stories and Roadmap	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Ajay Jagannathan	DBA/ARCH	Deep Dive into Azure SQL Database Managed Instance	Azure SQL Database, providing near 100% compatibility with SQL Server on-premises (Enterprise Edition). Managed Instance allows you to fast	Advanced
Break-Out	Ajay Jagannathan	ARCH/DBA	Azure SQL Database – the intelligent cloud database on autopilot that lets you focus on your business	service offering based on the world's best relational database management system, SQL Server. It consists of three deployment options -	Advanced
Break-Out	Ajay Jagannathan / Sudhakar Sannakkayala	ARCH/OSS	Relational Database Platform (Azure SQL Database, PostgreSQL, MySQL	database-as-a-service offering solves the demands of today's data estate involving omnipresence, heterogeneous and explosion. Built on world's	Intermediate
Break-Out	Allen O'Neill	DBA/DS	Add value & improve data quality with Data Science & Machine Learning (no math!)	attention to until it comes and bites us, and when it does, its usually a customer that notices it. As always, the poor beleaguered developer and	Intermediate
Chalk-Talk	Allen O'Neill	DEV/ARCH	How to build a robot and the data implications	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Amit Bansal	DEV/DBA	Intermediate - SQL Server Developer Tricks	your query is being treated inside the engine. Join this session to learn some real-world developer tricks that will improve performance of your	Intermediate
Break-Out	Anand Raman	DS/DEV	AI is the New Normal	Widespread use of mobile devices and powerful personal computing have driven a major shift among organizations of all types to adopt Artificial	Intermediate
Break-Out	Andy Leonard	BIA/DEV	Faster SSIS	Leonard as he runs test loads using sample and real-world data and shows you how to tune SQL Server 2016 Integration Services (SSIS 2016)	Intermediate
Break-Out	Andy Leonard	ARCH/BIA	Designing an SSIS Framework	consultant Andy Leonard explains the what, why, and how of a custom SSIS framework that delivers metadata-driven package execution.	Advanced
Open-Talk	Anupama Natarajan	DS/DEV	How Artificial Intelligence (AI) will transform modern workspace?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Anupama Natarajan	DS/DEV	Creating your first intelligent Bot	Why there is so much talk about bots? How do they work?	Basic
Break-Out	Anupama Natarajan	BIA/DEV	Deep Dive into SSRS2017 REST API	objects in a SQL Server 2017 Reporting Services report server catalog. The REST API exposes endpoints to navigate the folder hierarchy, discover the	Intermediate
Break-Out	Arun Khetarpal / Kapil Raja	DEV/ARCH	How to take advantage of scale out graph in Azure Cosmos DB	graph database applications on Azure Cosmos DB and explore the different solutions that it provides to common data scenarios in the	Basic

Break-Out	Avanish Panchal	DBA/ARCH	Azure ARM & Azure PowerShell for SQL Server Automation	Automation plays a key role in building a sustainable and repeatable framework for creating and managing SQL Server in Azure. Automation is	Intermediate
Break-Out	Ben Weissman	BIA/DBA	The Self-Tuning SSIS Package	parallel in SSIS. Regardless of whether your limit is on the source or destination side, you will eventually reach those limits.	Basic
Break-Out	Ben Weissman	BIA/ARCH	ETL and DWH Design with Metadata	Instead of reinventing the wheel every time you need to change or extend a package, let's talk about metadata models and how we can use them to	Intermediate
Chalk-Talk	Ben Weissman	BIA/DEV	What's in an ensemble? An intro to Data Vault	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Casey Karst	BD/DEV	Azure SQL Data Warehouse and Azure Databricks: Integration for the Modern Data Warehouse	that your implementation is future proof. In this session, you will learn how customers are using Azure SQL Data Warehouse and Azure Databricks	Intermediate
Chalk-Talk	Damian Widera	DBA/BIA	Row Level Security now and in the past - my working solutions	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Damian Widera	DBA/DEV	Travelling in time with SQL Server 2016	Temporal tables. You can make queries to historical data lot easier by using this feature. The mechanism is very simple however you all should	Intermediate
Break-Out	Damian Widera	BIA/DEV	U-SQL in great depth	Lake and the USQL. I would like to show how quickly you can do data analysis using traditional C# and a new language that is a bit similar to the	Intermediate
Break-Out	Denny Cherry	DBA/ARCH	Storage for the DBA	around storage. It's also one of the hardest places to troubleshoot performance issues because storage engineers and database	Intermediate
Break-Out	Denny Cherry	DBA/ARCH	How to Maintain the Same Level of utilities in Cloud Deployments - Securability, Reliability and Scalability".	Microsoft SQL Server in Microsoft Azure and on-premises from a Security, Reliability and Scalability perspective.	Intermediate
Break-Out	Devashish Salgaonkar	DBA/ARCH	Deploy and Manage SQL Server Configuration as a CODE	scale can be a daunting task even for the most experienced of DBA's. In today's cloud world, where speed and agility is a key driver, less manual	Intermediate
Open-Talk	Dinesh Priyankara	BIA/ARCH	Why should we focus more on logical data warehouse than physical data warehouse?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Dinesh Priyankara	BD/DEV	Making unstructured data analysis-ready using Data Lake Analytics	ignorable, it is becoming a part of almost all data-oriented solutions. Therefore, we just cannot ignore them and continue our journey, we	Intermediate
Chalk-Talk	Dmitri Korotkevitch	DBA/DEV	Tips and tricks for successful In-Memory OLTP implementation	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Dmitri Korotkevitch	DBA/DEV	Data Partitioning for Database Architects and Mere Mortals	simplify database administration and improve system availability and performance. Contrary to popular believe, it is not limited to partitioned	Intermediate
Break-Out	Dmitri Korotkevitch	DBA/DEV	Deep Dive into Blocking and Deadlocks Troubleshooting	and least understood part of SQL Server Internals. Blocking issues and deadlocks occur unexpectedly and negatively impact performance and	Intermediate
Break-Out	Edwin Sarmiento	DBA/DEV	Leveraging Microsoft PowerShell for Managing SQL Server VMs on Amazon AWS	if you simply rely on the graphical user interface. Imagine having to manage multiple SQL Server VMs as part of your day-to-day operations.	Intermediate
Break-Out	Edwin Sarmiento	DBA/OSS	Getting Started with Linux for the SQL Server DBA	escaping the fact that SQL Server DBAs need to be familiar with the Linux operating system. So, how do you start?	Intermediate
Open-Talk	Edwin Sarmiento	CG/DBA	Transitioning from a DBA to a Manager	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Erika Menezes	DS/DEV	Deep Learning for data geeks	around the use of artificial neural networks. Deep learning is being used in self-driving cars, medical diagnosis and	Basic
Chalk-Talk	Erika Menezes	DS/DEV	Babysitting for machine learning models – how do you know it's not working	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	
Break-Out	Erika Menezes	DS/DEV	Advanced - AI for Fashion	developed for one task is reused for other tasks and saves users a lot of time by not having to training deep learning models from scratch.	Advanced

Break-Out	Hamish Watson	DEV/ARCH	Building a Database DevOps Pipeline in the Cloud under 59 minutes	Continuous Delivery.	Basic
Chalk-Talk	Hamish Watson	DEV/DBA	Ways to get your database into Source Control for DevOps deploys	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Hamish Watson	DEV/DBA	Test Driven Development in SQL Server – how to deploy database code safer	enabled application developers to write cleaner code.	Intermediate
Break-Out	Harsh Raj Thakur / Rahulinder Singh	BIA/DEV	Evolution of the Enterprise Data Platform: Leveraging Azure to get the most out of your data	enterprise data platforms? From SQL Server to Azure Data Lake to Machine Learning, the storage and compute options have evolved over	Intermediate
Break-Out	Ilyas F	DS/ARCH	Moving Intelligence to IoT Edge	been here for many decades. Edge computing has a significant role in the IoT systems, the connected devices has a massive potential to generate	Basic
Break-Out	Janakiram MSV	DS/DEV	Machine Learning is not Magic - Getting Started with Azure ML	Intelligence can get overwhelming for developers. The sheer number of ML frameworks, tools, and more importantly the prerequisite of maths	Intermediate
Chalk-Talk	Janakiram MSV	DS/DEV	What is Edge Computing?	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	
Chalk-Talk	Joe Yong	BIA/DEV	How we designed Azure SQL DW to support PB-scale data warehouses	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Joe Yong	BD/ARCH	Azure SQL Data Warehouse: your PB scale cloud data warehouse and big data platform	scale cloud data warehouse that can be deployed in minutes. This session will walk through key architecture and design concepts that enable these	Basic
Break-Out	Joe Yong	BD/DBA	Azure SQL Data Warehouse: best practices and lessons learned	your enterprise data warehouse on it for some time, you may have spent time and effort trying to figure out how to load Terabytes of data per hour	Advanced
Break-Out	Joey Dántoni	ARCH/DBA	Introduction to Azure Infrastructure	components that make up the Microsoft Azure platform. When it comes to moving SQL Server systems into the Azure platform having a solid	Basic
Break-Out	Joey Dántoni	DBA/DEV	Cloud Application Development & Deployment	infrastructure, especially amongst developers and CI/CD practitioners. However, in the database space, container adoption has been lower. SQL	Intermediate
Open-Talk	Kane Conway / Khushboo Gupta	BIA/DEV	Analysis Services Troubleshooting Methodology	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Leila Etaati	DS/ARCH	End To End Data Science Solution: Azure ML Workbench	and advanced analytics solution. It enables data scientists to prepare data, develop experiments, and deploy models at cloud scale. In this session,	Intermediate
Chalk-Talk	Leila Etaati	DS/DEV	What Microsoft AI tools for what use case	architectural understanding, that too with only whiteboard and marker. No Laptops, no PPTs, no demos – only whiteboard-ing!	N/A
Break-Out	Leila Etaati	DS/DEV	Machine Learning Revolution with Azure Databricks	the Microsoft Azure cloud services platform and easy to set up with one-click. It has an interactive workspace that enables collaboration between	Intermediate
Break-Out	Luan Moreno	BD/DEV	Big Data on Azure: Processing, Wrangling & Analyzing Datasets using Scalable Programming Languages	enriching their capabilities even more, now there are a set of scalable languages that you can use in the Microsoft Azure environment to tackle	Basic
Chalk-Talk	Luan Moreno	DBA/DEV	Big Data Solutions on Azure, AWS & GCP in a Nutshell	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!	N/A

Break-Out	Luan Moreno	BD/ARCH	Use Cases & Best Practices of Big Data on Microsoft Azure	Abstract: The world is changing rapidly and the big data ecosystem is enriching their capabilities even more, now there are a set of scalable languages that you can use in the Microsoft Azure environment to tackle and solve different problems, in this training you will learn the best practices, common scenarios and use cases for Hive, Pig, PolyBase, Scala, Python, Spark-SQL and U-SQL, in the end of the day you will be able to understand the best fit for each one of those scalable languages and how this new way of work with data can enrich your way of thinking. Key Learning: 1. This course will explain Big Data Features from the very beginning with a short overview and intro about those. 2. At this course, we will touch base each one of the listed areas in a basic to an intermediate level. 3. The intention of this training is to show to DBA's / Developers about new ways to threat and work with data, we will explain different models and architectures in order to clarify how large companies work with data nowadays. Demos: Hive, Pig, PolyBase, Scala, PySpark, Spark-SQL, U-SQ	Basic
Break-Out	Madhan Gajendran	ARCH/DEV	Technical Overview of Azure Cosmos DB	Abstract: In this technical overview of Azure Cosmos DB you will learn how easy it is to get started building planet-scale applications with Azure Cosmos DB and benefit from the platform's turn key global distribution, guaranteed low latency access and elastic scale. We'll then take a closer look at important design aspects around global distribution, consistency, and server-side partitioning. How to model your data to fit your app's needs using tools and APIs you love.	Basic
Chalk-Talk	Manish Sharma	NoSQL/ARCH	How to do infinite Scale in Azure CosmosDB?	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!	N/A

Break-Out	Manish Sharma	ARCH/DEV	How to migrate your existing MongoDB/ SQL to Azure Cosmos DB	Abstract: Bring your Mongo DB applications to Azure Cosmos DB and benefit turnkey global distribution, guaranteed low latency for cloud scale. Learn how easy it is to migrate your existing NoSQL applications to Azure Cosmos DB by using the MongoDB API. In this sessions you will find out about the tools used in the process, see the code that leverages Azure Cosmos DB, and learn about techniques for working around the differences between the products.	Basic
Break-Out	Manohar Punna	BIA/DEV	Real-Time Analytics with Power BI	Abstract: Power BI helps you build appealing visualizations of your data. With ever growing footprint of data it is equally important to get real-time insights into your data. It can be as simple as monitoring a single metric or viewing real-time sales performance across multiple locations. Power BI real-time streaming enables you to stream data and update dashboards in real-time. Any time sensitive data can be a source of streaming data set like IoT sensor devices, social media sources, service usage metrics etc. In this session 1. I will dive deep into real-time analytics with Power BI. 2. How to build streaming tiles in Power BI Dashboards. 3. Demo using Power BI REST APIs. 4. Demo using Azure Event Hub and Azure Stream Analytics. Key Learning: Audience will learn and explore different approaches to create live tiles in Power BI Demos: 1. Demo using Power BI REST APIs. 2. Demo using Azure Event Hub and Azure Stream Analytics.	Intermediate
Break-Out	Nagaraj Venkatesan	DBA/DEV	Adaptive Query Processing and Automatic Tuning	Abstract: Automatic Tuning and Adaptive Query Processing are perhaps the two key features of SQL Server 2017 and Azure SQL Database which would influence query optimizer's plan choices. The following session will provide an overview of the features and explain how it can help to address performance issues like parameter sniffing, plan regressions, lack of indexes or too many indexes without any manual intervention and change of code. Key Learning: Understanding on Adaptive Query Processing scenarios -> Adaptive Joins, Memory Grant Feedback, Interleaved execution Understanding on Automatic Index Correction, Automatic Index Tuning Demos: Demo on Adaptive Joins, Memory Grant Feedback, Interleaved execution,Automatic Index Correction	Basic

Break-Out	Nagaraj Venkatesan	DBA/ARCH	"Stretching" using Polybase	Abstract: To archive older data and to seamlessly query it, SQL Server 2017 offers "Stretch" database as a solution. However, "Stretch" database is an expensive solution which not many can afford. Can "Polybase", the feature to query external data can be used as a cost-effective alternative for archival? Join this session to find out how "Polybase" can be used for archival and what are the pros and cons between "Stretch" Database and "Polybase" Key Learning: Overview on "Polybase" Understand how to archive using "Polybase" Compare Stretch database against "Polybase" Understanding the cost of both the solutions Demos: * Polybase configuration * Archival using Polybase * Using Stretch SQL Database to query archived data	Intermediate
Break-Out	Nikhil Gaekwad	ARCH/BIA	Distributing content with Power BI	Abstract: Effective ways of distributing content in Power BI using Apps, Mobile, Embed, and SharePoint. Key Learning: Learn about the different tools in Power BI that allow end users to consume content. Focused on Mobile, Embed, Apps, etc. Demos: Power BI Mobile, SharePoint Embed.	Basic
Break-Out	Nikhil Gaekwad	BIA/ARCH	Enterprise BI Deployments and Governance with Power BI	Abstract: Whether you're planning an enterprise-wide reporting deployment or providing structure to self-service BI activities within teams, Power BI has you covered. Learn about tools for developing, publishing, and managing your BI assets. This session will cover the data gateway, managing report lifecycle, publishing options, administration and governance contols, and end-user capabilities across devices and platforms. Key Learning: Managing and deploying Power BI within an organization. Demos: Power BI service and administration tools.	Intermediate
Open-Talk	Nikhil Gaekwad	BIA/DEV	Empowering end users in Power BI – getting the most of out your data	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Nilabja Ball	DS/DEV	Business Analytics with Microsoft R	Abstract: SQL Server Machine Learning Services has the power of proprietary packages to deliver advanced analytics at scale, and the ability to bring calculations and processing to where the data resides. In this session you will learn how to install and configure of Machine Learning Server and R language concepts. It also covers Machine Learning features in SQL Server such as Real Time scoring, in-database processing etc and how to run R scripts from Power BI to data visualization.	Intermediate
Open-Talk	Patrick Flynn	DBA/ARCH	Backup and Restore – Test your knowledge	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A

Break-Out	Patrick Flynn	DBA/DEV	Supercharge your Database Management with DBATools	Abstract: Need to constantly scale your work across more and more Servers? Required to monitor and maintain an increasing array of database instances and versions? Struggling with being able to accurately measure and provide the metrics and reports required by Management? In this session we will investigate how to supercharge the management of your Database Environment via the use of the open source module dbatools.. Used and endorsed many of the Worlds top DBAs , the dbatools collection of 400+ functions can be used to supercharge management of any sized environment and give you back time to concentrate on delivering true Business value! Key Learning: How to use PowerShell and the dbatools module to automate and scale management of Database environment from 1 to 100's of Servers Demos: Following Scenarios will be demonstrated (a) Migration of entire SQL Server instance to new Server and Instance with single line of PowerShell (b) Automated testing of Backups by restoring to new instance and running DBCC CheckDB (c) Validate of Database environment using	Intermediate

Break-Out	Patrick Flynn	DBA/DBA	Database Corruption - Advanced Recovery Techniques	Abstract: Being able to monitor, diagnose and recover from Database Corruption is a critical skill for any SQL Server DBA. In this session we will walk through the techniques required to detect and repair various forms of Corruption Using a number of example corrupted database we will explore how to determine and fix corruption while avoiding some of the common mistakes. This is 400+ level session that uses knowledge of the internal storage structures of SQL Server to Key Learning: Demonstrate Techniques to detect and recover data from damaged or corrupted databases Demonstrate how knowledge of SQL Server internals and physical storage structures can be used to perform data recovery from Damaged databases Demos: Session is based around multiple demos showing (a) Impact of Page Checksums on Backup and Restore Integrity (b) Use of Advanced restore options to recover with minimal data loss or downtime (c) Use of DBSS Page and DBCC Ind to restore data directly from Damaged Pages (d) Use of fn_dblog and log file internals to recover data directly from log file	Advanced

Break-Out	Peter Myers	BIA/ARCH	The A-Z of Power BI Dashboards	Abstract: Power BI dashboards enable the display of a consolidated view across the organization, regardless of where the data is stored. A dashboard consists of tiles, with each displaying a value or data visualization. At last count, there were 10 different ways to add a tile to a dashboard. In this session, everything you need to know about Power BI dashboards will be covered. This includes describing and demonstrating each of techniques used to add tiles, and how to ensure that dashboard data remains current. In addition, techniques to share and integrate dashboards, and to deliver real-time dashboard tiles will be introduced. Key Learning: In this session, you will learn: • How to create dashboards • How to add dashboard tiles • How to effectively share dashboards • Tips and tricks to produce and manage great dashboards Demos: A demonstration of each of the 10 “add tile” techniques.	Basic
Open-Talk	Peter Myers	BIA/ARCH	How Best to Share Power BI Content	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Peter Myers	BIA/DEV	Power BI Embedded Analytics	Abstract: Power BI supports embedding content into your apps. Content can consist of reports, specific report visuals, dashboards, specific dashboard tiles, and the Q&A experience. In this session, you will learn about the potential to embed Power BI analytics, and see a series of demonstrations that produce a single solution. This session is targeted at experienced application developers. Key Learning: In this session, you will learn: • How to embed Power BI content • How to leverage features supported by the Power BI JavaScript API • How to apply row-level security (RLS) • About Power BI licensing options for Power BI embedding Demos: A series of demonstrations will show embedding, implementing JavaScript API functionality, and securing data.	Intermediate
Break-Out	Philip Seamark	BIA/DEV	Practical DAX for Power BI	Abstract: A walkthrough showcasing a wide variety of practical tips and tricks using DAX with Power BI. This will cover data-generation, summarisation, pivots, complex joins, calculation optimisation and much more. Key Learning: Practical tips and tricks by well known author on how to improve your calculations Demos: Not sure what to add here	Intermediate

Break-Out	Philip Seamark	BIA/DEV	Data modelling for Power BI using brand new Analysis Services features	Abstract: A session covering some of the new data modelling features for Power BI . These will likely be public come the summit, so a session to demo and cover these in more detail. Features such as incremental refresh, aggregations and more. Key Learning: Do see how some of the fantastic new features look Demos: Not sure what to put here	Advanced
Open-Talk	Praveen Srivatsa	DS/DEV	Monetizing AI Services	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Praveen Srivatsa	DS/DEV	Machine Learning in Medical Rehabilitation	Abstract: One of the key areas that machine learning can make a big impact is in the area of medical rehabilitation. In this session, we will look at the complete life-cycle of a medical rehabilitation platform and how it uses IoT data, VisionAPI and sentiment analysis to provide a comprehensive patient view to the doctors. Key Learning: In this session, the audience will learn how technologies like IoT, VisionAPI and sentiment analysis can be put together to analyze a patient's rehabilitation in the real world.	Intermediate
Chalk-Talk	Raj Pochiraju / Mukesh Kumar	DBA/DEV	Deep dive to Azure Database Migrations (DMS) service architecture	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!	N/A

Break-Out	Raj Pochiraju / Nikhil Patel	DBA/ARCH	Migrating to Azure: Moving from on-premises SQL Server and Oracle databases to Azure	Abstract: Azure SQL Database and Managed Instances allows you 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 and show you how we have made it easy to migrate SQL Server and Oracle instances to Azure by using the Azure Database Migration Service and related tools. We will also cover most commonly seen migration blocking scenarios and demonstrate how our service can unblock your migration to Azure SQL Databases. We will give you a deep dive how to perform scale migrations using our CLI components. After completing the session, you will understand how easy it is to migrate from SQL Server and Oracle to Azure database platforms when empowered with the right tools, services, and best practices. We will also showcase lift and shift migrations from open source databases MySQL and PostgreSQL to our Azure Database platform. Key Learning: 1. You get proficient with our new DMS service, use in your upcoming upgrade/migration projects. 2. You will be able to pinpoint breaking and behavior changes and deprecated features upgrading to Azure SQL Database. 3. Execute the end-to-end database migration seamlessly, includes schema and data Demos: 1. Migrating databases from SQL Server and Oracle to Azure SQL Databases and Managed Instances using our DMS service 2. Migrating databases from Oracle to Azure SQL Databases. 3. Perform minimal downtime migrations to Azure using DMS service. 4. Perform minimal downtime migrations from MySQL to Azure Database for MySQL using DMS 5. Perform scale migrations using our DMS CLIs.	Advanced
Break-Out	Reza Rad	BIA/ARCH	Power BI On Premises - Report Server	Abstract: Come to this session to learn all about using Power BI in an on-premises solution named Power BI report server. you will learn about all limitations and advantages of Power BI report server through many live demos. Key Learning: Using Power BI on-premises Demos: many live demos	Intermediate
Open-Talk	Reza Rad	BIA/DEV	Starting Power BI? Come and have a chat	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A

Break-Out	Reza Rad	BIA/DEV	M and Power Query Beyond Limits	Abstract: Power Query is the transformation engine in Power BI. This is the engine that you do all data preparation before loading data into the model. This is the keystone of your Power BI solution. In this session, you will learn what things Power Query can do. M is the language behind the scene of Power Query, you will learn how M can be more beneficial than the Power Query Graphical interface. You will see demos of transformations that you can do with M Code. This session is full of live demos. Prepare to be amazed with what Power Query can do with heaps of demos in this session. Key Learning: Power Query Power BI Data Preparation with Power BI M Demos: heaps of live demos through the session	Advanced
Break-Out	Sandeep Alur	DS/DEV	Deep Neural Networks – Let us build a Image/Currency Classifier	Abstract: Data comes in various shapes and sizes, and so are expected outcomes from a business standpoint. Given the fact that 'Artificial Intelligence' has stood up to solving complex industry problems, computer vision(image) is one area where significant innovation is underway. This is where Deep Learning architectures such as Deep Neural Networks step in and give us techniques to build a model targeting a use case under consideration. Join this session to understand the basic building blocks of a Neural network and peek into building a model that helps one classify images/currency. Quite a few techniques are involved in the process and we will delve into technicalities of building such a solution. Critical component of this journey is the tool set and the compute environment required to run the neural network. You will get acquainted with what we have on Azure in building AI solutions. Key Learning: Below are the key take away for the audience - Understanding of Deep Learning - Peek into an image classifer	Advanced

Break-Out	Sandip Pani	DS/DEV	Build Predictive model using Azure Machine learning Studio	Abstract: If you are new to Machine learning, don;t know where to start, what is the difference between ML and AI, and you want to learn how to start, Then this session is for you. In this session I will explain what is machine learning, What is AI? Life cycle of a machine learning project. How to train a predictive model using Azure Machine learning studio. If you think you need deep understanding of statistics to start your Machine learning project then don't worry I will cover minimum important and fundamental key component of Statistics which is sufficient to build your first predictive model.	Basic
Break-Out	Sandy Winarko	BIA/ARCH	Embrace and Extend: First-Class Activity and 3rd Party Ecosystem for SSIS in ADF	Abstract: This session focuses on the deeper integration of SQL Server Integration Services (SSIS) in Azure Data Factory (ADF) and the broad extensibility of Azure-SSIS Integration Runtime (IR). We will first show you how to provision Azure-SSIS IR – dedicated ADF servers for lifting & shifting SSIS packages – and extend it with custom/3rd party components. 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 trigger/schedule SSIS package executions as first-class activities in ADF pipelines and combine/chain them with other activities, allowing you to inject/splice built-in data transformations in your ETL/ELT workflows, automatically provision Azure-SSIS IR on demand/just in time, etc. And finally, you will learn about the licensing model for ISVs to develop paid components/extensions and join the growing 3rd party ecosystem for SSIS in ADF. Key Learning: Learning how to lift & shift their traditional ETL workloads and create modern ETL/ELT workflows with SSIS in ADF Demos: - Provision Azure-SSIS IR dedicated for SSIS package executions using ADF app, join it to a VNet to enable data access on premises, and customize it by installing additional (un)licensed components/extensions - Using SSMS, deploy existing SSIS projects/packages into a catalog (SSISDB) that is hosted by Azure SQL Database/Managed Instance and attached to Azure-SSIS IR - Trigger/schedule SSIS package executions on Azure-SSIS IR as first-class activities in ADF pipelines using SSMS/ADF app - Chain/combine SSIS activities with other activities in ADF pipelines to create modern ETL/ELT workflows	Intermediate
Chalk-Talk	Sourabh Agarwal	DEV/DBA	Intelligent Database - Database Intelligence built In	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!	N/A

Break-Out	Sourabh Agarwal	DBA/DEV	Troubleshooting Always On Availability Group Failovers	Abstract: Unexpected failovers or Unsuccessful failovers are by far the most common problem surrounding Always On Availability Groups. Troubleshooting failover related issues can be a challenge and time consuming. In order to effectively troubleshooting these issues data needs to be collected and analyzed from different servers and replicas. Things get further complicated by the fact that there are multiple logs which needs to be looked at and that these logs are in different time zone formats. In this session we will look at how to easily identify the reasons for failover of an Availability Group and what new capabilities are being introduced to help with the analysis of AG failovers. Key Learning: Tools and Techniques to troubleshoot Always On Availability Groups. Demos: Tips to troubleshoot Always On Availability Groups.	Advanced
Break-Out	Sourabh Agarwal / Tejas Shah	DBA/ARCH	SQL Server vNext - Whats New	Abstract: In this session we will talk about the new capabilities which will be introduced in SQL Server vNext in the High Availability and Disaster Recovery space. Get a sneak peak into how HA/DR works when SQL Server in running in containers, in VMs or on physical nodes on either Windows or Linux. Key Learning: New HA and DR functionalities coming to SQL Server VNext. Design for a Highly Availabille SQL Server running in either containers, VM's or physical servers Demos: In this session we will demo 1) Setting up HA for SQL Server running on Containers 2) Setting up HA for SQL Server Running on Linux VM's	Intermediate
Chalk-Talk	Sravani Saluru	DBA/DEV	Temp db contention - Important things to optimize	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!	N/A

Break-Out	Sravani Saluru / Swetha Reddy	DBA/ARCH	Always On Availability Groups on SQL Server on Linux	Abstract: This session is targeted for high availability techniques on Linux for SQL Server. we will demonstrate how to enable and use availability groups for Linux. Manage the failover using pacemaker cluster. Troubleshooting scenarios like automatic failover issues , databases not synchronizing and performance issues. This session walks you through recommendations for configuration settings at pace maker cluster level , availability group resource , SQL Server settings to get the optimum performance . We will also talk about automatic and manual failover of Availability groups between replicas with and without external cluster . The demos will be based on Red hat Linux and SQL Server 2017. Key Learning: Configure and troubleshooting Always On Availability Groups on SQL Server on Linux Demos: Configure availability groups for SQL server on Linux	Intermediate
Break-Out	Steph Locke	DS/DBA	Probability & Statistics 1010	Abstract: We can be better at our jobs if we have a good grasp of basic statistics. It doesn't matter if you're a DBA looking to understand query plan performance, a data warehouse person needing to come up with ETL load time estimates, or an analyst needing to report figures to managers. Statistics can help you all. If only maths classes hadn't been so darn boring! Instead of going all mathsy, we'll be doing some real-time data capture and taking an intuitive and visual approach through summary statistics right up to understanding how to produce simple predictive models. By the end of the session, you'll understand concepts like sampling, error, regression, and outliers – important day-to-day stuff and a great base upon which to build. By the end of the session, you'll wonder how people could have it made seem so hard for so many years. Key Learning: - Understand and correctly utilise descriptive statistics - Understand simple regression concepts - Understand hypothesis tests Demos: - Interactive application with an associated survey that people will fill in realtime	Basic

Break-Out	Steph Locke	DEV/ARCH	Anchor Modelling: Agile databases	Abstract: Anchor Modelling is a fantastic database modelling paradigm that uses sixth normal form (6NF) to store data and provides third normal form (3NF) views for ease of use. This session deep dives into all the concepts behind Anchor Modelling (and indeed databases generally!) and then takes you through how Anchor Modelling uses these concepts to move away from the traditional data warehouse paradigm to deliver a purely additive, agile database. Key Learning: - Understand normalisation - A high-level understanding of anchor models - Able to make architectural decisions about the appropriateness of Anchor models Demos: - Normalising data - Building an anchor model - Adding to an anchor model - Generating deployment code	Intermediate
Break-Out	Sudhir Rawat	NoSQL/ARCH	Azure CosmosDB Underneath	Abstract: Azure Cosmos DB is Microsoft's proprietary globally-distributed, multi-model database service "for managing data at planet-scale". It's a fastest growing database in NoSQL world. This session will focus on some of the key areas of features provided by this database like partition, Consistency, globally distribution, throughput etc. This session will help you understand how to design database in CosmosDB, setting right knobs for designing application friendly database and how to reduce the cost but maintain good database standard. . Key Learning: > Design optimized DB > Understanding advance concept > Tips and tricks Demos: The demos will be around the proposed topics.	Intermediate

Break-Out	Sunil Agarwal	DBA/ARCH	Strategies to certify your SQL Server application in the Cloud era	Abstract: Application certification requires extensive testing for functional correctness, performance and scale that can take few weeks to months. Many IT organizations are struggling in managing legacy applications certified on older versions of SQL Server. Now, with the frequent releases of SQL Server and the advent of Azure SQL Database, many application vendors (ISV) and even custom developed application owners are challenged to minimize the effort required for application certification but still enabling customers to run the application on the latest SQL Server version or in SQL Azure Database. Come to this session to learn a new way to certify your application on SQL Server in a version agnostic way. Key Learning: Database compatibility level has been supported in SQL Server for a long time as a way to simplify database upgrade to new version of SQL Server. We will show how you can leverage DBCompat and the DMA tool to certify your application in SQL Server version agnostic way. Demos: Demo show casing (a) DMA tool on assessing the source database that can be upgraded to SQL Server 2017 (b) a source database that can't be upgraded to SQL 2017 due to functional incompatibilities	Intermediate
Open-Talk	Sunil Agarwal	DEV/IOT	Clustered Columnstore Index: How did I scale dataload for IOT scenario	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Sunil Agarwal	DBA/DEV	Maximizing the Query Performance with Columnstore Indexes in SQL Server	Abstract: Columnstore index can speed up the performance of analytics queries significantly but are you getting the best performance possible? Come to this session on to learn how to diagnose performance issues in queries accessing columnstore index and the steps you can take to troubleshoot. Some of the techniques we discuss here are rowgroup elimination, statistics, partitioning, improving the query plan quality, tweaking the schema, and creating one or more nonclustered btree indexes. Familiarly with columnstore index is required. Key Learning: Most people know that columnstore index can deliver significant query performance boost but many don't realize that they are not getting the optimal performance possible. This session will allow attendees to learn how to diagnose the query performance and take steps to fix it Demos: Demo heavy session. I present 5 common issues that lead to poorly performing queries on columnstore index and then show how you can diagnose and fix it. One demo per query performance scenario	Advanced
Open-Talk	Swetha Reddy	DEV/DBA	Why don't I see SSL certificate listed under Certificates Tab in SQL Server Configuration Manager?	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A

Open-Talk	Tejas Shah	DBA/DEV	Spinlocks, latches, locks – differentiation between them and how to plan troubleshooting them	Open-Talks are 30 minutes free-flowing discussion on a specific topic. No laptops, no PPTs, no demos – only discussion and Q & A	N/A
Break-Out	Tejas Shah	DBA/OSS	The Do's and Don'ts of Running SQL Server On Linux	Abstract: In this session we will cover the learnings from actual customer deployments about running SQL Server on Linux and ensuring that you get optimal performance from your SQL Server Deployments. We will share tips and tricks and demos on how certain settings on both SQL/Linux OS can impact performance. Key Learning: Deploying SQL Server on Linux Best practices for running SQL Server on Linux. Performance impact of OS/SQL properties when running SQL Server on Linux. Demos: Demos would include performance impact of because of incorrect or wrong settings on OS or SQL Server.	Advanced
Break-Out	William Durkin	DBA/DEV	Query Store without SQL 2016 = Open Query Store	Abstract: When SQL Server 2016 was released, it offered a fantastic new feature with the Query Store. Long term, statistics based, query tuning became a reality. But what about the thousands of servers that aren't upgrading to SQL 2016 or newer? The open source project Open Query Store is designed to fulfill that need. This session will give a short introduction to the Query Store feature in SQL 2016 and then dive into the Open Query Store (OQS) solution. William (co-creator of the OQS project) will explain the design of OQS and demonstrate the features. You will leave this session with an understanding of the features of Query Store and Open Query Store, and a desire to implement OQS in your systems when you return to the office. Key Learning: The audience will understand that the Query Store feature is and see how Open Query Store can provide a similar experience for SQL Server instances on versions where "real" Query Store are not available (versions below SQL Server 2016). This is a key administration tool for many thousands of SQL Server instances across the world. Demos: A full demo of how "real" Query Store works: collecting a workload and understanding how Query Store can be used to analyse this A second full demo of how Open Query Store works: collecting the same workload and understanding how Open Query Store can be used to analyse the data	Intermediate
Chalk-Talk	William Durkin	DBA/DEV	DBAtools – making DBA life simple	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!	N/A

Break-Out	William Durkin	DBA/ARCH	SQL Server Replication: What, How, Why	Abstract: Replication is one of the oldest data distribution technologies inside SQL Server (available since version 6.0). The age of replication shows in the management and troubleshooting tools (or lack thereof). In this session, we will dig into some real-world implementations and see how to manage deployments, performance problems and troubleshooting scenarios. We will look at: - Considerations for topology choices in a replication deployment - How to keep replication running smoothly - How to identify internal performance issues in replicating data - How to approach schema modifications in a replicated database - Approaches for troubleshooting errors - Uses for replication in modern SQL Server environments You will leave this session with a deeper understanding of the internals of replication. You will also be confident in identifying and triaging issues connected with replication systems. This session will cover features that are available in all versions of SQL Server from 2005 up to the latest and greatest release. Key Learning: Replication is powerful, but not well understood. This session will explain how it works and when it can/should be used. Key topics are - Considerations for topology choices in a replication deployment	Intermediate