

Data Platform Virtual Summit 2020 - Session List (Part 3)

Session Release Part 3. Release Date: October 6, 2020. Session Count: 155

Important Notes. Please read carefully.

1. This is the session release part 3. Final release will be in November first week. Total Sessions will be 200+
2. Final agenda/scheduling will be available two weeks prior to the conference.
3. With 200+ sessions, DPS 2020 is The Largest Online Conference On Microsoft Data Platform, Advanced Analytics & Artificial Intelligence. DPS 2020 brings you the finest & latest content from world's best educators.
4. Triple content compared to previous years. Half the price. Complimentary recordings. 50% OFF (Limited Time). Block your seat today.



To view, up-to-date sessions list ----->>

<https://dataplatfromgeeks.com/dps2020/sessions-agenda-schedule/>

Session_Type	Speaker	Track(s)	Session Title	Abstract	Level
Break Out	Kevin Feasel	Data Administration	Optimizing Backup Performance Using Data Science Techniques	<p>One of the most important tasks for a database administrator is taking (and testing!) backups. As databases get larger and larger, the amount of time it takes to perform a backup can grow as well, to the point where your backups take longer than your available backup window. There are several settings we can use to optimize backup performance, such as buffer counts, maximum transfer size, and the number of files, but trying every combination of settings on a single production-sized database could take weeks or even months. In this talk, we will apply data science techniques to the problem of backup settings optimization and look at different models for approaching the problem and analyzing data. Some statistics background would be helpful, but is not required; the big requirement is a desire to speed up backups.</p> <p>#SQLServer #Backups #datascience #R</p>	Advanced
Break Out	Niels Berglund	Development	Set Your SQL Server Data Free With The Use Of Apache Kafka	<p>In this session we look at how we can stream data from SQL Server to the de facto standard for streaming: Apache Kafka.</p> <p>We look at tools like Kafka Connect, and external languages, and after the session we should have a good understanding in various ways we can "set the data free".</p> <p>#sqlserver #apachekafka #streaming</p>	Advanced

Break Out	Dejan Sarka	Data Science (AI/ML)	Text Mining With TSQL	<p>Analyzing text is another modern hype. You have many tools, packages, and languages available. But did you know that you can do really a lot of text mining also in pure Transact-SQL? Yes, your SQL Server can also become a text mining engine. In this advanced session, you will learn how to analyze text in multiple languages with pure T-SQL, using also features from the full-text search (FTS).</p> <p>#textmining #TSQL</p>	Advanced
Break Out	Alexander Arvidsson	Professional Development	Talking To Myself(?) - Lessons Learned From Presenting Online	<p>If you had asked me in March if it was possible to speak or conduct training - with high quality - via the internet I would have said sure, it's doable, but it's different and it's going pale in comparison to an in-person delivery. I would avoid it."</p> <p>Then the world went off the rails.</p> <p>Suddenly everyone had to adapt to being in quarantine or lockdown and not going to any conferences or trainings for the foreseeable future. I was forced to figure out a way to conduct technical training using tools that I didn't believe in, by to people that I couldn't see - without it sucking.</p> <p>I had to figure out how to position myself to be able to use my hands, how to seamlessly switch between PowerPoint and the Azure portal and how to gauge my students' information retention - all in about a week. It turns out that there are a lot more hours in a week than the 40 most people are used to working.</p> <p>This session walks through what I've learned from speaking and conducting training over the internet the last few months - both what works well and what ... doesn't.</p> <p>#Training #Presenting #ProfessionalDevelopment</p>	Intermediate
Break Out	Armando Lacerda	Data Administration	Stretched Databases From Zero To Hero	<p>This is a deep dive session full of demos about SQL 2016 stretched databases feature. It will go from pre-reqs and blocking factor through deployment and disaster recovery. Everything you need to know and when to use this new cloud-based exciting capability.</p> <p>#Azure #AzureSQLDB #hybrid #cloud #database #mvpbuzz</p>	Intermediate

				<p>DevOps promotes self-service access to usable, production-like dev/test environments – but so many data folks are encumbered with shared development databases which are often inconsistent with production. It’s difficult to isolate the development of different features, so they are rolled together into big releases. It’s no wonder database deployments are often expensive, complicated and risky.</p> <p>What if it was possible for developers to provision themselves a disposable, masked copy of last night’s production 64 TB database on their regular workstation, with limited storage, in a matter of seconds? What if it was as easy as “git clone f5”?</p> <p>In this session attendees will learn to use GitHub, Docker, SSDT and PsDatabaseClone to unshackle themselves from “Wild West” dev databases and empower them to use modern branching techniques and Azure DevOps to continuously deliver their updates to production.</p>	
Break Out	Alex Yates & Sander Stad	Development	Solving the dev database problem with GitHub, Docker and dbaclone	The result? Empowered developers who can deliver value in	Intermediate
Break Out	Javier Villegas	Data Administration	Monitoring And Troubleshooting SQL Server Environments Using Free Community Tools	<p>In this session we will see how to monitor and troubleshoot different SQL Server environments using free tools developed by community experts</p> <p>We will cover First Responder Kit , Diagnostics Queries, BPCCheck, dbatools / dbachecks and more</p>	Intermediate

Break Out	Martin Cairney	Architecture	My Top 5 Omissions From Azure SQL Database Applications And How To Fix Them	<p>As a consultant, I am often called upon to help troubleshoot Azure SQL Database applications. Most of these could have been addressed with better design and planning.</p> <p>Come along to this session and discover the common omissions from a robust Azure SQL Database implementation and how you can go about fixing these in your own environment before they bite you.</p>	Basic
Break Out	Edwin M Sarmiento	Data Administration	Building A Custom SQL Server On Containers Image For Deployment	<p>You have been given new responsibilities to automate deployment, scaling and management of containerized applications not just SQL Server databases. And while Microsoft provided publicly available images for you to use, you want to create your own custom images to meet your corporate policy requirements.</p> <p>In this session, learn how to build a custom SQL Server on Containers image solution that you can incorporate in your DepOps workflow. You will also learn how to leverage containers for faster SQL Server deployments - be it in development or production environments.</p> <p>#SQLServer #Docker #Containers #Linux #DevOps #Tag6 etc..</p>	Intermediate
Break Out	Radu Vunvulea	Architecture	The Monster Under The Bed - Overengineering The cloud	<p>Nowdays, every new project that we start is around microservices. You cannot build anymore a simple solution that goes in production in 2 week time. Many times we grow the project complexity because of the current IT trends without thinking about what we need. The focus of the session is around microservices and what is the impact of it during development, automation and operation phase. The purpose is to identify the balance between the size, number and complexity of the services taking into account the delivery time and budget, especially during the first iterations.</p>	Advanced

Break Out	Warwick Rudd	Data Administration	Azure Data Studio Above And Beyond	<p>The data platform environment is exploding so quickly and as a data platform professional how can you be productive and able to support your environments easily?</p> <p>In this session Warwick will introduce you to Azure Data Studio and take you on a demo rich journey on how, when, where and why Azure Data Studio can make your life easier in supporting your ever expanding data estate.</p> <p>If you have not heard of or seen Azure Data Studio before, Warwick will introduce the tool, it's history and introduce you to its continual enhancements to assist you.</p> <p>#AzureDataStudio #DataPlatform #Tooling #SQLServer #Learning</p>	Basic
-----------	--------------	---------------------	------------------------------------	---	-------

Break Out	David Pless & Brian Carrig	Data Administration	SQL Server 2019: Building A Foundation With Persistent Memory	<p>Data volumes have exploded. There is much more data at rest than ever before with only a small amount of the ingested data ever processed. Intel estimates that this is less than 1% of the volume ingested. Yet, we continue to produce data at record levels and systems can barely keep up to meet the demands.</p> <p>In this session we will cover some of the benefits of a hyper-converged infrastructure, Storage Spaces and Storage Spaces Direct in Windows Server 2019, the capabilities of Intel Optane DC Persistent Memory, and how SQL Server 2019 can build upon these technologies.</p> <p>Intel Optane DC Persistent Memory is an innovative memory technology that delivers a combination of large capacity storage and support for data persistence. Persistent Memory can help with increased capacity needs and unique memory modes, lower the overall total cost of ownership while maximizing virtualization densities, and increase memory security with automatic hardware-level encryption.</p> <p>Persistent Memory can be a powerful technology for larger systems such as SQL Server, SharePoint, and private cloud deployments.</p> <p>We will cover how SQL Server can take advantage of Persistent Memory today and where the opportunities lie for all supported versions.</p>	Advanced
-----------	----------------------------	---------------------	---	---	----------

Break Out	Hasan Savran	Development	What Do Graph Tables Bring To The Table	<p>There are many ways to represent data. Most of us stuck with Relational Database Data Models and we don't feel comfortable when data gets represented in different ways.</p> <p>Graph Processing Tables is yet another way to represent data, it can change the way you look at the data. Relational Databases don't adapt well to changes. Adding a new domain or relation to your data model should not be that difficult. You want to create flexible data models that evolve with your business process.</p> <p>Join me to learn how to create flexible data models with Graph Database tables.</p> <p>#SQL Server, #Graph Tables</p>	Intermediate
Break Out	Magnus Ahlkvist	Development	Statistics, An Unreliable Friend	<p>Statistics are essential for SQL Server's optimizer to generate a good query execution plan. A common practice is to update the statistics regularly. Although recent version of SQL Server has improved automatic statistics update with a lower threshold before automatic update is performed the larger an index grows, it is still a common problem that the statistics are not accurate when that one important and time critical query needs to run.</p> <p>In this demo-packed session, I will demonstrate how statistics are used by the Query Optimizer, introduce the Ascending Key problem and talk about how SQL Server 2016 offers a great improvement compared to previous versions of SQL Server. I will also show some examples of missing statistics problems that have NOT been addressed by Microsoft, and a few different ways to address them in your code.</p> <p>#statistics #tsql #performancetuning #ascendingkeyproblem</p>	Intermediate
Break Out	Damir Matešić	Development	MS SQL New functions, Syntaxes, Tips & Tricks	<p>In this session we will see what bring us some newly introduced MS SQL functions and syntax's like COMPRESS, STRING_SPLIT, STRING_AGG, DROP IF EXISTS (a.k.a. DIE), UTF8 Support and many more. How are they performing compared to some old solutions?</p> <p>Accompanied with multiple examples, this session will show you many interesting topics.</p>	Intermediate

Break Out	Pedro Lopes & Joe Sack	Architecture	Azure SQL: Getting Started With An Intelligent Database	<p>Azure SQL encompasses Azure SQL Database, Azure SQL Managed Instance, and SQL Server.</p> <p>As part of the Intelligent Query Processing family, the latest SQL Server 2019 and Azure SQL Database introduce ground-breaking enhancements.</p> <p>These enhancements make up the latest in Microsoft's mission to improve critical parallel workloads when running at scale while remaining adaptive to the ever-changing landscape of data.</p> <p>Join us to know more about the Intelligent Database world that's available to you today, and to know more about the journey onwards.</p>	Expert
Break Out	Markus Ehrenmueller-Jensen	Development	A Game Of Hierarchies: Advanced Graph Algorithms In T-SQL	<p>Even when there is a big need from a business perspective, the solutions in relational databases are not straightforward. We will use this 60min to walk through solutions for typical questions we can ask a graph: centrality, PageRank, transitive closure, and shortest path.</p> <p>Join this session for a journey through best practices to transform your hierarchies and graphs into useful information. We will have fun playing around with a sample database based on G. R. R. Martin's famous Game of Thrones</p> <p>#TSQL #SQL #SQLServer #GraphDB #Graph #Hierarchy #AGameOfThrones #AGameOfHierarchies</p>	Advanced
Break Out	Gilbert Quevauvilliers	Business Intelligence & Advanced Analytics	How I Reduced My Power BI Dataset By 60%	<p>How to optimize Power BI Datasets to ensure that they can run as fast as possible with the smallest amount of memory possible.</p> <p>The session will cover the following topics:</p> <ul style="list-style-type: none"> Data Modelling with the Star Schema Looking at columns in the dataset Data Types and how this affects the size of your dataset Using DAX Studio for analysis of datasets Real world optimizations that I put into practice <p>Questions</p>	Basic

Break Out	Uwe Ricken	Development	Advanced Insides Into System Versioned Temporal Tables	<p>SQL Server 2016 introduces support for system-versioned temporal tables as a database feature that brings built-in support for providing information about data stored in the table at any point in time rather than only the data that is correct at the current moment in time.</p> <p>This session takes you behind the basics of Temporal Tables and answers - after a short introduction - the following questions:</p> <ul style="list-style-type: none"> - How to handle the change of NULL constraints? - Can Temporal Tables be renamed and what happens to the System Versioned table? - What will happen if you change meta data (data length, data type, ...)? - What happens if you drop a column or add another column to the Temporal Table? - what about security for the system versioned table? - what about calculated columns in the temporal table? 	Advanced
Break Out	Jose Manuel Jurado & Roberto Cavalcanti	Data Administration	Azure SQL Database - Troubleshooting real world scenarios with Microsoft Support Engineers	<p>In this session you will learn the best practices, tips and tricks on how to successfully use Azure SQL Database on production environments. You will learn how to monitor and improve Azure SQL Database query performance. I will cover how Microsoft CSS has been using Query Store, Extended Events, DMVs to help customers monitor and improve query response times when running their databases in the Microsoft Azure cloud. These learnings are fruit of Microsoft CSS support cases, and customer field engagements. This session includes several demos</p> <p>#AzureSQL #QDS #Monitoring #Improving</p>	Advanced

Break Out	Torsten Strauss	Architecture	Microsoft SQL Server - In-Memory OLTP Design Principles	<p>In this session we will look at the design principals of the in-memory OLTP engine. We will understand how the in-memory engine optimizes data storage for main memory, eliminates latches and locks, and uses native compilation to reduce the CPU overhead.</p> <p>For this we will compare the traditional on-disk engine with the in-memory engine to decide when it makes sense to use in-memory OLTP.</p>	Expert
Break Out	Kevin Chant & Sander Stad	Industry Solution	Implementing SQL Server DevOps Using Azure DevOps	<p>Welcome to Azure DevOps Duet, a tale about how a development team and an operations team have to bond together and start using Azure DevOps for SQL Server related deployments.</p> <p>This session will cover the process of developing a CI/CD process starting at getting the team on board and ending with making an actual release.</p> <p>We will discuss:</p> <ul style="list-style-type: none"> the challenges of implementing DevOps setting up a database solution project getting started with source control testing your database releases using tSQLt setting up your Azure Devops pipelines <p>As part of the session we will cover the default services you get with Azure DevOps including:</p> <ul style="list-style-type: none"> Azure Repos Azure Boards Azure Pipelines Azure Test Plans Azure Artifacts <p>After this session you will have the tools and knowledge to get started with DevOps and get your development process to the next level. In addition, some ideas to get you developers and</p>	Intermediate

Break Out	Will Velida	Development	Introducing Graph Databases With Azure Cosmos DB	<p>Graph databases are an absolute game changer when it comes to storing our data! Data in the wild the more naturally connected than we might force it to be in a traditional relational database. However, you may not even know what a graph database is and how you can get started building one!</p> <p>In this session, I'll cover the basics of what graph databases are and how we can build a graph database in the cloud thanks to Azure Cosmos DB's Gremlin API. Then I'll discuss how we can model our graph data in Cosmos DB, how we can query our data that's stored in a graph database and how we might apply graph databases to a variety of different use cases.</p> <p>By the end of this session, you'll leave with an understanding of what Graph databases are, what you can use graph databases for and how we can use Azure Cosmos DB to build a graph database.</p> <p>#azure #cosmosdb #graphdb</p>	Intermediate
Break Out	Benjamin Kettner	Development	Cloud-Based ETL With Azure Data Factory And Azure Functions	<p>Today, many ETL workloads and Data Warehouses move to the cloud. But also the data sources and requirements have changed. How can you process streaming data for a cloud-bourne DWH? How can you pull data from REST APIs in batches? What tools do you have for processing data in the cloud? In my session I will discuss scenarios that can be solved using the Azure Data Factory, scenarios that are beyond ADF capabilities and how Azure Functions can be employed in both scenarios. I will show real-life examples utilizing Storage queues for asynchronous processing, durable entities for streaming data and orchestrator functions.</p> <p>#ETL, #Azure, #ADF, #AzureFunctions, #ModernDWH, #DWH, #Cloud</p>	Intermediate

Break Out	David Goad	Architecture	A Guide To Developing Your Ideal Internet Of Things (IoT) Architecture: Nine (9) Key Questions To Ask Yourself BEFORE Developing Your IoT Solution	<p>Unlike the Internet which has been around in some shape or form for almost 30 years now and has a relatively stable architecture based on a handful of well known and recognized industry standards, the IoT has really only been around for the last 5 or 10 years (depending on your definition of IoT) and has a quickly evolving landscape of vendors, protocols, standards, and technologies with perhaps hundreds of options and combinations available to the practitioner to use in terms of technical architecture. Come to this session to learn more about the key questions you need to ask yourself when designing your IoT architecture.</p> <p>#architecture #iot</p>	Expert
Break Out	Stefano Tempesta	Data Science (AI/ML)	AI-Powered SharePoint Intranets	<p>This session combines the agility of building pages and web parts in SharePoint Framework, with the power of the Microsoft AI platform. Specifically, I'll present a dashboard in SharePoint that displays Machine Learning-powered sentiment analysis of your intranet contributions; an automatic document and image classification with Cognitive Service, and a content filtering engine that learns from new entries and improves accuracy of detection over time.</p> <p>#MachineLearning #SharePoint</p>	Advanced

Break Out	Haniel Croitoru	Architecture	Manage Your Power Automate Governance Like A Rockstar	<p>You've managed to get Power Automate introduced and even adopted within your organization, which is awesome. But are you sure that everyone is using it the way they should? Are you up at night wondering whether</p> <ul style="list-style-type: none"> - Do my flow makers know the confined in which they can and should build flows? - Are there flows that are beyond the rights and capabilities of my flow makers? - Is any of my organizational business data being compromised? - Do my flow makers able to leverage Power Automate for their business needs without negatively impacting themselves or their peers? - Are my flow makers following organizational best practices recommended by my IT department? - Is my IT department being taxed with additional support for flow makers beyond their means? <p>If you've answered yes to any of these questions, then you shouldn't miss this session, where I will introduce and demonstrate topics such as</p> <ul style="list-style-type: none"> - Data Loss Prevention in Power Automate - Automating Power Automate onboarding and training for new flow makers - Building a governance matrix to establish what flows should be owned by whom - Leveraging the Power Apps Center of Excellence to audit and 	Advanced
Break Out	Tomaz Kastrun	Data Science (AI/ML)	Applied Data Science With Azure DataBricks	<p>Azure Databricks is an Apache Spark based analytics service for big data and data analytics on top. In this session we will create Databricks scenarios for useful business scenarios.</p> <p>Data engineers and business analysts (data scientists) can now work on RDD structured files using workbooks for collaborative projects, using ANSI SQL, R, Python or Scala, easily covering both analytical and machine learning solutions on one hand, and also giving the capabilities to use it as a datawarehouse.</p> <p>#data #Analytics #Azure #Databricks #datascience</p>	Advanced

Break Out	Monica Rathbun	Data Administration	Mastering Tempdb	<p>Have you experienced performance problems caused by contention in TempDB? Have you ever wondered why your TempDB is suddenly 3 TB? In this session, you will learn about all the various components of SQL Server that use TempDB. Whether it be AlwaysOn Availability Groups, Read Committed Snapshot version stores, spills, or simply temporary tables, learn about how to identify what SQL Server or your applications are doing in TempDB. Once you understand all the ways SQL Server uses this critical resource, and how to properly configure it, you'll be better prepared for your workloads whether it be an Azure VM, a physical server, or a container.</p>	Intermediate
Break Out	Paul Ou Yang	Development	SQL Code Testing Pipeline	<p>Imagine you write SQL code that must be compatible with the 5 latest versions of SQL Server: 2012, 2014, 2016, 2017, and 2019. Manually testing your SQL code in each version is going to be time consuming and perhaps neglected. This is where automation comes in. Using ARM templates, Azure DevOps, Pester and PowerShell, test environments can be easily created and destroyed when new versions are released. This session will help you automate the testing and ensure that your SQL code will run efficiently in any version. Demos planned include:</p> <ul style="list-style-type: none"> - Using ARM templates to create multiple SQL VMs - Working with PowerShell to deploy SQL scripts - Testing SQL execution with Pester - Orchestrating the tasks with Azure DevOps pipelines <p>#DevOps #PowerShell #Pester #SQL #Azure</p> <p>Note: Complete code from start to finish is in GitHub repo https://github.com/paulouyang2/SQLCodeTestingDemo</p>	Advanced
Break Out	Steve Jones	Professional Development	Blogging for the Tech Professional	<p>When a company posts a job opening, they are often inundated with dozens, if not hundreds, of applicants. How do they decide which ones are worth pursuing with an interview? Standing out from other candidates is important to ensure you receive the consideration that you deserve. Building a blog and including those links in your CV can help. Come learn how to build a blog, choose topics, and share your knowledge in a way that boosts your career.</p> <p>#blogging #career #prodev</p>	Basic

Break Out	Steve Jones	Development	Adopting A DevOps Process For Your Database	<p>When your application developers want a fast moving, DevOps software development process, the database often becomes a hindrance. It doesn't have to be. While database development can be challenging when the system needs to rapidly evolve, many best practice software development practices still apply. Come learn about the ways you can mitigate the challenges of performing database development in a DevOps environment.</p> <p>#DevOps #CICD #ContinuousIntegration #Testing</p>	Intermediate
Break Out	Deborah Melkin	Development	Single Statement, Many Changes: How One Statement Can Modify Multiple Tables	<p>You can only insert, update, or delete from one table at a time. At least that's what they tell us when we first learn to write SQL statements. However, that one statement could modify multiple tables, and we may or may not even realize it is happening.</p> <p>In this session, we will examine how a single data manipulation (DML) statement could change data for many tables. We will approach these from two different angles: implicit database design & explicit SQL code and objects. Syntax, performance gains, and gotchas of these different methodologies will be discussed. Finally, we will explore often overlooked changes that occur further downstream as a result of our DML statement.</p> <p>When you leave, you will understand and appreciate how a DML statement against one table affects not only that table but how it can have a ripple effect of changes throughout your entire database.</p>	Intermediate
Break Out	Andy Yun	Development	Demystifying SQL Server I/O To Improve Your T-SQL Performance	<p>Do you understand how SQL Server handles I/O? If not, you should. Understanding I/O is more important than ever with increasing cloud adoption, due to cloud costing models. But even on-premises, I/O has a direct impact on T-SQL performance.</p> <p>Join me for a deep dive into the storage engine to understand how SQL Server handles I/O. Concepts will be paired with T-SQL examples to illustrate the practical impacts.</p> <p>When you leave, you'll be able to optimize your T-SQL, reduce I/O and improve performance across the board.</p> <p>#sqlserver, #tsql, #storageengine</p>	Advanced

Break Out	Tom Martens	Business Intelligence & Advanced Analytics	DAX - Time Series Analysis, Sequences, And Some Other Fancy Stuff	<p>Here Time Series Analysis does not mean calculating the YTD value of a numeric value. This is about discovering the clients who have been buying 10 days in a row and how concepts of time series analysis can be used with DAX.</p> <p>These are the requirements to have fun attending this session: a basic understanding of the Filter Context is required, this article has to be read at least once https://mdxdax.blogspot.com/2011/03/logic-behind-magic-of-dax-cross-table.html and it's also helpful to have a good understanding of extended tables.</p> <p>#DAX #PowerBI</p>	Brain-Hurting
Break Out	Marco Russo	Business Intelligence & Advanced Analytics	Inside The VertiPaq Engine	<p>The VertiPaq engine used by SQL Server Analysis Services Tabular, Power BI, and Power Pivot, is a columnar database capable of incredible performances, both in speed and compression ratio. In this session, we will perform a deep dive in the internals of the database architecture, discovering how Vertipaq stores information, in order to gain better insights into the engine and understand the best way to model your data warehouse to leverage the features of VertiPaq. We will show common and useful techniques to increase the compression ratio and obtain better performances from your Tabular data model.</p>	Advanced

Break Out	Andrew Brust	Business Intelligence & Advanced Analytics	AI And Analytics With Apache Spark And Azure Databricks	<p>Open source technology Apache Spark is the analytics and machine learning platform of choice for many companies. While Spark has manifested in numerous parts of the Microsoft stack, including HDInsight, Synapse Analytics and even SQL Server 2019, Microsoft's go-to Spark service is Azure Databricks.</p> <p>The service, from Microsoft and Databricks (the company founded by Spark's creators), is a versatile one, geared towards data lake management, analytics, data engineering and data science. Azure Databricks lets developers work in notebooks, offline, interactively with running clusters, or scheduled as production jobs that provision Spark clusters on-demand.</p> <p>This session will cover the concepts, service mechanics, and code necessary for you to do analytics and machine learning on Azure Databricks, and integrate it with other Microsoft cloud services and on-premises technologies.</p> <p>You will learn:</p> <p>About the fundamentals of Apache Spark, Spark SQL and Spark MLlib</p> <p>How to use Databricks notebooks and manage clusters</p> <p>The rigors of integrating Databricks with Azure Storage, Azure SQL Database and Power BI</p> <p>How to write Python code for both analytics and machine</p>	Intermediate
Break Out	John Q. Martin	Data Administration	SQL Server And Network Security	<p>The network is often forgotten when securing SQL Server is completed. However, this is a primary attack vector which needs to be designed and configured properly to help add the layers of protection needed.</p> <p>In this session we will explore the network architecture you should look to implement as well as how to leverage Operating System Firewalls as well as Azure Network security configurations. When combined this will add more depth to the defence of your SQL Server security and help you meet compliance and regulatory requirements.</p> <p>#Security #SQLServer #Azure</p>	Advanced

Break Out	Andreas Erben	Data Science (AI/ML)	Cognition For Applications - Cognitive Services And Mixed Reality In Action	<p>Cognition is defined as acquiring knowledge and understanding through thought, experiences, and the senses. This summarizes very well what can be accomplished when using Microsoft Azure Cognitive Services and Mixed Reality capabilities together.</p> <p>In this demo-driven session, the presenter uses those services together enabling advanced scenarios. Cognitive Services allow understanding the environment, combined with Mixed Reality services, this understanding maps to locations in the real world.</p> <p>This allows to create completely new spatially aware experiences.</p> <p>Example code and demos will show both decision makers and software developers of all types what is possible, and how to start implementing in their organizations.</p>	Advanced
Break Out	Greg Low	Development	Really Understanding Character Data Types In SQL Server	<p>If you still think that varchar(10) means up to 10 characters, or you don't know what that N is in front of some strings, and why it matters, or if you don't really know your collations from your code pages, this session is for you. Let Greg explain it to you.</p>	Intermediate

Break Out	Andreas Erben	Professional Development	Hackathons - unleashing hidden potential in organizations and empower individuals	<p>Hackathons are a hot topic to drive innovation. Organizations across the globe have discovered that it enables them to help their digital transformation endeavors through unleashing hidden potential within the organization or across an industry. For individuals, hackathons can help to dive into a new subject matter, empowering to work on subjects they always wanted to work on, and in some cases hackathons have changed career paths.</p> <p>This session will discuss topics around hackathons: What different types of hackathons exist? What value do hackathons have for organizations and for participants? How does one best prepare for a hackathon?</p> <p>Most importantly, hackathons are about culture and in some cases helping to change culture in a positive way. The best hackathons are inclusive and empower people to succeed.</p> <p>Attendees will learn about the value of hackathons for organizations and individuals, and tips will be shared how to run or participate in a successful hackathon. This session is for everyone who cares about innovation, digital transformation, and how to disrupt and reinvent themselves.</p>	Basic
Break Out	Martin Catherall	Development	T-SQL Window Functions	<p>Window Functions first emerged with the 2005 release and were enhanced in 2012.</p> <p>They provide a window over the result set and the ability to access data other than the current row.</p> <p>This means that somethings that may - at first - seem like they need row-by-row processing can be done in one single set-based operation.</p> <p>The set of window functions available also includes server analytical functions that give the ability to perform statistical analysis.</p> <p>Come along to this session and see a deep dive into window functions and learn how their use will add to your ability to process data and write more efficient and cleaner queries.</p>	Intermediate

Break Out	Hamish Watson	Data Administration	KQL Will Be The Next Query Language You Will Learn To Manage Azure SQL Databases	<p>Azure is the leading platform for many companies, and to manage your databases and other infrastructure you need to have insights into what is happening.</p> <p>You need to be able to query the Azure platform, to not only understand your infrastructure but also to leverage monitoring and analytics to react to operational changes occurring in it.</p> <p>This session will introduce you to the Kusto Query Language (KQL) which will allow you to query a variety of Azure resources. The session will also show how machine learning and time series analytics can help advance your insights into you Azure services.</p> <p>We will also look at how we can export our results out to csv files and PowerBI for full-scale reporting of our Azure platform.</p> <p>#Azure #Database #Cloud #DevOps</p>	Intermediate
Break Out	Rainer Stropek	Architecture	Privacy In The Public Cloud - Azure SQL Private Endpoints	<p>Azure SQL Database has been designed so that it can run connected to the public internet. It has proper authentication and encryption built-in. However, administrators are humans and humans sometimes make mistakes. So it cannot harm to add another layer of security and limit access to cloud databases using network security mechanisms. Recently, Microsoft has been adding support for Azure Private Endpoints to more and more PaaS offerings including Azure SQL DB. In this session, long-time Azure MVP and RD Rainer Stropek will demonstrate how to use Private Endpoints to limit access to databases on a network level. He will show how you can connect your databases to Azure Virtual Networks and securely access them from applications running in the cloud and from on-premise systems using Azure VPN Gateways.</p> <p>#Azure #Security #Network</p>	Advanced

Break Out	Andre Melancia	Data Science (AI/ML)	I'm not sorry, Dave... ML.Net can do that!	<p>Machine Learning is trending and many do it with Python or R libraries... But what if you prefer .Net languages? Do Machine Learning in your .Net world in an easy way, by using the ML.Net framework. Don't know how Machine Learning works? This session covers that too.</p> <p>Disclaimer: No monoliths were harmed in the making of this session.</p>	Intermediate
Break Out	Thomas Grohser	Architecture	Building High End SQL Server Infrastructure	<p>While the overall goal is to utilize cost effective commodity hardware for the majority of the database systems sometimes it is necessary to have an efficient high end server to get the job done.</p> <p>This session explains how to scale servers to hundreds of cores, terabytes of RAM and petabytes of storage in an optimized way for SQL Server. We will cover edge technologies like NVDIMM for the tail of the log buffering and talk about CPU and IO affinity and NUMA and how to expand the concept it from memory to your disks and network.</p> <p>The most amazing part of this presentation is that you will learn that these high end servers are no longer as expensive as they one were.</p> <p>This is the session for you if you need (or just want to know) how to handle millions of transactions per second or multi petabyte tables.</p> <p>#SQLInfrastructure #HighEndSQL #CostEffective</p>	Brain-Hurting

Break Out	Anthony Nocentino	Architecture	Containers - What's Next?	<p>You've been working with containers in development for a while, benefiting from the ease and speed of the deployments. Now it's time to extend your container-based data platform's capabilities for your production scenarios.</p> <p>In this session, we'll look at how to build custom containers, enabling you to craft a container image for your production system's needs. We'll also dive deeper into operationalizing your container-based data platform and learn how to provision advanced disk topologies, seed larger databases, implement resource control and understand performance concepts.</p> <p>By the end of this session, you will learn what it takes to build containers and make them production ready for your environment.</p>	Intermediate
Break Out	Richard Conway	Business Intelligence & Advanced Analytics	Building an advanced practice using Azure Data Explorer	<p>In this talk Richard will talk through how to build a scalable, interactive data platform using Azure Data Explorer (ADX) and Kusto Query Language (KQL) based on work he's been doing with one of his clients.</p> <p>In this deep dive he'll go through how to develop and productionise with examples using advanced concepts such as table permissioning, ingesting data, update policies, caching, row level security and other key enterprise features.</p> <p>He's the author of 30 lightning facts in 60 days which parade short blog posts on ADX which you can read more about here https://aizoo.info</p> <p>And an admin command line tool dexcmd, which you can see here:</p> <p>https://github.com/elastacloud/dexcmdline</p>	Expert

Break Out	Sandip Pani	Data Science (AI/ML)	Data Leakage- Most Ignored Problem In Machine learning	<p>When you share information while training model which you shouldn't is referred as data leakage. Beginners to experts many do this mistake of sharing information while training model. Sometime it happens accidentally and sometimes users not even aware about his mistake.</p> <p>In this session we will discuss what is data leakage?, how it impact my model and how will you detect it? and how to prevent data leakage? I will demonstrate with various examples.</p>	Intermediate
Break Out	Christopher Adkin	Business Intelligence & Advanced Analytics	Kubernetes And OpenShift Infrastructure For SQL Server 2019 Big Data Clusters	<p>Planning on building out an infrastructure for a SQL Server 2019 Big Data Cluster and you don't know where to start ?, then this session is for you. It will cover everything you need to know in terms of:</p> <ul style="list-style-type: none"> - Planning your infrastructure - Critical components such as load balancers and what to do for storage - Managing the life cycle of your Kubernetes cluster - Deployment tools and automation - How to backup and protect your environment - How to dedicate critical SQL Server 2019 Big Data Cluster components to specific Kubernetes nodes - OpenShift, what is it, why you should care about it and how it relates to Kubernetes <p>. . . and much much more.</p> <p>This is not just a vanilla Kubernetes 101 session, but a session on what you need to do in order to successfully build a production grade environment to run your SQL Server 2019 Big Data Cluster on.</p>	Intermediate

Break Out	Mihail Mateev	Data Science (AI/ML)	Anomaly Detection, Powered By Azure AI And Digital Twins	<p>Detecting the Onset of System Failure using Anomaly Detection Techniques is one of the key demands in Industry 4.0. Nowadays implementation of this functionality is often related to two concepts: Digital Twins and Anomaly Detection with AI</p> <p>Anomaly detection is the identification of rare items, events or observations which raise suspicions by differing significantly from the majority of the data. Typically, anomalous data can be connected to some kind of problem or rare event such as e.g. bank fraud, medical problems, structural defects, malfunctioning equipment etc.</p> <p>Digital twins have been around for several decades, the rapid rise of the internet of things (IoT) is that they have become more widely considered as a tool of the future. Digital twins are getting attention because they also integrate things like artificial intelligence (AI) and machine learning (ML) to bring data, algorithms, and context together</p> <p>In this presentation will be discussed various techniques that can be used to detect the onset of failure occurring in systems in the context of Microsoft Azure, using Azure Digital Twins Service and Azure Cognitive Service Anomaly Detector API.</p> <p>Azure Digital Twins is a game-changer in the modern IoT and AI solutions. This a SaaS offering easy to build digital models of complex systems.</p> <p>The Anomaly Detector API enables you to monitor and detect abnormalities in your time series data with machine learning. This feature adapts by automatically identifying and applying the best-fitting models to your data, regardless of industry, scenario, or data volume.</p>	Advanced
Break Out	Alex Whittles	Data Science (AI/ML)	Machine Learning In Power BI	<p>Two of the hottest topics in data are Machine Learning and Power BI, but how to the work together? In this session we'll look at what ML tools and features are available in the Power BI world, and how you go about using these in your data and reports. This will include using both built in and 3rd party Power BI visuals, consuming Azure Cognitive Services, as well as the new ML features available in data flows.</p> <p>#PowerBI #ML #AI #CognitiveServices</p>	Advanced

Break Out	Jorge Maia	Industry Solution	IoT Edge - Overview And Case Study For Smart Buildings Focused On Data	<p>This session is a merge of some clients' use cases in the last five years on the smart buildings, facilities, and office automation projects that CrazyTechLabs be leading inside them.</p> <p>Many examples of IoT presents a single sensor or a couple of them and a chart with their data. This session will cover how IoT can lead digital transformation in several buildings and present all Azure Environment that can be used, featuring Azure IoT Edge, Time Series Insights, and Azure Digital Twins to reduce energy waste, count people in rooms, and enhancing the user experience on the buildings.</p> <p>The main idea behind the scenes on almost all of them on this kind of project was automation connected with authorization systems, to track the energy utilization by floors and sectors, based on occupancy and measurements. All data acquired must be stored, and after some months, the data extracted from the application promotes insights about rooms, illumination of individual suites, and common halls. It can turn HVAC, lighting, and access control systems automation with intelligence and control based on predictions and real-time data.</p>	Basic
Break Out	John Morehouse	Data Administration	Optimizing Query Performance In Azure SQL Database	<p>Many think that moving to the cloud will not only help brighten your teeth but also solve all of your bad coding practices that give you poorly performing queries. If it's done correctly, implementing Azure SQL Database can help with one of those two and while it can mask things well, the best solution is to fix the bad code. In this hour-long session, we'll examine several different methods that you can utilize to help fix bad performance starting with the underlying service tier. Next, we'll investigate what options are available directly from the Azure portal to determine where the bottlenecks might reside along with possible ways to fix them. Lastly, we will interrogate which native SQL Server tools exist within Azure SQL Database that can really help solve any performance issues you might be having. You'll leave this session with a solid understanding of how to trouble shoot performance issues in Azure SQL Database and what you might be able to do to help fix them.</p>	Basic

Break Out	Parikshit Savjani & Sunil Agarwal	Development	Modernize your data on Azure Database for Postgres and MySQL Flexible Server	Are you looking to migrate it to the cloud, gain operational efficiency, and yet more importantly, not give up control of your database? Look no further. In this session you will learn how the new Flexible Server for Azure Postgres or Azure MySQL gives you full control for choosing your configuration, optimize cost, manage automatic maintenance, with flexible high availability options for your databases. For each of these capabilities, we will cover the 'why' and how using with simple demos to make your learning real.	Advanced
Break Out	Sunil Agarwal & Arun Kumar	Data Administration	Best Practices For Migrating Your Workloads To Azure Database For PostgreSQL And Azure Database For MySQL	For open source relational database workloads, customers are increasingly adopting Azure Database for MySQL and Azure Database for PostgreSQL to take advantage of a fully managed database service technical teams innovate faster. In this session, we will cover the best practices to migrate existing MySQL/PostgreSQL/Oracle workloads to Azure PaaS. We will also dive deeper into various migration tools available for offline (one-time migrations and online (ongoing replication) migrations with a live demo of a MySQL to Azure DB for MySQL online migration.	Advanced
Break Out	Marsha Pierce	Data Administration	Migrate Your Database Like A Rockstar	We often have to migrate databases between servers. It could be for a SQL Upgrade or it could be for a Hardware Upgrade. When you move servers there are a ton of things that can cause your move to fail. Learn from someone who has moved 1000s of instances. I will show you how to move your databases to a new server with just a few minutes of downtime and how to reduce the risk of needing to rollback. We will account for everything except an IP change including replication and not having to reseed it,.	Brain-Hurting

Break Out	Alicia Moniz	Data Science (AI/ML)	Data Stewardship In An AI-Driven Ecosystem: InterpretML, FairLearn, WhiteNoise	<p>At the core of Microsoft's AI are the principles of fairness, reliability & safety, privacy & security, inclusiveness, transparency & accountability. As AI capabilities increase along with adoption, it is important that we also leverage tools that enable us to practice AI responsibly.</p> <p>Responsible ML provides us with tools to ensure that as practitioners we</p> <p>Understand machine learning models - Are we able to interpret and explain model behavior? Are we able to assess and mitigate model unfairness</p> <p>Protect people and their data - Are we actively working to prevent data exposure with differential privacy?</p> <p>Control the end-to-end machine learning process - Are we documenting the machine learning life cycle?</p> <p>Announced at Build this year were multiple Responsible ML open source packages. The accessibility of these freely available tools enables every machine learning developer to consider incorporating Responsible ML into the development cycle of their AI projects.</p> <p>InterpretML - An open source package that enables developers to understand their models behavior and the reasons behind individual predictions.</p>	Advanced
-----------	--------------	----------------------	--	--	----------

Break Out	Ronen Ariely	Data Administration	SQL Internals - Physical Table Structure Under The Hood, And Implementation On Real Case Scenarios	<p>Understanding what we have under the hood is not only done for the sake of learning theory, but it directly impacts your daily work, and it can help us to improve performance and reduce resources dramatically.</p> <p>During this session, we will go over the structure of the tables behind the scenes. I will show the changes in the table structure behind the scenes, related to actions which we execute on the table. Using this information, I will demonstrate several real case scenarios and the dramatic impact of understanding the internals in choosing our solution.</p> <p>Is the order of the columns in the Create Table statement important? Can a specific order of columns affect performance or lead to an error in the future use of the table? Is dropping a column that is no longer needed is the right solution? Questions like these are very common in forums, but usually do not get the right answers even from experts.</p> <p>Note: This session is a live demo, level 500 internals. It based on undocumented tools.</p>	Brain-Hurting
-----------	--------------	---------------------	--	--	---------------

Break Out	Mohit Batra	Business Intelligence & Advanced Analytics	Building Streaming Pipeline With Azure Databricks Using Structured Streaming And Delta Lake	<p>Modern data pipelines not just work with batch processing of data, but it often includes streaming data that needs to be processed in real-time. However, many a times, requirements go much beyond that. The processed data may then be consumed by downstream batch, as well as another streaming pipelines. There are several challenges that you need to address here.</p> <p>First, to use a common platform to build unified batch and streaming pipelines. Second, to store the processed data into a Data Lake, and still make sure that downstream systems can consume it reliably, without any consistency issues. And third, to have a unified environment for development and deployment. And this is what we are going to address here.</p> <p>This session will take you through three components.</p> <p>Structured Streaming stream processing component of Apache Spark to build reliable streaming pipelines</p> <p>Delta Lake an open-source storage layer that brings ACID transactions and reliability to your data lakes.</p> <p>And Azure Databricks a unified analytics platform that runs on Azure, handles the infrastructure, and provides a collaborative environment for doing development.</p> <p>We will first talk about the end-to-end architecture. Then we</p>	Advanced
-----------	-------------	--	---	--	----------

Break Out	Ferenc Csonka	Business Intelligence & Advanced Analytics	XMLA Read-Write Endpoint: The Cornerstone For Power BI As An Enterprise BI Solution	<p>Since the release of Power BI Premium in 2017, Microsoft has been continuously expanding the list of features that differentiate its solution running on a dedicated capacity from the shared capacity version. A game-changer step in this process is the release of the XMLA read-write endpoint, which makes Analysis Services Tabular models of Power BI reports deployed on Power BI Premium available for both reading and writing. This makes a myriad of applications and techniques previously developed for on-premises Analysis Services Tabular and Azure Analysis Services models available for developing and operating Power BI reports, taking Power BI Enterprise BI solutions to the next level. In this session, we present several practical use cases for XMLA read-write endpoint: migrating existing Analysis Services Tabular models to Power BI Service, significantly improved application lifecycle management for Power BI artifacts, debugging, monitoring and tracing functionalities for Power BI reports, much more sophisticated incremental refresh solutions, using such Analysis Services Tabular-specific features which are not available in Power BI, etc.</p> <p>Power BI Premium, XMLA, Analysis Services, monitoring, ALM, incremental refresh</p>	Advanced
Break Out	Tillmann Eitelberg	Architecture	Structuring The Data Lake	<p>The Data Lake, infinite storage space, infinite possibilities to collect and hoard data. But is that the goal? Especially in a data lake there are rules how to store data, developers have to deal with naming conventions, business rules and policies.</p> <p>In this session we will show which different concepts exist for structuring a data lake. Which naming conventions have become established or should be avoided. The Azure Data Lake also offers some support on the infrastructure side in the form of policies that can help to build a clean data lake.</p>	Intermediate

Break Out	Melody Zacharias	Architecture	The Download On Azure Arc For Data Services	<p>Ever wondered what Azure Arc was and how it relates to the data professional. This session gives you the full download on Azure Arc for data services. This session will review the architecture as well as the business use cases. We will go into how this service can help your business or clients as well as what to avoid. This session looks at how Arc can be used to extend to a hybrid infrastructure in both a connected and disconnected scenario. A demo of the environment and management options will showcase the latest features of Arc for data services.</p> <p>#ARC #Containers #Hybrid #Azure</p>	Basic
Break Out	Paresh Motiwala	Professional Development	How to Get and Nail Your Interviews	<p>Workforce reduction can strike anyone, anywhere and anytime. Are you prepared for that moment?</p> <p>Let's look at some real life examples through this session that draws full rooms wherever I have presented it!</p> <p>Also, do you feel you need to get out, try either a new company or a new role or even a new boss? Then, this very entertaining, informative and highly interactive session is perfect for you. I assure you that you'll walk out of it totally reassured and enlightened.</p> <p>You'll learn techniques for finding new opportunities in a competitive job market. I'll also share how to exploit the social media to your advantage. (They are not bad if used judiciously) What can you do once you get that rare and elusive interview? There are several things we need to take care of like, overall appearance, food habits, arrival, sitting posture, small talk or ice-breakers and humor. Finally you will learn about the follow-up and importance of keeping a log.</p> <p>#Networking #Interviewing #JobHunting</p>	Basic

Break Out	Damian Widera	Data Science (AI/ML)	Quantum Computing Explained With Demos!	<p>This talk is inspired by the quantum revolution we will see in a moment.</p> <p>I would like to introduce the topics:</p> <ul style="list-style-type: none"> - reminder why basic linear algebra (matrices and vectors) is important - qbits, superposition, and quantum logic gates - Show the simplest problem where a quantum computer outperforms classical methods <p>If time permits - quantum entanglement</p> <p>I will do demos implemented in Q# with the Microsoft Quantum Development Kit.</p> <p>I do not assume any special knowledge, but if you know basic linear algebra (matrices, vectors, matrix multiplication) than the topic will be more easily understand!</p>	Intermediate
Break Out	Ben J Miller	Data Administration	SQL Server Encryption Unplugged	<p>There are so many ways to encrypt data today, so you need to get your game on. From TDE, Always Encrypted, cell encryption and Backup encryption, there are some key pieces of information and methods that will help you get these all right. This session covers a lot of ground and will be over 60% demos. I will take you through each one of these and show you how it works. As an added bonus, we will hook TDE on premises up to Azure Key Vault too. We will also take you through rotation of the keys for TDE. Join me for a fast and furious demo packed session to get your data encrypted like a pro.</p> <p>#TDE #SQLEncryption</p>	Advanced

Break Out	Davide Mauri	Development	Create Secure API With .NET, Dapper And Azure SQL	<p>Creating backend REST API is a very common duty in today's work. But what about creating an API that is also secure and only shows the authorized data to each user invoking it? Thanks to Azure SQL's Row Level Security feature, it's very easy to create an API, even on an existing data, without any change to the database schema, that is secure and fast. In this session we'll see how we can create a full REST API in an hour, using Azure SQL, the micro-ORM Dapper and .NET Core.</p> <p>You'll be amazed to see how simple it is, even while you are not giving up on performances and security!</p>	Advanced
Break Out	Manish Kumar	Data Administration	Fully Secure And Control Your Data In Cloud With Azure Database For MySQL Service	<p>With rapid digital transformation, Azure database for MySQL continues to attract increasing number of business-critical applications handling financial, healthcare, government, and other sensitive data. With these critical workloads moving to the Azure Database for MySQL, the need for stronger security capabilities has never been greater. Join us in this session, to see the latest security innovations for Azure Database for MySQL and learn how to secure your MySQL environments in Azure. Learn about the new cutting-edge capabilities in network security, authentication, data protection, security management etc. and about some of the best practices to secure your MySQL workloads on Azure.</p>	Advanced
Break Out	Amol Bhatnagar	Development	Turbocharge Your MySQL Engine On Azure To Make It Fast, Resilient And Scalable	<p>Looking to unleash the full power of the cloud and build fast, resilient and internet scale applications with MySQL on Azure? Are you looking to extract more performance and scale out of your Azure Database for MySQL service? Join this action packed deep technical session to understand how do you measure and monitor the right metrics with Azure Database for MySQL service, configure your Azure DB for MySQL server for best performance and how to scale your application seamlessly from 0 to thousands of QPS without any application changes. Whether you are a DBA, developer or devops engineer, you'll leave the session with a list of great tips and tricks for optimizing performance of your MySQL server without any application side changes.</p>	Advanced

Break Out	Sunitha Muthukrishna	Development	Build Modern Cloud Native Applications With Kubernetes And OSS Databases (MySQL & Postgres)	<p>With changing customer requirements, you need to think about modernizing your approach to application development to be cloud-ready. In this session, we will show you how to modernize your application with Azure Kubernetes Service (AKS) and MySQL on Azure. We will walk through on how to migrate an application that uses PostgreSQL or MySQL to Kubernetes using AKS and Azure Database for MySQL. In addition to the migration, you will also learn how to optimize AKS to work best with Postgres or MySQL on Azure and to simplify migration and operational efficiency resulting in increased scalability, improved performance, and reduced costs.</p>	Intermediate
Break Out	Guy Glantser	Development	Analyzing Execution Plans Like A Pro	<p>The most powerful tool when tuning a query is the execution plan. In most cases, by analyzing the execution plan of a query, you can learn all you need to know about how the query processor has chosen to execute the query, what is causing the execution to take so long, and what you can do in order to make it run faster.</p> <p>But there are common pitfalls that you should be aware of when analyzing an execution plan, such as the fact that the operator cost values are always estimated, or the true meaning of the cardinality numbers within some of the plan operators. These pitfalls can be very misleading, and they might lead you to spend your tuning effort in the wrong direction. By understanding these pitfalls and avoiding them, you can save a lot of time and effort. You'll become a performance tuning hero by getting the slow queries to run faster faster</p>	Intermediate

Break Out	Rob Sewell	Development	Notebooks, PowerShell And Excel Automation	<p>Everyone loves Excel. A common request that I hear is can you put that into Excel for me, please. I love automation and so have combined two of my favourite tools PowerShell and Notebooks. Using Jupyter Notebooks to accomplish regular tasks and run-books has become common over the past year. A particular use-case that my clients have found very useful is to quickly create Excel sheets of information using Notebooks, PowerShell, and the ImportExcel module. Let me share my experience with you and show you how to create your own Notebooks for your users and the tactics I have learned to enable users to self-service and create their own Excel Workbooks.</p> <p>#Automation #Excel #PowerShell #SQL Server #dbatools</p>	Basic
Break Out	Melissa Coates	Data Administration	What You Can Learn From The Power BI Activity Log And REST APIs	<p>The Power BI Activity Log and the Power BI REST APIs are a goldmine of information for understanding usage patterns and activities in your Power BI environment. In this demo-packed session, we will look at examples of the type of data you can obtain and how this data can make a significant contribution to your Power BI governance, security, management, and adoption efforts.</p>	Intermediate
Break Out	Louis Davidson	Development	Matters Of Concurrency	<p>OLTP databases can be constantly written to and reporting databases are written to at least periodically. In order to ensure consistent results, connections must be isolated from one another while executing, ideally with the lowest possible cost to concurrency. How this isolation is handled is based on the isolation level, whether the classic lock based or the newer optimistic scheme of the in-memory OLTP engine is used, or even if both engines are enlisted in the same transaction. In this session we will look at examples of how SQL Server isolates reading and writing operations from other writing operations to explore how this may affect your application through error messages and performance hits.</p> <p>Tags: #Concurrency #MemoryOptimizedTables #SQL Server 2019</p>	Advanced
Break Out	Amit R S Bansal	Development	Developer Best Practices In SQL Server	<p>In this session, you will learn how to write an optimized T-SQL code for SQL Server. You will see real-world examples related to Sarg-ability, Re-Writing T-SQL for better performance, SQL Injection, Parameterization, and if time permits, we will also see the good and bad of Dynamic SQL.</p>	Intermediate

Break Out	Grant Fritchey	Development	Learn How to Deploy Databases Using AWS DevOps Tools	<p>DevOps itself is all about a cultural change in how you manage your IT resources. However, the biggest challenge when implementing DevOps is learning all the necessary tooling in support of automation. This issue is even stronger when it comes to automating database deployments. This session will show you the tools and mechanisms needed to get your databases deploying through AWS. We'll slowly add additional steps and tools to expand the complexity and functionality of the deployment process. Understanding how to build out the bare bones of an AWS deployment process will make it easier to expand your deployments into the complex undertakings that they will ultimately become.</p>	Intermediate
Break Out	Matt Gordon	Industry Solution	Crushing a Cloud Migration: Moving Mountains While Migrating Data	<p>Are you an architect or DBA who has been tasked with executing a cloud migration? Are you a consultant who is trying to wrestle with the scale and scope of a cloud migration for a client? While there are certainly technical issues to solve and tooling selections to understand and execute, what often gets lost in a migration are the political and organizational mountains that need to be moved for the project to be successful. Technical success will not happen with our organizational success - they work hand in hand to lead to a successful conclusion.</p> <p>As we walk through this session we will discuss and demo the technologies and tools that are useful for a cloud migration - and we'll review some that are not useful and talk about why they should be avoided. We will also mix in discussions of successful migrations and how the team worked together to move those organizational and political mountains out of the way in order to execute a successful cloud migration to make clients, customers, and managers happy!</p> <p>Tags: #Cloud #Migration #DBA #Architect #Azure #AWS</p>	Intermediate

Break Out	Tracy Boggiano	Professional Development	Mental Health and Wellness in IT: Safeguarding our most precious resource	<p>One in four Americans suffers from a mental health challenge each year. In the tech community, this number rises to 42%. The stigma associated with discussions around mental health, as well as how it can negatively impact your work and home life, make people reluctant to discuss the topic. Now is the time to start discussing mental health in tech and how to take advantage of the opportunity to safeguard our most precious resource: people. We'll discuss how mental health affects job quality, including sharing examples from my professional life. We'll discuss the four primary workplace factors which play a pivotal role in mental health challenges in the tech community. I'll provide tips on what is best to say and not say to people you know whose lives are being affected by their mental health. Lastly, we'll cover how to make your workplace a safe place where mental health can be talked about and how to encourage your company to offer resources to help maintain everyone's mental health.</p>	Basic
Break Out	Daron Yondem	Development	Cosmos DB: Jack of All Trades, Master of Many	<p>How would it look like if you were building a wishlist for an imaginary database; turnkey global distribution, 5 well-defined consistency levels, SLAs across 4 dimensions (High Availability, Performance Latency, Performance Throughput, and Data Consistency), Schema-agnostic automatic indexing, Key-Value, Tabular, Document, Graph, Relational Data Models, MongoDB, Gremlin, T-SQL API support... Did we go too far? Not really. So far we have just listed some of the features of CosmosDB. Come and join me to take a closer look and see if it is all blue skies, or maybe not?</p>	Intermediate

Break Out	Mike Nelson	Data Administration	Cross-Platform Management and Fun with PowerShell, the WSL, and Windows Terminal!	<p>Let's face it - Windows and Linux are becoming very comparable with cross-platform databases and the interoperability between them. It is becoming a time when data administrators must embrace both worlds and dive into understanding things they have not had to before. For instance, did you know that PowerShell is available for other platforms than just Windows? Did you also know that there is an exceptionally easy way to combine all of those platforms into a single pane of glass with Windows Terminal, and run them simultaneously with the Windows Subsystem for Linux (WSL)? This session will unravel all of the mystery and complexity around how to effectively manage these environments using these tools and some tricks to help guide your way to becoming cross-platform gurus!</p> <p>Tags: #WSL #management #PowerShell #Windows Terminal #Linux #Windows</p>	Intermediate
Break Out	Sray Agarwal	Data Science (AI/ML)	How to Remove Bias and Make Machine Learning Models Fair & Discrimination-Free at the Example of Credit Risk Data	<p>There have been multiple instances when a machine learning model was found to discriminate against a particular section of society, be it rejecting female candidates during hiring, systemically disapproving loans to working women, or having a high rejection rate for darker color candidates. Recently, it was found that facial recognition algorithms that are available as open-source have lower accuracy on female faces with darker skin color than vice versa. In another instance, research by CMU showed how a Google ad showed an ad for high-income jobs to men more often than women. Using credit risk data where we wanted to predict the probability of someone defaulting on a loan, we were able to shortlist features that were discriminatory in nature.</p>	Advanced

Break Out	Eva Pardi	Data Science (AI/ML)	Reinforcement Learning in gaming	<p>Microsoft uses reinforcement learning in many ways to improve our products and services. AI and machine learning can help accelerate game development by providing more realistic worlds and challenges as well as support automation and live operations.</p> <p>At Game Stack Live, Microsoft Research announced Project Paidia, a research effort aimed at exploring new opportunities created with AI based reinforcement learning in gaming. During this session, attendees will be able to get started with Reinforcement Learning for gaming with the use of Azure Machine Learning.</p> <p>Tags: #AzureML #Microsoft #ReinforcementLearning #ProjectPaidia</p>	Expert
Break Out	Ginger Grant	Business Intelligence & Advanced Analytics	Machine Learning with Spark and Azure Synapse	<p>Azure Synapse Analytics contains a number of different features, and this session will focus on how to use the integrated Spark clusters to analyze data using Auto ML and other coding examples. See how you can use Azure Synapse to create data pipelines and do exploratory analysis without creating a SQL Pool. This session will provide real world examples showing you how you can use Azure Synapse to create solutions with your data and the demos will explain how to make it possible.</p> <p>Tags: #AzureSynapse #ApacheSpark #MachineLearning</p>	Intermediate
Break Out	Neil Hambly	Data Science (AI/ML)	Data Science - Feature Engineering	<p>Data is the life (blood) of Analytics. Without it we have nothing. In this session we will focus on a core aspect of the TDSP "Feature engineering". Feature Engineering involves aggregation and transformation of data variables to create features which are used in the analysis. We need to understand how features relate to each other and then select an optimum algorithms for these features. Feature Selection is where we define a subset of features to reduce the dimensionality of the training (model) and improve performance & costs of the ML Algorithms. One way we can understand these features better is to perform a PCA (Principal Component Analysis). For this important step using an example, we step through the process in a Jupyter (python) notebook to identify the features of a dataset which conveys important information about the structure of the dataset.</p>	Advanced

Break Out	Lindsey Allen	Data Science (AI/ML)	Azure AI, Power new possibilities for every organization	AI and data is at the center of the digital feedback loops. The quintessential characteristic of every application going forward will have AI, we have invested in a comprehensive portfolio of AI tools, infrastructure and services. We are also pushing the bounds of how computers and AI can generalize learning beyond narrow domains. We work with our partners on our collective pursuit to democratize AI and its benefits for everyone. Come to this session to get an update of Azure AI with demos.	Intermediate
Break Out	Reid Havens	Business Intelligence & Advanced Analytics	Unlocking New Visualizations and Features in Power BI	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. Tags: #PowerBI #Visualizations #Storytelling #CustomVisualizations #Reporting #KPI	Basic
Break Out	Argenis Fernandez	Data Administration	Advanced Storage Troubleshooting for SQL Server	In this session we will go over the tools and tricks used to track down nasty storage issues as they affect SQL Server databases on hypervisor plus storage array or hyper-converged configurations. We will go ver the different places in the stack where I/O could be stuck, queued, or just flat out slow. This is an advanced session and as such you can expect tons of bits and bytes being shown!	Brain-Hurting

Break Out	Sebastiano Galazzo	Data Science (AI/ML)	Machine Learning - Best Practices and Vulnerabilities	<p>Artificial Intelligence and Machine Learning are a must nowadays.</p> <p>For projects carrying a simple or well-known problem we can find a lot of ready-made solutions, but the game changes when facing with specific custom problems.</p> <p>The first part of this session is a deep down on techniques approaches and best practices in configuring ML algorithms but much more, do we really need it always?</p> <p>The second part will cover vulnerabilities of ML, discovering how easy could be to fool and to hack a neural network by some techniques (Like pixel attack) and their implication in (our) security.</p> <p>A demonstration will focus on a case of e-commerce using cloud ML (Cognitive) services, breaking them then possible solutions and workarounds.</p> <p>#AI #MachineLearning</p>	Intermediate
Break Out	Leila Etaati	Data Science (AI/ML)	AI Builder: AI in Power Apps and Power Automate	<p>There are many ways to extend the AI capabilities inside Power Apps and Power Automate. In this session, you will see an overview of different capabilities with a brief demo on create an application to scan a business card and store back the result into Microsoft Outlook with help of AI builder and Power Automate.</p>	Advanced
Break Out	Sam Nasr	Development	Storing and Searching Files in SQL Server FileTables	<p>FileTables is a SQL Server features that allows easy storage and access to documents dynamically. After some initial setup, users can simply drag files into a folder and then be able to access their meta-data and contents through SQL Server, seamlessly. Due to the way FileTable uses the FileStream feature, it stores files on the file system but maintains the metadata in the database. This allows full control over the file while not bloating the database.</p> <p>Tags: #FileTables, #Search</p>	Intermediate
Break Out	Peter Myers	Business Intelligence & Advanced Analytics	Working with different Power BI data model architectures	<p>Power BI provides you with different data model architectures. It's all controlled by setting the storage mode of model tables, as either Import, DirectQuery, or Dual. In this presentation, learn why and how to develop the model that best fits your data and circumstances. Also, learn how you can extend an existing Power BI data model with new data and calculations.</p>	Intermediate

Break Out	Kevin Kline	Data Administration	Query Tuning Internals for the Advanced SQL Developer	<p>Skilled SQL developers know that the SQL Server query optimizer uses a multi-step process to produce execution plans and that SQL Server uses a cost-based optimizer to construct query execution plans. But how does the optimizer decide the cost of a given operator over another equally valid operator, say between a nested loop join and a merge join? This session will teach you advanced techniques using undocumented trace flags to deeply explore the query optimization process. These choices can have a dramatic impact on performance, so we will pay special attention to the algebraizer, including associative, commutative, and transitive transformations. Together, we will explore the SQL Server internal memo structure to see how SQL Server uses the heuristics of the algebraizer and query optimizer. We will examine a variety of everyday queries whose performance can be greatly improved by applying a deeper understanding of these internal behaviors. Lots of examples and demos!</p> <p>Tags: #optimizer #algebraizer #parsing #binding #JoinOrder #MemoStructure #ParseTree</p>	Expert
Break Out	Jeremy Frye	Business Intelligence & Advanced Analytics	Power Apps and AI - A Flow of Location Events	<p>The Microsoft Power Platform has become an innovative bridge between data interaction, visualization and analytics. The possibilities have become almost endless for businesses to engage with their data in amazing ways.</p> <p>In this session, I will show you an application using Power Apps and location services to track latitude and longitude coordinates on Bing maps. Attendees will learn how to configure and use the application on a mobile device. I will also show two approaches using AI within the application by using Azure SQL Database, Cognitive Services and Power Automate.</p> <p>An end-user will be able to leverage Bing Search for location-aware web searches while an administrator will be able to leverage the Azure SQL Database and Power Automate to send alerts based on predefined metrics.</p>	Intermediate

Break Out	Harsh Chawla	Professional Development	DBA to Data engineer	<p>How transition from monolithic to Microservices applications brought this radical change</p> <p>Evolution of Data engineering and data scientist roles</p> <p>Transition from Data warehouse to Modern data warehouses</p> <p>Transition from Data mining to Advanced data analytics</p> <p>Real time and batch mode data processing</p> <p>Lambda and Kappa Architecture</p> <p>Anatomy of Modern data warehouse and Advanced data analytics solutions</p>	Intermediate
Break Out	Bob Pusateri	Data Administration	Minimizing User Impact with Advanced Restore Methods	<p>We all know that backups are only half the battle - restores are what really matter when disaster strikes. Standard restores, while effective, may require additional downtime and further affect the business. This session will demonstrate three advanced restore methods you should know: point-in-time restores, piecemeal restores, and page restores, and will discuss when each method is appropriate. Attend this session to learn how to be a better DBA by minimizing downtime and user impact after disaster has struck!</p>	Intermediate
Break Out	Reza Rad	Business Intelligence & Advanced Analytics	Power BI for Ticketing and Subscription: A design pattern review	<p>This session explains one of the common design patterns for Power BI implementation. A design pattern that can be used for subscription or ticket systems. Through this session you will learn the structure of the data model, the data preparation steps to get there, DAX expressions to support analytics.</p> <p>This is not a beginner-level session, you should have some understanding of Power BI, Power Query, and DAX already.</p> <p>Tags: #PowerBI, #PowerQuery, #DataModeling, #StarSchema, #DAX #Analytics</p>	Expert
Break Out	Reza Rad	Professional Development	Stop looking for job, let the job finds you	<p>Come to this session to learn tips and tricks from Reza himself, things that he learned through his professional career. Tips and tricks that help you to find your dream job. This is not a technical session, it is all about professional development and soft skills.</p> <p>Tags: #ProfessionalDevelopment, #SoftSkills, #JobHunting, #CareerSuccess</p>	Basic

Break Out	Paul Turley	Business Intelligence & Advanced Analytics	Paginated Reports: the New Old Operational Reporting Platform	<p>Power BI Paginated Reports (aka SQL Server Reporting Services) was old but now it's new again. Available on-premises or in the Power BI service with flexible licensing, you have multiple options to implement operational reports. This session will briefly cover the differences between analytic and operational reports; and help you understand the advantages and trade-offs using Power BI Paginated Reports, Power BI Report Server and SQL Server Reporting Services. Material from our forthcoming book: Paginated Report Recipes.</p>	Intermediate
Break Out	Joshua Higginbotham	Business Intelligence & Advanced Analytics	Functional Python for the ETL Developer	<p>Python has been around for close to 2 decades now, with it's popularity and adoption growing larger each day. As Data Professionals moving towards a cloud approach, we will find Python as an offering in the majority of areas in the cloud environment. To be able to support this language, we can begin the adoption early on while still on-premises, replacing some of our traditional ETL's while increasing efficiency and accuracy of our existing solutions.</p> <p>In this session:</p> <p>We will leverage commonly used packages across the open source community that lowers your cycle time from start to completion. Building out reusable and scalable functions that facilitate commonly used ETL processes across your estate.</p> <p>Tags: #DataEngineering #Analytics #ETL #Python</p>	Intermediate

Break Out	Thomas Leblanc	Business Intelligence & Advanced Analytics	Building an Agile Data Warehouse	<p>The Data Vault method can be used in agile development. If you follow 2.0 structure the hubs, links and satellites only have inserts. The difficult or more time consuming part is the Information Mart. This is where views for reporting can be used for reporting or converted into tables. The view code can be used, with some date manipulation, to schedule updates for structured dimensional model tables. The dimensional Data Mart is what most reporting or visualization tools (Power BI or Tableau) are programmed optimally.</p> <p>It still gives those power query writers (T-SQL experts) the ability to see the archival data in the Data Vault if they so desire. What really becomes apparent is the integration of new applications that replace the existing applications. Satellites are great at structuring the hard rules of the data to tables and relationships. Again, the Information Mart views or tables is where the soft rules will be applied to help reporting queries.</p>	Intermediate
Break Out	Warner Chaves	Architecture	Global Analytics with Azure Cosmos Db and Synapse Analytics	<p>Cosmos Db is Azure's NoSQL Database as a Service, born in the cloud and designed to take advantage of the flexibility, elasticity and global reach of cloud computing.</p> <p>Synapse Analytics is Azures data analytics services that integrates on-demand SQL querying, Spark big data, Data Lake Store Gen2 as well as an integrated authoring experience.</p> <p>Together these two services can be used to develop solutions in a simple and elegant way that would have been incredibly complex before. The most ambitious is the capability of doing Global Analytics, being able to do analytical queries over your live operational data coming from anywhere in the planet. All without having to handle one piece of infrastructure yourself.</p> <p>In this demo-heavy session we will look at the C# code, features and configuration of Cosmos Db and Synapse and see the Global Analytics in action live.</p> <p>#CosmosDB #NoSQL #Synapse #BigData #ADLS2 #Analytics</p>	Intermediate

Break Out	Satya Ramesh	Development	Understanding NULLs in SQL Server	<p>SQL Server follows three-valued logic (3VL) (instead of Boolean logic) to evaluate any condition to True or False or Unknown. Unknown value in SQL Server is represented as NULL. NULLs in SQL Server are often misunderstood and causes data inconsistency issues. In this session, we will try to go through different test cases to understand NULLs properly.</p> <p>#3VL #NULL #SQL #SQLServer #TSQL</p>	Intermediate
Break Out	Manohar Punna	Development	Cloud Infrastructure as Code (IaC) using Terraform	<p>In recent years, the cloud has been adopted by most companies. It offers greater flexibility in terms of infrastructure and resource management at a greater scale. Although, if not noticed the prices of maintaining this infrastructure and resources will blow through the roof. DevOps and deployability of resources as needed have become even more critical with the cloud.</p> <p>Infrastructure as code is a necessity than a luxury. Terraform offer great IaC solution across multiple cloud vendors. In this session, I will explain the importance of infrastructure as code for cloud resources and introduce how Terraform can be used for building your IaC solutions. I will dive deep into using Terraform for deploying cloud resources for an end-to-end application platform in Azure.</p> <p>#IaC #Terraform</p>	Intermediate
Break Out	Michelle Gutzait	Architecture	Hey! I really need to save on SQL Server license costs, would you help me consolidate?	<p>Abstract In this session, Michelle will talk about consolidation options for SQL Server environments, discuss license costs reduction, Cloud alternatives and provide examples from real customers.</p> <p>Tags: #sqlserver #consolidation #azure #sqlserverconsolidation #sqlazuredatabase</p>	Intermediate

Break Out	William Durkin	Data Administration	Upgrading/Migrating a Microsoft Data Platform Environment: Experiences from the Field	<p>In this session we will go over some experiences, thoughts, and pitfalls around upgrading and migrating SQL Server solutions. The audience will hear how different approaches are necessary depending upon business needs and that the most obvious solution isn't always the right one! Along the way, we will gain insight into the tooling and methods that can be used to help during the different phases of these projects. How to prepare, how to execute and how to complete a migration or upgrade. Attendees will leave the session with important information to help design and execute their future projects, including questions to ask and options to consider before embarking on the upgrade journey.</p> <p>#Migration #Azure #SQLServer</p>	Intermediate
Break Out	Charley Hanania	Business Intelligence & Advanced Analytics	Querying best practices for Azure Synapse Analytics	<p>Data warehousing is evolving into the Cloud so it's time to get on board. With Azure Synapse Analytics there are now more ways to transform, manage and present your data pipelines with scale and performance in mind. Recent experiences from clients on the platform have shown that if you just try to dump data into the platform and use it without understanding the underlying architecture and capabilities, you'll leave fairly quickly as you see your costs skyrocket and the benefits decline. Join this session to get an understanding of what options are available to reduce your query times and lower your costs while allowing more users to access the data that's been loaded into the platform. In this session we'll cover: 1. Introduction to Azure Synapse Analytics Components 2. Azure Synapse Analytics Architecture Overview (querying-focused) 3. Best Practices & Tips for Querying data from Azure Synapse Analytics</p> <p>Tags: #Azure #Synapse #Analytics #Loading #Polybase</p>	Intermediate

Break Out	Erland Sommarskog	Development	Don't Bite Off More than You Can Chew - Take it in Cunks	Any SQL programmer with some experience knows that loops are bad and that you should work with all data at once in set-based statements. However, you may have experienced situations where this strategy did not work out well and you ran into problems like out-growing the transaction log or blocking other users. You can solve this by operating on the data in chunks. Implementing chunking is not that difficult, but there are still pitfalls you can run into. In this session, I will discuss in what situations you may want to use chunking. I will give some best practices for how to implement chunking for good performance and I will highlight some things you need keep in mind, for instance, recovering from interruptions. I will also look at how you can use chunking for error handling: You want process many rows and you want the good rows to succeed even if some rows yield errors – something you cannot do in a single set-based statement.	Advanced
Break Out	Erik Darling	Data Administration	Defeating Parameter Sniffing With Dynamic SQL	Parameter sniffing can be a huge performance problem, and fixing it can feel like an exercise in futility. If you're ready to go beyond recompiles and unknowns and solve problems for real, come to this session. I'll teach you techniques to fix parameter sniffing problems for good. #SQLServer #Performance #Query #Index #Query-Plan #Tuning	Intermediate
Break Out	Leonard Lobel	Development	Building Event-Driven Microservices with the Azure Cosmos DB Change Feed	The change feed in Azure Cosmos DB is one of the most overlooked features of Microsoft's globally distributed, massively scalable, multi-model database service. Similar to the transaction log of a relational database, the change feed gives you a continuous record of changes as they occur. It therefore serves as an excellent event source for a wide range of cloud-based microservices. In this demo-filled session aimed at developers and data professionals alike, you'll learn how to leverage the change feed and build event-driven microservices with Azure Functions to achieve replication, denormalization, notifications, materialized views, and data archival.	Expert
Break Out	Prashant G Bhoyar	Industry Solution	Applied AI Practical Use Case : Content Classification of Network drives	Most of the organizations have humongous file shares/network drives where the contents are not organized and labeled as per the organization's liking. This poses a high problem when organizations want to move the legacy contents in the public cloud. In this demo-driven talk, we will cover the case study of a recent file share migration project we did where we used Microsoft Cognitive Services (Applied AI) for intelligent content classification.	Intermediate

Break Out	Louis Imershein	Architecture	Securing SQL Server on Linux	Securing database workloads involves more than just configuring the database and a firewall, the operating system security is deeply involved. Linux has built-in security features that can be leveraged by SQL Server database admins to protect the entire environment. These include features around access control, advanced vulnerability analysis and patching, non-root container support, identity and key management, as well as cryptography and encryption. Most DBAs don't realize that Microsoft has built SQL Server on Linux to leverage all of these features. In this talk, we'll dive deeper into the interaction between SQL Server security and the security features of Linux and show you how to leverage Linux security features to architect hardened database environments. We also take a look at some specific live demos that show SQL Server interacting with these features.	Advanced
Break Out	Anupama Natarajan	Data Science (AI/ML)	Design Patterns for Machine Learning Projects	Machine Learning (ML) Design Patterns are proven techniques to solve common ML problems. In this session you will learn the techniques of applying these design patterns to solve real world problems. You will learn each design pattern with a problem statement followed by the solution provided by the pattern. We will be mainly looking at design patterns relating to Supervised, Unsupervised and Semi-supervised Machine Learning and each Design Pattern will be explained through a real-world use case. #MachineLearning #DesignPattern #AI #ML #bestpractices	Advanced
Break Out	Miguel Martinez & Peter Myers	Business Intelligence & Advanced Analytics	Monitor Your Data in Real-time with Microsoft Power BI	Why make decision based on data from yesterday when you can learn what is happening now? Wouldn't it be better to look forward instead of the rear-view mirror? Learn about the real-time data visualization capabilities for datasets, reports, and dashboards In this session you will learn: How to create dashboards, reports, and visuals with data that's always up to date. Different options for real-time datasets: streaming, push, et. How to monitor critical events with automatic page refres. #real-time #dataviz #streaming	Intermediate

Break Out	Vin Yu	Data Administration	What is Azure Arc Enabled Managed Instance?	<p>Have you checked out Azure Arc yet? Azure Arc enabled Managed Instance enables you to run SQL Server compatible workloads with the added management capabilities anywhere. In this session, we'll bring you quickly up to speed on Azure Arc enabled Managed Instance, discuss use cases, and demo this product. We'll also cover how you can get easily get started with this product.</p> <p>#AzureArc</p>	Intermediate
Break Out	Sasha Nosov	Architecture	Azure Arc Enabled SQL Server	<p>Even if you cannot migrate or modify your SQL Server application, you can leverage Azure. Whether your existing SQL server instances deployed to your private infrastructure, AWS or GCP, you can use Azure Arc to manage your global inventory, protect SQL Server instances with Azure Security Center or periodically assess and tune the health of your SQL Server configurations.</p> <p>#AzureArc</p>	Intermediate

Break Out	Jean-Yves Devant & Nikhil Patel	Data Administration	What is Azure Arc enabled PostgreSQL Hyperscale?	<p>You would like to modernize to the cloud but you can't migrate everything overnight? For regulatory/compliance reasons you need to keep some workloads on your premises while you move other applications to the cloud? You need a database engine that is able to scale dynamically, with no downtime to match the growth of your multi-tenant/SaaS workloads or your real-time analytics applications? You are already using the Postgres database engine or you are planning to migrate to it? You would like to use the same Postgres based solution both as a managed service in the cloud and in your data center? If you answered yes to any of these questions, join us in this session to learn about Azure Arc enabled PostgreSQL Hyperscale. This is a new hybrid Azure data service that runs on any physical infrastructure, on premises, at the edge or in the cloud (Azure, AWS, GCP). It is the same technology as the Azure Database for PostgreSQL Hyperscale (Citus) managed service and is now available on the infrastructure of your choice with Azure Arc. Like its sibling in Azure PaaS, Azure Arc enabled PostgreSQL Hyperscale uses the open source Citus extension to scale horizontally, transform Postgres into a distributed database by distributing your data and your queries across all the nodes in a cluster. In the cloud and on your own infrastructure, we have it for you. Let's connect.</p> <p>#Azure #PostgreSQL #AzureDBPostgres</p>	Intermediate
Break Out	Amit Banerjee	Data Administration	SQL Server Licensing: Demystified	<p>Learn about the new licensing benefits that were introduced with Software Assurance and how it can help you drive lower TCO for SQL Server deployments. In this session you will also understand how you can derive additional benefits by leveraging Azure SQL as an extension of your existing datacenter using Azure Hybrid Benefit.</p> <p>#AzureArc</p>	Advanced

Break Out	Bob Ward	Data Administration	Inside Waits, Latches, and Spinlocks Returns	<p>This session marks the return of a popular session dive into the internals of waits in SQL Server including latches and spinlocks. In this session, you will learn how SQL Server implements waits, how you can monitor and troubleshoot waits, and a deep dive into specific common wait types. This session will include new wait types specific to Azure SQL. The session will include plenty of demos and back by popular demand the use of the Windows Debugger to peek inside how waits are truly implemented in SQL Server.</p> <p>#sqlserver #azuresql</p>	Brain-Hurting
Break Out	Bob Ward	Data Administration	Inside SQL Server on Kubernetes	<p>Containers are the new "virtual machines" and one of the hottest technologies in the industry. While containers compliment virtual machines and provide consistency, portability, performance, and availability, how do you deploy containers at scale? For many, Kubernetes, an open-source system for automating deployment, scaling, and management of containerized applications, is the answer. In this session, we will dive deep into the Kubernetes platform and how SQL Server containers can be deployed, managed, and configured. You will see how SQL Server can take advantage of Kubernetes for "built-in" high availability, easy to deploy patching, and integration with Always On Availability groups. We will discuss deploying SQL Server in various popular Kubernetes distributions including but not limited to Azure Kubernetes Service, OpenShift, and kubeadm. This session assumes a fundamental knowledge of both SQL Server and containers. We will move fast, and go deep using demos as a method to learn SQL Server and k8s.</p> <p>#sqlserver</p>	Advanced

Break Out	Daniel Coelho & Rahul Ajmera	Data Science (AI/ML)	Effective Data Engineering and Data Science on SQL Server Big Data Clusters	<p>We are on a mission to show you how to effectively use SQL Server BDC to deliver Data Engineering and Data Science use cases. Let's go through an end-to-end implementation focusing on most used big data engineering patterns and how to use all BDC components in conjunction to achieve a sound solution architecture. We will double down on how to effectively enable and execute Data Science at scale while keeping things organized and secure. At the end, you will walk away with all the generalizable artifacts do make it your own and solve your organization Big Data modernization challenges.</p> <p>#sqlserver #BigDataClusters</p>	Intermediate
Break Out	Mihaela Blendea	Architecture	Gain insights with SQL Server Big Data Clusters on Red Hat OpenShift	<p>Big Data Clusters (BDC) is a set of capabilities introduced in SQL Server 2019 to help achieve data-driven business insights from ever increasing sources and amounts of data. With BDC, organizations can run containerized Apache Spark and Hadoop Data File System (HDFS) in an integrated matter side by side with SQL Server, in a single, secure and unified data platform. BDC requires Linux containers and Kubernetes, and recently Red Hat OpenShift was added as a commercially supported Kubernetes platform for BDC. OpenShift helps achieve the desired agility, consistency, flexibility, scalability, portability, etc. for data science and application development workflows. In this session, you will learn about BDC on OpenShift architecture, key capabilities and use cases, best practices, and technical resources to get you up to speed.</p> <p>#SQLServer #BigDataClusters</p>	Advanced

Break Out	Tejas Shah	Data Administration	Deployment, High Availability and Performance guidance for SQL Server on Linux in Azure IaaS ecosystem	<p>SQL Server on Linux has been one of the fastest growing database on Linux platform over last couple of years. At the same time many customers are looking to migrate their SQL Server instances and SQL based applications to the cloud to take advantage of scalability, flexibility, advanced high availability and disaster recovery options, optimized licensing scenarios, access to Azure monitoring & Azure security and advisory technologies. If you are considering SQL Server on Linux in cloud, look no further than Azure. In this session you will learn about how to quickly deploy production grade SQL Server on Linux in Azure ecosystem, including new fully supported High Availability configuration. You'll learn how to configure AD authentication reliably and quickly with new and improved experience. You'll also see suggestions on how to get best performance out of SQL Server on Linux in Azure. We look forward to seeing you there!</p> <p>#SQLServeronLinux</p>	Intermediate
Break Out	Julie Koesmarno & Alan Yu	Development	Azure Data Studio Notebooks Power Hour	<p>Join this session to learn more about Notebooks in Azure Data Studio. We'll demo new features in notebooks in Azure Data Studio, through the different use case lenses. Learn how to author notebooks, convert + organize your scripts to more manageable and shareable notebooks, all the way to fun + productive tricks that you can try at home! DBAs, Data Engineers and all data professionals are welcomed.</p> <p>#AzureDataStudio, #Notebook, #AzureSQL, #ClientTooling</p>	Intermediate
Break Out	Drew Skwiers-Koballa & Udeesha Gautam	Development	SQL Database Projects	<p>Azure Data Studio has recently added project functionality, supporting both compatibility with SQL Server Data Tools (SSDT) and new features. With the growth of cross platform environments and DevOps for databases, now is a great time to learn about database development in Azure Data Studio. We will dive into the SQL Database Projects extension for Azure Data Studio through an end to end example with Azure SQL Edge as well as a review of the overall functionality.</p> <p>#AzureDataStudio, #AzureSQLEdge</p>	Beginner

Break Out	Vicky Harp & Ken Van Hying	Development/Database Administration	State of the SQL Tools	Learn about what's new and what's coming soon for SQL Server tools in this demo-heavy session from the SQL Server product team, highlighting Jupiter notebooks, sql database projects, and more. #AzureDataStudio #OSS	Beginner
Break Out	Amit Khandelwal	Development/Database Administration	SQL containers are ready, are you? - For DBA and Developers	In this session, learn how as a Data Engineer/DBA or a developer can you take your favorite SQL Server to the container world. In this session we will also show you how easy it is for you to implement SQL Server in containers and get started with it. We will talk about new features available, how you can implement the containers in a secure way for production environments, customize it as per your requirement, we will also talk about the scenarios that best suit for SQL containers. #containers #sqlcontainers	Intermediate
Break Out	Hannah Qin & Drew Skwiers-Koballa	Development/Database Administration	How to Become an Azure Data Studio Contributor	Learn how to make an impact on Azure Data Studio in this session, where we'll discuss how you can leverage the GitHub repository, extensibility points, and open-source community to maximize your value in the application. New features are added to Azure Data Studio every month, and the team is always listening for feedback. We will walk through the various ways the community is empowered to shape changes to the application. Users are invited to engage by voting on feature requests, discussing use cases, and testing the nightly Insiders build. The Azure Data Studio code base is both open and extensible – meaning that contributors are welcome to build an independent extension or submit changes to the core of the application. We will walk through the available extension APIs and newly improved extension generator, as well as demonstrate how to build the full application. Come to this session and become a part of the Azure Data Studio community. #AzureDataStudio #OSS	Beginner
Break Out	Kevin Farlee	Data Administration	New ways to keep your SQL databases on Azure VMs protected and available	Keeping your SQL database available and protected is a critical mission of any DBA. In this session, we will discuss new ways and best practices for keeping your SQL Server databases available on Azure VMs. #SQLHA	Intermediate

Break Out	Melony Qin	Data Administration	Administrating Big Data Clusters (BDC)	<p>SQL Server Big Data Clusters (BDC) is a cloud-native, platform-agnostic, open data platform for analytics at any scale orchestrated by Kubernetes, it unites SQL Server with Apache Spark to deliver the best data analytics and machine learning experience. Join this session if you are interested in the administration scenarios of the big data clusters, including tooling for monitoring, how to deploy and secure the environment, and to learn about the latest improvements.</p> <p>#SQLServer #BigDataClusters #Kubernetes</p>	Intermediate
Break Out	Vasiya Krishnan & Sourabh Agarwal	Data Administration	Real-time data analysis using Azure SQL Edge	<p>Join us to learn how Azure SQL Edge provides the flexibility and capabilities to develop IoT/Edge solutions that leverages machine learning & real-time analytics to drive continuous optimization and efficiency.</p> <p>#AzureSQLEdge #SQL #IoTolutions #AzureSQL #SQLServer #Edge</p>	Intermediate
Break Out	Rajesh Setlem & Mohammed Kabiruddin	Data Administration	Understanding your SQL Server readiness to migrate databases to Azure SQL using database assessment and migration tools	<p>In this session we will walk you through database assessment and migration tools that will assist in seamlessly migrating your SQL Server databases to Azure SQL. The session will guide you through migration process on how to determine readiness of your SQL Server instances, identify a recommended target, and complete the migration to Azure SQL.</p> <p>#AzureSQL</p>	Intermediate
Break Out	Raj Pochiraju & Mukesh Kumar	Data Administration	Database modernization best practices and lessons learned through customer engagements	<p>In this session you'll learn the real-world ins and outs of how we successfully migrated customers to Azure data platform. You will get a detailed view of Microsoft investments accelerating Azure SQL migrations, the migration lifecycle phases, and specific tools/ services help you in the migration journey.</p> <p>You will listen to real world customer migration case studies, the speciifc migration strategies adopted and best practices implemented.</p> <p>#AzureSQL #DMS #SQLMigration</p>	Intermediate

Break Out	Raj Pochiraju, Balmukund Lakhani & Ajay Jagannathan	Data Administration	App Modernization and Migration from End to end, using data migration tools and Azure SQL	<p>In this session, you'll learn about an end-to-end scenario that takes customers through the application and database migration process.</p> <p>This demo heavy session will help you understand the Microsoft recommendations, tools, & services that help you migrating your .NET application and databases to Azure.</p> <p>#AzureSQL #DMS #SQLMigration</p>	Advanced
Break Out	Balmukund Lakhani & Dimitri Furman	Data Administration	Fine Tuning Clouds: Azure SQL Database Performance Tuning Tips and Tricks	<p>Performance tuning might be seen as magic, and those who can tune might look like magicians. There are no “go-faster” buttons and in reality, performance tuning is more art than magic. In this session we share common performance issues seen by various customers while using Azure SQL Database, and possible ways to fix them. You will learn database tuning techniques to gain the best performance for your application. While there are many tips and tricks available for traditional SQL Server installations, there are additional tuning best practices specific to Azure SQL Database. This session covers efficient data loading, dealing with high resource consumption, practical usage of various DMVs and various diagnostic telemetry available to look back in time.</p> <p>#AzureSQL</p>	Advanced
Break Out	Borko Novakovic, Vladimir Ivanovic, Srdan Bozovic	Data Administration	Modernize your SQL applications with the recently enhanced version of Azure SQL Managed Instance	<p>Azure SQL Database Managed Instance offers the best combination of a fully managed cloud service in Azure and compatibility with SQL Server on-premises product. This makes it the recommended destination for SQL Server workloads migrating to Azure cloud.</p> <p>If you think about digital transformation in Azure, start exploiting the increased productivity and cost savings that the managed instance can provide to you.</p> <p>Managed Instance remains to be one of the strategic investments in Azure SQL product portfolio. Over the last year it has been greatly enhanced in terms of new programmability surface area and increased application compatibility, improved price-performance, networking security and compliance, and better migration and manageability experience.</p> <p>Come and hear about the best practices of migration to managed instance, the latest product enhancements, and get a sneak peek at the future product roadmap.</p> <p>#AzureSQL</p>	Intermediate

Break Out	Sanjay Mishra	Data Administration	The Future of Cloud Relational Databases	<p>Relational databases have been a key foundation for applications for over 40 years. How does the future look for relational databases, especially in the cloud? This session will also bust several myths that have come to surround relational databases in the last few years through demos.</p> <p>#AzureSQL</p>	Intermediate
Break Out	Arvind Shyamsundar	Development	DevOps for AzureSQL	<p>"But... it worked on my computer!" A good DevOps practice and implementation will help eliminate that unwelcome phrase! Come to this session to learn how to use tools and platforms like Visual Studio Code, Git, Azure DevOps, GitHub Actions for a modern, secure and easy-to-use database DevOps experience for Azure SQL. We will use demos to show and evaluate two approaches: first, a code-centric migration-based approach (using DbUp and GitHub Actions), and second, a more complex project/state-based approach (using Visual Studio SQL Projects and Azure DevOps and a locked-down Azure SQL DB) to understand some of the many possibilities for implementing DevOps for Azure SQL.</p> <p>#AzureSQL</p>	Intermediate
Break Out	Denzil Ribeiro	Data Administration	Azure SQL Hyperscale Deep Dive	<p>Azure SQL is a world class database and it's new tier Hyperscale enables separation of compute and storage and gives you the scale and availability needed in the cloud. In this session, we will go through the Hyperscale architecture and all its components such as Page servers, read scale replicase, Xlog as well as multiple cache layers. We will also cover analyzing performance and the diagnostics unique to hyperscale and discuss best practices to migrate to Hyperscale. We see Hyperscale as the future, come learn about it!</p> <p>#AzureSQL</p>	Advanced
Break Out	Ajay Jagannathan	Data Administration	Azure SQL: What to use when and product updates	<p>Come learn about the latest capabilities in the Azure SQL family (VM, SQL Managed Instance, SQL Database) in the past year, along with the latest "game changers" that Azure SQL brings to the table for organizations, including hyperscale, serverless, intelligence, and more.</p> <p>#AzureSQL</p>	Intermediate

Break Out	Emily Lisa & Shreya Verma	Data Administration	Azure SQL High Availability and Disaster Recovery	<p>If you've ever had to deploy an Availability Group, Failover Cluster Instance, or backups you'll appreciate learning about the availability architectures and disaster recovery solutions that Azure SQL provides automatically when you deploy Azure SQL Database or Azure SQL Managed Instance. You'll learn about the latest innovation in availability, including topics around Availability Zones and Monitoring. You'll also learn about options, recommendations, and tutorials for recovering from a disruptive event that could cause data loss.</p> <p>#AzureSQL</p>	Advanced
Break Out	Alain Dormehl & Mara Steiu	Data Administration	360-degree view of Azure SQL Monitoring	<p>Take a trip with the product group through all the monitoring options for Azure SQL. Learn more about some of the new upcoming features and capabilities in Azure to monitor Database workloads. You can come and see more on a new solution for gathering and ingesting telemetry, which allows you access to new telemetry for your database, with full granular control of the collection and ingesting of data. This session will help you understand all your options to adequately monitor your SQL workloads in Azure.</p> <p>#AzureSQL</p>	Intermediate
Break Out	Silvano Coriani	Development	Leverage Azure SQL Database for Hybrid Transactional and Analytical Processing (HTAP)	<p>In real world, applications can rarely be classified as "pure OLTP" or "pure OLAP", but more often they will be present a mix of these two workloads. Learn how Azure SQL Database features and capabilities can nicely cover customer requirements and data processing scenarios.</p> <p>#AzureSQL</p>	Intermediate
Break Out	Aditya Badramraju & Pam Lahoud	Data Administration	SQL Server in Azure Virtual Machines Reimagined	<p>Most of the enterprise workloads would want both control and Manageability of their databae solutions. For such scenarios, deploying SQL Server in VM is a best choice. By deploying SQL Server in Azure Virtual Machines customers can take advantage of many features around administration, high availability and performance. In this session we will look at different best practices,advantages we have when we deploy SQL Server on Azure Virtual Machines.</p> <p>#AzureSQL #SQLonVM #SQLIAAS</p>	Intermediate

Break Out	Andreas Wolter	Data Administration	Securing your data in Azure SQL Database	<p>In this session we will guide you through the steps on how to secure your databases. Where to start, what to use, and what to watch out for. While the focus will be on Azure SQL Database, Microsoft's PaaS offering, you will also learn about the few potential differences when it comes to Managed Instance or SQL Server in a VM (IaaS) and on-prem. The aim is to give first-time users or adopters a checklist to make sure all areas of data security are covered.</p> <p>#SQLSecurity #AzureSQL</p>	Beginner
Break Out	Rodrigo Souza	Business Intelligence & Advanced Analytics	Self-service analytics and BI using Azure Synapse Link for Azure Cosmos DB	<p>In this session, we will cover how you can build real time BI dashboards with deep granularity using Azure Synapse and Azure Cosmos DB. Learn how to enable no-ETL analytical processing on real time operational data with no performance impact on mission critical operational workloads.</p> <p>Azure Synapse Link for Azure Cosmos DB is Microsoft's cloud-native implementation of hybrid transactional/analytical processing (HTAP).</p> <p>#Azure #AzureCosmosDB</p>	Expert
Break Out	Andy Feldman	Development	Best practices for deploying Azure Cosmos DB in enterprise mission-critical scenarios	<p>In this session, we will cover the best practices for deploying a mission-critical Azure Cosmos DB workload. We will begin by walking through how to secure, stress test, and diagnose a pre-production environment ahead of your go-big launch. Next, we will discuss how to monitor livesite, as well as, further optimize the Azure Cosmos DB production environment using various knobs including autoscale, indexing, and high availability. By the end of the session – you will be able to safely and confidently go big with your Azure Cosmos DB workload – just like how existing Fortune 500 customers like Walmart, ExxonMobil, American Cancer Society, ASOS and Symantec have done before.</p> <p>#Azure #AzureCosmosDB</p>	Advanced

Break Out	Mark Brown	Development	Data modelling and partitioning in Azure Cosmos DB: What every relational database user needs to know	<p>If you are new to Azure Cosmos DB and have only ever used relational databases, this session is for you! Azure Cosmos DB is unique in that a database that it can provide the same exact performance whether it is Megabytes in size or Petabytes in size. But the only way to get that kind of performance is understand how to model and partition data. Even if you're not trying to design a database with ultimate scale, you may have other questions How should I structure my data? When should I co-locate data in a single container? Should I de-normalize or normalize properties? What's the best partition key for my model?</p> <p>In this demo-filled session, we discuss the strategies and thought process one should adopt for modeling and partitioning data effectively in Azure Cosmos DB. Using a real-world example, we explore Azure Cosmos DB key concepts—request units (RU), partitioning, and data modeling—and how their understanding guides the path to a data model that yields best performance and scalability. If you're familiar with relational databases, and want to dive into the non-relational world, this is the session for you.</p>	#Azure #AzureCosmosDB	Advanced
Break Out	Kyle Weller	Business Intelligence & Advanced Analytics	How Delta Lake with Azure Databricks can accelerate your big data workloads	Delta Lake empowers you to build reliable Data Lakes at scale. Come learn how you can leverage the new and advanced features of Delta Lake on Azure Databricks to easily transform your big data analytics and machine learning workloads. We will discuss Apache Spark optimizations like file compaction, Z-Order partitioning, schema evolution, unified batch and streaming, time travel, and how you can set expectations on data quality.	#AzureDatabricks #Spark #BigData #Analytics	Advanced

Break Out	Arindam Chatterjee	Business Intelligence & Advanced Analytics	Migrating & Operating Business Critical Open Source Analytics Applications with Azure HDInsight	<p>The world today runs on data. And enterprises need to gain lightning fast insights from increasingly sensitive data while meeting the strictest security & compliance requirements. And they need to do so while controlling costs and ensuring high levels of productivity amongst their employees. Azure customers can do just that and more with Azure HDInsight – now with new cost control, performance and enterprise grade security features all based upon Apache analytics frameworks built and fully supported by Microsoft. Learn how to use the Migrate, Optimize and Modernize strategy to easily migrate your big data applications from on-premises and transform them to run optimally on Azure.</p> <p>#HDInsight #BigData #Spark #Analytics</p>	Intermediate
Break Out	Wee Hyong Tok & Nellie Gustafsson	Data Science (AI/ML)	How Data Engineers and Data Scientist can get started with Azure Synapse Analytics, Azure Data Factory and Azure Machine Learning	<p>Data is the fuel of AI. Data engineers and data scientists spend significant time performing data preparation, and feature engineering. Often, the data that sits in diverse data sources spread across the organization, and lots of time is spent on figuring out data acquisition. In this session, you will learn how Azure Data Factory and Azure Synapse Analytics can empower data engineers and data scientists to connect to various data sources, perform data transformation, and operationalize their data models, all at elastic scale. You will learn how data scientists can leverage this new found superpower to get to the data they need to develop and train and deploy the machine learning models using Azure Machine Learning, orchestrate it using Data Factory, and scale their models using Azure Synapse.</p> <p>Don't miss this demo-packed sessions as you learn how Azure Data Factory, Azure Synapse, and Azure Machine Learning can come together to help you achieve more!</p> <p>#MachineLearning #AzureSynapse #AzureDataFactory #BigData #Analytics</p>	Intermediate

Break Out	Arnaud Comet	Business Intelligence & Advanced Analytics	Building a connected analytics hub with Azure Synapse Link	<p>In this session, we will cover how you can build real time BI dashboards with deep granularity using Azure Synapse and Azure Cosmos DB. Learn how to enable no-ETL analytical processing on real time operational data with no performance impact on mission critical operational workloads. Azure Synapse Link for Azure Cosmos DB is Microsoft's cloud-native implementation of hybrid transactional/analytical processing (HTAP).</p> <p>#Analytics #AzureSynapse #CosmosDB #BigData</p>	Intermediate
Break Out	Arun Ulagaratchagan & Amir Netz	Data Administration	Microsoft Power BI: Business Intelligence Strategy, Vision, and Roadmap Update	<p>Our main goal for Power BI is to help our customers drive a data culture – helping everyone make better decisions based on data. Join Microsoft's Business Intelligence leadership team to learn about our strategy, experience the newest capabilities available in Power BI, and see what is coming next.</p> <p>#PowerBI</p>	Intermediate
Break Out	Priya Sathy	Business Intelligence & Advanced Analytics	Building systems of insights for enterprise scale with Power BI and Azure	<p>Bringing business intelligence to thousands of users in the enterprise requires scaling for the largest datasets in the organization without compromising performance or security. Join us to learn about how Power BI deeply integrates with Azure Synapse Analytics and how you can easily leverage your investments in Azure to bring insights to those who need it the most.</p> <p>#PowerBI</p>	Intermediate

Break Out	Dmitri Korotkevitch	Architecture	Architecture and Maintenance of VLDBs (Part 1)	<p>The modern Data Management systems deal with enormous amount of information. The business and compliance requirements force them to collect and process large amount of data and retain them for prolonged amount of time. As result, the size of the databases grows exponentially year after year, which leads to database administration and performance challenges, and increases hardware cost. This, 2-part session, discusses various aspects of architecture, implementation and maintenance of Very Large SQL Server Databases (VLDB) working under mixed (OLTP + DW/Reporting) workload. It will: Provide an overview of three different technologies supported by SQL Server - row-based, column-based and In-Memory OLTP and discuss their benefits, downsides and the best use-cases</p> <p>Show how to design the systems in the way, that would benefit from all those technologies</p> <p>Discuss typical challenges related to VLDB maintenance, such as how to reduce database backup size and hardware cost, improve RTO, maintain indexes, among with several other topics.</p>	Architecture
Break Out	Dmitri Korotkevitch	Architecture	Architecture and Maintenance of VLDBs (Part 2)	<p>The modern Data Management systems deal with enormous amount of information. The business and compliance requirements force them to collect and process large amount of data and retain them for prolonged amount of time. As result, the size of the databases grows exponentially year after year, which leads to database administration and performance challenges, and increases hardware cost. This, 2-part session, discusses various aspects of architecture, implementation and maintenance of Very Large SQL Server Databases (VLDB) working under mixed (OLTP + DW/Reporting) workload. It will: Provide an overview of three different technologies supported by SQL Server - row-based, column-based and In-Memory OLTP and discuss their benefits, downsides and the best use-cases</p> <p>Show how to design the systems in the way, that would benefit from all those technologies</p> <p>Discuss typical challenges related to VLDB maintenance, such as how to reduce database backup size and hardware cost, improve RTO, maintain indexes, among with several other topics.</p>	Architecture