

Special Topics:

# Self-Driving Database Management Systems

## Autonomous Systems V

@Karthik R // 15-799 // Spring 2022

Lecture #22

# LAST CLASS

---

## Series on End-to-End Autonomous Systems

- UDO
- openGauss
- Database Learning
- Transferable Model for Learning



# TODAY'S AGENDA

---



- Database Offerings at Microsoft
- Auto-Indexing Architecture
- Index Recommendations
- Experiments
- Parting Thoughts



# AZURE SQL PRODUCT-LINE



## SQL Server on Azure Virtual Machine

- Vanilla DB-on-VM offering
- Pay for virtual machine + software license
- Access to underlying OS



## Azure SQL Managed Instance

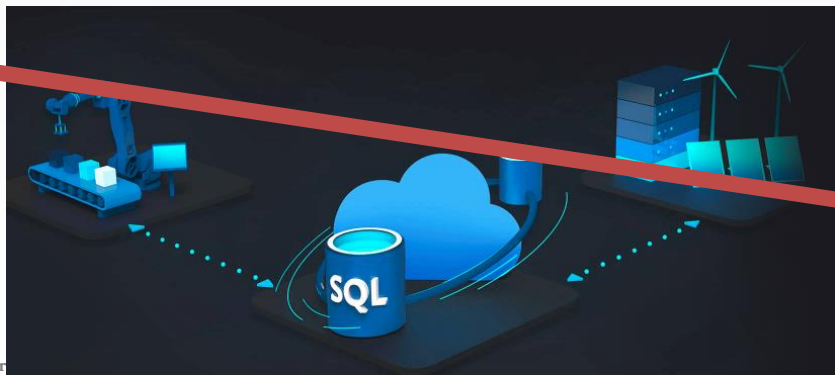
- Fully “managed” instance
- Pay for (compute, storage)
- Auto-scales
- No access to underlying box
- Auto-upgrades

# AZURE SQL PRODUCT-LINE



## Azure SQL Database

- Flagship offering
- All the features of SQL Managed Instances + “Intelligence”
- Currently available for storage < 100 TB
- Server-ful and serverless variants



## Azure SQL Edge

- Streaming, Time-series and in-database Machine Learning capabilities

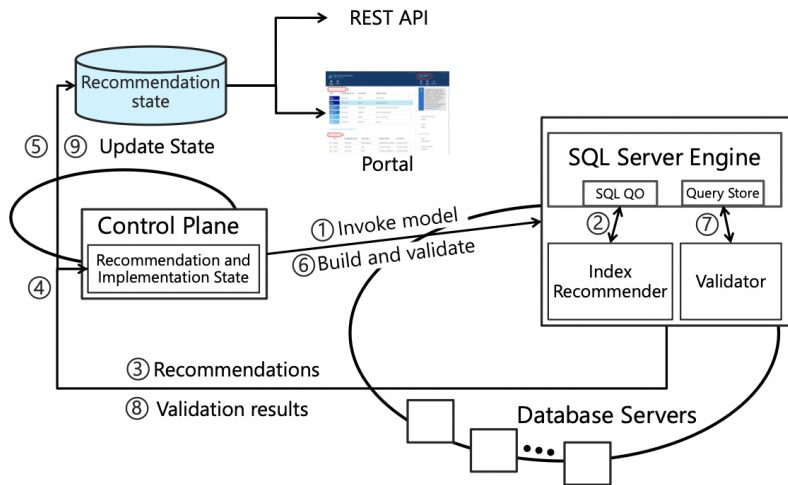
# CHALLENGES AT SCALE

---

Even if you could recommend the “best” indexes, how do you:

- Identify a representative workload?
  - Analyze prod instances ensuring QoS compliance?
  - Implement the indexes (the WHEN)?
  - Ensure no regressions?
  - Detect drifts in workload and data distributions?
- ... while still being profitable?

# AUTO-INDEXING ARCHITECTURE



**Figure 4: Conceptual architecture of auto-indexing service within a single Azure region.**

## Moving Parts:

### Central:

- Control Plane aka “The Brain”

### Per Database Instance:

- Index Recommender
- Validator

### Within SQL Server:

- Query Optimizer
- Query Store

### Everywhere:

- Telemetry

# CONTROL PLANE

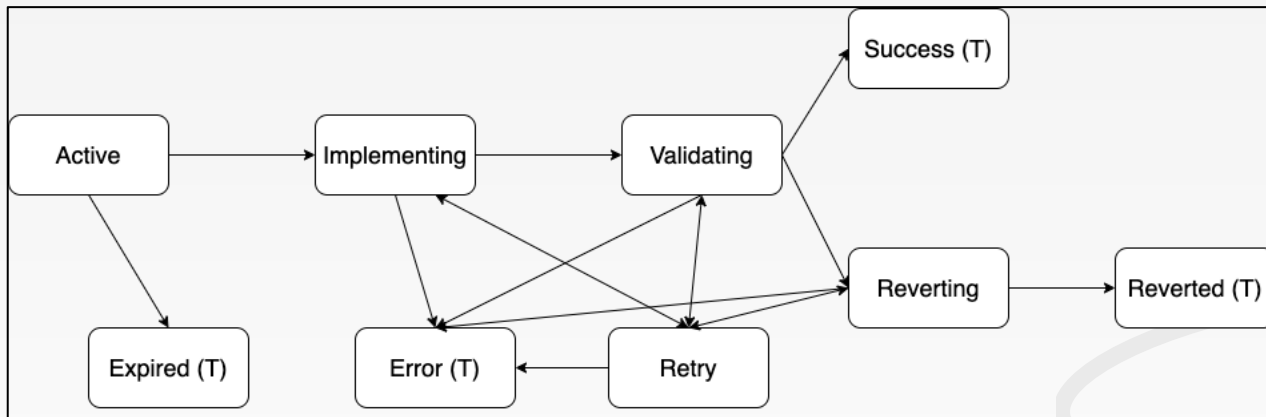
---

- Highly Available and Fault Tolerant service
- One auto-indexing service per Azure region
- Manages a persistent, highly available data store
- Collection of micro-services
  - Invoke Index Recommendation service
  - Implement recommendations
  - Validate recommendations
  - Detect anomalies in the system



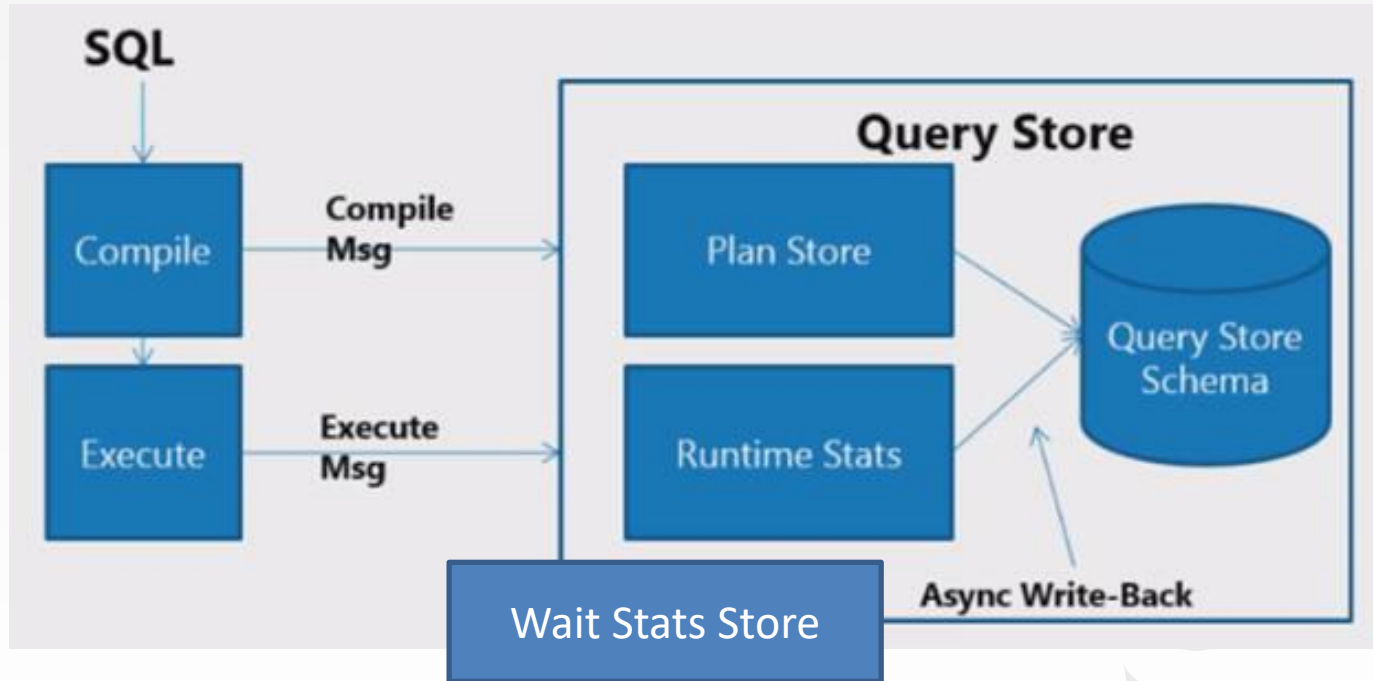


# CONTROL PLANE



Every index recommendation is persisted as a state machine.

# QUERY STORE



# QUERY STORE

---

Stores the following:

- **Plans:** Execution plans for a given query
  - Useful for forcing query plans and detecting impact of indexes
- **Runtime stats:** Resource consumption per query
  - Can answer questions like “Top N queries for resource A in the last X hrs”
- **General stats:** audit information
- **Wait stats:** Locking, resource limits, system degradation



# QUERY STORE

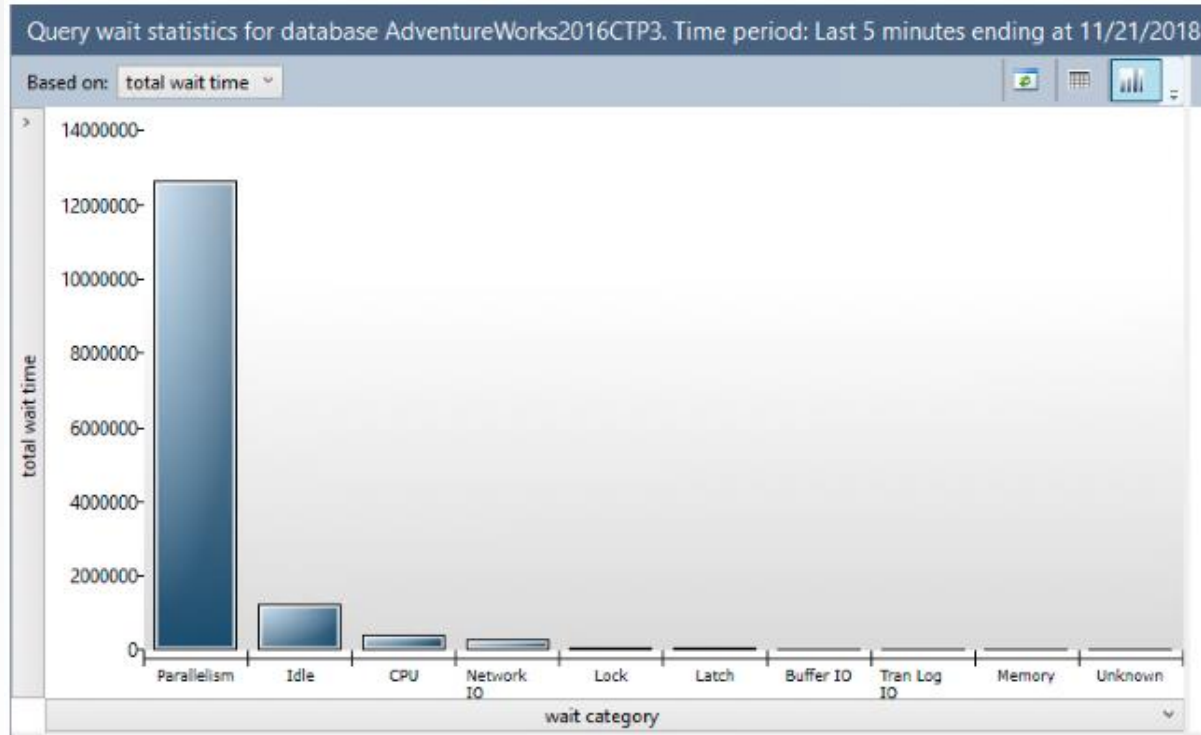
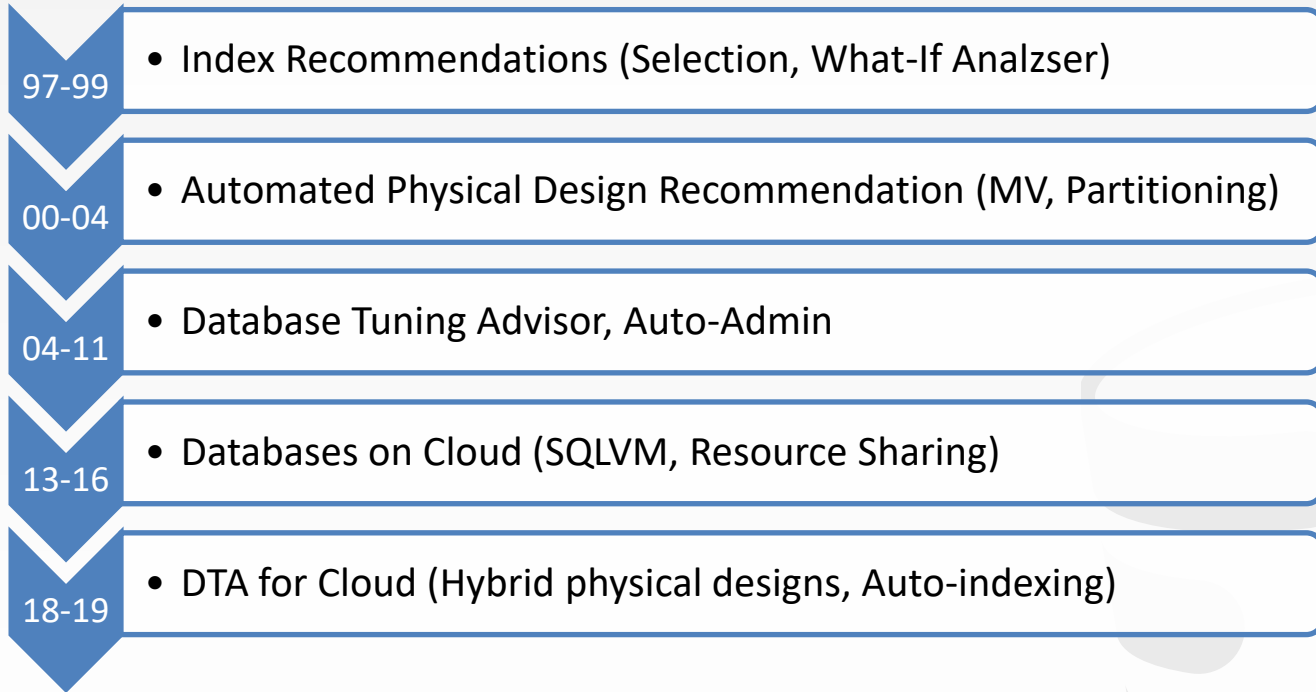


Fig: Query wait categories

# HISTORY OF INDEX RECOMMENDATIONS



# INDEX RECOMMENDER

**Create index**  
dbo.orders

+ Apply Discard View script

Recommended action	Status	Last update	Initial
Create index	Active	10/23/2018 8:14:47 AM	N/A

[Learn more](#)

**Estimated impact**

Impact	Medium
Disk space needed	10.00 MB

**Details**

Index name	_dta_index_orders_5_1285579618_K1_5_6
Index type	NONCLUSTERED
Schema	dbo
Table	orders
Index key columns	[O_OrderDate]
Included columns	[O_OrderDate], [O_OrderPriority]

**Execution plan**

```

SELECT LastAccessDate, DisplayName, Reputation
FROM StackOverflow2013.dbo.Users WHERE DisplayName LIKE 'Brent Ozar%';

```

Query 1: Query cost (relative to the batch): 100%

```

SELECT LastAccessDate, DisplayName, Reputation FROM StackOverflow2013.dbo.Users WHERE DisplayName LIKE 'Brent
Missing Index (Impact 84.0089): CREATE NONCLUSTERED INDEX [<Name of Missing Index, sysname,>] ON [dbo].[Users]

```

Execution plan details:

- SELECT (Cost: 0%)
- Parallelism (Gather Streams) (Cost: 3%)
- Clustered Index Scan (Clustered) [Users].[PK\_Users\_Id] (Cost: 97%)

**Figure 3: A snapshot of the UI where customers see the recommendation details.**

Two components:

- Missing Index Recommender
- Database Tuning Advisor

# INDEX RECOMMENDER

---

## Missing Indexes (MI) Recommendation

- Lightweight, always ON tool
- Analyzes the leaf nodes of plans (query-level)
- Suggests “good-to-have” indexes

```
SELECT column_a, column_b
FROM table_1
WHERE column_a < 10000
AND column_b < 10000
ORDER BY column_b, column_a
```

- Doesn't benefit JOINS, ORDER/GROUP BYs
- Does not account for index maintenance overheads
- Primarily memory resident



# MISSING INDEXES RECOMMENDATION

---

For each query:

- Step 1: Identify candidate columns
- Step 2: Leverage statistics to determine optimizer-estimated benefit on a per query/statement basis
- Step 3: Filter out candidates with few query executions
- Step 4: Assess the impact of candidates over a period of time
- Step 5: Merge of index candidates



# INDEX RECOMMENDER

---

## Database Tuning Advisor (DTA)

“DTA is a comprehensive physical design tool that given a workload  $W$ , finds the physical design that minimizes the optimizer estimated cost of  $W$ .”

Problem: How to identify a representative workload  $W$  automatically?

Leaving workload-capture ON 24\*7 is resource intensive!

# IDENTIFYING WORKLOADS

---

## How about leveraging the Query Store?

- Query are sampled before logging in Query Store (QS)
- Conditional, Loops, Variables etc. are not logged
- A single query isn't always the right granularity

## Solution:

- Query Text in QS is augmented with system metadata
- Batch definitions are pulled from plan cache of QS
- Queries incompatible with “What-If” API are rewritten

# DATABASE TUNING ADVISOR (DTA)

---

**Problem**: How to ensure that database is always online?

**Solution**:

- Resource governance for “system tasks” in a multi-tenant setup
- Low-priority locks for modifying hypothetical indexes
- Change to DTA to reduce number of invocations




# RESOURCE SHARING

---

## In a multi-tenant environment:

- Resource guarantees must be absolute
- Services must be cost-effective for provider
- Independent accountability from provider
- No assumptions about workloads, usage timelines
- Support for policies such as elasticity



# RESOURCE SHARING

---

## Solution

- Concept of an SQLVM
  - A logical construct with a promise of a set of resources for a tenant inside a database server (ie. VM)
  - Resources: CPU, I/O, Memory, (potentially network bandwidth, etc)
  - Isolates system/mgmt tasks from the tenant's resource reservation
- Independent metering
- Scheduling based on Largest Deficit First (LDF)
  - Can support elasticity, and heuristics such as max-ing provider's revenue

# DATABASE TUNING ADVISOR (DTA)

---

## Putting it all together,

- Service-ification – invoked by the control plane
  - Analyses the execution stats of top  $k$  queries over last  $N$  hours by resource  $X$
  - Recommends ONLY non-clustered B+ tree indexes
- (Remember that DTA also supports clustered indexes, MV, partitions)
- Considers JOINS, ORDER BY, GROUP BY columns
  - Works under imposed constraints (such as storage budget)

# DATABASE TUNING ADVISOR (DTA)

## Dropping Indexes:

- Manually created indexes for infrequent but imp. queries
- Query hints and forcing query plans
- Duplicate indexes; which one to drop?

SQL

```
CREATE PROCEDURE Sales.GetSalesOrderByCountry (@Country_region nvarchar(60))
AS
BEGIN
    SELECT *
    FROM Sales.SalesOrderHeader AS h, Sales.Customer AS c,
         Sales.SalesTerritory AS t
    WHERE h.CustomerID = c.CustomerID
         AND c.TerritoryID = t.TerritoryID
         AND CountryRegionCode = @Country_region
END;
```

```
sp_create_plan_guide
@name = N'Guide1',
@stmt = N'SELECT *FROM Sales.SalesOrderHeader AS h,
         Sales.Customer AS c,
         Sales.SalesTerritory AS t
         WHERE h.CustomerID = c.CustomerID
              AND c.TerritoryID = t.TerritoryID
              AND CountryRegionCode = @Country_region',
@type = N'OBJECT',
@module_or_batch = N'Sales.GetSalesOrderByCountry',
@params = NULL,
@hints = N'OPTION (OPTIMIZE FOR (@Country_region = N''US''))';
```

# INDEX IMPLEMENTATION

---

**Problem**: How to schedule indexing operations?

- Index creation is resource hungry
- Impacts concurrently running workloads

**Solution**:

- Resource governance for “system tasks” in a multi-tenant setup
- Forecasting periods of low activity



# INDEX VALIDATION

---

**Problem:** How to detect implemented indexes are *actually* good?

- Optimizer costs can be inaccurate
- Collected data can be noisy due to concurrency
- If the index is bad, what do you do?

**Solution:**

- Leverage ONLY logical metrics for goodness of query execution
- Consider only executions with a plan change
- Use of statistical measures to determine regression

# EXPERIMENTATION AT SCALE

---

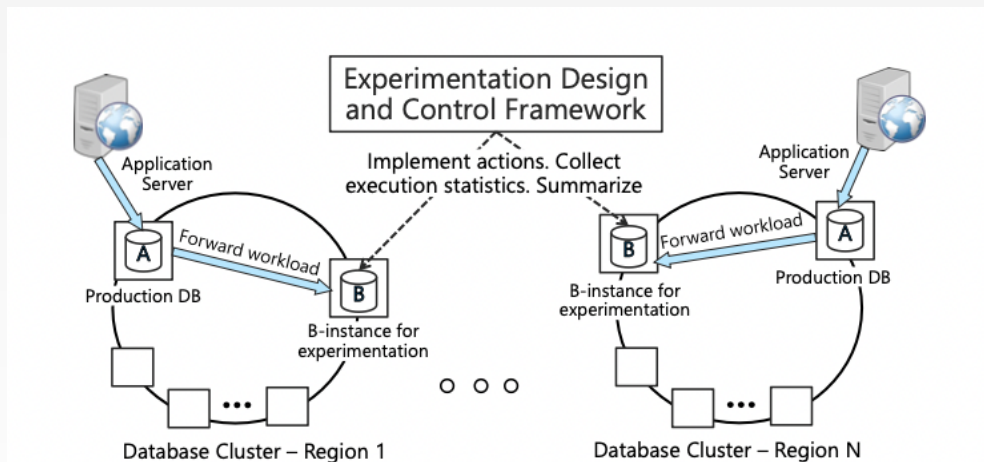
**Problem**: How are changes to the system tested?

- Changes cannot be run in production
- What constitutes a representative set of instances?

**Solution**: “B-instances” to the rescue!

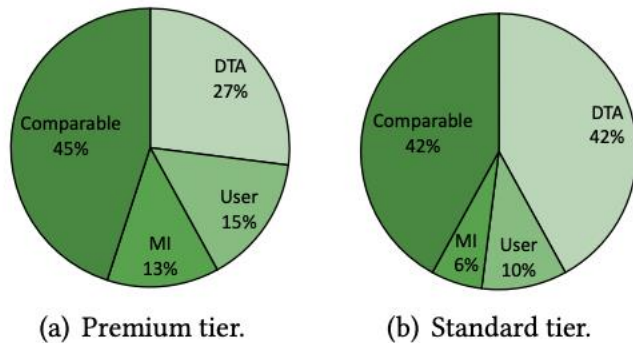
- Not a secondary replica
- An entirely different instance with a snapshot of the database
- Asynchronous to the primary, a best-effort clone
- Allows running different binaries, having different resources

# EXPERIMENTATION AT SCALE




**Figure 5: A conceptual architecture for experimentation using B-instances in Azure SQL Database.**

# EXPERIMENTATION AT SCALE



**Figure 6: Experimentation at scale with production databases in the premium and standard tier.**

# WHAT ABOUT PERF ISOLATION?




**Figure 5:** Differentiated service and optimizing for provider revenue.

# SOME NUMBERS...

---

As of October 2018,

- 250K recommendations for CREATE INDEX
- 3.4M recommendations to DROP INDEX
- Most databases reach a steady state
- DTA results in ~82% CPU time improvement, MI ~72%
- 11% regression rate



# USE CASE

---

## Tailored indexes for each out of 28K databases

SnelStart is a company from Netherlands that uses Azure and Azure SQL Database to run their software as a service. Over the last few years, SnelStart has worked closely with the SQL Server product team to leverage the Azure SQL Database platform to improve performance and reduce DevOps costs. In 2017, SnelStart received the Microsoft Country Partner Netherlands award, proving their heavy investment in Azure and collaboration with Microsoft.

SnelStart provides an invoicing and bookkeeping application to small and medium-sized businesses. By moving from desktop software to a hybrid software-as-a-service offering built on Azure, SnelStart has drastically decreased time to market, increased feature velocity and met new demands from their customers. By using the Azure SQL Database platform, SnelStart became a SaaS provider without incurring the major IT overhead that an infrastructure-as-a-service solution requires.

SnelStart uses a database per tenant architecture. A new database is provided for each business administration and each database starts off with the same schema. However, each of these thousands of customers have specific scenarios and specific queries. Before automatic tuning, it was infeasible to tune every database to its specific usage scenario. The result was over indexing from trying to optimize for every usage scenario in one fixed schema. Individual databases did not get the attention they needed to be tuned and this resulted in less than optimum performance for each database workload.



# PARTING THOUGHTS

---

## Contributions of the paper:

- Identification of a representative workload
- Guaranteeing resource isolation for tenant

## What's missing?

- Decision-making framework

## What's next for Microsoft?

- Physical design automation (knobs maybe?)





# NEXT CLASS

---

- Project Presentations
- Guest speakers next week!



