dbWatch report 1-2017

Azure database performance

Azure performance measurements February 2017

Marek Jablonski, CTO

dbWatch AS

Azure database performance

Introduction

The popular image of cloud services is both attractive and enticing. You get access to an unlimited amount of hardware at a reasonable price. The cloud provider ensures stability and redundancy, leaving you without any worries. Your systems can run seamlessly and at peak performance anywhere in the world, and It is natural to assume that the performance of cloud services is identical and stable, regardless of whether it's located in Asia, Europe or North America.

In this article, we take a closer look at the Microsoft Azure technology with our focus being on the performance of the Azure SQL databases, a service that runs on the cloud platform. With the Azure technology, an implementation of a database is just a few clicks away. Both the performance, measured in DTU (Database Transaction Unit) and the geographical location can be easily adapted to your needs. The systems can later be scaled up or down as needed, so you only pay for the resources you use. This concept seems to be the ideal solution for both small and large systems.

In the physical hardware world, a CPU unit normally has a stable and predictable performance, but in the virtual world, the resources are shared and can vary if the underlying hardware is for example overloaded. In a virtual system that you control yourself, it's easy to check the underlying infrastructure for overloads and performance bottlenecks. This is significantly harder in the Azure world were you simply are left assuming that the DTU resources you pay for are really there.

So what is a DTU really? How does this compare to the database systems you already have, both virtual and physical? How can you calculate what DTU strengths you need for the Azure database to take over your production database, if you plan to move into the cloud?

To answer these questions, we have set up a series of tests on three different Azure DTU Tiers at three different geographical locations. From the results we can conclude that there are wide variations in service comparing geographical locations. At the same time, we measured considerable variations in strength for the same Azure database Tier at various times, and we are surprised by the weak performance of Standard S3 and Basic Service Tiers.

In this article we describe the test procedure and the results of the selected Azure DTU Tiers, giving you some insight into the dark clouds of the Azure world.

The performance test procedure

The test results presented in the report are based on repeated and identical test sequences performed on Azure databases with three different DTU Tiers. Basic (5 DTU), Standard S3 (100 DTU) and Premium P2 (250 DTU). All three Azure databases are duplicated to three different locations: West Europe (WE), West Japan (WJ) and West US (WU).

All test sequences were performed once a day for a period of one week (7 test sequences per database per location). All tests were performed in December 2016. The test consisted of executing a locally defined Transact-SQL procedure that performed a specified number of insert, select, update and delete statements towards a locally defined table. After every test run the table was truncated. The procedure dynamically generated each row consisting of two values, a number and an 80-character text string.

Every test sequence was performed on 5K, 10K, 20K and 40K rows, and for Azure S3 and P2 the test runs were also performed on 80K and 160K rows.

Since all tests are based on executing a procedure inside each database, no network traffic was involved. Thus, the database location should not affect any of the results. There was no other activity on the databases when we conducted the tests. All tests and measurements are performed by the database monitoring, management and reporting tool dbWatch from dbWatch AS.

During every test sequence the following DTU resources were monitored:

- Data I/O
- Log write
- CPU

The percent usage for every resource is registered internally every 15 seconds in an Azure database and the values are available in the sys.dm_db_resource_stats dynamic view. This allows us to find what percentage of each resource type is used in relation to the max (100%) available in this DTU Tier. In Figure 1 we show a test sequence performed on an Azure S3 (100 DTU) Tier database. The top two graphs ("DTU % usage" and "DTU (CPU) % usage") display the DTU percentage used for the three DTU resources. The time interval for all graphs is 15 minutes and values are plotted every 15 seconds. In this test sequence, there are 6 test runs on 5K, 10K, 20K, 40K, 80K and 160K rows.

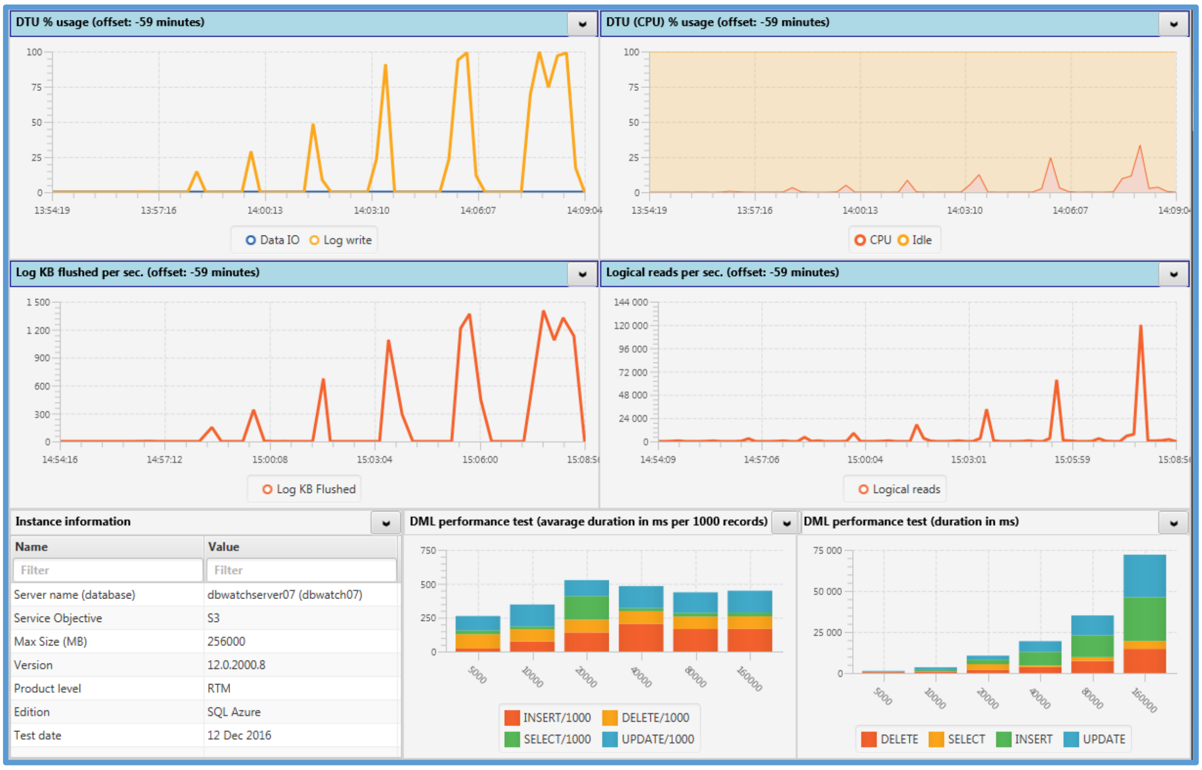

Figure 1: Example test sequence for an Azure S3 Tier database (100 DTU)

Simultaneously, while we monitor the values in dm_db_resource_stats view, to metrics from the sysperfinfo table are gathered:

- Log Bytes Flushed/sec
- Page lookups/sec

Values from these metrics are registered with the same time interval every 15 seconds and visualized in the two middle graphs in Figure 1. The graphs show "Log KB Flushed" per second, and "Logical reads" per second. The reason why the "Logical reads" was selected as the monitor metric is that it is not possible to measure directly the CPU consumption in an Azure database, and that the "Logical reads" reflects the pattern of CPU usage reasonably well. It’s also important that both metrics in the sysperfinfo table are available for both platforms, Azure and Microsoft SQL Server, which allows us to compare the two platforms.

In graphs related to writing transaction log data (redo), we observe a “double-top” (or split top) for some of the test runs. Those test runs usually lasts for a considerably longer time than 15 seconds (our time resolution). Every test run starts with an “INSERT” statement, followed by a “SELECT” and then “UPDATE” and “DELETE” operation. Since the “SELECT” statement is not creating any redo, it causes a dip between “INSERT” and “UPDATE” statement in a test runs that lasts for considerably longer time. For test runs that last less than 15 seconds all four operations are blended together.

The bottom part of Figure 1 shows the database information that is being tested, such as database-name, DTU strength, test date etc. In addition, there are two graphs plotting the results of the test sequence. The first graph shows average speed per 1000 rows per DML statement. (“INSERT”, “SELECT”, “UPDATE” and “DELETE”), while the last graph shows the total time per test run per DML statement. The results for all the test sequences will be presented and interpreted later in this article.

Background noise

Several tests/ measurements were performed to determine the “background noise” to find out how large a degree of influence it had on the test results. The measurements were performed both before and after the tests. In figure 2, 3 and 4 we present the “background noise” measurements for all 3 Azure database Tiers with the same time interval as the performance test sequences. The top two graphs show 0% activity measured in DTU unit. The two bottom graphs refer to the data gathered from the sysperfinfo table described in the section “The performance test procedure”.

All locations show a similar pattern with some minor activity every minute, and some larger activity every 5 minutes. This is mainly observed in the “Logical reads” per second graphs. Only a few sporadic and insignificant values were registered for “Log KB flushed”.

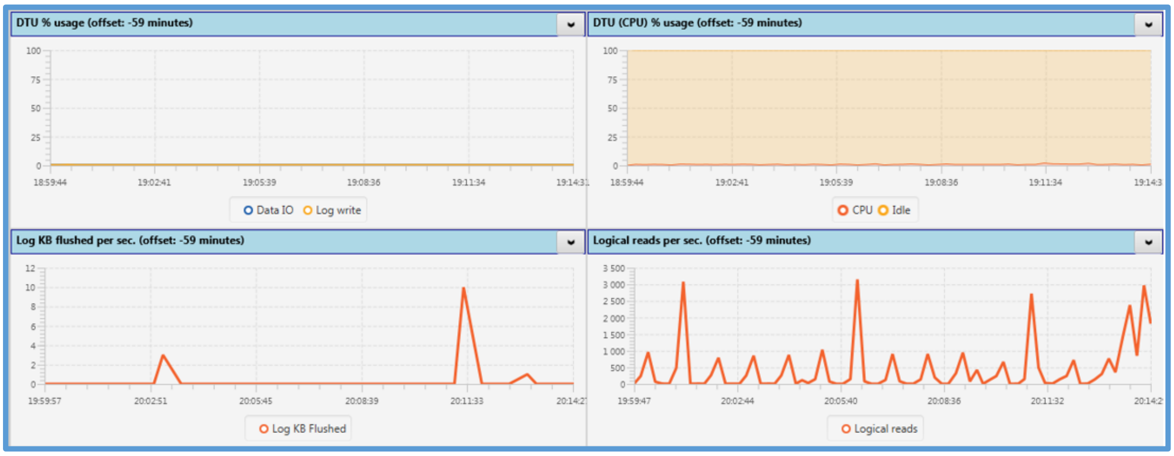

Figure 2. Background noise, Azure Basic (5 DTU), location: West Japan.

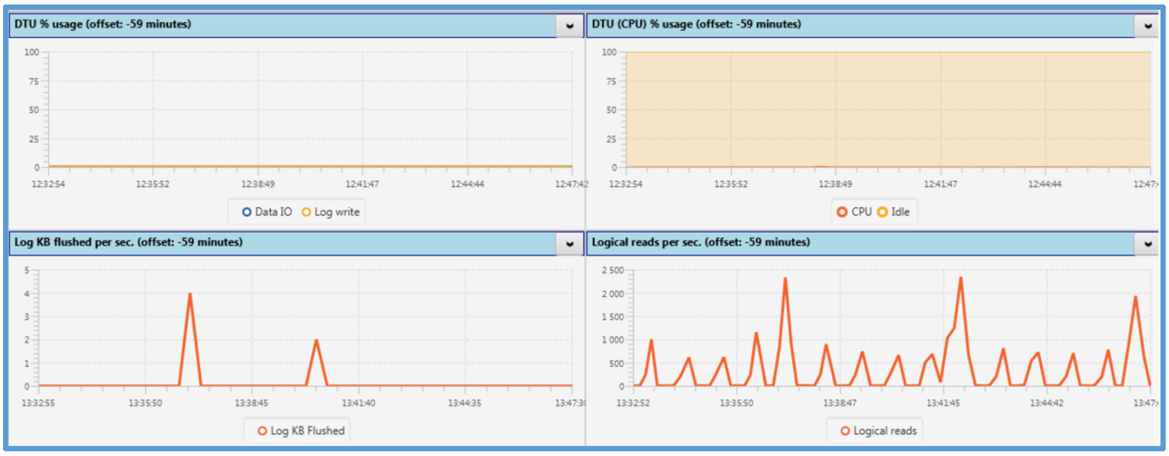

Figure 3. Background noise, Azure S3 (100 DTU), location: West US.

Figure 4. Background noise, Azure P2 (250 DTU), location: West Europe.

The largest values measured for “background noise” (for “Logical reads” per second graphs) are comparable to values registered during test runs performed on 5 000 rows, causing the highest impact on test results for the test runs with the lowest number of rows.

Test of Azure Basic (5 DTU) database

Based on Microsoft Azure documentation, the Azure Basic (5 DTU) is a good Tier for small databases, such as test and development databases, with only one active database user per system.

To create a good reference point to verify the strength and stability at this Tier, 3 databases at 3 different locations where tested: West Europe (WE), West Japan (WJ) and West US (WU). The test procedure was as described in “The performance test procedure”.

After few test sequences we noticed that the performance of the databases varied significantly between locations. In figure 5 we present a test sequence performed on an Azure database located in West Europe (WE). Independent of the number of rows of each test run, we measured the average execution time per 1000 rows (for all four DBL operations) to be 2.4 seconds.

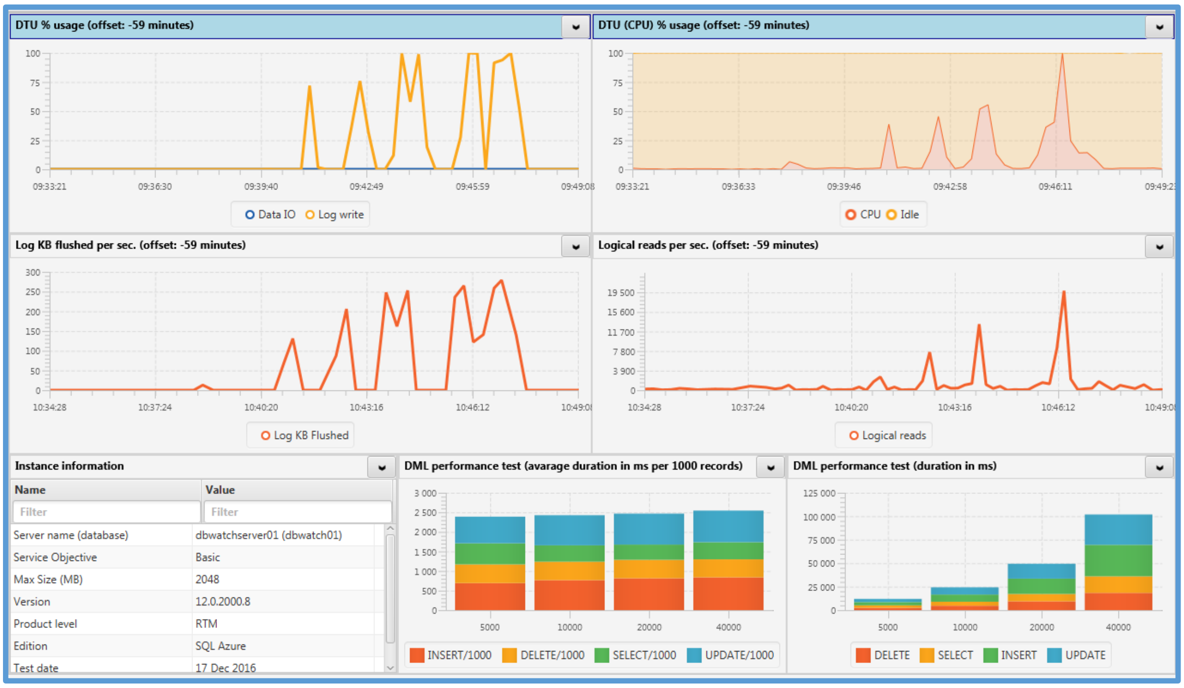

Figure 5. Test sequence on Azure Basic (5 DTU), location West Europe.

During the test run for 40 000 rows, we observe saturation for two DTU resources, the “Log write” and “CPU”. This corresponds to an approximate maximum value of 300 “Log KB flushed” per second and 20 000 “Logical reads” per second.

In Figure 6 we present a test sequence on a database located in West Japan (WJ). Here the average execution time for all 4 DML operations per 1000 rows is 3.0 seconds, which is 20% higher than the test performed on the database located in West Europe. When looking more closely at the statistics presented in Table 1, we see that the “SELECT” operation, on average, takes twice as long time in West Japan, than at the other two locations, West Europe and West US.

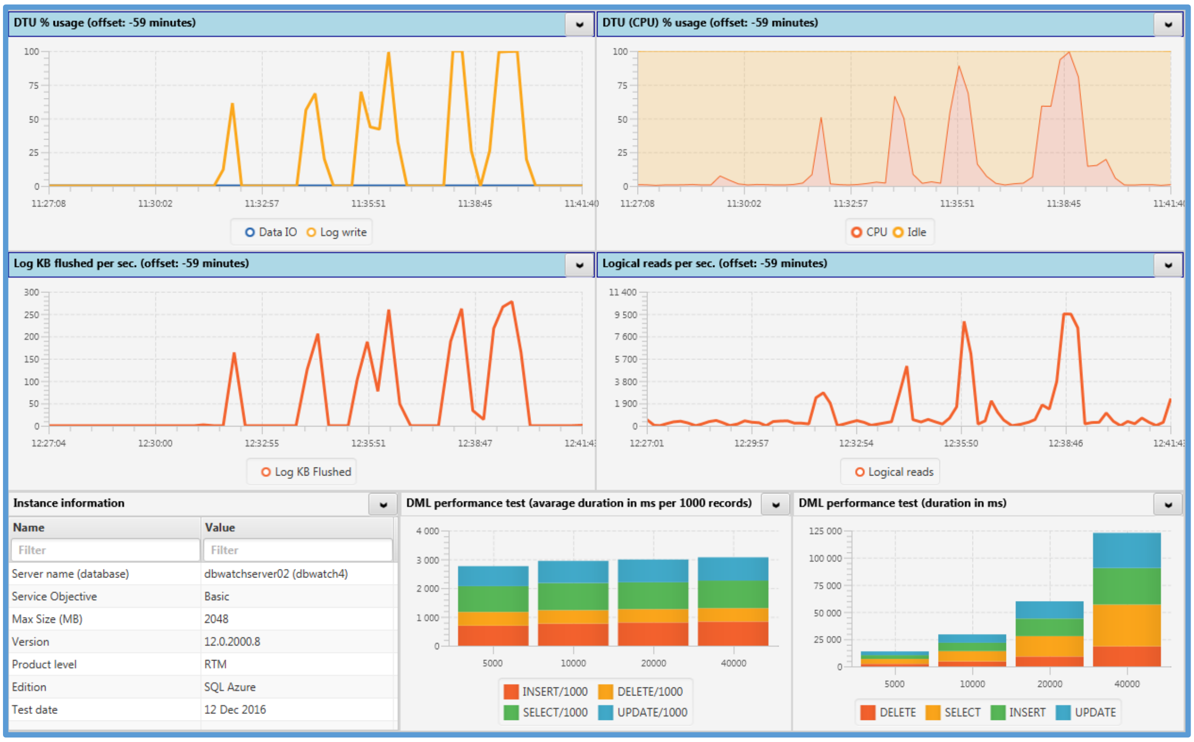

Figure 6. Test sequence on Azure Basic (5 DTU), location West Japan.

Test sequences on all three locations reach a maximum DTU resources for "Log write" at the same level, 300 "Log KB flushed" per second. While the maximum value for "Logical reads" per second registered in West Japan was half of what we measured at the two other locations (about 9 000 "Logical reads" per second).

In Table 1 we present the average values for all test sequences for each DML call per 1000 rows. All values are in milliseconds. Last column shows the total time for all 4 DML operations per 1000 rows. There is a clear difference in the "SELECT" operation on the database located in West Japan, which in test runs on 40 000 rows it consumed about 25 seconds more than the test runs in the other two locations.

Table 1. Average values (in milliseconds) for 7 test sequences for each DML call per 1000 rows on Azure Basic (5 DTU) databases.

Location	INSERT	SELECT	UPDATE	DELETE	Total (ms)
West Europe	773	410	759	473	2415
West Japan	781	980	775	490	3025
West US	772	432	758	477	2439

We performed 28 test runs at each location. All test values were surprisingly high (which means the performance is surprisingly low). The same test sequence performed for MS SQL Server running on a standard laptop, executes 100 times faster than on a 5 DTU Azure database. Originally we planned to include 80K and 160K rows test runs for Azure Basic Tier, but those test runs fully saturated "Log write" and "CPU" DTU resources not giving any meaningful results.

Test of Azure Standard S3 (100 DTU) databases.

Based on the Microsoft Azure documentation, the Azure Standard (Tier S0 - 10 DTU, S1 - 20 DTU, S2 - 50 DTU, and S3 - 100DTU) is well suited for cloud applications with a low to medium amount of I/O operation and with support for multiple concurrent queries.

We have chosen to run the test sequences on the highest Tier level, S3, at the same locations as for the test sequences conducted on Azure Basic Tier (5 DTU).

In Figure 7 we present a test sequence, from 17 December 2016, on an Azure S3 database located in West Europe. The test sequence consists of 6 test runs on 5K, 10K, 20K, 40K, 80K and 160K rows. All 7 test sequences show similar results, with average values presented in Table 2. Within all test sequences, we reach 100% DTU saturation during 160 000 rows test run for the "Log write" resource. For some test runs the saturation was also reached during the 80 000 rows test run. This corresponds to an approximate maximum value of 1400 "Log KB flushed" per second, almost 5 times higher than for the Basic Tier. The maximum DTU percentage value reached for CPU DTU resource was close to 40%.

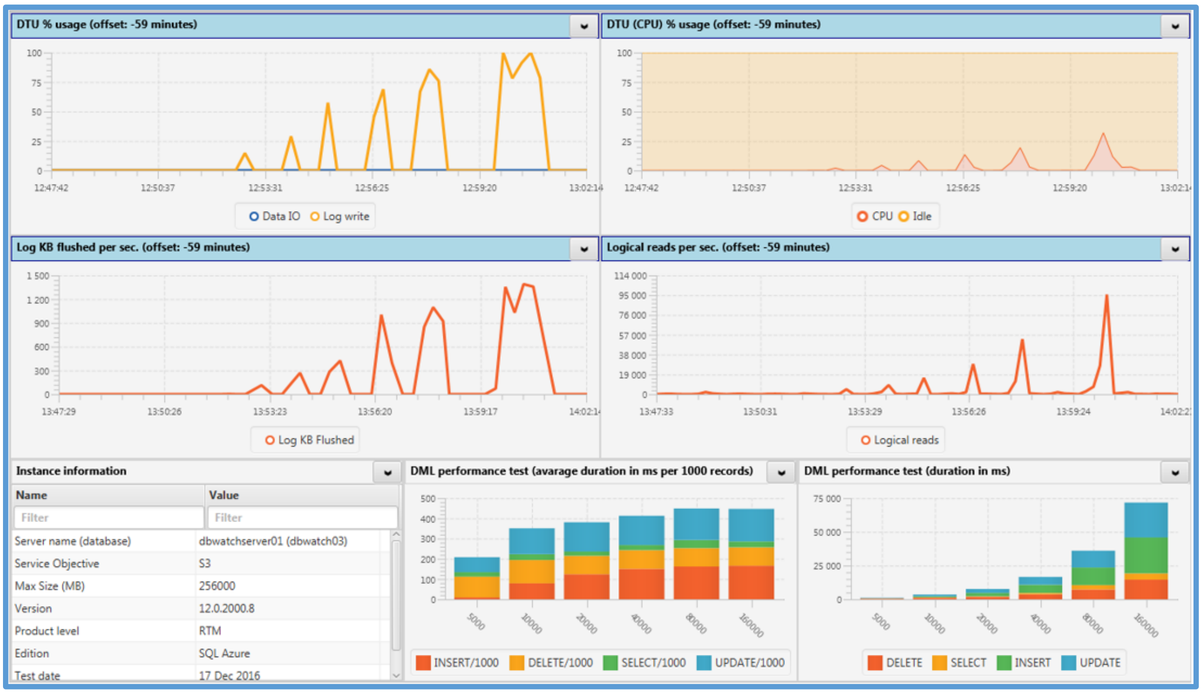

Figure 7. Test sequence for Azure S3 (100 DTU), location: West Europe.

In contrast to the results from the test sequences for the Azure Basic databases, we discover that during a 5 000 rows test runs, the"INSERT" operations, per 1000 rows, are significantly faster than for the remaining test runs. The execution time values for "INSERT" operations for the initial test run, varies between 9 and 50 milliseconds per 1000 rows. For test runs on higher values of rows, we see a gradual increase in execution time. For the test run on 10 000 rows, the execution time varies between 79 and 93 milliseconds, and for test runs with higher values of records the execution time stabilizes at around 160 milliseconds per 1000 rows. This may suggest some sort of caching is involved initially. An average execution time value for all measurements is 118 milliseconds (see Table 2).

Table 2. Average values (in milliseconds) for 7 test sequences for each DML operation per 1000 rows on Azure S3 (100 DTU) databases.

Location	INSERT	SELECT	UPDATE	DELETE	Total (ms)
West Europe	118	27	137	92	374
West Japan	175	101	122	95	494
West US	133	60	136	93	422

By examining the results for the same type of database (Azure S3) in a different location, we see that the differences between locations are even greater than we observed for Azure Basic databases. In Figure 8 we show a test sequence from 18 December 2016 on an Azure S3 (100 DTU) database located in West Japan.

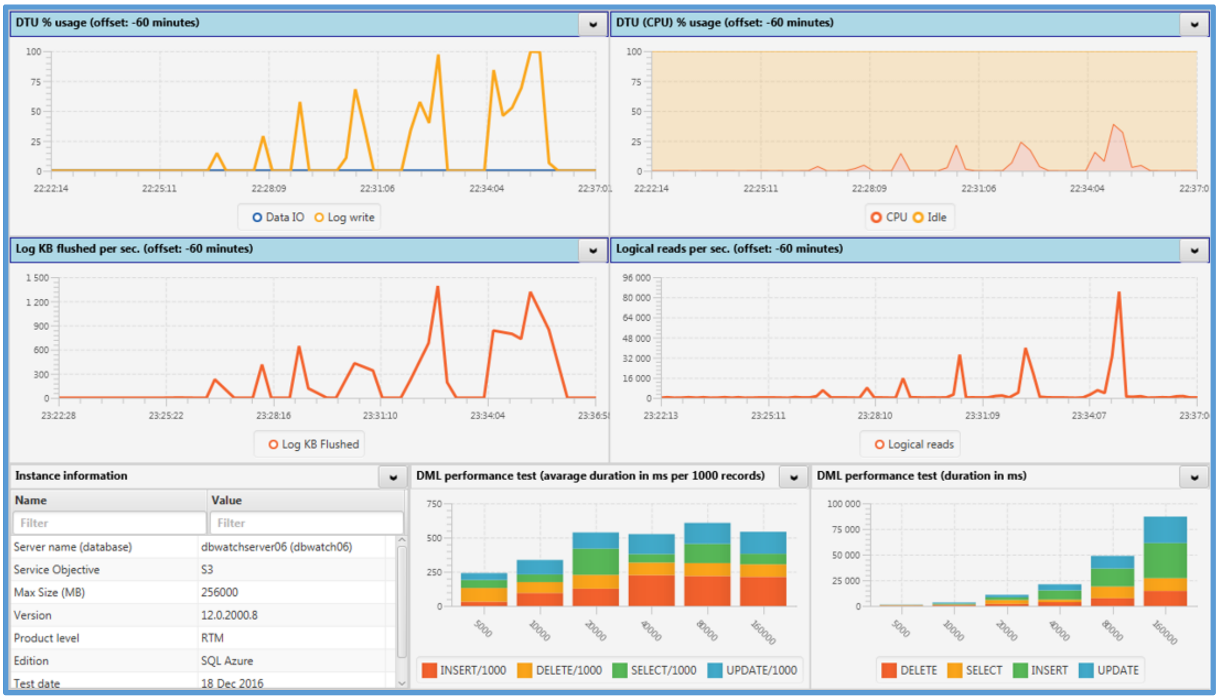

Figure 1: Test sequence on an Azure S3 (100 DTU), location: West Japan.

We observe a similar increase in execution time for “INSERT” operations in relation to the number of rows as seen in the other locations. For the initial test of 5 000 rows, the measurements vary between 28 and 59 milliseconds per 1000 rows. The execution time, for test runs on higher values of rows, increase to around 200 milliseconds, but the measurements vary significantly between the test runs. For the test runs on 40K, 80K and 160K rows the execution time values vary between 181 and 344 milliseconds per 1000 rows. The corresponding values for West Europe location were between 147 and 167 milliseconds per 1000 rows, which indicates significantly higher stability. An average execution time value for the “INSERT” operation, for all measurements, was 175 milliseconds for the database located in West Japan (see Table 2).

The “INSERT” operation of all the test runs showed great variation, but the largest difference in time execution was recorded for “SELECT” operation. In table 2 we see that the average time execution value for “SELECT” operation in West US is twice as high as for West Europe, while for West Japan the value is almost 4 times as high (last column in Table 3). In addition, we registered large variations in the measurements from West Japan and West US compared to West Europe. The Table 3 shows the shortest, longest and average time for the "SELECT" operation per 1000 rows. The last column shows the difference in operation time for the average time consumed between West Japan or West US and West Europe.

Table 3. Test measurement of "SELECT" operation for 1000 rows on Azure S3 (100 DTU). All values are in milliseconds.

Location	SELECT (min)	SELECT (max)	SELECT (avg)	Ratio
West Europe (WE)	20	47	27	100% WE
West Japan	52	378	101	374% WE
West US	20	174	60	222% WE

When we look at the two remaining operations, "UPDATE" and "DELETE", we see similar values across all three locations. In Table 4 and 5 we present the corresponding data from all test sequences performed on 10 000 rows and higher.

Table 4. Test measurements of "UPDATE" operation for 1000 rows on Azure S3 (100 DTU). All values are in milliseconds.

Location	UPDATE (min)	UPDATE (max)	UPDATE (avg)	Ratio
West Europe	128	162	137	112% WJ
West Japan (WJ)	76	180	122	100% WJ
West US	118	162	136	111% WJ

The "UPDATE" operation was performed fastest on the database located in West Japan, while the two other locations performed almost identically with a speed approximately 10% lower. However, we also see the largest variations in speed in the database located in West Japan.

Table 5. Test measurements of "DELETE" operation for 1000 rows on Azure S3 (100 DTU). All values are in milliseconds.

Location	DELETE (min)	DELETE (max)	DELETE (avg)	Ratio
West Europe (WE)	78	115	92	100% WE
West Japan	79	111	95	103% WE
West US	75	106	93	101% WE

The "DELETE" operations were almost identical across all locations, and the variation in measurements were less than the general uncertainty of the measurements.

Test of Azure P2 (250 DTU) databases

There are multiple strength Tiers in the Premium level of Azure databases. You can choose between 6 levels: P1 (125 DTU), P2 (250 DTU), P4 (500 DTU), P6 (1000 DTU), P11 (1750 DTU) and P15 (4000 DTU). Even though there is a large difference in strength at this level, Microsoft describes them all as well suited for databases with high transaction volumes, that demands a high I/O performance while supporting many users.

We have chosen to conduct all the test sequences on the Azure databases at strength Tier P2 (250 DTU) in the same regions as our test sequences for S3 and Basic Tiers.

In Figure 9 we present a test sequence performed on the database located in West Europe. None of the test runs could saturate any of the DTU resources. The highest percentage value of DTU usage was measured for “Log write” at about 45%, while for the “CPU” resource the maximum value was measured at around 25% DTU usage. This corresponds to approximately 3500 “Log KB flushed” per second, and 130 000 “Logical reads” per second.

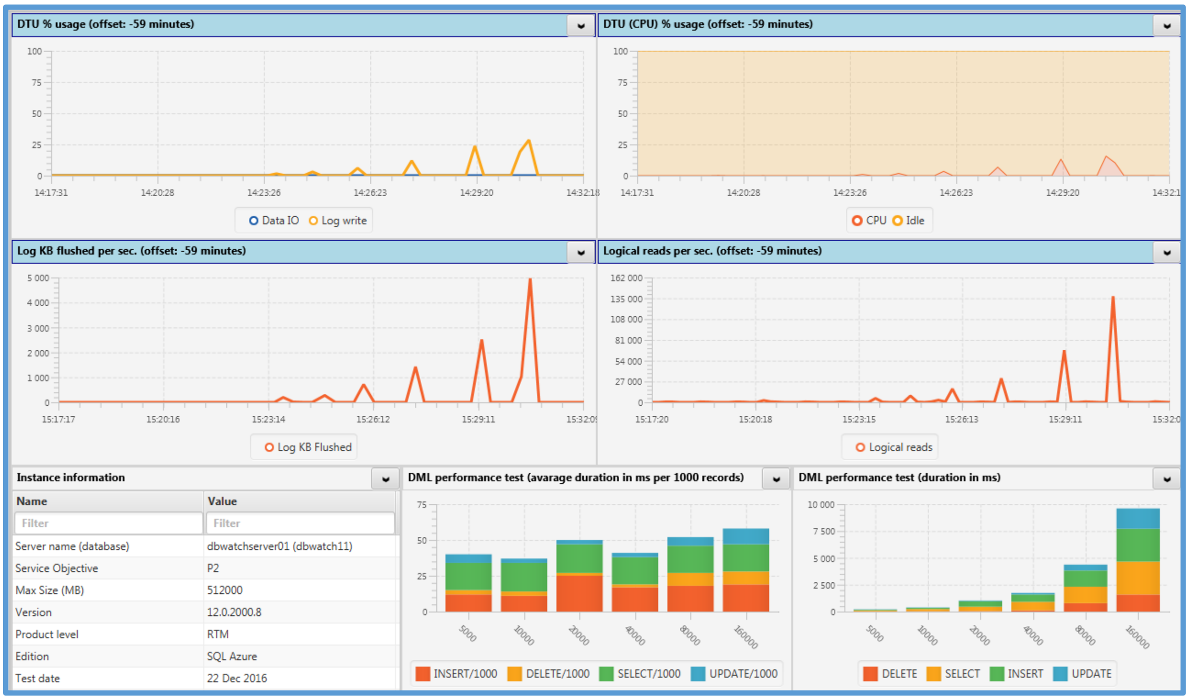

Figure 9. Test sequence on Azure Premium P2 (250 DTU), location West Europe.

The test results from the database located in West Japan shows significant higher execution values for the test runs. In Figure 10 we show one of the test sequences for this location performed on 27 December 2016. Although the percentage values measured for the DTU resources are relatively similar compared to the database located in West Europe, the maximum values, measured for “Log KB Flushed” and “Logical reads” per second are respectively 35 % and 20% lower.

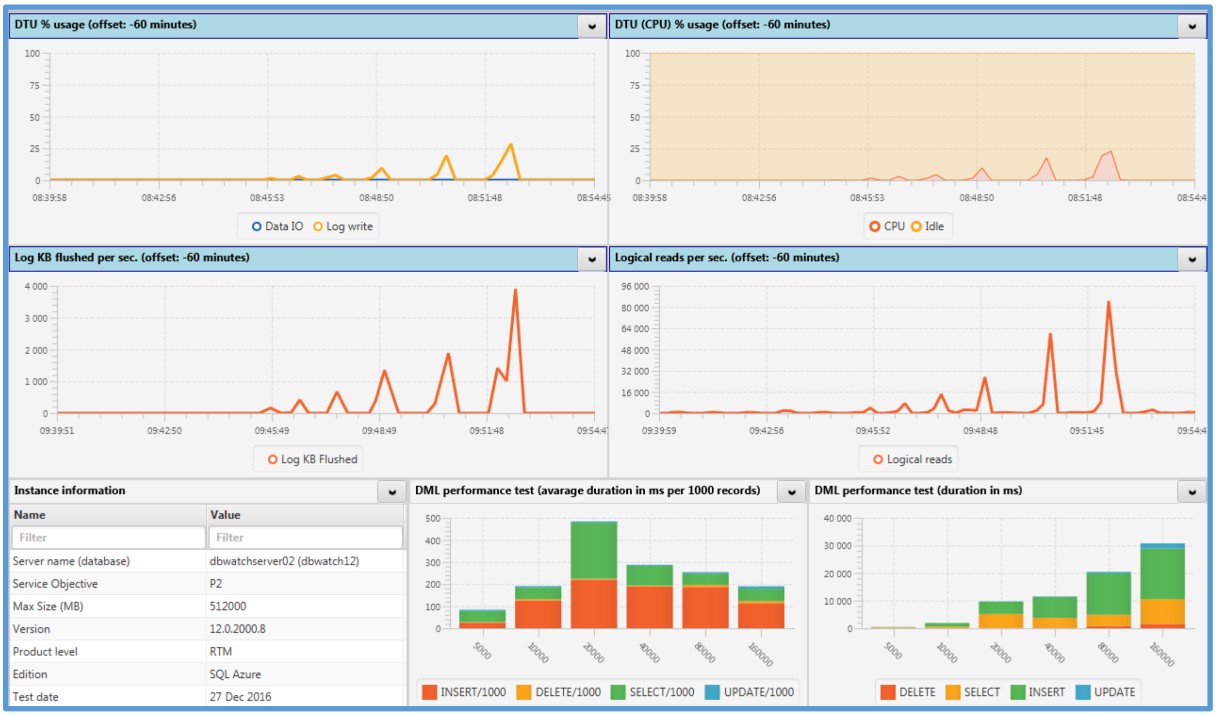

Figure 10. Test sequence on Azure Premium P2 (250 DTU), location: West Japan.

The difference is clearly visible when we compare the values for average execution time for all DML operations per 1000 rows. In Table 6 we see that there are significant differences between the execution times for "INSERT" and "SELECT" operations for all three locations. The execution time in West Europe is 5 times faster than in West Japan and 4 times faster than in West US.

Table 6. Average values in milliseconds for 7 test sequences for each DML call per 1000 rows on Azure Premium P2 (250 DTU) databases.

Location	INSERT	SELECT	UPDATE	DELETE	Total (ms)
West Europe	15	21	5	5	46
West Japan	116	101	7	6	230
West US	90	58	12	12	173

In addition to large differences in execution time between locations, we registered great variations in individual measurements for each site, and considerably more in West Japan and West US, than in West Europe. In Table 7 and 8 we show the shortest, longest and average time for "INSERT" and "SELECT" operations per 1000 rows. The last column shows the ratio between locations compared to the lowest average time value.

Table 7. Test measurements of "INSERT" operation per 1000 rows for Azure Premium P2 (250 DTU). All values are in milliseconds.

Location	INSERT (min)	INSERT (max)	INSERT (avg)	Ratio
West Europe (WE)	8	42	15	100% WE
West Japan	21	278	116	773% WE
West US	9	220	90	600% WE

Table 8. Test measurements of “SELECT” operation per 1000 rows for Azure Premium P2 (250 DTU). All values are in milliseconds.

Location	SELECT (min)	SELECT (max)	SELECT (avg)	Ratio
West Europe (WE)	18	43	21	100% WE
West Japan	50	356	101	481% WE
West US	17	169	58	276% WE

The “INSERT” operation take almost 8 times longer to execute in the database located in West Japan, and 6 times longer to execute in the database located in West US, compared to the database located in West Europe. For the “SELECT” operation the execution takes 5 times longer in West Japan and almost 3 times as long in West US compared to the database located in West Europe.

The fastest operations are by far the “UPDATE” and “DELETE” operations. For West Europe and West Japan the values are comparable, while for West US location both operations take about twice as long (see Table 6).

Comparative test of MS SQL Servers on physical and virtual environments

To get a better idea on how fast a MS SQL Server platform is compared to the Azure platform we have performed a series of test sequences on different MS SQL Servers used both internally and at some customer locations. The test sequences were conducted in the same fashion as they were on Azure platform using the dbWatch software suite. In Figure 11 we present one of the test results performed on a 2 year old HP Elitebook laptop running a MS SQL 2012 Server.

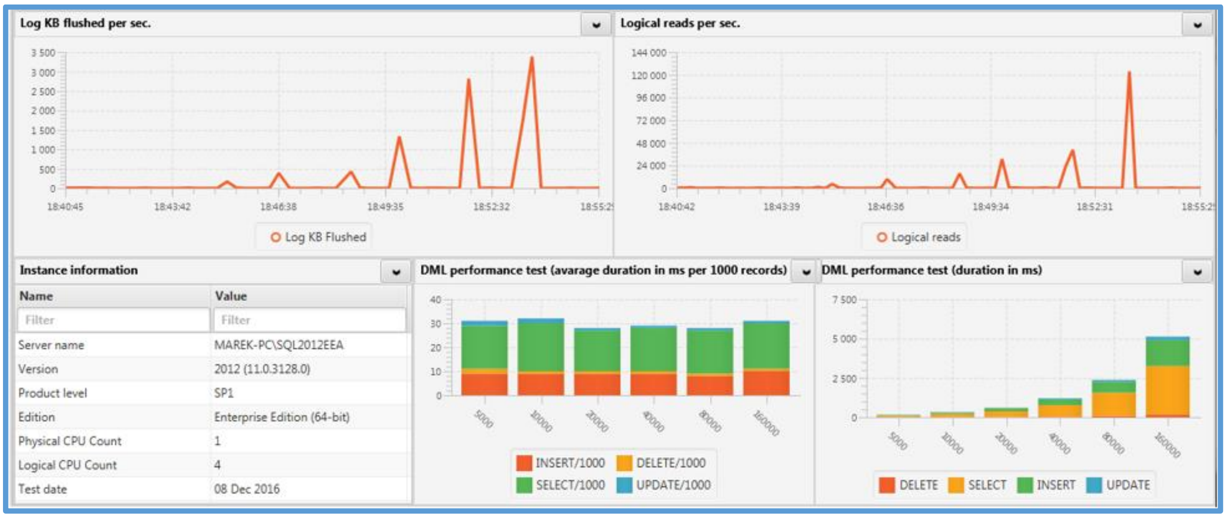

Figure 11. Test sequence on the HP Elitebook Laptop.

In addition, we have tested 2 physical and 3 virtual servers in 4 different customer database environments. In Table 9 we compare the values of average execution time for all DML operations per 1000 rows from all environments (the Azure and MS SQL Server platforms).

Table 9. Average values in milliseconds for each DML statement per 1000 rows for all Azure databases (Basic, S2, P2) and all databases running MS SQL Server on physical and virtual hardware.

Server/Location	INSERT	SELECT	UPDATE	DELETE	Total (ms)
Basic, WE	773	410	759	473	2415
Basic, WJ	781	980	775	490	3025
Basic, WU	772	432	758	477	2439
S3, WE	118	27	137	92	374
S3, WJ	175	101	122	95	494
S3, WU	133	60	136	93	422
P2, WE	15	21	5	5	46
P2, WJ	116	101	7	6	230
P2, WU	90	58	12	12	173
Laptop #1	9	18	2	1	30
Laptop #2	7	8	3	2	20
Server #1	5	10	1	1	17
Server #2	4	9	1	1	15
VM Server #1	6	11	2	1	20
VM Server #2	8	16	2	1	27
VM Server #3	8	11	2	1	22

The main criteria for the selected servers was a low load during the test period to minimize the impact on the test measurements. A short description of the physical and virtual environments is presented in Table 10.

Table 10. Technical information for the servers involved in the tests.

Server	Description
Laptop #1	HP EliteBook 8470p, Intel i5-3230M CPU 2.60 GHz
Laptop #2	HP Zbook B&O, Intel Xeon iS05M CPU 2.80 GHz
Server #1	HP ProLiant BL460c, Xeon(R) E5-2623 v3 CPU 3.00GHz
Server #2	HP ProLiant DL 380, Xeon E5-2637 v4 CPU 2.50GHz
VM Server #1	Intel(R) Xeon(R) E5-2640 v4 CPU 2.4gHz
VM Server #2	HP ProLiant BL460c, Xeon E5-2637 v3 CPU 3.50 GHz
VM Server #3	Intel Xeon E5-2699 v3 CPU 2.30 GHz

The results for both, the physical and virtual machines, shows that the execution time is significantly lower/faster than for the Azure Premium P2 Tier. Only the execution time for Azure Premium P2 database, located in West Europe, is comparable to results from physical and virtual machines, but still 50 % slower than the slowest servers. It's important to point out that the performance of both physical and virtual MS SQL Servers, show a much more even test results with much less variations compared to the test results from the Azure platform.

Conclusions

The Azure platform offers databases with different strength Tiers measured in DTU. The strength Tiers vary from Basic (5 DTU), Standard S0-S3 (10 – 100 DTU) up to Premium P1 - P15 (125 – 4000

DTU). In this article, we have focused on three strengths at the bottom half of the scale, with Basic (5 DTU), Standard S3 (100 DTU) and Premium P2 (250 DTU). Based on our measurements we have seen:

- Significant variations in performance between the different geographical locations
- Significant variations in performance for the same database at different times
- Surprisingly poor performance for the Standard S3 and Basic Service

Based on our measurements we can conclude that an Azure database with a specific DTU strength does not necessarily provide the same performance at all cloud locations. For the Azure Basic databases, there was a 20% performance difference between the three tested locations (West Europe, West Japan and West US). For Standard S3 databases the difference had grown to 25%, while on Azure Premium P2 databases the performance difference was a staggering 500%. When the same test takes 5 times as longer on one database than the other, given the same DTU strength, it's surprising that it can be sold as the same product. So in other words, ordering the same service at different locations you may end up with completely different underlying hardware. It addition there are large variations in the performance at the same location, which is indicative of an immature and unstable service performance.

It's also surprising how poor the execution time was on the Azure database, compared to a relatively old PC. On the Azure basic (5DTU) Tier the test took 100 times longer than on a 2 year old laptop, the HP Elitebook, while on the Azure Standard S3 it took 10 times longer. Even the test on Azure Premium P2 (250 DTU) took 50% longer. A similar test on 5 different database systems at client locations, performed twice as fast as the fastest Azure Premium P2 database.

In our opinion, the performance of these products is surprisingly weak compared to their pricing. The list price for an Azure Premium P2 database in our market (Norway) is just above US$10000/year, and that's not a full database server. It's only one database. The next level, a Premium P4 (500 DTU), will cost you around twice as much. Even the Azure Standard S3 will cost more than US$1500/year.

Based on collected statistics for all test sequences, we have calculated/estimated the maximum values of two resource on the Azure databases. In Table 11 we show the maximum values for "Logical reads" (per second) and "Log KB flushed" (per second) which the different Tiers of Azure (Basic, Standard S3 and Premium P2) can support.

Table 11. Maximum values supported per resource and per Azure database Tier.

Resource	Basic (5 DTU)	S3 (100 DTU)	P2 (250 DTU)
Maximum "Logical reads" per second	15 000 - 20 000	300 000 - 350 000	400 000 -700 000
Maximum "Log KB flushed" per second	300 - 350	1 300 - 1 500	10 000 - 15 000

Since we have only measured three Azure database Tiers, it is difficult to estimate equivalent values for other Azure database Tiers with a higher DTU strength. Nevertheless, we can see that the growth rate for the "Logical reads" per DTU is close to linear, which would indicate maximum values at 1.5 million "Logical reads" (per second) for P4 (500 DTU), and 15 million "Logical reads" (per second) for P15 (4000 DTU). The values for "Log KB flushed" (per second) per DTU strength, is far from linear, making it impossible to estimate any reasonable values for the other DTU Tiers.

There are many positive properties in the Azure platform, but the fuzzy and unclear definition of the DTU unit make it difficult to calculate the correct Tier level for your application. However, if you are considering moving your databases to the Azure platform, it's important to perform thorough performance measurements of your existing database solutions. Given the great variations in performance we have measured, it is probably necessary to aim for higher Azure Tier to compensate for the unstable resources, so that your service das not periodically degrade under your own minimum requirements.

Moving Forward

To establish a better understanding of the Azure database cloud service, it's important that we perform additional tests on other Azure Tiers and across other locations. We anticipate conducting further surveys in the near future to discover more about the cloud service.

Meanwhile, if you would like to conduct your own tests using out methods, we would recommend to download you the dbWatch tool from http://dbwatch.com/azure, where you will find all necessary information on how to perform the same tests.

Good luck.

Marek Jablonski / dbWatch AS