

---

## Virtual Machine Cloud Infrastructure Management: Determining Differences between Azure Regions within the Same VM Size

---

Submitted 08/09/24, 1st revision 20/10/24, 2nd revision 06/11/24, accepted 25/11/24

Michał Zagrabą<sup>1</sup>

**Abstract:**

**Purpose:** This article delves into the problem of selecting the most suitable IT infrastructure while opting for one of the leading cloud providers, specifically Microsoft Azure. The primary goal is to tackle a couple of business problems: are certain regions more advantageous for hosting infrastructure than others within the same VM sizes? Are there noticeable variations with the same sizes of VMs? Can we enhance the process of choosing infrastructure?

**Design/Methodology/Approach:** The study employs CoreMark benchmarking to assess CPU performance across Microsoft Azure virtual machines, focusing on configurations with upto 16 CPUs due to subscription and cost constraints. It acknowledges potential regional performance variability not captured by standard benchmarks and uses Infrastructure as Code (IaC) with Bicep templates to automate VM deployment and testing. This approach provides a comprehensive understanding of VM performance and pricing within a specific Azure region.

**Findings:** The results involve the automation of tests for different Virtual Machines (VMs) in Microsoft Azure using native Infrastructure as Code (IaC), all within the context of business decision-making. We will simplify the challenges related to VM series pricing and discuss how Infrastructure Management and IaC serves as efficient tools driving the testing process seamlessly across diverse Azure Regions.

**Practical Implications:** The study aids businesses in optimizing resource allocation and choosing appropriate Azure regions by analyzing VM performance and cost trade-offs while considering regional variability. It uses Infrastructure as Code (IaC) for efficient VM deployment and benchmarking, enhancing cloud resource management and operational efficiency.

**Originality/Value:** The study offers unique insights by integrating CoreMark benchmarking with regional performance considerations, highlighting potential discrepancies not captured by standard benchmarks. Its use of Infrastructure as Code (IaC) for automation provides a novel approach to efficiently manage and deploy cloud resources, adding significant value to cloud optimization strategies.

**Keywords:** Azure Virtual Machines, CoreMark benchmark, Cloud efficiency.

**JEL codes:** C88.

**Paper type:** Research article.

---

<sup>1</sup>Mgr., Katedra Informatyki, Wydział Zarządzania, Uniwersytet Łódzki, Poland,  
e-mail: [michal.zagrabam@uni.lodz.pl](mailto:michal.zagrabam@uni.lodz.pl);

## **1. Introduction**

In the ever-evolving realm of cloud computing, understanding how to effectively deploy and optimize Virtual Machines (VMs) within Microsoft Azure can be quite a puzzle. All running software on Azure VM's utilizes its CPU, memory and disk resources so in short software performance and process effectiveness is based on underlying VM performance. Azure offers a variety of VM series (Copeland *et al.*, 2015) each designed for different tasks. These series, like General Purpose VMs or Memory Optimized VMs, come with a wide range of configurations and pricing options (Microsoft official documentation "Azure Virtual Machines Pricing.", <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/>).

This diversity can make choosing the right VM a confusing task, even for experienced users. Naturally, opting for a 4-CPU VM configuration may offer improved performance compared to a 2-CPU VM, but it doesn't necessarily equate to more cost-effectiveness. Complicating matters further, identical VM types come in diverse versions, making the decision-making process even more intricate. On top of that each VM series may be built on different CPUs which leads to performance discrepancies between identical VMs.

To unravel this complexity, we reference "Microsoft Azure Virtual Machines Pricing", an official Microsoft resource. This source provides clear information on the various virtual machines series, and the pricing models associated with them. As an illustrative example, let's consider the D-series VMs, which are available in various versions, including Ds, Dds, Das, Dads, Dps, Dpds, Dpls, and Dplds. It's worth noting that these versions not only exhibit differing price points but also possess distinct configurations, further complicating the decision-making process. Adding to the complexity, the availability of these versions may vary across regions, and their pricing structures can also fluctuate.

In the dynamic realm of cloud computing, businesses and enterprises are presented with a multitude of choices when it comes to selecting virtual machines within the Azure ecosystem. The central question that arises is whether significant performance disparities exist among various virtual machines versions and types within a given series, and if these discrepancies are accurately mirrored in Microsoft's official pricing structure. Beyond this, we delve into the intriguing landscape of Azure regions, questioning whether disparities in performance persist across these diverse geographical locations.

As we embark on this exploration, we extend our inquiry to the intricate decision-making process surrounding datacenter architecture. With the advent of new technologies, a critical consideration arises: When constructing a new datacenter, should it be equipped with the latest CPUs for optimal performance, or does the existing datacenter's infrastructure take precedence for modernization? Unveiling the answers to these questions promises insights into potential discrepancies and unveils

---

the intricate balance between technological advancements and practical infrastructure decisions. Join us on this journey as we unravel the intricacies of Azure's performance dynamics and ponder the implications for strategic decision-making in the ever-evolving landscape of cloud computing.

## 2. Literature review

Navigating the of cloud infrastructure selection within Microsoft Azure has been a subject of exploration in recent research. While the field offers valuable insights, it's essential to note that the existing body of literature tends to be somewhat limited, with a predominant focus on specific aspects such as the choice of operating systems (Ristov and Gusev, 2013) not the virtual machines themselves.

## 3. Materials and Methods

### 3.1 Coremark Test

To assess the performance of CPUs across different virtual machine instances, a standard benchmarking tool commonly employed by cloud providers is CoreMark. CoreMark is an industry-standard benchmark that measures the performance of central processing units (CPU) and embedded microcontrollers (MCU) (Coremark project, <https://github.com/eembc/coremark>).

Google Cloud provides CoreMark scores (Google Cloud, "CoreMark Scores of VM Instances,"

<https://cloud.google.com/compute/docs/coremark-scores-of-vm-instances/>)

for various VM instances, offering a comparative metric to evaluate CPU capabilities. Similarly, Microsoft Azure acknowledges the significance of compute benchmarking and publishes benchmark scores for its virtual machines, including those on Linux instances (Microsoft Azure. "Compute Benchmark Scores." <https://learn.microsoft.com/en-us/azure/virtual-machines/linux/compute-benchmark-scores>)

While these standardized benchmarks contribute valuable insights into the intrinsic performance of CPUs within a cloud provider's infrastructure, it's crucial to note that they may not capture regional differences in performance. The geographical location of datacenters can introduce latency and variability in performance, which may not be fully reflected in these standardized tests. As we explore the nuances of performance in Azure regions, it becomes apparent that testing methodologies need to extend beyond these benchmarks to provide a comprehensive understanding of how regional dynamics may influence overall performance metrics.

### 3.2 Virtual Machines Series and Pricing in Microsoft Azure

Let's take a closer look at virtual machines series and pricing in Microsoft Azure.

Virtual machines series are like different toolkits, each best suited for specific tasks. They come in various sizes and capabilities. However, understanding how these series relate to pricing can be a head-scratcher. It's crucial to know how to pick the right virtual machine without breaking the bank.

In our comprehensive analysis of Virtual Machines within a specific Azure region, we meticulously considered all possible virtual machines configurations with CPU counts ranging from 1 to 16. This approach ensured a thorough examination of a diverse array of virtual machine types, resulting in a curated list of 250 virtual machines representative of the region's offerings. It is crucial to note that virtual machines configurations exceeding 16 CPUs were intentionally excluded from testing. This decision was driven by both Microsoft subscription quotas (Microsoft Azure. "Quotas, limits, and throttling in Azure." <https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits>) and the associated costs incurred for virtual machines with higher CPU counts.

Microsoft Azure imposes subscription-specific limitations on resource creation, and the constraints on virtual machines with 20 or more CPUs aligned with these subscription quotas. The quota limit can always be increased however this requires manual raising the ticket to Microsoft support (Microsoft Azure "Raise a quota increase" <https://learn.microsoft.com/en-us/azure/quotas/quickstart-increase-quota>).

This is fine and reasonable approach when using Azure for hosting your application, so in fact, using it as it is meant to be used, but our case was different. It is hard to imagine raising that many tickets. High CPUs virtual machines can be a subject of separate study.

As a result, our focus on virtual machines within the 1 to 16 CPU range aimed to provide valuable insights while navigating within the bounds of subscription limitations and cost considerations.

### **3.3 Infrastructure as Code: The Key to Automation**

To tackle the complexities of VM benchmarking, automation is the key. That's where Infrastructure as Code (IaC) comes into play. IaC lets you describe your infrastructure as code, making it easy to automate VM deployment and benchmarking. In this context, bicep templates, a default IaC framework, empowers the study to automate these processes.

For a deeper understanding of IaC's role, we refer to "Infrastructure as Code: Managing Servers in the Cloud" by Kief Morris (Infrastructure as Code: Managing Servers in the Cloud, pp. 15-25. O'Reilly Media, 2016). This book is a valuable resource for grasping IaC principles and its impact on optimizing cloud infrastructure.

### 3.4 A Tool for Automating Deployment

The testing process employed in this study was meticulously automated and characterized by its repetitive nature to ensure comprehensive coverage of virtual machine (VM) instances. The initial phase involved the execution of a script capable of analyzing approximately 250 different VM sizes, constrained to configurations with up to 20 CPUs due to cost limitations. Regions such as West Europe and North Europe were scrutinized within development and test subscriptions, providing valuable insights into performance dynamics.

Given that the Poland Central region was not available for testing in development and test subscriptions, the researcher opted for a pay-as-you-go subscription to explore this specific region, albeit incurring additional costs. Once the scope was defined, the script seamlessly deployed VMs using Bicep templates. Subsequently, a testing script was uploaded to the deployed VM, initiating the CoreMark benchmarking process.

The testing script, when activated, executed the CoreMark benchmark with the specified compiler flags: "-O2 -DMULTITHREAD=<number of vm's CPU> -DUSE\_FORK -DPERFORMANCE\_RUN=1 -lrt". The benchmark results, encompassing vital information such as hostname, date, model name, iterations per second, CoreMark size, total ticks, total time, compiler version, and compiler flags, were meticulously captured and uploaded to a designated blob storage.

An example of the result file structure is as follows:

```
{ "hostname": "StandardF1",  
  "date": "20231210_075054",  
  "model_name": "Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz",  
  "iterations_sec": "14820.303816",  
  "coremark_size": "666",  
  "total_ticks": "13495",  
  "total_time": "13.495000",  
  "compiler_version": "GCC5.4.0 20160609",  
  "compiler_flags": "-O2 -DMULTITHREAD=1 -DUSE_FORK -  
DPERFORMANCE_RUN=1 -lrt" }
```

Once the results were successfully uploaded, the corresponding VM was systematically deleted, paving the way for the initiation of the next testing cycle. This automated and iterative approach facilitated a thorough examination of VM performance across various sizes, regions, and configurations within the Azure environment.

The script itself and the IaC code used for automating are available under public GitHub repository: <https://github.com/mzgrb/vm-testing/>.

### **3.5 Costs, Time and Regions**

The typical cycle for the entire testing process involves the creation of a new resource group, the deployment of a virtual machine (VM) with essential components such as a network interface, public IP, virtual network (VNet), and disk. Following this, the script is uploaded, executed, and subsequently removed. On average, this entire cycle consumes approximately 4 minutes, as indicated by sample data derived from deploying 20 VMs consecutively.

This implies that the maximum number of results generated in a day is slightly below 370. It's noteworthy that the script execution was occasionally interrupted by VM restarts/crashes which slowed down the testing process.

Two regions under comparison were Western Europe and Poland Central. Western Europe stands as one of the longstanding data centers in Europe, while Poland Central represents one of the most recently established data centers on the continent (Microsoft "Azure Regions." <https://azure.microsoft.com/en-us/global/infrastructure/regions/>; Microsoft Azure Blog. "Microsoft expands cloud services in Europe and into the Middle East." <https://blogs.microsoft.com/blog/2019/03/14/microsoft-expands-cloud-services-in-europe-and-into-the-middle-east>).

## **4. Research Results**

Between December 1, 2023, and December 15, 2023, a total of 4,772 distinct tests were conducted. To expedite the research process, parallel execution across different Azure subscriptions was employed.

The analysis encompassed precisely 240 VM sizes, spanning from 1 to 16 CPUs. VMs with more than 16 CPUs were deliberately excluded from this study due to constraints related to subscription quotas, intricate setup requirements, and escalated costs. Given the availability constraints across regions, 167 VM results could not be directly compared, as they were exclusively accessible in the West Europe region. When script tries to deploy those VMs in Poland central region it returns an error.

Consequently, the research focused on 72 VMs that were comparable across both regions. The top 10 VMs exhibiting the most significant differences in favor of Poland Central are detailed in the table below, with the complete dataset provided in the appendix.

When analyzing the down 10 results we observe only minor discrepancies in the outcomes. There is no statistical difference between regions in this case.

**Table 1.** Average Coremark results – TOP 10 results

VM Type	Region						U_test	signif.
	Poland Central			West Europe				
	count	var	avg	count	var	avg		
StandardF1	14	428,95	14792,85	21	752496,13	12889,88	0	0,0000
StandardD11v2	14	56845,46	30724,90	20	1908960,09	26215,83	0	0,0000
StandardF8	13	336655,26	122170,11	20	70872193,93	104186,42	0	0,0000
StandardD12v2	14	167595,85	60785,29	19	15297656,34	52445,84	1	0,0000
StandardF4	14	26539,60	60864,35	20	7476468,34	50932,69	14	0,0000
StandardD4v3	13	187768,16	41134,73	14	4706265,90	35521,04	4	0,0000
StandardD3v2	13	5587,97	60924,97	10	3650743,12	51420,32	0	0,0000
StandardF2	14	129019,16	30714,55	20	4023667,06	26976,35	21	0,0000
StandardD8v3	12	57306,54	82678,01	16	25997975,58	71625,24	12	0,0000
StandardD13v2	13	15839070,28	120298,33	9	59532666,03	104794,10	4	0,0000

**Note:** *U\_test* - Mann-Whitney *U* test, *Signif* - significance level  $P < 0.05$ , *var* – variance, *avg* – average.

**Source:** Own study.

**Table 2.** Average Coremark results – DOWN 10 results

VM Type	Region						U_test	signif.
	Poland Central			West Europe				
	count	var	avg	count	var	avg		
StandardD2dv4	13	1102619,4	23342,98	17	497,82	23619,88	103	0,7729
StandardD8asv4	13	614373,38	125366,12	19	722692,17	125288,42	117,5	0,8206
StandardE4-2dsv4	13	165960,01	23778,77	18	154823,83	23795,20	111,5	0,8283
StandardE8-4asv4	9	70289,63	63243,95	20	1020228,96	63053,24	85	0,8351
StandardD2dsv4	13	253758,22	23478,08	17	206705,34	23849,22	106,5	0,8691
StandardF2s	13	16418,73	30786,12	20	3306851,79	30541,61	135	0,8702
StandardF1s	14	1337,16	14787,05	20	1183824,18	14538,71	137,5	0,9313
StandardDS2v2	13	32829,09	30724,41	19	3357284,94	30374,57	124	1,0000
StandardE2asv4	13	194316,52	31988,22	20	236455,72	31793,28	130,5	1,0000
StandardE2sv3	13	2845,88	20888,19	18	210136,51	20934,21	130,5	1,0000

**Note:** *U\_test* - Mann-Whitney *U* test, *Signif* - significance level  $P < 0.05$ , *var* – variance, *avg* – average

**Source:** Own study.

It is noteworthy that all significant advantages favor the Azure region Poland

Central. This suggests the possibility of newer CPUs being deployed in this data center, prioritizing the development of new capabilities and the expansion of cloud data centers over the modernization of older ones.

Let's analyze underlying CPU's for first 3 virtual machines series with biggest difference.

**Table 3. Underlying CPU for TOP 3 VM's**

VM-series and CPU	Poland Central	West Europe
<b>StandardD3v2</b>	<b>13</b>	<b>10</b>
Intel(R) Xeon(R) CPU E5-2673 v4 @ 2.30GHz		2
Intel(R) Xeon(R) Platinum 8171M CPU @ 2.60GHz		8
Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz	13	
<b>StandardF4</b>	<b>14</b>	<b>20</b>
Intel(R) Xeon(R) CPU E5-2673 v4 @ 2.30GHz		3
Intel(R) Xeon(R) Platinum 8171M CPU @ 2.60GHz		17
Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz	14	
<b>StandardF8</b>	<b>13</b>	<b>20</b>
Intel(R) Xeon(R) CPU E5-2673 v4 @ 2.30GHz		9
Intel(R) Xeon(R) Platinum 8171M CPU @ 2.60GHz		11
Intel(R) Xeon(R) Platinum 8272CL CPU @ 2.60GHz	13	

Source: Own study.

In each instance, VMs in Poland Central utilized newer and faster CPUs, clarifying the divergence in Coremark results. This not only elucidates the performance differential but also partially addresses the quandary of region selection. When deciding between Poland Central and West Europe for VM placement, the choice leans towards the former for certain VMs, disregarding cost considerations. Generally, West Europe tends to be more cost-effective than Poland Central (Average Price for Azure Regions, <https://cloudprice.net/regions>).

**Table 4. Average VM cost Region – Microsoft notation for region names**

Region Name	Region	Hourly average price per VM per hour
<b>Central India</b>	centralindia	1,05
<b>North Central US</b>	northcentralus	1,11
<b>West Europe</b>	westeurope	1,46
<b>Poland Central</b>	polandcentral	1,56
<b>Switzerland West</b>	switzerlandwest	2,09
<b>Brazil Southeast</b>	brazilsoutheast	2,5

Source: Own study.

## 5. Discussion - Potential for Further Research

While the current study provides valuable insights into the nuances of cloud

---

infrastructure selection within Microsoft Azure, there remains ample opportunity for further research to refine and extend our understanding. One avenue for future investigations involves conducting additional runs of the experiments to produce more accurate and robust results. By increasing the sample size, researchers can enhance the reliability and generalizability of the findings, offering a more comprehensive perspective on the performance dynamics of Virtual Machines (VMs) across various regions.

Moreover, the scope of this research can be broadened to encompass a comprehensive analysis of all regions within the Microsoft Azure ecosystem. While such an extension would undoubtedly enhance the depth of the study, it is essential to acknowledge the associated increase in costs and time.

However, the trade-off lies in obtaining a holistic view of the entire estate, providing a more nuanced understanding of the performance variations and factors influencing cloud infrastructure decisions.

Beyond Microsoft Azure, there exists a rich landscape of other cloud providers, including industry giants such as Google Cloud Platform and Amazon Web Services (AWS), as well as emerging players like Alibaba Cloud and OVH. A natural progression for future research would involve extending the investigation to include these diverse cloud platforms.

This comparative analysis across multiple providers would yield valuable insights into the generalizability of findings and the unique nuances of each provider's infrastructure, contributing to a more comprehensive understanding of cloud performance dynamics in the broader cloud computing ecosystem.

## 6. Conclusion

In the culmination of this research, our findings underscore the significant impact of underlying CPU variations on the performance of Virtual Machines (VMs) within Microsoft Azure. The study revealed noteworthy discrepancies of up to 20% in CoreMark results among VMs of the same size, emphasizing the crucial role of the CPU in determining overall performance.

Of particular interest was the performance comparison between the Western Europe and Poland Central regions. The Poland Central datacenter, being notably newer, exhibited superior performance in nearly 25% of cases, translating to every fourth VM deployed. Furthermore, Poland Central demonstrated a higher level of consistency in CPU performance, attributed to the deployment of a limited set of CPUs—a characteristic likely stemming from the datacenter's modern infrastructure built with a standardized line of CPUs.

This prompts a practical consideration for users deploying VMs in older

datacenters—particularly those using various CPUs. The observed variability suggests that redeploying VMs in such cases, if feasible, could be akin to "winning a lottery" of CPUs, potentially yielding substantial performance improvements.

As organizations continue to navigate the cloud landscape, understanding and leveraging these insights can contribute to more informed decision-making in optimizing cloud infrastructure within the Microsoft Azure framework.

**References:**

Copeland, M., Soh, J., Puca, A., Manning, M., Gollob, D., Copeland, M., ... , Gollob, D. 2015. Getting started with Azure virtual machines. Microsoft Azure: Planning, Deploying, and Managing Your Data center in the Cloud, 147-176.

Microsoft official documentation “Azure Virtual Machines Pricing.”  
<https://azure.microsoft.com/en-us/pricing/details/virtual-machines/>.

Ristov, S., Gusev, M. 2013. Performance vs cost for windows and linux platforms in Windows Azure cloud. 2013 IEEE 2nd International Conference on Cloud Networking (CloudNet), San Francisco, CA, USA, 2013, pp. 214-218. doi: 10.1109/CloudNet.2013.6710581.

Coremark project. <https://github.com/eembc/coremark>.

Google Cloud, CoreMark Scores of VM Instances.  
<https://cloud.google.com/compute/docs/coremark-scores-of-vm-instances/>.

Microsoft Azure. Compute Benchmark Scores. <https://learn.microsoft.com/en-us/azure/virtual-machines/linux/compute-benchmark-scores>

Microsoft Azure. Quotas, limits, and throttling in Azure. <https://docs.microsoft.com/en-us/azure/azure-resource-manager/management/azure-subscription-service-limits/>.

Microsoft Azure. Raise a quota increase. <https://learn.microsoft.com/en-us/azure/quotas/quickstart-increase-quota-portal/>.

Morris, K. 2016. Infrastructure as Code: Managing Servers in the Cloud, pp. 15-25. O'Reilly Media.

Microsoft. Azure Regions. <https://azure.microsoft.com/en-us/global-infrastructure/regions/>.

Microsoft Azure Blog. Microsoft expands cloud services in Europe and into the Middle East. <https://blogs.microsoft.com/blog/2019/03/14/microsoft-expands-cloud-services-in-europe-and-into-the-middle-east/>.

Average Price for Azure Regions. <https://cloudprice.net/regions>.

**Appendix A:**

	Region						U_test	signif.
	Poland Central			West Europe				
VM Type	count	variance	average	count	variance	average		
StandardF1	14	428,95	14792,85	21	752496,13	12889,88	0	0,0000
StandardD11v2	14	56845,46	30724,90	20	1908960,09	26215,83	0	0,0000
StandardF8	13	336655,26	122170,11	20	70872193,93	104186,42	0	0,0000

StandardD12v2	14	167595,85	60785,29	19	15297656,34	52445,84	1	0,0000
StandardF4	14	26539,60	60864,35	20	7476468,34	50932,69	14	0,0000
StandardD4v3	13	187768,16	41134,73	14	4706265,90	35521,04	4	0,0000
StandardD3v2	13	5587,97	60924,97	10	3650743,12	51420,32	0	0,0000
StandardF2	14	129019,16	30714,55	20	4023667,06	26976,35	21	0,0000
StandardD8v3	12	57306,54	82678,01	16	25997975,58	71625,24	12	0,0000
StandardD13v2	13	15839070,28	120298,33	9	59532666,03	104794,10	4	0,0000
StandardD2v3	14	2421,56	20874,05	18	3267579,00	18546,35	28	0,0001
StandardD2v2	14	492056,27	30534,78	18	4716014,50	26504,74	29	0,0001
StandardE2v3	13	2104,84	20857,99	17	2412917,88	18171,25	26	0,0002
StandardE4v3	13	63312,65	41156,30	17	9142324,58	35906,09	26	0,0002
StandardD1v2	14	1353,56	14787,83	19	1559643,98	13394,05	57	0,0048
StandardDS1v2	13	3271,91	14764,21	20	590026,40	14645,01	190	0,0269
StandardD4dsv4	13	136850,54	46805,86	16	838274,87	47384,58	146	0,0683
StandardE4v4	13	15025,84	46816,50	16	7565,43	46762,81	62,6	0,0683
StandardE2v4	13	7025,87	23673,22	17	3184,38	23627,15	67	0,0715
StandardD4sv4	13	88782,17	46794,10	17	951401,32	47547,01	151,5	0,0863
StandardD4sv3	13	180633,02	41202,40	17	129511,63	41531,91	151	0,0945
StandardE4sv3	13	130341,99	41142,69	17	2342213,65	41012,65	149,5	0,1033
StandardE4-2sv3	13	4473,33	20861,87	17	2112,37	20826,28	72,5	0,1127
StandardB1ms	14	419,42	14812,10	22	449925,60	14747,50	106,5	0,1243
StandardDS13-2v2	13	29690,83	30742,12	17	6198662,00	30158,63	147	0,1334
StandardB4ms	13	180057,82	60823,64	16	21776094,91	59332,34	71,5	0,1563
StandardB1ls	13	270,74	14811,29	14	535038,08	14695,62	65	0,2199
StandardE8sv3	13	56133,25	82684,35	16	3737800,85	82404,73	129	0,2875
StandardE8v4	12	18384,67	93704,12	17	27988,83	93811,72	129	0,2875
StandardDS12-2v2	13	32244,93	30676,55	19	3307282,66	30521,32	150,5	0,3050
StandardD2v4	13	2156,58	23633,94	17	334,43	23617,96	86,5	0,3203
StandardE8-2sv3	13	12823,33	20823,56	17	171983,94	20776,13	86	0,3203
StandardDS12-1v2	13	1926,06	14776,21	18	2560457,71	14775,20	140,5	0,3521
StandardE4asv4	13	37933,99	63188,89	19	51578,51	63256,97	148	0,3624
StandardD2sv3	13	2230,96	20848,97	17	2557713,84	19824,80	89,5	0,3851
StandardE8-4sv3	13	143653,92	41145,59	17	2935801,36	40994,59	132	0,3851
StandardD8sv4	13	357985,25	93635,98	16	54184,25	93837,60	124	0,3983

StandardF2sv2	13	12013,07	23636,90	15	35219,61	23646,43	80,5	0,4397
StandardB8ms	13	274492,16	122085,02	16	43332909,35	120316,39	86	0,4488
StandardD4asv4	13	612418,57	62850,61	20	365842,78	63073,95	151	0,4555
StandardE8-2asv4	13	389586,16	32124,71	20	9411,22	31932,84	109	0,4555
StandardF4s	13	2034057,84	60452,84	19	23828223,18	60045,28	142,5	0,4715
StandardD4v4	13	31979,68	46845,51	17	11929,16	46803,62	94,5	0,5088
StandardB2s	13	9109,14	30802,49	17	2061395,99	30655,31	95,5	0,5358
StandardF4sv2	13	3000445,60	46340,39	16	777449,12	47334,75	118	0,5595
StandardD4dv4	13	2823716,73	46375,11	17	14108,25	46836,22	124	0,5917
StandardD8dv4	13	19528541,60	92020,47	17	94956,42	93697,14	123,5	0,5917
StandardD2sv4	13	13419,38	23662,91	18	141442,97	23827,30	131	0,5941
StandardDS3v2	13	388007,96	60747,00	18	14886190,79	60107,83	131	0,5941
StandardDS11v2	13	878436,84	30550,00	20	2098922,62	30729,80	144,5	0,5981
StandardB2ms	13	7970,51	30811,04	16	3671277,06	30283,96	92	0,6193
StandardE2dsv4	13	8073,11	23654,46	17	147928,95	23799,35	98	0,6206
StandardDS11-1v2	13	1507,14	14793,25	20	312280,10	15024,12	143,5	0,6237
StandardF8s	13	394531,93	121927,06	19	51530914,43	121187,95	135,5	0,6498
StandardE4dsv4	13	165198,94	46790,41	17	96190,29	46926,29	122	0,6501
StandardD8dsv4	13	857655,72	93677,90	17	3463432,78	94818,57	100,5	0,6801
StandardE4-2asv4	13	9601,10	31926,54	21	11694,67	31930,69	147,5	0,7006
StandardB1s	13	1193,27	14785,07	21	906740,70	14746,85	146	0,7532
StandardD2asv4	13	13518,01	31914,11	20	20123,17	31911,51	139	0,7575
StandardD2dv4	13	1102619,42	23342,98	17	497,82	23619,88	103	0,7729
StandardD8asv4	13	614373,38	125366,12	19	722692,17	125288,42	117,5	0,8206
StandardE4-2dsv4	13	165960,01	23778,77	18	154823,83	23795,20	111,5	0,8283
StandardE8-4asv4	9	70289,63	63243,95	20	1020228,96	63053,24	85	0,8351
StandardD2dsv4	13	253758,22	23478,08	17	206705,34	23849,22	106,5	0,8691
StandardF2s	13	16418,73	30786,12	20	3306851,79	30541,61	135	0,8702
StandardF1s	14	1337,16	14787,05	20	1183824,18	14538,71	137,5	0,9313
StandardDS2v2	13	32829,09	30724,41	19	3357284,94	30374,57	124	1,0000
StandardE2asv4	13	194316,52	31988,22	20	236455,72	31793,28	130,5	1,0000
StandardE2sv3	13	2845,88	20888,19	18	210136,51	20934,21	130,5	1,0000

*Source: Own study.*