JOURNAL of COMPUTER SCIENCES INSTITUTE

JCSI 32 (2024) 157–162
Received: 17 May 2024
Accepted: 10 June 2024

Performance evaluation of designated containerization and virtualization solutions using a synthetic benchmark

Ocena wydajności wskazanych rozwiązań konteneryzacyjnych i wirtualizacyjnych przy użyciu benchmarku syntetycznego

Sebastian Tarasiuk*, Damian Traczuk, Kacper Szczepaniuk, Patryk Stoń, Jakub Smołka

Department of Computer Science, Lublin University of Technology, Nadbystrzycka 36B, 20-618 Lublin, Poland

Abstract

The aim of this paper is the analysis of the performance of virtualization and containerization technologies in the context of IT infrastructure. The following virtualization technologies were selected for the study: VirtualBox, VMware and QEMU, as well as containerization technologies: Docker, Podman and LXD. In addition, microVM technologies such as QEMU and Firecracker, which are increasingly important in the context of virtualization, were also included. The comparative criteria on which the analysis is based include aspects of computing and memory performance. Tests conducted on selected technologies included testing CPU performance, RAM efficiency and disk speed for reading and writing data. A SysBench, synthetic benchmark, was used to conduct the tests.

Keywords: virtualization; containerization; microVM; synthetic benchmark

Streszczenie

Celem niniejszego artykułu jest analiza wydajności technologii wirtualizacji i konteneryzacji w kontekście infrastruktury IT. Do badania wybrano następujące technologie wirtualizacji: VirtualBox, Vmware i QEMU, a także technologie konteneryzacji: Docker, Podman i LXD. Ponadto uwzględniono również technologie microVM, takie jak QEMU i Firecracker, które mają coraz większe znaczenie w kontekście wirtualizacji. Kryteria porównawcze, na których opiera się analiza, obejmują aspekty wydajności obliczeniowej i pamięciowej. Testy przeprowadzone na wybranych technologiach obejmowały sprawdzenie wydajności procesora, wydajności pamięci RAM oraz szybkości dysku do odczytu i zapisu danych. Do przeprowadzenia testów wykorzystano SysBench, czyli syntetyczny benchmark.

Słowa kluczowe: wirtualizacja; konteneryzacja; microVM; benchmark syntetyczny

*Corresponding author

Email address: sebastian.tarasiuk@pollub.edu.pl (S. Tarasiuk)

1. Introduction

In the era of rapid technology advances, there is an increasing emphasis on optimal and full utilization of hardware resources, as well as on reducing maintenance costs. The rapid increase in data transfer and the growing number of new applications on the market are forcing companies to implement solutions that will result in the most beneficial use of hardware. A topic that is also often discussed is the isolation of application operation from the physical layer of machines.

These issues are closely related to virtualization techniques. They can be divided into two main categories: container-based virtualization (so-called lightweight virtualization) and hypervisor-based virtualization. The latter is further divided into relevant subcategories, however, only one of them will be considered in this paper: type 2 hypervisor. Each category is marked by unique features: in the case of containerization, we are dealing with an engine installed on the host operating system, which is what allows individual containers to use physical resources. In addition, this solution uses host operating system resources such as libraries or other structures. In the case of virtualization, we are dealing with a virtual machine that has its own hardware specification and operating system. These two different but related solutions are a frequent topic of discussion as to which one is better to use and from what angle.

The purpose of this work is to perform an analysis and comparison of the performance of selected containerization solutions and type two hypervisors. For this purpose, a synthetic benchmark will be used - SysBench, which will check the performance of the processor, RAM, and disk operations (writing and reading). This will allow to determine in which applications which solution is a more favourable choice.

2. Purpose of research

Aim of this research was to find best performing virtualization and containerization tools. This comparison will serve in finding the best tool to use for specific tasks. Additionally, it would allow for checking if containerization is more suitable to use in businesses. Also, as part of research microVM will be evaluated as an alternative to currently used solution.

3. Literature review

This section describes scientific publications that deal with the fast-growing technologies of recent years: virtualization and containerization. It also presents research on a novel solution for lightweight virtual machines called microVM.

The authors of the article [1] compared the performance in hardware aspects of three virtualization and containerization platforms, namely VMware Workstation, Hyper-V, and Docker with different operating systems: Ubuntu 20.04 and Windows 10. More specifically, they examined the performance of the CPU, hard drive, and RAM. For the study, they used several test applications such SysBench, fio, mbw, Fritz Chess Benchmark and Crystal Disk Mark. In this study, the host was Windows 10.

The researchers of the article [2] compared the performance of a virtual machine that was based on the KVM kernel with a containerization solution like Docker. They examined CPU utilization, memory usage, performance memory and disk I/O performance, using benchmarks such as mpstatcommand, freecommand and cpcommand. In addition, they performed a benchmark test with web server. For this purpose, they used JMeter. In the end, the results showed that Docker more efficient than a virtualization solution such as KVM.

In the conference paper [3], the authors analysed the performance of virtualization and containerization solutions: KVM, Docker, LXC. The results obtained in the tests were also additionally compared with native computer hardware. The researchers performed CPU, disk, RAM, and network performance tests. To do this, they used various tools, namely Y-cruncher, NBENCH, Linpack, Bonnie++, STREAM, Netperf.

The authors, in a scientific publication [4], compared the performance of virtual machines with containers. They used the SysBench tool to examine performance. Using this tool, they performed the oltp benchmark test on a MySQL database. The researchers also compared the results here with native computer hardware.

In addition to virtualization and containerization, a similar solution, microVM, has appeared on the market. It is an innovative technology that is remarkably like classic virtual machines. The main difference is that they are more optimized for resource consumption. This is confirmed, among other things, by studies in scientific publications [5,6]. One popular lightweight VM solution is Firecracker.

In several conference papers [7-9], the authors used Filebench to examine the performance of virtualization solutions. This is a tool that is used to examine disk and file system performance. On the plus side, it already includes several predefined workloads, among others: webserver.f, varmail.f, fileserver.f.

The Filebench tool was also used in conference papers [10, 11]. This time, however, the authors approached the performance testing of virtualization solutions a little differently. Benchmarks were performed while simultaneously running one, two and three virtual machines.

Summarizing the literature review, each researcher took a different approach to analysing the performance of containerization and virtualization solutions. They used different research tools, host operating systems and guest operating systems. Considering the above the topic of evaluating the performance of virtual machines as well as containers can be approached in separate ways, there are no rigidly defined rules or principles here.

4. Research method

The purpose of the research conducted is to determine which virtualization or containerization solution is the best in terms of performance. Performance is defined as a component of individual component efficiencies: CPU, RAM, and disk operation performance. This means that the winner is determined by ranking. In individual tests, each virtualization and containerization solution are ranked. Then, it will be possible to determine which solution performed best.

4.1. Object and scope of research

The objects of research are the following solutions:

- virtualization: VMware Workstation Pro, VirtualBox, QEMU (Quick Emulator),
- containerization: Docker, Podman, LXD,
- microVM: QEMU microVM, Firecracker.

The scope of research includes the following performance comparisons:

- processor,
- RAM memory,
- disk operations.

4.2. Research bench

The research was conducted on the same computer workstation. Its specifications are shown in Table 1.

Table 1: Specification of equipment - research bench

Processor	Intel i5-9300H 2.40 GHz
Graphics	Intel UHD Graphics 630
Memory	SK Hynix DDR4 1333MHz 24 GB
Hard Drive	SSD Toshiba KXG50ZNV256G 256GB
Operating System	Ubuntu 22.04.4 LTS

Each virtual machine, regardless of the hypervisor, had the same hardware configuration. Details are shown in Table 2.

Table 2: Virtual machine specifications

Processor	2 virtual cores
Memory	8192 MB
Hard Drive	25 GB
Storage controller	LSI Logic
Operating System	Ubuntu 22.04.4 LTS

Appropriate limits have also been imposed on individual containers (see Table 3) to maintain test consistency.

Table 3: Containers limits

Processor	2 cores
Memory	8192 MB

4.3. Research group

The research group was selected based on an article [12], which presents the most popular virtualization and containerization solutions. The Type 2 hypervisors analyzed were: VirtualBox, VMware and QEMU, while the containerization solutions were Docker, Podman and LXD. In addition, performance tests were performed for microVMs: Firecracker and QEMU microVM.

4.4. Research procedure

Both the host and guest operating systems were a Linux distribution: Ubuntu 22.04.4 LTS. Performance tests of the CPU, RAM and disk input-output operations were repeated 10 times using the SysBench tool. The same configuration shown in Table 1 was used in each testing scenario.

5. Results analysis

The results obtained with the SysBench tool were visualized into bar charts and tables so that they could be analysed. Not only were the main differences between virtualization and containerization analysed, but also between the specific technologies.

5.1. Processor efficiency analysis

Analysis of CPU performance has shown that there are significant differences in the number of events processed per second not only between containerization, virtualization and microVM solutions, but also between specific technologies within a solution type. The best performance was observed for LXD and Podman containerization solutions, with 545.01 and 543.28 events per second, respectively, and QEMU microVM technology, with 544.41 events per second. Slightly worse results were achieved by microVM Firecracker technology - 537.66 events per second and QEMU virtualization technology - 536.61 events per second. Weaker results were observed for VirtualBox virtualization technology and Docker containerization technology - 534.18 and 533.47 events per second, respectively. The weakest and significantly outlier in CPU performance was demonstrated by VMware virtualization technology, with 528.69 events per second.

The analysis showed that containerization technologies have, on average, a slight 2% better CPU performance than virtualization technologies. MicroVM technologies achieved results almost identical to containerization technologies.

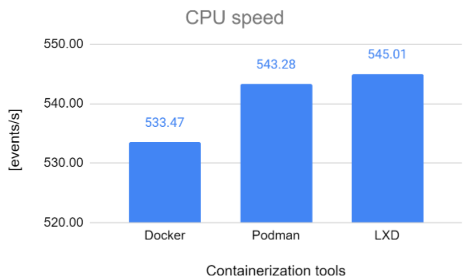

Figure 1: CPU speed for containerization technologies.

A comparative analysis of CPU performance was made by visualization showing averaged test results in the form of bar graphs. Figure 1 shows the results obtained by containerization technologies.

Figure 2 shows the results obtained by virtualization technologies.

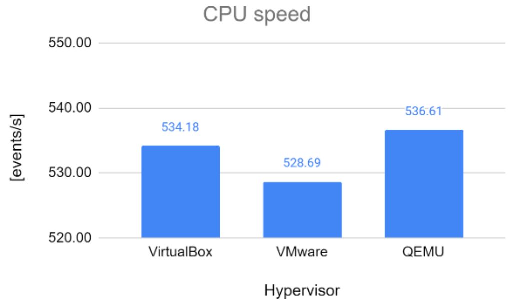

Figure 2: CPU speed for virtualization technologies.

The results obtained by microVM technologies are shown in Figure 3.

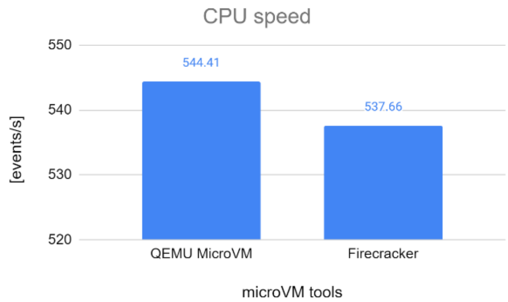

Figure 3: CPU speed for microVM technologies.

5.2. RAM efficiency analysis

Analysis of RAM performance showed that there are significant differences in the amount of data processed per second not only between containerization, virtualization and microVM solutions, but also between specific technologies within a solution type. The greatest variation was observed for virtualization technologies. For containerization solutions, the results are similar while for microVM solutions, the results are almost identical. The best performance was observed for LXD and Podman containerization solutions, with 7301.8 and 7282.98 mebibytes of data processed per second, respectively, and VMware virtualization technology, with 7083.28 mebibytes of data processed per second. The result of less than 7,000 and more than 6,000 mebibytes of data was observed for four technologies: microVM Firecracker and QEMU technologies, Docker containerization technology and QEMU virtualization technology, and these results were respectively: 6771.39, 6711.37, 6579.97 and 6288.32 mebibytes per second. The lowest and the greatest outlier was observed for VirtualBox virtualization technology - it was the result of 4812.81 mebibytes of data processed in one second.

The analysis showed that containerization technologies performed on average 3% better than virtualization technologies (not including VirtualBox technology which scored significantly lower, deviating from all other technologies studied). MicroVM technologies achieved results slightly lower than those of the most efficient technologies in the virtualization and containerization categories, but near to the average among these technologies.

A comparative analysis of RAM performance was made by visualization showing averaged test results in the form of bar graphs. Figure 4 shows the results obtained by containerization technologies.

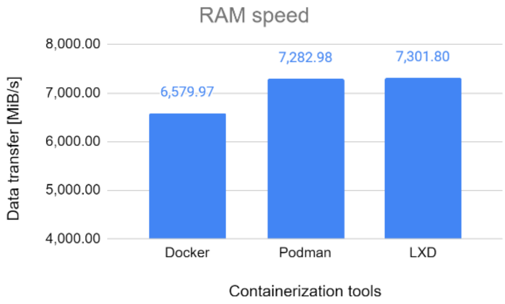

Figure 4: RAM speed for containerization technologies.

The results obtained by virtualization technologies are shown in Figure 5.

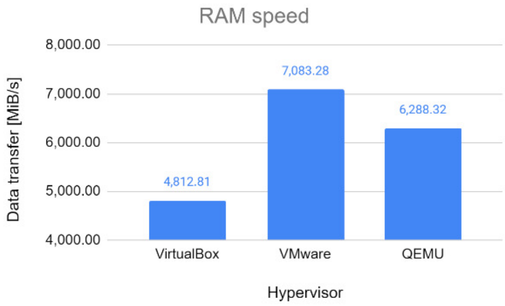

Figure 5: RAM speed for virtualization technologies.

The results obtained by microVM technologies are shown in Figure 6.

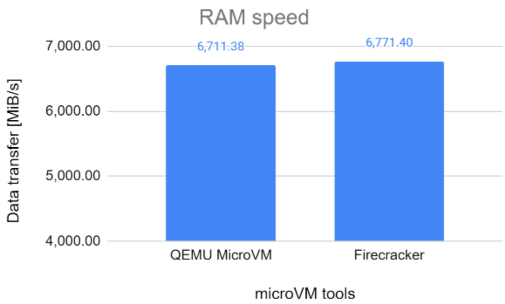

Figure 6: RAM performance for microVM technologies.

5.3. Disk read/write efficiency analysis

Analysis of disk read and write performance showed that there are huge differences in the amount of data read and written to disk per second not only between containerization, virtualization and microVM solutions, but also between specific technologies within a solution type. The biggest differences were observed for virtualization and microVM technologies. For containerization solutions, the results are similar. The best performance was observed for the microVM Firecracker solution, with a score of 316.90 mebibytes of data read from disk per second, and 210.91 mebibytes of data written to disk per second. Scores significantly lower, but still above the average for all technologies evaluated, were observed for the two virtualization technologies VMware and VirtualBox, with scores of: 235.81 and 203.81 mebibytes of data read from disk per second, and 156.94 and 135.64 mebibytes of data written to disk per second. Significantly lower results were observed for all three containerization technologies evaluated (Docker, Podman, LXD), as well as for the virtualization technology QEMU and QEMU as microVM. The results for reading data from disk in one second were 6.92 mebibytes of data for Docker, 11.84 mebibytes of data for Podman, 8.69 mebibytes of data for LXD, 5.09 mebibytes of data for QEMU and 4.24 mebibytes of data for QEMU microVM, respectively. The results for writing data to disk in one second were 4.60 mebibytes of data for Docker, 7.88 mebibytes of data for Podman, 5.78 mebibytes of data for LXD, 3.39 mebibytes of data for QEMU and 2.82 mebibytes of data for QEMU microVM, respectively.

The analysis showed that virtualization technologies (not including QEMU, whose score stands out significantly from the other virtualization technologies evaluated) performed twenty-seven times better than containerization technologies on average in terms of both writing and reading data from disk. A large discrepancy was observed for microVM technologies. Firecracker proved to be unrivalled when evaluated for both writing and reading data and performed on average more than 45% better for both writing and reading than VMware and VirtualBox technologies. QEMU microVM, on the other hand, achieved the lowest score of all the technologies evaluated, about half of containerization's result for both writing and reading data, and about 20% lower than QEMU Virtualizer.

A comparative analysis of the performance of writing data to disk and reading data from disk was made by visualization showing the averaged test results in the form of bar graphs. Figure 7 shows the results obtained by containerization technologies when reading data from disk.

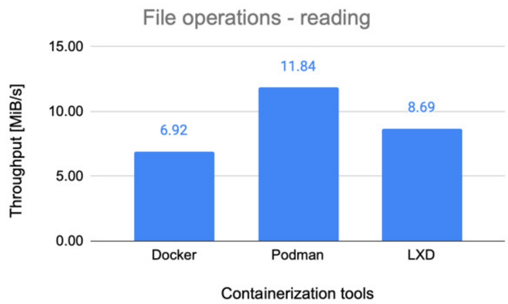

Figure 7: Reading from disk efficiency for containerization technologies.

Figure 8 shows the results obtained by containerization technologies when writing data to disk.

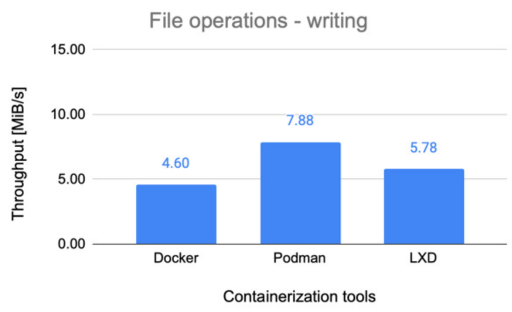

Figure 8: Writing data to disk efficiency for containerization technologies.

Figure 9 shows the results obtained by virtualization technologies when reading data from disk.

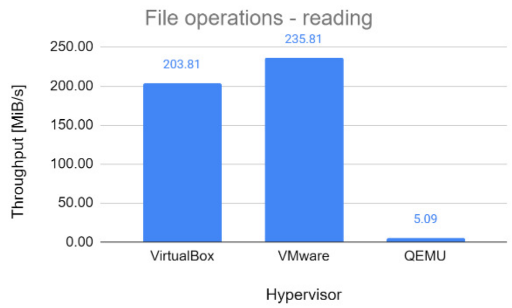

Figure 9: Reading from disk efficiency for virtualization technologies.

Figure 10 shows the results obtained by virtualization technologies when writing data to disk.

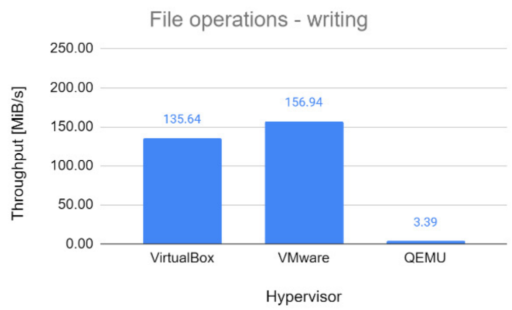

Figure 10: Writing data to disk efficiency for virtualization technologies.

Figure 11 shows the results obtained by microVM technologies when reading data from disk.

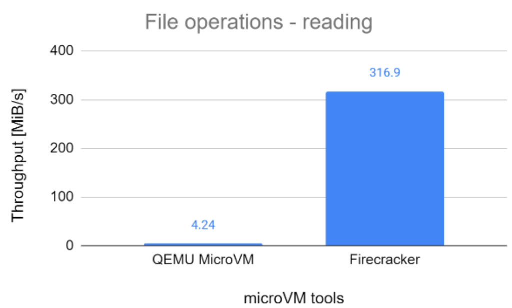

Figure 11: Reading from disk efficiency for microVM technologies.

Figure 12 shows the results obtained by microVM technologies when writing data to disk.

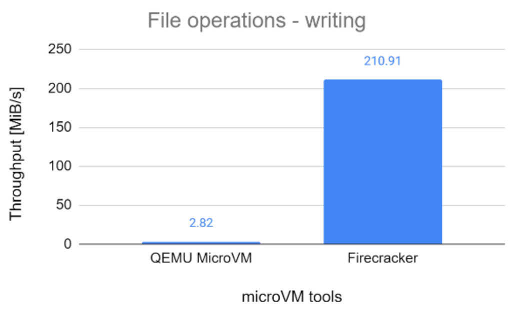

Figure 12: Writing data to disk efficiency for microVM technologies.

6. Conclusions

Based on research performed it is possible to notice that containerization solutions performed better in most aspect examined. The best achieving one were LXD and Podman. Hypervisors gained edge only during testing operations input and output on disk achieving vastly higher results than other analysed tools. Best achieving one was VMware, whose result does not deviate in meaningful way from those achieved by containerization solutions. This, among other things such as ease of scalability and their lightweightness, explain why containers was so widely adopted in industry. Additionally, microVM solutions seems like viable competitors to containerization in regards of performance. Additionally, they offer isolation like traditional virtualization solutions and lightweightness native to containers. The main thing stopping this solution from being more widespread is lack of supporting tools which would make using it easier.

References

[1] Z. Li, Comparison between common virtualization solutions: VMware Workstation, Hyper-V and Docker, In 2021 IEEE 3rd International Conference on Frontiers Technology of Information and Computer (ICFTIC) (2021) 701–707, https://doi.org/10.1109/icftic54370.2021.9647226.

[2] M. Chae, H. Lee, K. Lee, A performance comparison of linux containers and virtual machines using Docker and KVM, Cluster Computing 22 (2019) 1765–1775, https://doi.org/10.1007/s10586-017-1511-2.

[3] R. Morabito, J. Kjällman, M. Komu, Hypervisors vs. Lightweight Virtualization: A Performance Comparison, In 2015 IEEE International Conference on Cloud Engineering (2015) 386–393, https://doi.org/10.1109/IC2E.2015.74.

[4] W. Felter, A. Ferreira, R. Rajamony, J. Rubio, An updated performance comparison of virtual machines and Linux containers, In 2015 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS) (2015) 171–172, https://doi.org/10.1109/ispass.2015.7095802.

[5] F. Manco, C. Lupu, F. Schmidt, J. Mendes, S. Kuenzer, S. Sati, K. Yasukata, C. Raiciu, F. Huici, My VM is Lighter (and Safer) than your Container, Proceedings of the 26th Symposium on Operating Systems Principles (2017) 218–233, https://doi.org/10.1145/3132747.3132763.

[6] G. Li, K. Takahashi, K. Ichikawa, H. Iida, P. Thiengburanathum, P. Phannachitta, Comparative Performance Study of Lightweight Hypervisors Used in Container Environment, Proceedings of the 11th International Conference on Cloud Computing and Services Science (2021) 215–223, https://doi.org/10.5220/0010440502150223.

[7] B. Djordjević, V. Timčenko, N. Kraljević, I. Jovičić, N. Davidović, Comparison of VMware Workstation, VirtualBox and MS Hyper-V hypervisor performance with MS Windows OS based guests, In 2023 22nd International Symposium INFOTEH-JAHORINA (INFOTEH) (2023) 1–5, https://doi.org/10.1109/INFOTEH57020.2023.10094080.

[8] B. Đorđević, N. Kraljević, N. Davidović, Guest file system behavior for type-2 hypervisor-based virtualization in VMware Workstation, In 2023 22nd International Symposium INFOTEH-JAHORINA (INFOTEH) (2023) 1–6, https://doi.org/10.1109/INFOTEH57020.2023.10094170.

[9] B. Đorđević, S. Gucunja, N. Kraljević, N. Davidović, Performance comparison of different hypervisor versions of the type-2 hypervisor VirtualBox, In 2023 10th International Conference on Electrical, Electronic and Computing Engineering (IcETRAN) (2023) 1–6, https://doi.org/10.1109/IcETRAN59631.2023.10192192.

[10] D. T. Vojnak, B. S. Đorđević, V. V. Timčenko, S. M. Štrbac, Performance Comparison of the type-2 hypervisor VirtualBox and VMWare Workstation, In 2019 27th Telecommunications Forum (TELFOR) (2019) 1–4, https://doi.org/10.1109/TELFOR48224.2019.8971213.

[11] B. Đorđević, V. Timčenko, O. Pavlović, N. Davidović, Performance comparison of native host and hyper-based virtualization VirtualBox, In 2021 20th International Symposium INFOTEH-JAHORINA (INFOTEH) (2021) 1–4, https://doi.org/10.1109/INFOTEH51037.2021.9400684.

[12] S. Giallorenzo, J. Mauro, M. G. Poulsen, F. Siroky, Virtualization Costs: Benchmarking Containers and Virtual Machines Against Bare-Metal, SN Computer Science 2 (2021) 404, https://doi.org/10.1007/s42979-021-00781-8.