Virtualization

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Bronte Sundstrom

Spring 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

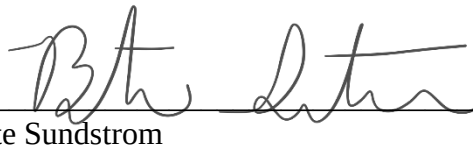

Bronte Sundstrom Date: April 24, 2023

_____ Date:

Richard D. Jacques
Department of Engineering and Society

STS Research Paper

Abstract

Virtualization started as a concept in the late 1960s when IBM was researching time-sharing. Their research led them to the first operating system that supported virtualization through their use of a hypervisor, which was later improved upon by VMware. Virtualization today can be comprised of hypervisors or containers. Hypervisors are software that allows for multiple virtual machines on a single physical system. Containers are independent images that run an application but use the same operating system as the host. Virtualization can be further divided into six main types: server, storage, network, data, application, and desktop. While all of these are similar, they all have their own purposes and vary from one another. Virtualization has many benefits such as the cutting of costs, recovery, IT management, and some security increases. However, it also presents the consequences of virtual machine sprawl, some security decreases, and compatibility issues. Overall, virtualization is more beneficial for society, mainly because of its increased efficiency.

Introduction

Virtualization is growing, but many people do not quite understand what it means. Virtualization, in its simplest form, is the use of software to simulate physical hardware. This allows for more efficient hardware, more storage, and less physical hardware to be needed. Many companies use virtualization to make their work more efficient while using fewer physical resources and therefore less energy. The internet has also become more virtualized as cloud computing becomes more popular.

Virtualization dates back as a concept to the late 1960s and has been on the rise ever since. There are six main types of virtualizations, each virtualizing a different aspect of physical hardware. Like all technologies, there are benefits and consequences that are present and must be addressed. This leads to the questions that need to be asked before using it. Do the benefits outweigh the consequences? How does virtualization affect society? Is this something our society should be working towards?

History

The concept of virtualization can be traced back to the late 1960s with IBM's investment in research on time-sharing. (“Brief History of Virtualization,” n.d.) At the time, computing resources were rare and inefficient. Only one person could use them at a time, and they required an operator to get the results. Therefore, engineers determined that if more than one user could use a computer at one time, efficiency could increase as the computer would be working on other tasks while receiving inputs or giving outputs. This was first implemented in a literal sense, where tasks would simply switch from one to another based on "quantums" of time. While this was an increase in efficiency from before, it was far from multitasking or true virtualization.

This then led to IBM creating the first operating system that supported virtualization. They achieved this through their development of a hypervisor, which simply manages the multiple virtual machines. The first advantage of hypervisors as opposed to time-sharing was that all the physical resources were being shared rather than used one at a time, which led to higher efficiency. Next, multiple operating systems could be run at the same time. And lastly, due to various operating systems, the entire system became more dependable. (Fedoseenko, 2019)

However, during the 1980s, the drive for virtualization slowed to a halt as inexpensive x86 computers led to distributed computing, a system where multiple computers can run as one. (IBM, n.d.) This new process was more easily scalable and less likely to fail due to its redundant nature. Yet even with this increased use, the operational and management costs began to show themselves. VMware attempted to remedy this through virtualization in 1999 with its hypervisor. (Probrand, n.d.)

Key Technologies

1. **Hypervisors**

Hypervisors are simply the layer of software that manages multiple virtual machines that are on the same physical system. They reallocate physical resources such as CPU, storage, and memory between the VMs. With hypervisors, not only can different applications be running on the same software, but multiple operating systems can be as well, as shown in Figure 1 below. They separate the virtual machines from each other in such a way that they do not interfere with each other. If one gets ruined, the others will still survive without any impact. (IBM, n.d.)

Hypervisors can be categorized into Type 1 and Type 2. Type 1 is directly on the host hardware and manages all the virtual machines from there. The resources are therefore given directly from the hardware. This makes them useful for data centers and servers, as a hypervisor can be hosted on one and the VMs can be accessed remotely with no user interface needed. Type 2, however, is run on an existing operating system, and the resources are scheduled through the host OS. This type is more useful for individuals, as multiple operating systems can easily be run on a personal computer. (Red Hat, n.d.)

2. **Containers**

Containers are another way to host different applications on the same physical machine. Unlike hypervisors, containers all use the same operating system, as shown in Figure 1. Each container has its own image with its own configurations, directories, and files. These images can run their own applications, which would not interfere with the dependencies of the other containers. (Red Hat, n.d.) These are often used in application development, as it allows developers to write code based on a container configuration rather than having to worry about the dependencies and environments of different computers. Furthermore, containers only use the resources that they need so that the applications run efficiently. (Google, n.d.) Due to their smaller size and portability, containers are particularly useful for applications on the cloud.

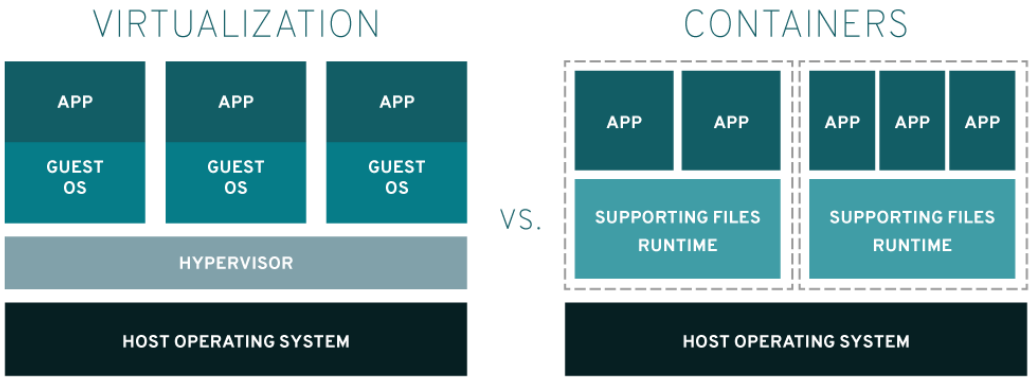

Figure 1. Virtualization with Hypervisors versus Containers (Red Hat, n.d.)

Types

Virtualization has become so innovative over the past couple of decades that society can now virtualize many distinct types of physical infrastructure that could not be done before.

1. **Server**

Servers are often underutilized, as only a small part of their processing power is used. This causes large racks of servers to sit idle as the workload is distributed to only a small portion of the server. Yet at the same time, more servers are needed because the parts of them that are being used are becoming full. This costs businesses money, space, and energy to host servers that are largely unused.

However, server virtualization mitigates these consequences, as each server can be broken up into several smaller virtual servers that can function as their own unique devices. These VMs, with their own applications and operating systems, can more efficiently divide up the work presented to the servers in such a way that more of the servers' processing power is used. This allows for fewer servers, saving space and energy. This is shown in Figure 2. (VMware, 2022)

Figure 2. Decreasing number of servers using Virtualization (JMA Wireless, 2022)

2. **Storage**

Ordinarily, data is stored on one server or in one data center, all in one place. While this seems practical, it falls apart when data that should be together is stored in places across the country. However, this is not a problem for storage virtualization. It is simply "the pooling of multiple physical storage arrays from [storage area networks] and making them appear as a single virtual storage device." (DataCore Software, 2023) This does not care about different networks or data centers when storing the data, allowing for ease of view and management from a single computer.

Similar to most virtualization, the largest advantages of this type of storage are the scalability of storage resources and the ability to store substantial amounts of data in a small amount of space. As our society progresses toward using exponentially more data every year, this virtualized storage has become more of a necessity. This is often used by IT organizations as storage is more efficiently and effectively used without the need for constant monitoring and data reallocation. Therefore, there is more control from a single point when management is required. Storage devices can be added and removed as needed without having to adjust any physical hardware.

3. **Network**

As virtual as networks seem to many people, there is a large physical network infrastructure that every network runs off. These physical elements include switches, routers, and firewalls. While this work well, there is the management aspect that requires maintenance in the physical locations where these elements reside. However, virtualization allows all this management to be done in one central location virtually by administrators for a more simplified management

system. (The University, 1978) The multiple physical networks can be combined into one virtual network, or one physical network can be broken into multiple, independent virtual networks. Further, networks can be moved and altered across domains without the need to reconfigure the network. This allows networks to be updated quickly, so new applications can be hosted without having to create new infrastructure. (VMware, 2022)

4. **Data**

Data can be stored on servers for easy access, but what if related data are stored in different databases or servers? This would require the physical movement or copying of data. This is expensive and inefficient. Data virtualization is the layer between the physical data and the applications that use the data. This allows users to access, combine, transform, and deliver data from diverse sources without messing with the physical. This unified view can then be used by the application with increased speed and cost-effectiveness and can be used similarly to a traditional database. Furthermore, it can also lead to less duplicated data as copies do not need to be created. Like all virtualization, this layer is also easily scalable, allowing for databases that can alter size depending on changing business requirements. (TIBCO Software, n.d.)

5. **Application**

Traditionally, application deployment is time-consuming, complicated, and often requires significant IT resources. This is because applications may have a potential conflict with existing software that is on the same machine. So, to ensure that a company's application is compatible with any machine that it could run on, developers need to use extra time and resources that could be going towards other projects. This, however, can be solved by application virtualization.

Simply put, application virtualization is the act of deploying an application within a container so that everything that is needed for it is deployed with it and there are no dependencies that should not be there. This allows for ease of deployment to multiple devices that host varying operating systems, as there is a layer of insulation between the application and the underlying hardware and system. Many companies that release applications that will be used on a range of devices and operating systems can use this technique to ensure it is easily compatible with all that use it. (VMware, 2022)

6. **Desktop**

Desktop virtualization allows users to access desktop environments using any device from anywhere. Within a company, employees can work virtually anywhere with a network connection, regardless of the device or operating system that they are using. This is a large security benefit for the companies as the data being accessed is not present on the physical device but rather in a remote location accessed only by password, key, or whatever other security the company implements. Furthermore, it gives a consistent desktop environment for all who use it; no matter what device is used, all the same applications and files will be accessible. (VMware, 2022)

In addition to providing security and ease for employees, desktop virtualization also allows IT departments to easily manage desktop environments. This being that all the desktops are hosted on a centralized server. Hardware upgrades would only have to apply to the central server, and software and security upgrades would be the same. (VMware, 2022)

There are two kinds of desktop virtualization: virtual desktops and local desktops. The first of which is what most companies would use. This is when all the virtual desktops are hosted on a

remote server and users can access them using client devices. The second of which can still be used in a company setting, but it is more commercial as anyone with a desktop computer can achieve this. Local desktop is when the hypervisor is run on the local machine and the virtual computer, with the same or a different operating system, is run on top. The local and virtual desktops can be switched between in the same way that someone would switch applications. Two common examples of this are VirtualBox and NoMachine. (The University, 1978)

Benefits

While each type of virtualization has its own benefits, collectively, the list of benefits is lengthy. The first and largest of them is the cutting of costs. This shows itself when looking at server virtualization. When running multiple physical servers, there is a large energy cost to keep them running, a cost for maintenance, lots of space is needed, and there is a cost for cooling. When there are virtual servers, fewer physical servers must be running to have the same amount of computing power. This allows for less energy, less maintenance, and less space. (IBM, n.d.)

Another advantage is the ability to recover data in the event of a disaster. As screenshots can be taken of virtual machines without hassle, if a machine fails at any point, it is easy to move an instance to a new virtual machine. This can help with data, storage, and server virtualization. This can also prevent tests from having to be restarted because progress made in a previous snapshot can be resumed. This would traditionally take days or weeks to re-setup physical hardware, as opposed to a virtual machine taking a few minutes. (Ayuya et al., 2022)

From a business standpoint, virtualization simplifies IT management as departments can manage multiple servers or devices from their desk rather than having to individually check

multiple physical devices. This lets IT teams reallocate resources in a more efficient way, as they can monitor and adjust based on real-time usage. It also allows for easy updates, security patches, and fixes from one place.

Lastly, security is greatly improved when using virtualization. Keeping data and applications in a virtualized form keeps them off the physical hardware. This prevents the spread of viruses to the actual device as well as keeping the data in a secure area that allows for fewer security breaches. Moreover, for application virtualization, applications can work on the same network as each other but stay isolated so that data can be protected from being shared between them.

Challenges

As companies begin to use virtualization more, it is easy to think of them as not taking up memory or CPU and making as many as desired. However, this leads to the first consequence of virtualization: virtual machine sprawl. This phrase refers to the situation in which a person or organization has an excessive number of virtual machines relative to their available memory, storage, and CPU. While virtualization makes these things more efficient by eliminating idle time, the other extreme is when too many virtual machines cause the hardware to overload and become inefficient again. This inefficiency can have major hardware and energy costs from the overuse of hardware, as well as security risks from virtual machines that are not responsibly managed. It can also lead to decreased performance as the overhead of the hypervisor has extra virtual machines to manage. This can also lead to one virtual machine disproportionally using resources, taking them away from others. (“Virtual machine (VM) Sprawl Control Guide,” 2020)

Security risks do not only come from an overuse of hardware, though. A security breach in one virtual machine could result in security breaches in other virtual machines or even the host system because so many virtual machines share the same hardware and physical resources. Also, since virtual machines can be easily cloned or moved, virtualization can make it more difficult to manage and enforce security policies. Further, this copying and moving can give unauthorized access to sensitive data, risking data loss or theft. This is further propagated by virtual machines being able to be on any cloud service provider and being able to be accessed anywhere. (Seget, n.d.)

As virtualization becomes more common, our society is moving away from traditional hardware and applications. However, these practices are not going away. They, too, are ingrained in our society. This is another consequence of virtualization—that there are compatibility issues with legacy hardware and applications. A prime example of this is that some applications may rely on certain hardware that is not available in virtualized environments, while others will not work correctly because of how the hardware is being virtualized. Due to this, applications need to be evaluated more to make sure they are not only compatible with the physical hardware but also the virtual hardware. (Business News Daily, n.d.)

How This Affects Society

While there are valid consequences of virtualization, overall, the benefits are moving society in a better direction. Due to the increased efficiency, businesses that virtualize their infrastructure can reduce maintenance costs and streamline their IT departments. In addition to the reduction of space, this results in cost savings that allow for investment in more research. The reduction in

physical hardware leads to reduced energy consumption, which helps companies with sustainability. This is particularly applicable as so many businesses care more about their impact on the environment.

Recently, the government-instituted lockdowns forced everyone to stay in their homes. As time went on and the government did not reduce its hold, businesses adapted and made it possible for people to work from home. Virtualization was particularly helpful as it allowed them to access their work environments from anywhere in the world. Even since the tyrannical lockdowns have been lifted, companies have still allowed people to work from home and have looked around the world for talent rather than locally, as employees can easily access their work while at home or traveling. This has both increased productivity and provided a better work-life balance for many.

Lastly, virtualization can allow educational institutions to access more powerful computing capabilities than they otherwise would be able to. New opportunities for research and practical hands-on experience are available to those who did not have them before. While virtualized equipment is not particularly cheap, it is an investment that would benefit colleges and high schools in preparing their students for the workforce.

Conclusion

Virtualization is not going away any time soon. While it is a broad topic, this advancement has increased the efficiency and ease of too many things for it to be thrown out right now. Further, the consequences are all remediable by taking extra precautions for security and by doing a little bit of extra work to make legacy hardware available. And since the benefits in

themselves are positive, the impact that they have on society overall is also positive. Companies can accomplish more, people have more flexibility, and it gives them more opportunities.

I have used virtualization in my time at UVA for classes as well as for my internship at JMA Wireless during the summer of 2022. While virtual machines can be confusing at first, they are great for educational purposes as well as practical in the workforce. Particularly in my experience, it allows for users to use alternate operating systems that they would not otherwise have access to. Virtualization, and by extension, cloud computing, are in the future, and many people use them without even knowing it. There are things that need to be kept updated, and security measures need to be put in place, but we should be excited about the future prospect of virtualization.

References

1.1.1. Brief History of Virtualization. 1.1.1. brief history of virtualization. (n.d.). Retrieved April 13, 2023, from https://docs.oracle.com/cd/E26996_01/E18549/html/VMUSG1010.html

What is distributed computing. IBM. (n.d.). Retrieved April 13, 2023, from https://ibm.com/docs/en/txseries/8.2?topic=overview-what-is-distributed-computing

History of virtualisation. Probrand. (n.d.). Retrieved April 13, 2023, from https://www.probrand.co.uk/it-services/vmware-solutions/history-of-virtualisation

What is a hypervisor? Red Hat - We make open source technologies for the enterprise. (n.d.). Retrieved April 13, 2023, from https://www.redhat.com/en/topics/virtualization/what-is-a-hypervisor

Fedoseenko, V. (2019, November 25). *ISPSystem*. ISPsystem. Retrieved April 13, 2023, from https://www.ispsystem.com/news/brief-history-of-virtualization

What are hypervisors? IBM. (n.d.). Retrieved April 13, 2023, from https://www.ibm.com/topics/hypervisors

Google. (n.d.). *What are containers? | google cloud*. Google. Retrieved April 13, 2023, from https://cloud.google.com/learn/what-are-containers

What's a linux container? Red Hat - We make open source technologies for the enterprise. (n.d.). Retrieved April 13, 2023, from https://www.redhat.com/en/topics/containers/whats-a-linux-container

The University. (1978). *What is ...* Amazon. Retrieved April 13, 2023, from https://aws.amazon.com/what-is/virtualization/

What is server virtualization?: Vmware glossary. VMware. (2022, December 19). Retrieved April 13, 2023, from https://www.vmware.com/topics/glossary/content/server-virtualization.html#:~:text=Server%20virtualization%20is%20the%20process,its%20own%20operating%20systems%20independently.

What are containers? IBM. (n.d.). Retrieved April 13, 2023, from https://www.ibm.com/topics/containers

Virtualization 101: What is a hypervisor? Pluralsight. (2022, November 3). Retrieved April 13, 2023, from https://www.pluralsight.com/blog/it-ops/what-is-hypervisor

Storage virtualization: A deep dive. DataCore Software. (2023, January 25). Retrieved April 13, 2023, from https://www.datacore.com/storage-virtualization/

XRAN® - JMA Wireless. JMA Wireless -. (2022, January 3). Retrieved April 13, 2023, from https://jmawireless.com/xran/

What is network virtualization?: Vmware glossary. VMware. (2022, October 21). Retrieved April 13, 2023, from https://www.vmware.com/topics/glossary/content/network-virtualization.html#:~:text=Network%20Virtualization%20(NV)%20refers%20to,into%20separate%2C%20independent%20virtual%20networks.

What is data virtualization? TIBCO Software. (n.d.). Retrieved April 13, 2023, from https://www.tibco.com/reference-center/what-is-data-virtualization

What is application virtualization?: Vmware glossary. VMware. (2022, December 19). Retrieved April 13, 2023, from https://www.vmware.com/topics/glossary/content/application-virtualization.html

What is desktop virtualization?: Vmware glossary. VMware. (2022, December 19). Retrieved April 13, 2023, from https://www.vmware.com/topics/glossary/content/desktop-virtualization.html

5 benefits of virtualization. IBM. (n.d.). Retrieved April 13, 2023, from https://www.ibm.com/cloud/blog/5-benefits-of-virtualization

Ayuya, C., Staff, T. R., Crouse, M., Abdullahi, A., Greenberg, K., Okeke, F., & Wallen, J. (2022, October 7). *Virtualization: Benefits, drawbacks and defining features*. TechRepublic. Retrieved April 13, 2023, from https://www.techrepublic.com/article/virtualization-benefits-drawbacks-defining-features/

Contributor, S. (2020, July 22). *Virtual machine (VM) Sprawl Control Guide*. DNSstuff. Retrieved April 13, 2023, from https://www.dnsstuff.com/what-is-virtual-machine-server-vm-sprawl

Seget, V. (n.d.). *The basics of virtualization security*. Netwrix Blog. Retrieved April 13, 2023, from https://blog.netwrix.com/2020/01/09/virtualization-security/

Virtualization Pros and Cons. Business News Daily. (n.d.). Retrieved April 13, 2023, from https://www.businessnewsdaily.com/6014-pros-cons-virtualization.html