IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

Mitigating Security Threats Arising from Server Sprawl in Virtualized IT Environments

Nirjhor Anjum[1], Md Rubel Chowdhury[2], Md Anwarul Kabir[3]

Assistant Professor (P.T.), CIS Department, Daffodil International University, Bangladesh[1]

Web Developer (Associate), SuperbNexus Limited, Bangladesh[2]

Assistant Professor, Department of Computer Science, American International University Bangladesh, Bangladesh[3]

Abstract: Virtualization is now a central part of modern IT infrastructure, offering better resource utilization, flexibility, and cost savings. However, one serious issue that comes with it is server sprawl. Server sprawl happens when too many virtual machines (VMs) are created without strong policies or proper management. Over time, this leads to a large number of unused or outdated VMs, many of which go unpatched or misconfigured. These neglected virtual machines increase the attack surface, making the entire system more vulnerable to cyberattacks. Server sprawl is not just a problem of wasted resources. It is a growing security threat. This paper explores the root causes of server sprawl in virtualized environments, the kinds of security risks it introduces, and practical ways to reduce these threats. Key strategies such as centralized VM governance, RBAC, hypervisor security, and automated patching are discussed. The paper also reviews past research, presents observations from real-world practices, and proposes directions for future improvements in managing server sprawl securely.

Keywords: Server Sprawl, Virtual Machines, Virtualization Security, IT Infrastructure, RBAC, Hypervisor Security, VM Lifecycle Management, Threat Mitigation.

I. INTRODUCTION

Server sprawl has become a major concern in modern virtualized IT environments. As organizations embrace virtualization to reduce hardware dependency and improve operational flexibility, the number of virtual machines (VMs) in use grows rapidly. At first, this sounds like a good thing. Creating new VMs is easy, fast, and cheap. But over time, without strong policies or tracking systems, many VMs are left running without purpose. Some are forgotten, some are underused, and others are not updated at all. This uncontrolled growth is called server sprawl.

The core issue with server sprawl is that it increases the attack surface. More virtual machines mean more entry points for attackers. And the problem gets worse when these machines are misconfigured, outdated, or completely abandoned. These forgotten VMs may still hold sensitive data or connect to internal networks, making them weak spots for cybercriminals to exploit. Even if only one VM is compromised, it could lead to much bigger problems for the organization.

Virtualization offers many advantages, like saving money, simplifying disaster recovery, and allowing easier scaling of resources. But these benefits come with responsibilities. Every virtual machine that's created needs to be tracked, managed, and secured. If that doesn't happen, the very system that was designed to help can become a liability.

Several common causes contribute to server sprawl. Sometimes, different teams create VMs without telling central IT. Sometimes, old test environments are never deleted. And sometimes, machines are cloned and reused without clear documentation. These habits may seem harmless, but in a large organization, they lead to hundreds or even thousands of unmanaged virtual machines.

Security risks connected to server sprawl are serious. Attackers are constantly looking for overlooked systems. VMs that haven't been patched or are running outdated software become easy targets. Even worse, when these machines are not documented properly, it becomes hard to detect if they've been breached. This makes incident response more difficult and increases the chance of long-term damage.

So, organizations must treat server sprawl as more than just a resource management issue. It is also a cybersecurity issue. And it needs a strong, focused approach. This paper aims to understand the root causes of server sprawl, the types of security threats it brings, and the most effective ways to deal with it in real-world environments. It is not just about using more tools - it is about adopting a culture of visibility, accountability, and continuous improvement in virtualized systems.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

II. BACKGROUND RESEARCH

Virtualization became popular as a solution to reduce hardware costs and improve IT flexibility. Instead of running one operating system on one physical server, virtualization allows many virtual machines (VMs) to run on a single physical system. Each VM acts like a separate computer, with its own operating system, applications, and configurations. This makes resource use more efficient and supports better scalability.

Many researchers and industry experts have talked about the benefits of virtualization. It helps organizations lower power consumption, reduce physical space requirements, and recover faster during system failures. Virtual machines can be moved, copied, or restored much more quickly than traditional servers. These advantages have made virtualization a common choice in modern data centers (Pogarcic et al., 2012).

But with these benefits came new challenges. As more virtual machines are created, managing them becomes harder. Without proper planning, virtual environments can become overcrowded and disorganized. This leads to what is called server sprawl. According to Tank et al. (2019), server sprawl occurs when virtual machines are created faster than they can be tracked or managed. Over time, this results in many inactive, outdated, or unpatched VMs that are still running or stored in the environment.

Security issues caused by server sprawl have also been discussed in past research. Some VMs might still have default credentials or missing patches, making them easy targets for cybercriminals. These machines may also contain sensitive information from previous use. If not properly decommissioned, they remain a risk to the organization (Wei et al., 2009).

Researchers have recommended several strategies to control server sprawl. Role-Based Access Control (RBAC) is one way to prevent unauthorized creation of VMs (Hirano et al., 2008). This method allows only specific users to create, modify, or delete virtual machines based on their role in the organization. Automated tools have also been suggested to help track VM inventory, apply updates, and manage the VM lifecycle.

Another suggestion from Kolahi et al. (2020) is to use centralized management systems. These platforms give IT administrators a complete view of all running and inactive virtual machines. They can apply policies, monitor performance, and shut down unused resources. This kind of control is important in keeping the virtual environment clean and secure.

Earlier studies have also emphasized the importance of securing the hypervisor. The hypervisor is the software layer that runs and manages all virtual machines on a host server. If an attacker gains control of the hypervisor, they can access all connected VMs. So, protecting the hypervisor through hardware-assisted security and isolation is critical (Lombardi & Di Pietro, 2010).

In summary, past research has clearly shown that virtualization brings many advantages, but also new risks. Server sprawl is not just a technical mess, it is a growing security threat. The background work by researchers between 2008 and 2020 lays a strong foundation to understand how these risks grow and what can be done to reduce them. The next chapter will explain the research method used in this study to explore these issues further.

III. RESEARCH METHOD

This research is based mainly on reviewing existing literature and analyzing real-world practices related to virtualization and server sprawl. A qualitative approach has been used. The goal is to understand how server sprawl becomes a security risk in virtualized IT environments, and what strategies are effective in managing that risk.

The first step of the research involved collecting secondary data. This includes academic papers, industry reports, and case studies published before or during 2020. Only peer-reviewed journals, conference papers, and technical whitepapers were used. The selection of documents was focused on virtualization management, security threats, VM lifecycle, and best practices for infrastructure governance.

To keep the study focused, documents were filtered based on a few key terms. These included “server sprawl,” “virtual machine security,” “hypervisor management,” “RBAC in virtualization,” “VM lifecycle management,” and “virtual infrastructure risks.” Only those studies that directly addressed challenges or solutions related to server sprawl in virtual environments were included in the final review.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

After selecting the materials, a thematic analysis was carried out. This means the content was read and broken down into themes. These themes included causes of server sprawl, types of security threats, warning signs of unmanaged virtual environments, and proven mitigation techniques. This helped in identifying patterns and repeated suggestions from multiple researchers.

In addition to literature review, industry case observations were included. These are not original interviews or surveys, but examples and summaries from companies that have shared their virtualization experiences publicly. These real-world situations helped support the findings from the academic papers. They added more depth and context, especially when discussing what works and what fails in controlling server sprawl.

This study does not use numerical data or statistical models. Instead, it relies on descriptive insights to explain the problem clearly and offer solutions that are practical and realistic. The goal is not to build a new theory but to make sense of existing knowledge and apply it to the growing security issue caused by server sprawl.

The next chapters will use this method as a base. First, the specific challenges of server sprawl will be explained in more detail, followed by the types of security threats it creates.

IV. CHALLENGES OF SERVER SPRAWL

Server sprawl creates several challenges in virtualized IT environments. These problems are not always noticed at first. In the beginning, creating virtual machines is easy and quick. That is one of the biggest reasons organizations start using virtualization. But the more machines are created without a plan, the more difficult it becomes to manage them later.

One of the first challenges is lack of visibility. When too many VMs are deployed across different departments or teams, it becomes hard to keep track of them all. Some teams may create virtual machines for short-term projects and forget to delete them afterward. Others may duplicate existing VMs just to test something. Over time, these machines pile up. Without a centralized inventory or a tracking system, IT administrators don't even know how many VMs exist, which ones are active, and which ones are abandoned.

Another challenge is wasted resources. Even idle or forgotten VMs continue to use memory, storage, and processing power. This can lead to performance issues for the entire environment. Some systems may slow down, while others may run out of capacity. In some cases, organizations keep buying new hardware without realizing that many VMs could have been removed or optimized. This goes against the original purpose of virtualization, which was to save resources and money.

VM sprawl also makes maintenance harder. Regular patching, updates, and security monitoring become difficult when the number of machines is too high. IT teams may not have enough time or tools to patch every VM. Some machines may run outdated software for months or years. Others may be misconfigured or run unnecessary services. These unmanaged machines can easily become entry points for malware, ransomware, or unauthorized access.

Documentation issues are also common. In large organizations, VMs are often created without proper naming conventions or records. It becomes unclear who owns the machine, what it is used for, or whether it is still needed. When an incident happens, like a network breach, it takes a long time to figure out if an unknown VM is part of the company's infrastructure or not.

Another serious problem is lack of ownership and responsibility. When teams work independently, virtual machines are sometimes created and left behind when the project ends or when team members leave the organization. Without a clear policy on who manages VM cleanup, no one takes responsibility. These orphaned VMs stay in the system for years, quietly consuming resources and introducing risks.

Sometimes server sprawl happens because the organization is growing too fast and virtualization makes it easy to expand. But growth without governance is risky. A culture of "create first, manage later" leads to disorganization. Even if a security team exists, they can't protect what they don't know exists. At present days, the challenge has grown even more complex with many organizations running hybrid environments that combine on-premise virtualization with cloud-hosted virtual machines, making visibility and control across platforms a critical concern.

In summary, server sprawl challenges start small but grow fast. They include hidden machines, wasted resources, difficulty in updates, weak documentation, and confusion over ownership. These are not just technical problems. They affect security, efficiency, and cost. Without a solid strategy to control VM creation, track assets, and remove what is no longer needed, organizations end up with a cluttered and vulnerable virtual environment.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

V. SECURITY THREATS IN VIRTUALIZED ENVIRONMENTS

Virtualization was designed to make IT environments more efficient, but it also brings new types of security threats. When server sprawl happens, the number of virtual machines increases fast, and each new machine becomes a potential target for cyberattacks. The more machines that exist, the harder it is to monitor and protect them all.

One of the most serious threats is unpatched virtual machines. In many organizations, some VMs get ignored after they are created. These machines may still be running outdated software or operating systems that no longer receive updates. If they are exposed to a network, they become easy targets for attackers who are scanning for known vulnerabilities. According to Wei et al. (2009), unpatched VMs can act as silent doors that attackers can walk through without resistance.

Another major issue is unauthorized access. When access to virtualization platforms is not properly restricted, users might create or modify VMs without approval. This can lead to machines that are not secured or machines that are used for unsafe purposes. Without Role-Based Access Control (RBAC), it is difficult to limit who can do what inside the virtual environment. Hirano et al. (2008) explained that RBAC is important in preventing misuse and in making sure only authorized users can perform actions like creating, modifying, or deleting virtual machines.

There is also the problem of hypervisor attacks. The hypervisor is the layer of software that manages all the virtual machines running on a single physical server. If the hypervisor is compromised, all connected virtual machines can be affected. Lombardi and Di Pietro (2010) mentioned that hypervisors must be carefully protected using isolation techniques and secure configurations. If a hacker gains access to the hypervisor, they can control or spy on every VM that runs through it.

Notably, security researchers in recent years have highlighted the risks of side-channel attacks and privilege escalation vulnerabilities in hypervisors like VMware ESXi and Microsoft Hyper-V, emphasizing the importance of applying security patches and using hardware-level protections like Secure Boot and virtualization-based security.

Data leaks can also happen because of server sprawl. Sometimes, virtual machines are copied or cloned without cleaning the data inside. If these machines are left in the system or accidentally shared, sensitive data can be exposed. Old VMs that are not properly deleted may still contain customer information, passwords, or confidential files.

Zombie VMs are another threat. These are virtual machines that are still active in the system but serve no current purpose. Because they are not in active use, they are often not monitored or updated. However, they are still connected to the network. Attackers look for these forgotten machines because they are easier to break into. Once inside, attackers may use these machines to move through the network or launch attacks from inside the organization's infrastructure.

Another concern is resource exhaustion. Too many virtual machines running at the same time can overload the physical hardware. When CPU, memory, or disk space is stretched too thin, the system becomes slower and less stable. At that point, even security tools and monitoring systems may fail to work properly. Kolahi et al. (2020) showed how overloaded environments can cause delays in threat detection and lead to higher risk of security failure.

Lastly, poor visibility makes everything worse. When there are too many VMs without proper documentation or monitoring, security teams don't know what they need to protect. Attackers, on the other hand, only need to find one weak point. If they find an old VM with weak credentials or exposed ports, they can use it as a starting point to move deeper into the system.

To sum it up, server sprawl increases the risk of unpatched systems, unauthorized access, hypervisor attacks, data leaks, and network misuse. It also causes problems with monitoring and performance. These are not small issues. They can lead to serious breaches, loss of data, and financial damage. The next chapter will explore real-world insights and observations that support these concerns.

VI. EMPIRICAL INSIGHTS OR STUDY

Although this research does not include original interviews or experiments, it relies on careful observation of real-world practices, published case studies, and industry reports shared before or by 2020. These examples show how different companies faced server sprawl and what they did to solve or prevent the problem. They also help to confirm that the threats described in earlier chapters are not just theoretical, but actually happening in real IT environments.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

In one commonly discussed scenario, a large enterprise deployed hundreds of virtual machines across various departments. Each department had access to create their own VMs. At first, this approach worked well. It allowed teams to set up testing environments and launch new services quickly. However, within two years, the IT department found that over 40 percent of the VMs were either inactive, unpatched, or had not been logged into for several months. These virtual machines were still connected to the network, still consuming resources, and still posing risks. The organization had no central tracking system at that time, so they didn't even know who had created many of those machines. This is a classic case of server sprawl.

Another real-world case involved a medium-sized financial company that used VMware to virtualize its infrastructure. They had proper hypervisor security in place but had not implemented any role-based access control for virtual machine creation. As a result, many junior staff created VMs for training and testing and then left them running. Some machines used default usernames and passwords. A later security audit discovered that one of these test VMs had been accessed from an external IP, but since no one was monitoring it, the event went unnoticed for weeks. This situation supports the concern shared by Hirano et al. (2008) and others regarding the need for access controls and VM usage monitoring.

There are also examples from government IT departments that showed similar patterns. A report from one public-sector project found that after moving to a virtualized environment, the number of VMs had doubled within a year, but there was no VM lifecycle policy. Machines were created but never deleted. Some were used just once for software demonstrations. Others were restored from old images but never shut down. The result was a bloated system, with high storage use and declining performance. Security risks were also increasing because many of those old machines had outdated operating systems with known vulnerabilities.

What these examples show is that the core problems of server sprawl are not unique to one industry or type of organization. Whether in finance, education, healthcare, or public service, the problem shows up the same way. It starts small, with a few extra machines created for convenience. Then it grows quietly, until the organization realizes that it has hundreds of unmanaged virtual systems, and some of them may already be compromised.

Many IT professionals interviewed in published studies noted that fixing this problem after it has grown is much harder than preventing it in the first place. Once a virtual environment becomes messy, it takes time, tools, and dedicated effort to clean it up. And if there is no good documentation or tagging system, teams might even delete a VM that is still critical to operations.

These industry insights help validate the challenges and threats already discussed. They also support the need for strong mitigation strategies, which will be explained in the next chapter.

VII. MITIGATION STRATEGIES

Solving the problem of server sprawl in virtualized environments is not just about reducing the number of virtual machines. It requires a full plan that includes governance, automation, security controls, and continuous monitoring. These strategies must be practical and long-term. Otherwise, the same problems will come back again.

Many of these mitigation strategies are also recommended by well-known cybersecurity frameworks. For example, the NIST SP 800-53 guidelines emphasize proper access control and continuous monitoring of assets. The CIS Benchmarks provide hardening rules for virtual machines and hypervisors. Similarly, ISO/IEC 27001 outlines requirements for secure configuration management and lifecycle policies. Organizations can align their virtualization governance with these frameworks to improve both security and compliance.

One of the most effective strategies is to use centralized virtual machine management tools. Platforms like VMware vCenter or Microsoft System Center Virtual Machine Manager help administrators keep track of all VMs across different departments and hosts. These tools provide a dashboard where every virtual machine is listed along with its status, location, and resource usage. Kolahi et al. (2020) pointed out that centralized visibility is the first step in controlling virtual machine sprawl. Without it, no other strategy works properly.

Another important step is setting clear policies for VM lifecycle management. Every VM should go through stages: request, approval, creation, active use, and finally, decommissioning. This process needs to be written down and followed. If a machine is created, there should be a reason, a responsible owner, and a plan for when it should be deleted. Some companies even set expiration dates on VMs that are meant for temporary use. After the date passes, the machine is flagged or shut down unless someone extends it.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

Role-Based Access Control (RBAC) is also critical. As Hirano et al. (2008) suggested, access to create, modify, or delete VMs should be based on job roles. Not everyone in IT should have full access. For example, a junior tester might only be allowed to start and stop pre-approved machines, while a senior engineer can create or manage them. This reduces accidental creation and makes people more responsible for the systems they control.

Automated patching systems should be used to keep all active VMs up to date. It is easy to forget one or two machines during manual updates. Over time, those few machines become the weak points in the network. Using tools like Red Hat Satellite or Windows Server Update Services helps ensure that every VM receives updates at the right time. This lowers the risk of leaving known vulnerabilities open to attackers (Wei et al., 2009).

Regular cleanup audits are also necessary. Every few weeks or months, the IT team should review the list of all running and idle VMs. Machines that haven't been used in a long time should be flagged for review. If no one claims them, they can be archived or deleted. This keeps the environment clean and prevents buildup. Some organizations set up automated scripts that scan for inactive VMs and send reports to system administrators.

Protecting the hypervisor is another must. The hypervisor is the core of the virtual environment. If it is weak, all the VMs on it are exposed. Lombardi and Di Pietro (2010) recommended using hardware-level security like Intel TXT and enabling features like Secure Boot. It is also important to limit access to the hypervisor interface and monitor all changes closely.

Tagging and documentation may sound simple, but they are often ignored. Every VM should have a proper name, purpose description, and owner listed. This helps during audits, handovers, and security reviews. When teams can quickly tell what a machine does and who is responsible, it reduces confusion and errors.

Finally, organizations should invest in staff training and awareness. Sometimes server sprawl happens because people do not realize its impact. They create VMs and forget about them. If employees understand the cost, risk, and responsibility behind every virtual machine, they are more likely to follow the rules.

Mitigation is not a one-time fix. It is a habit. The more disciplined the organization becomes, the easier it is to keep server sprawl under control. The next chapter will summarize the main findings observed through this research.

VIII. RESEARCH FINDINGS

This research clearly shows that server sprawl is not just a technical inconvenience. It is a serious problem that creates both operational inefficiencies and security risks. From reviewing past studies and real-world cases, a few key findings stand out. These findings help explain how server sprawl grows, why it becomes dangerous, and what actually works in controlling it.

The first finding is that server sprawl is often invisible at the beginning. Most organizations do not notice the problem until it has already grown. Virtual machines are easy to create, and without strong policies, they keep piling up. In many cases, more than 30 to 40 percent of all VMs in an environment are unused or forgotten. These machines still use hardware resources and still pose a risk if not properly patched.

Another finding is that lack of governance is a common root cause. Organizations that allow teams or individuals to create VMs without approval or tracking end up with a mess. Machines are created for testing, demos, training, or personal projects, but no one deletes them. In some companies, there are even duplicate VMs running the same services because the old ones were never shut down.

The third major insight is that many security threats come directly from unmanaged virtual machines. As discussed in earlier chapters, unpatched VMs, weak credentials, old software, and leftover data can all become entry points for attackers. Machines that are not actively monitored or updated are often the easiest targets. This has been highlighted in research by Tank et al. (2019) and confirmed in several industry reports before 2020.

Also, it was found that technical solutions alone are not enough. Tools like patch management, hypervisor protection, or centralized dashboards are helpful, but they must be combined with people and process. Without clear policies, access controls, and regular audits, even the best software will not solve the problem. Some companies had good tools but still suffered from server sprawl because they lacked discipline and structure.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

One more important finding is that automation helps, but human review is still required. Automated scripts and patching systems can reduce manual workload, but they can't always decide which VMs are important and which are not. For that, human knowledge is still needed. That is why a good mitigation strategy needs both technical tools and responsible staff.

Finally, this study confirms that early prevention is more effective than late clean-up. Once a virtual environment grows without control, fixing it takes much more time and effort. It is better to start with strong governance, tagging, and access control from the beginning. The organizations that followed this approach had less sprawl and fewer security issues.

These findings support the need for continuous improvement in how virtualized systems are managed. The next chapter will offer future research directions for those who want to explore this topic further.

IX. FUTURE RESEARCH DIRECTIONS

Server sprawl in virtualized environments is a growing problem, and while this paper has covered the current challenges and solutions, there is still room for more research. The topic is closely connected to cybersecurity, cloud computing, and infrastructure management. So, future studies can go in several important directions to make these systems even safer and more efficient.

One of the areas that needs more attention is automation of VM lifecycle policies. While many tools already support VM creation and deletion, there is limited research on how these tools can be connected to business rules. For example, a virtual machine created for testing should automatically expire after a set time unless it is extended by an authorized person. Future work could look into intelligent policies that adjust themselves based on machine usage, role, and risk level.

Another direction is behavior-based monitoring of virtual machines. Instead of just checking if a VM is running, future systems could analyze how the machine is behaving. Is it sending large amounts of data suddenly? Is it using more CPU than usual? These changes might signal that the machine has been compromised. Research could explore how AI or rule-based engines can help detect these behaviors in real time, especially in large virtual environments where manual checking is not practical.

Hypervisor security is another area that still needs improvement. While current practices involve using trusted computing technologies and secure configurations, there is room to explore stronger isolation methods or backup hypervisors that can take over in case the main one fails or is attacked. Future studies could focus on lightweight secondary hypervisors or virtual sandboxing as extra protection layers.

Integration between cloud and on-premise virtualization management is also a growing need. Many organizations now use hybrid environments. Some virtual machines are hosted in local data centers, while others run in the cloud. Managing both together is difficult, and research could help in designing unified platforms that can enforce security rules across both types of systems without creating conflicts or blind spots.

There is also a chance to study the human side of virtualization governance. Most failures in server sprawl come from human habits—creating machines without cleanup, skipping documentation, or ignoring security alerts. Future work could look into behavior change, training strategies, and simple tools that make responsible VM usage easier for non-technical staff. Additionally, future research should explore the emerging issue of container sprawl, particularly in environments using Docker and Kubernetes. While containers differ from traditional VMs, they can accumulate similarly if not governed properly. Comparing VM sprawl and container sprawl - along with evaluating tools and practices that manage both - would offer more comprehensive insights into infrastructure hygiene.

Finally, more case-based studies and practical frameworks would be helpful. Every organization is different. A strategy that works for a university might not work for a government office or a private bank. So, researchers can collect more real-world data and propose flexible models that organizations can adapt based on their size, risk level, and resources.

In summary, this topic is still developing. Server sprawl will not go away on its own. As virtualization grows and IT systems become more complex, there is a need for smarter tools, stronger controls, and better understanding of both technical and human factors. Future research can play a big role in shaping safer virtual environments for the years to come.

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

X. CONCLUSION

Server sprawl is a serious issue in virtualized IT environments. While virtualization brings many benefits like reduced hardware cost, better flexibility, and faster deployment, it also opens the door to risks if not managed properly. One of the biggest risks is the uncontrolled growth of virtual machines. When too many VMs are created without planning or policies, the system becomes hard to manage and easier to attack.

This paper explored the causes of server sprawl and how it creates new security threats. It was found that unpatched VMs, poor documentation, abandoned machines, and weak access controls all add to the problem. These machines increase the attack surface and often go unnoticed until something bad happens. Once a threat gets inside through one of these weak points, it can affect the rest of the environment.

Several studies, published before 2021, helped confirm that server sprawl is not only common but also dangerous. They showed that many organizations suffer from this problem without fully realizing it. Real-world examples proved that even large companies with advanced virtualization platforms face the same issues when there is no strict governance or lifecycle policy.

To reduce these risks, this paper highlighted a set of clear and practical mitigation strategies. These include centralized VM management, role-based access control, automated patching, hypervisor protection, tagging, regular audits, and strong documentation. Training and awareness also play an important role. No tool can replace human responsibility. When teams understand the risks, they become more careful with how they create and manage virtual machines.

In the end, controlling server sprawl is about discipline. It takes a mix of technology, policy, and teamwork. The problem grows silently, but with the right plan and mindset, it can be prevented. Organizations that invest in strong virtualization governance now will be more secure, more efficient, and more prepared for the future.

REFERENCES

[1]. Hirano, M., Shinagawa, T., Eiraku, H., Hasegawa, S., Omote, K., Tanimoto, K., Horie, T., Kato, K., Okuda, T., Kawai, E., & Yamaguchi, S. (2008). Introducing Role-Based Access Control to a Secure Virtual Machine Monitor: Security Policy Enforcement Mechanism for Distributed Computers. Proceedings of the 3rd IEEE Asia-Pacific Services Computing Conference. https://doi.org/10.1109/apscc.2008.14

[2]. Kolahi, S. S., Hora, V. S., Singh, A. P., Bhatti, S., & Yeeda, S. R. (2020). Performance Comparison of Cloud Computing/IoT Virtualization Software, Hyper-V vs vSphere. 2020 Advances in Science and Engineering Technology International Conferences (ASET). https://doi.org/10.1109/aset48392.2020.9118185

[3]. Lombardi, F., & Di Pietro, R. (2010). Secure virtualization for cloud computing. Journal of Network and Computer Applications, 34(4), 1113–1122. https://doi.org/10.1016/j.jnca.2010.06.008

[4]. Pogarcic, I., Krnjak, D., & Ozanic, D. (2012). Business Benefits from the Virtualization of an ICT Infrastructure. International Journal of Engineering Business Management, 4, 42. https://doi.org/10.5772/51603

[5]. Rekha, G. S. (2018). A Study On Virtualization And Virtual Machines. International Journal of Engineering Science Invention (IJESI), 7(5), 51–55. https://www.ijesi.org/papers/Vol(7)i5/Version-3/I0705035155.pdf

[6]. Shital, M., Bahale, V., & Gupta, S. (2014). Virtualizing Disaster Recovery Management Based On Cloud Computing. International Journal of Research in Advent Technology, 2(5). https://ijrat.org/downloads/Vol-2/may-2014/paper%20ID-25201464.pdf

[7]. Tank, D., Aggarwal, A., & Chaubey, N. (2019). Virtualization vulnerabilities, security issues, and solutions: a critical study and comparison. International Journal of Information Technology. https://doi.org/10.1007/s41870-019-00294-x

[8]. Wei, J., Zhang, X., Ammons, G., Bala, V., & Ning, P. (2009). Managing security of virtual machine images in a cloud environment. Proceedings of the 2009 ACM Workshop on Cloud Computing Security. https://doi.org/10.1145/1655008.1655021

IARJSET

ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Vol. 8, Issue 9, September 2021

DOI: 10.17148/IARJSET.2021.8980

BIOGRAPHY

Nirjhor Anjum is a cybersecurity expert, eGovernance strategist, and digital services specialist with a strong academic and industry background. As a National Cybersecurity Consultant for the Ministry of Planning (Bangladesh), he is contributing to eGovernance projects and their security. As an Assistant Professor at DIU and an active researcher, he has contributed to cybersecurity, AI-driven automation, and software engineering through peer-reviewed publications. He has taught more than 2,000 students as an academician. In executive roles in the IT industry as CTO, CBO, and CAO, he has led cybersecurity, eGovernance, FinTech, enterprise automation, business intelligence, and various other enterprise digital transformation projects. With expertise in large-scale IT infrastructures, SaaS solutions, and public-sector digitalization, he is a thought leader connecting research with real-world innovation.

Md Rubel Chowdhury is a growing researcher specializing in computer networking, cybersecurity, and web development. He is working as a Web Developer and is involved in building secure web applications at SuperbNexus Limited. For his performance in 2020, he received the Rising Star Award in 2021 at SuperbNexus. His expertise includes vulnerability assessment, cloud server management, PHP development, CMS management, QA, Laravel, OpenCart, and more. He is a top graduate with the highest CGPA of 3.92 from Bangladesh University. His work focuses on web security and digital infrastructure resilience.

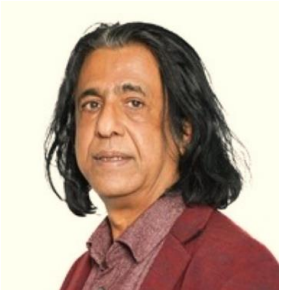

Md Anwarul Kabir is an Assistant Professor of Computer Science at American International University-Bangladesh (AIUB). With an academic background in Physics from the University of Dhaka and an MSc in Computer Science from the University of Wales, Swansea (UK), he teaches both undergraduate and postgraduate courses. His areas of teaching include Software Engineering, Systems Analysis and Design, Software Project Management, e-Governance, and Ethics in Computer Science. He has supervised numerous final-year projects, contributing to developments such as an expert system for child leukemia, a Bangla C compiler, and a preprocessor for Bangla OCR. His research focuses on software requirement analysis, open-source process models, computer security and computer ethics.