VIDDER

Segmentation for Security

Do It Right Or Don't Do It At All

Executive Summary

During the last 30 years, enterprises have deployed large open (flat) networks to provide convenient connectivity to everyone in the enterprise. Today, those networks make it easy for adversaries to quickly compromise servers containing intellectual property, financial data, and the personal information of employees, partners, and customers.

The adversary's entry point into the network can be achieved by compromising a device attached to the corporate network. From that foothold, adversaries can move laterally to compromise internal servers, including "shared service" applications, which will then provide them direct access to databases and file servers of interest.

Security professionals have been advised to segment their networks to defeat lateral movement. But traditional network-based segmentation approaches have failed. Data Center segmentation is only effective if combined with a method to control user access to Data Center partitions, which is difficult-to-impossible using traditional network segmentation techniques.

By deploying a "trust-aware" boundary between the corporate access network and the data center (or other areas where servers are deployed), *zero-trust partitions* can be deployed economically to insulate critical applications from compromises and attempted breaches that might be occurring throughout other areas of the corporate network.

This paper describes how to use Trusted Access Control to properly segment your network to limit the impact of any successful device compromise or credential theft.

TABLE OF CONTENTS

Adversaries Have Many Entry Options

Some of the many ways adversaries can gain a foothold in networks are by stealing credentials, compromising a device on the corporate network, or by physical access.

Credential theft may allow adversaries remote access VPN connectivity to the LAN, or access to an internet-facing server the attacker can then exploit to gain LAN access.

Another way for adversaries to gain a presence is to compromise an employee's PC via a phishing attack or some other form of social engineering. There are lots of different types of devices and lots of ways to attack them.

Or adversaries can just walk into a company office anywhere in the world and plug a "pen testing device" into an Ethernet port and gain global access. And that's just a few of the ways. In summary, corporate networks have huge attack surfaces which are growing in size and diversity.

Flat Networks Make for an Easy First Hop

Today's corporate networks are built to allow any device to potentially access any other resource on the network, including access from all branch offices and remote user PCs.

Unfortunately, these networks provide the same "easy connectivity" for adversaries who successfully compromise any device on the global network.

Once the foothold is established, any device on the network with any sort of vulnerability - user devices, printers, networked cameras, IoT devices - can be compromised in the attempt to access or at least move closer to the sensitive data that attackers desire.

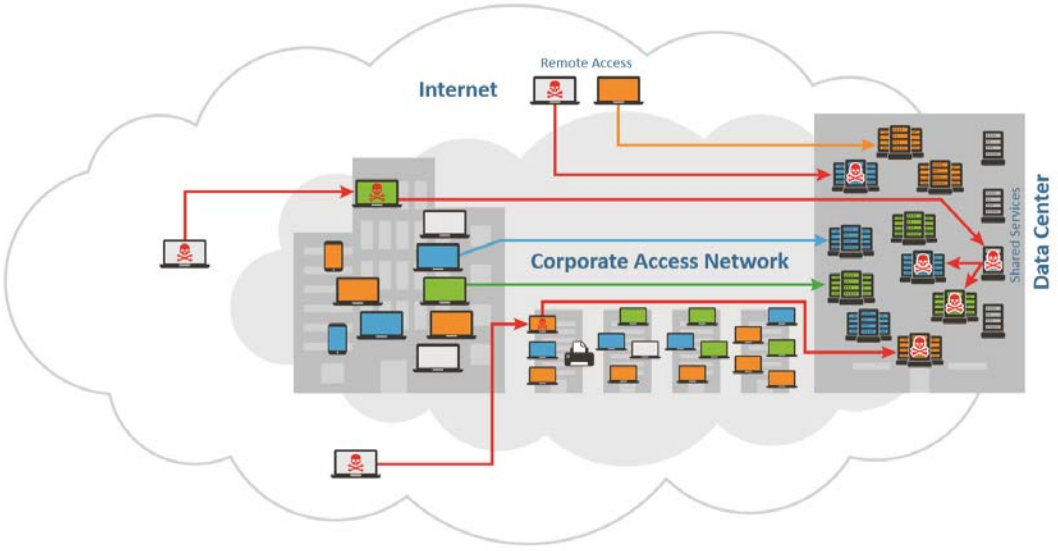

Most of the sensitive apps and data are accessed via servers. In flat networks, adversaries can move through the network and eventually directly connect to any server to find vulnerabilities or configuration errors that allow those servers to be compromised.
According to a recent Rapid7 research report[i], 96% of their pen tests of internal networks found at least one server with vulnerabilities that could be exploited to compromise the server.

A NOTE ABOUT SERVERS

One tends to think servers have just a couple of services running on them – services such as http/s and RDP. However, Windows Server 2016 can have over 100 different services listening on almost that many ports.

Every one of the 100+ services can have a vulnerability or configuration error that would enable an adversary to compromise it.

Network Segmentation Doesn't Solve the Problem

In the past, to try to reduce the openness of such environments, almost everybody segments users from data center servers by putting each into different subnets and putting firewalls in between. Beyond that, sometimes the data center is segmented so that the servers of one application cannot connect to the servers of other applications. This is sometimes combined with network segmentation in the access network to put users/devices into different subnets than servers, and then configure rules to manage the flow of traffic between the subnets.

But configuring static and complex firewall rules based on network segmentation at the subnet level is a blunt instrument that is time consuming to create and maintain. Worse than that, it is ineffective. There are multiple reasons for this.

First, when doing segmentation in the data center, servers running shared services cannot be isolated from other server subnets. They must be able to connect. In many cases, the servers running shared services are also reachable from the user network. Any compromise of such servers can easily traverse network segmentation boundaries to access any other server on any data center subnet. And some of those servers may even have Domain Admin credentials. If the adversary pops one of those servers, it is game over. For adversaries, such "shared services servers" are a key prize to find and compromise because they provide a direct path to so many other internal servers.

YOU CAN'T AVOID SHARED SERVICES

Many servers need to connect to many other servers.

For example, the web server needs to access the database, but so does the business intelligence application that needs to analyze the contents of the database, and so does the scanner that needs to access the list of running processes in the server, add to that back-up servers, patch servers, log servers, boot servers, etc.

These servers are often part of a group of servers called "shared services" because they have a legitimate business reason to connect to the servers of many different enterprise applications. This makes it impossible to completely partition servers into subnets with full traffic blocking between them.

Secondly, the corporate access network cannot be segmented in a manner where there is a clear mapping of subnets in the corporate network to subnets in the data center. IP addresses *might* indicate where on the network a user or device connected, but they rarely can be correlated to user *authorization*, as users with very different access privileges are often operating on the same floor or the same branch office, and, therefore, are usually configured within the same subnet.

Furthermore, source IP addresses have no value at all in assessing whether the user is who he or she claims to be, or whether the device the user is operating on is "clean". Compromised devices have IP addresses that "pass the test" from a firewall's point of view.

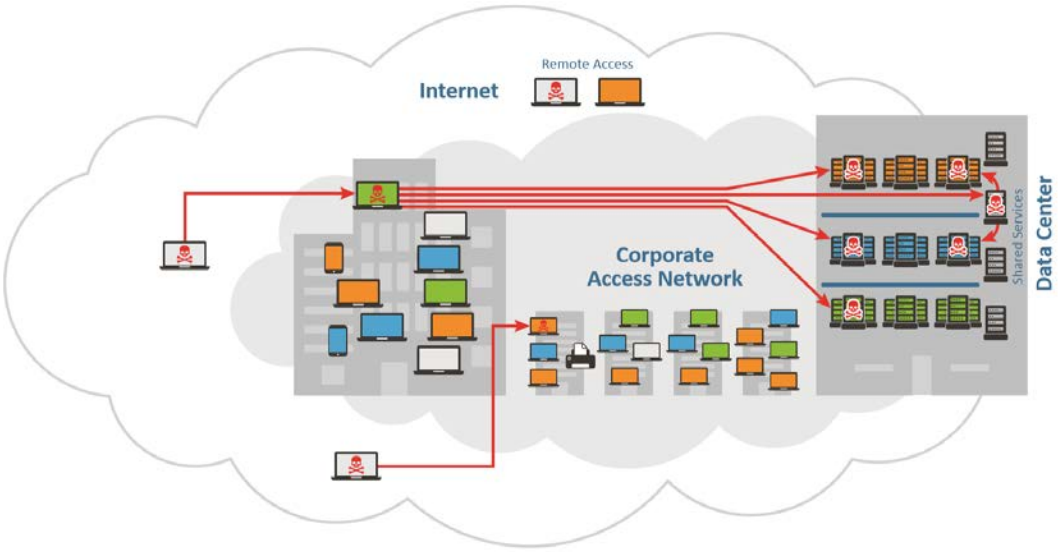

As a result, traditional network segmentation, both in the data center and the access network, is ineffective at thwarting adversaries' ability to move laterally through the network to access valuable data, once they gain an internal foothold.

Summarizing the Problem

1) The adversary's foothold is any device on the corporate network.
2) Enterprise servers have 10's of open ports that may have vulnerabilities or misconfigurations.
3) The network provides an easy path to enterprise applications and data once the adversary establishes a foothold.

A New Approach to Segmentation

Make the First Hop the Hardest Hop

Network-based firewalls at the interface between users and servers creates, at best, a speed-bump to protect against compromises moving from the corporate access network into servers and applications. This is an interface where you don't want a *speed-bump*. You want a *barrier*. The first step in executing an effective segmentation strategy should be to create a strong barrier between users and servers that can execute *trust-aware* policies for controlling access to applications.

Trust-aware means the access control system should act based on deep and extensive knowledge about the user, the device being used, it's location, and the sanctity of the software on that device.

A trust-aware access control barrier verifies that the user is truly who they claim to be via strong multifactor methods and confirms their authorization to use an application *before* they can interact with, or even see, protected servers and applications.

A trust-aware access control barrier verifies that the software running on an authorized user's device is running the expected client security software, that the entire software on the device does not have critical vulnerabilities, and that the device has not been successfully compromised.

A trust-aware access control barrier creates a *zero-trust partition*, which reduces the attack surface of servers to just the presently active authorized users, and then further reduces the attack surface so that even compromised devices of authorized users can't interact with or see protected applications and servers.

Such a barrier between corporate access networks and servers enables **Trusted Access Control**, which prevents adversaries who gain a foothold from proceeding any further. They can't turn that first success into what they really want – which is access to servers, applications, and data.

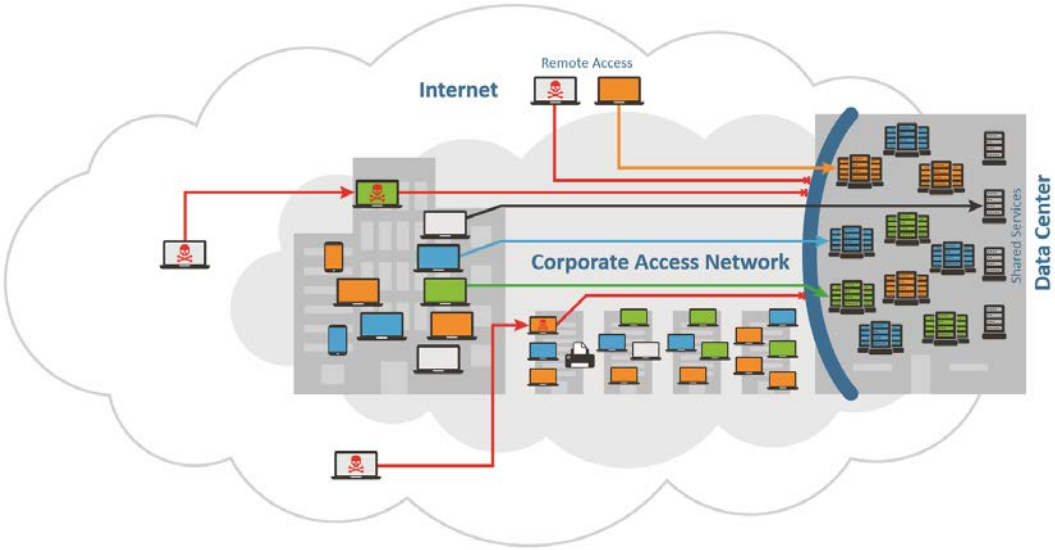

Creating such a barrier is what Vidder's PrecisionAccess does.

First, it completely isolates servers from any users and devices that are not authorized to access them. Uniquely, PrecisionAccess uses application-layer tunnels to reach the servers, which means that authorized users do not have access to a network of servers, they don't even have access to the whole front-end server. Rather they only have access to the authorized port on the authorized server – not the 10's of open ports that server is likely to have. This defeats server exploitation. That is, it mitigates server vulnerabilities and configuration errors by removing access to them by unauthorized users.

Second, PrecisionAccess has integrated multifactor authentication to defeat credential theft. Uniquely, it uses "the ideal MFA" that not only defeats all forms of compromised passwords, but is transparent to users, easy to deploy to all applications, and proven to work (see the white paper, "The Ideal MFA").

Third, PrecisionAccess has Trust Assessment to prevent compromised PC's from moving laterally. It is the only solution to combine posture assessment, vulnerability assessment, and threat assessment to defeat the adversary's movement from the initial entry point to any protected server.

By using PrecisionAccess to execute Trusted Access Control between users and servers, adversaries who establish a foothold on any device in the corporate network get "stuck" there. They cannot find a path to access data delivered by protected servers and applications.

Is There Any Value to Network Segmentation?

One might ask – if I have such a powerful boundary between my user devices and my servers, do I need to do any traditional network segmentation at all?

The answer is – probably not in the corporate access network. But traditional network segmentation between servers in the Data Center can be a useful complement to add a layer of security in the Data Center.

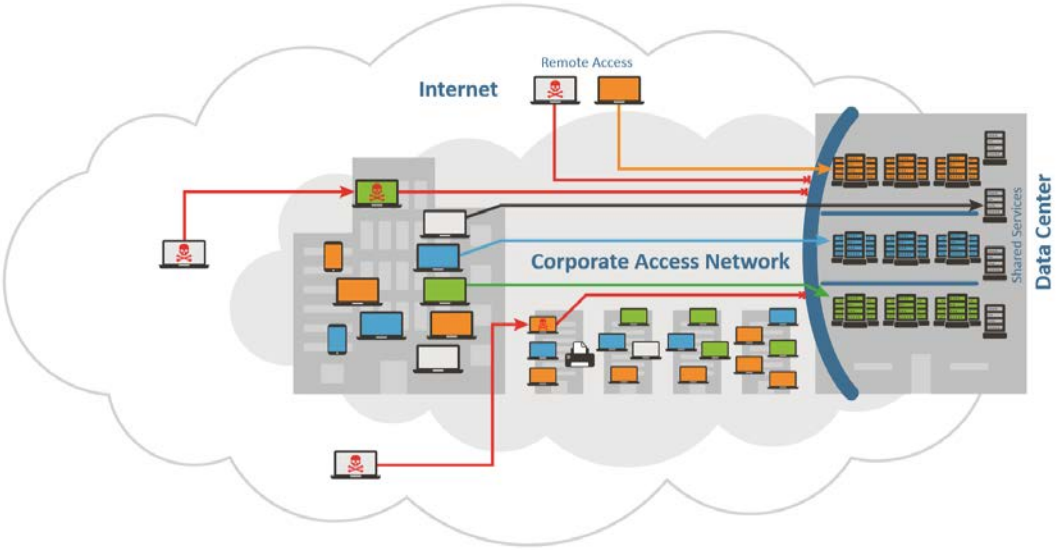

In the diagram above, user access to the green, blue, and red servers is protected using Trusted Access Control. And, then, traditional network data center network segmentation is used to eliminate lateral paths between the groups of servers. If additional segmentation is needed within the server partitions described above, and/or to further hone the access of the shared servers to the rest of the data center, Data Center micro-segmentation solutions can be deployed to achieve more granular control over the data center topology.

Another situation where data center network segmentation augments Trusted Access Control is when Trusted Access Control is used to secure access to specific critical business applications. Such applications might be associated with a software development team, financial operations, HR and PII information, critical Intellectual Property, workgroup operations, divisional or BU specific applications, or a variety of other reasons why it might be highly desirable to raise the security posture of a subset of corporate applications by creating a *Secure Enclave*. In these cases, the data center components of such applications need to be protected from the lateral movement of other servers in the data center that aren't as well protected on the front end.

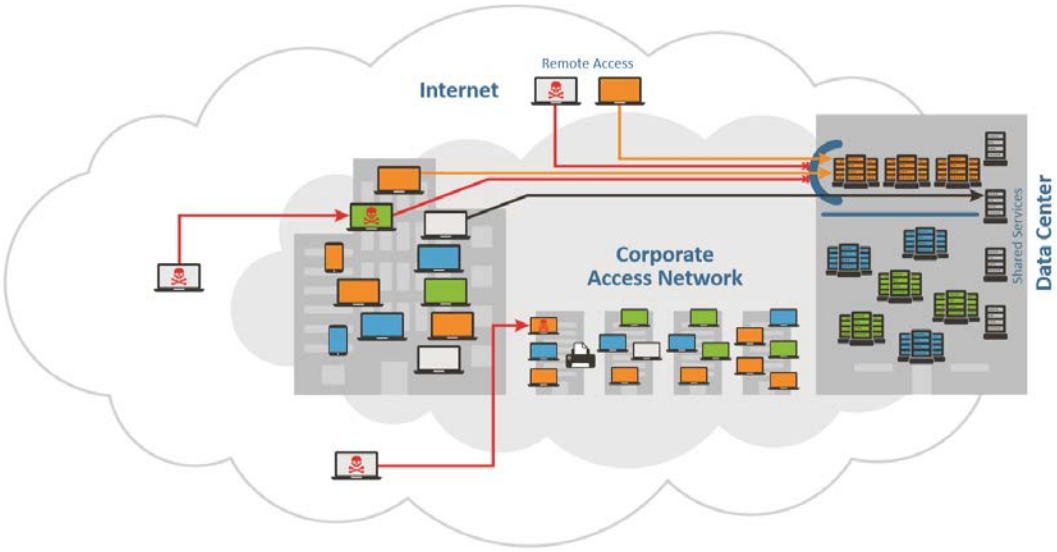

In the diagram above, user access to the orange servers is protected by Trusted Access Control. Data Center network segmentation is used to eliminate lateral paths between the orange servers and all other servers except the shared services servers. This approach is effective only if the shared services servers in the data center are protected by Trusted Access Control, because those servers, by design, need to be able to connect to the orange servers, and, therefore, any compromise of them could spread into the orange servers.

Thus, there are many examples where data center segmentation itself may not be effective, but when combined Trusted Access Control becomes a very important component of an overall security-by-segmentation strategy.

Extending Segmentation to the Cloud

As networks expand to include cloud infrastructure some of the problems with traditional approaches to segmentation introduce new challenges. For hardware-bound solutions, elaborate hairpins from the user to the data center to the cloud add cost and latency. And other approaches, including network access control, are not extensible to the cloud.

Trusted Access Control, however, is easily extensible to the cloud, allowing your teams to enforce a Zero Trust Partition consistently across data center and cloud environments. No other solution offers an equivalent hybrid cloud-ready trust-aware control barrier.

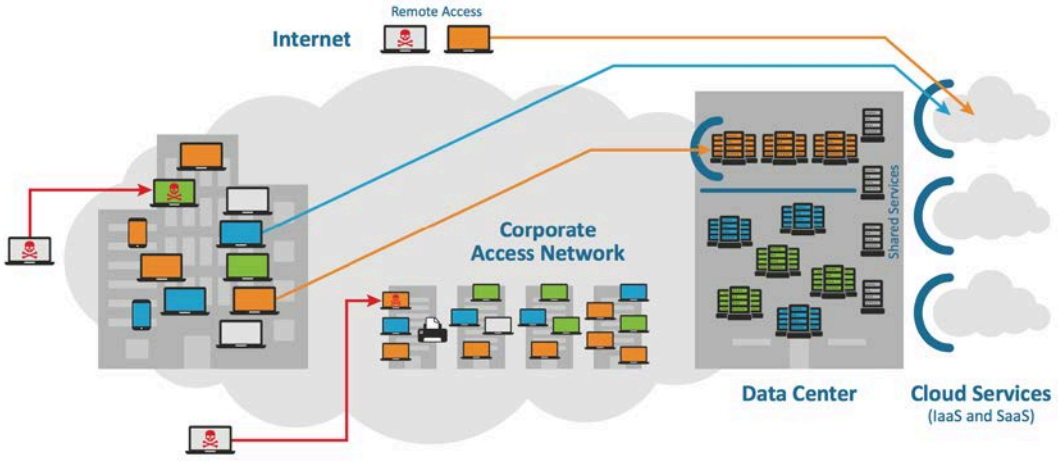

Are There Other Options for User-to-Server Segmentation?

Windows Authentication

In theory, the authentication provided by NTLM and Kerberos, coupled with the user-to-server authorization provided by the Domain Controller, should provide strong access control. But we know it doesn't. The problem is that to execute these functions requires that servers be in a "listen all" mode for all users and devices that might want to connect to them. Server vulnerabilities can be exploited. NTLM hashes can be stolen, impersonated, and relayed; tokens can be scraped from RAM; and PCs can be compromised with backdoors added. Bottom line, there are many ways to defeat the access control of Windows Authentication.

Network Access Control (NAC)

NAC exerts control at the wrong place. It attempts to keep rogue devices off the network but does not segment unauthorized users from authorized users per application. NAC is also time consuming and expensive to implement, and complex to manage.

The Next Gen Firewall (NGFW)

NGFW as defined by industry analysts, is the combination of a stateful firewall and an IPS. This can be useful in segmenting the servers in the data center by CIDR blocks but will not segment the users from the data center applications because different groups of users are not segmented by CIDR blocks.

Application-Aware Firewalls

Application-Aware Firewalls integrate with Active Directory or other identity systems to determine the user associated with a Source IP address and exert policy based on that knowledge. This is closer to the desired capability at the user-to-server boundary than anything else available to an enterprise in the past. But such firewalls don't perform authentication and therefore anything that spoofs the authentication system spoofs them as well. They are also not trust-aware in that they have no knowledge of the security state of the software operating on user devices, and they are generally very expensive when running at the scale of the boundary between the corporate access network and the data center.

Summary

Many devices on the corporate network can be successfully attacked and compromised by today's sophisticated adversaries – creating a foothold to further the attack. In a flat network, it is nearly impossible to stop the adversary's movement.

Traditional network segmentation controls can be easily bypassed by adversaries, resulting in a broad and multi-faceted attack surface that is almost impossible to defend. However, inserting a trust-aware boundary (Trusted Access Control) between corporate access networks and servers creates a zero-trust partition that strands adversaries before they can reach critical assets.

As enterprises embrace hybrid cloud operating models, the enterprise is exposed to new attack vectors and traditional approaches to segmentation and access control introduce even more cost and complexity while becoming even less effective. Because Trusted Access Control is topology independent, it is the only solution for protecting hybrid environments with highly secure and consistent control.

About the Author

Brent Bilger was at Cisco in the late '80s and most of the '90s as the product manager of routers, which can use access control lists to create network segmentation. He was VP of product marketing at Altor Networks, inventor of the first virtual firewall and was the acting VP of product marketing at Nicira, which became the NSX product at VMware. Brent was also the 3rd employee and VP of product marketing at CloudPassage, the first host-based firewall company that supported elasticity.

[i] Rapid7 research report, Under the Hoodie: Actionable Research from Penetration Testing Engagements, Feb 8, 2017.