VECTRA

Six critical attack vectors to detect in your data center and private cloud

Move beyond segmentation

Data center and virtualized security are built in the image of traditional campus network security. Using the network perimeter as its model, the industry first focused on recreating firewall-like abilities to segment and enforce rules on the flow of traffic in the virtual data center.

This included simply porting traditional firewalls to run as virtual machines, and then progressed into more agent-based segmentation models that were closely integrated with the virtualization platform software itself. Both approaches are largely focused on how to enforce policies within the cloud data center.

However, creating and enforcing rules is not the same as catching cyber attackers. At the perimeter, firewalling is complemented with a variety of threat detection and prevention technologies, such as IDS/IPS, anti-malware solutions and web filtering.

And like their firewall brethren, many of these perimeter threat-prevention technologies have been simply ported over to run on virtual machines to replicate the campus network security architecture.

The problem is that cloud data centers are not simply perimeter security 2.0. Cloud data centers often encounter cyber threats in the more advanced phases of attack than the perimeter, and likewise, will experience different types of threats and attack techniques.

Specifically, perimeter threat prevention is overwhelmingly focused on detecting the initial compromise or infection (e.g. exploits and malware). Cloud data center cybersecurity must focus on detecting attackers who have already compromised the perimeter and have moved on to more advanced attack phases, such as internal reconnaissance, lateral movement, and data exfiltration.

What's needed is a unique approach to cybersecurity that concentrates on the inside of the network. This approach must employ artificial intelligence detection models based on machine learning and behavioral analytics to reveal attackers who have gained a position of trust inside the network.

This inside approach to attack detection is even more important in the cloud data center. Long before attackers reach a virtual workload, they will have already compromised the perimeter, an end-user device, and stolen administrative credentials.

Instead of firing exploits or malicious payloads directly at data center resources, data center cyber attackers are far more likely to use their position of trust to access or damage critical assets.

The practice of building cloud data center security in the image of campus security has led to a cybersecurity vacuum in the data center. While a great deal of effort has been spent to bring native policy enforcement and segmentation to the virtualized network, the same amount of innovation has not been invested in threat detection in the virtual network.

Worse still, the traditional models of perimeter intrusion prevention are not designed to find the mature attacks that face today's data centers. A new model is required that brings cybersecurity into the native virtualized environment of the cloud data center.

Native integration with virtualized environment

In addition to detecting advanced phases of an attack, data center cybersecurity must be natively integrated with the virtualization platform. An analysis of cloud data centers shows that 80% of traffic stays inside the data center. At the most basic level, a security solution must be within the virtual platform to have visibility into potential threats.

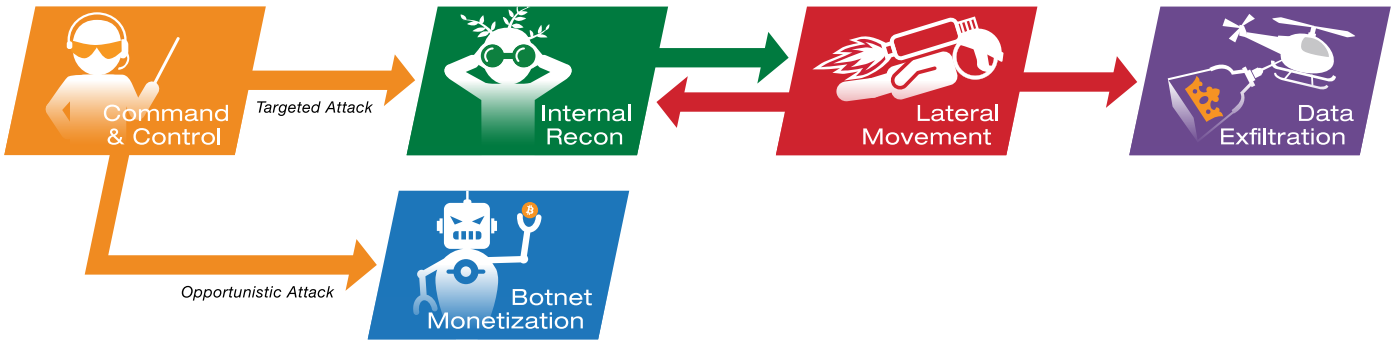

The phases of a cyber attack inside the kill chain

However, simply inserting into the virtualization platform is not sufficient. The virtual environment is always in flux. The dynamic and agile nature of virtualization is one of its most attractive qualities.

Developers can quickly spin up new applications. As needs change, applications can easily move to new locations, potentially on a completely separate physical host.

Any security solution that hopes to find attacker behavior and the progression of an attack must be able to retain context and visibility throughout all of these virtual environment changes. As a result, it isn't enough to just be another host in the virtual environment running security as an application.

Security must have the native visibility and context of the virtualized platform itself. Instead of simply being a piece on the virtual chess board, security must be able to retain an understanding of all pieces on the board, and retain that context over time. Without this context, behavioral modeling becomes all but impossible.

Unified visibility for all teams

In addition to detecting active attacks, it is necessary to have unified visibility into data center security that spans operational teams. By nature, cloud data centers involve the work of multiple teams, each with their own priorities and timelines.

Developers are typically driven to build applications quickly and the virtualization team often wants to deploy and support them as quickly as they can. Consequently, the security team is not always aware of changes that are made in the virtual environment.

The critical attack vectors

Data centers and the wealth of information they contain represent the ultimate prize for attackers. But unless the attacker gets lucky and finds an Internet-facing vulnerability, compromising a data center takes a significant amount of effort and planning.

As a result, cyber attacks that target data centers tend to be patient, mature operations that emphasize persistence and require flying below the radar of security teams.

This section examines the critical attack vectors and techniques that sophisticated cyber attackers use against data centers.

Co-opting administrative access

Administrators have unparalleled access to the data center and as a result are natural targets for attackers. Administrative protocols can give attackers backdoor access into the data center without the need to directly exploit an application vulnerability. And by using standard admin tools such as SSH, Telnet or RDP, attackers can easily blend in with normal admin traffic.

Since these phases of attack use allowed protocols and don't rely on malicious payloads, it is important to use behavioral models to detect cyber threats. If possible, perform behavioral modeling against actual network traffic because logs are not often available for the protocol being used.

Closing the local authentication loophole

In addition to the standard paths utilized by administrators, many data centers rely on local authentication options that can be used in an emergency. For example, if a domain controller or other authentication infrastructure fails, admins still need the ability to manage the data center.

In these cases, admins rely on local authentication to access the hosts and workloads they need to manage. However, these local authentication options are not logged and the same login credentials are often shared across hosts and workloads for the sake of simplicity.

While essential, these local authentication channels present a serious risk to the security of the data center. When attackers find the credentials by compromising an administrator, they can silently access the data center without fear of their activity being logged.

The administrative hardware backdoor

Local authentication offers an example of a backdoor that administrators – and attackers – can use to gain access to a data center. However, there are other examples that take the same approach and extend it deeper into the hardware.

While the data center is synonymous with virtualization, the virtualized environments and resources still need to run on physical hardware. Virtual disks are ultimately dependent on physical disks, and the physical disks run in physical servers.

Physical servers likewise have their own management planes designed for lights-out and out-of-band management. The management planes have their own management protocols, power, processors, and memory, which allow admins to mount disks and re-image servers even when the main server is powered off.

These actions are often performed via protocols such as the Intelligent Platform Management Interface (IPMI). While many hardware vendors have their own branded versions of IPMI, such as Dell iDRAC or HPE Integrated Lights-Out (ILO), they are all based on IPMI and perform the same functions.

These functions are independent of the data plane of the server. They effectively sit beneath the virtualization layers, below all host operating systems, and even below the BIOS of the main board.

IPMI and its related protocols have well-documented security weaknesses and are often slow to receive updates and fixes. The combination of IPMI vulnerabilities and its immense power make it a major weak point for attackers who are trying to subvert the security of the data center.

Advanced attackers aim low

Unfortunately, hardware problems in the data center don't end with IPMI. Advanced attackers, including nation-states, increasingly target physical servers, routers, switches, and even firewalls.

Tools, such as Synful Knock, have shown how attackers can burrow beneath the operating system to gain complete administrative control over a router and subsequently launch attacks against other systems and routers in the same network.

At a fundamental level, these tools represent rootkits that sit below the level of the operating system, making them extremely difficult to detect using traditional methods.

Disclosures of attack tools used by the Equation Group provide insight into the threat arsenal and techniques used by nation-states. This includes a variety of techniques and tools that implant software and firmware in a wide range of firewalls and security appliances.

These techniques allow attackers to infect the very devices charged with protecting the network, and then use those devices to launch attacks deeper into the network. Again, the strategic nature of these devices give attackers with the ability to monitor or reroute traffic and launch attacks from a position of trust.

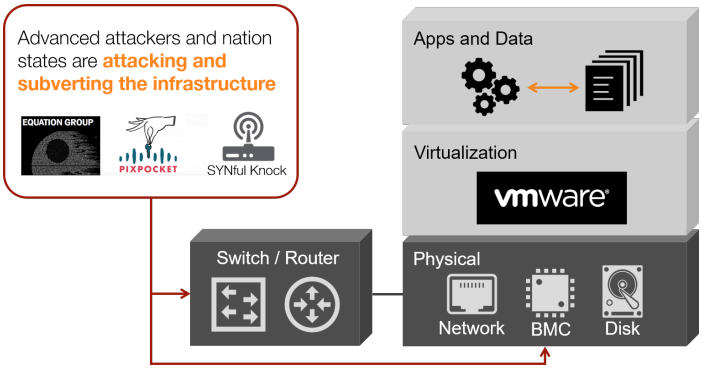

Data center attacks focus on the underlying physical infrastructure

Keeping an eye on data

The ultimate goal of most attacks is to steal data. Therefore, it must be the ultimate goal of security teams to always to identify attacks well before data is accessed, including attacks at the exfiltration phase.

Depending on their needs and skill level, attackers can use a variety of approaches to smuggle data out of the data center. The most obvious approach involves moving data in bulk out of the data center, either directly to the Internet or to an intermediate staging area in the campus network.

Subtle attackers may attempt to stay low-and-slow by patiently exfiltrating data at rates that are less likely to be noticed or arouse suspicion. Efforts can also be made to obscure data exfiltration in hidden tunnels within allowed traffic, such as Web or DNS traffic.

Blending physical and virtual context

Data centers are unique to their own organizations and vary based on applications and how users interact with them. The most common type of data center today is the private enterprise data center. Attacks against these data centers are typically extensions of attacks against the larger enterprise.

For example, attackers may have initially compromised an employee laptop via a phishing email or social engineering. Next, attackers typically look to establish persistence within the network by spreading from the initial victim to other hosts or devices.

To control the ongoing attack, attackers will plant backdoors or hidden tunnels to communicate back and forth from inside the network. Over time, attackers will map out the internal network, identify valuable resources, and compromise devices and user credentials along the way.

The most coveted stolen asset for an attacker is administrator credentials because they ensure near autonomy inside the victim's network. Administrator credentials are particularly essential for data center attacks, since administrators are often the only individuals who can access data en masse.

The key point is that an attack is typically at a mature stage by the time it reaches a private data center. The hidden command-and-control traffic, the reconnaissance, the lateral movement, the compromise of user and admin credentials are all prerequisites that lead up to the intrusion into the data center.

Each of these phases represents an opportunity to detect an attack and it is important for security teams to see as much of this context as possible before the attack reaches the data center.

This is why a consistent approach to cybersecurity – from the campus to remote sites to the data center – is important. Cyber attacks are complex, interconnected events and treating data center security as a separate silo only helps the attackers.

Conclusion

With their wealth of data and applications, today's data centers are the ultimate prize for cyber attackers. Yet while most data center security has focused on protecting the virtualized layers of the data center, real-world attackers are increasingly subverting the physical infrastructure that the data center depends on.

It is imperative to have the ability to identify cyber attacks that target data centers. With advanced detection models that expose attacks against application, data and virtualization layers in the data center, as well as the underlying physical infrastructure, security teams will be able to address critical vulnerabilities at every layer of the virtualized data center.

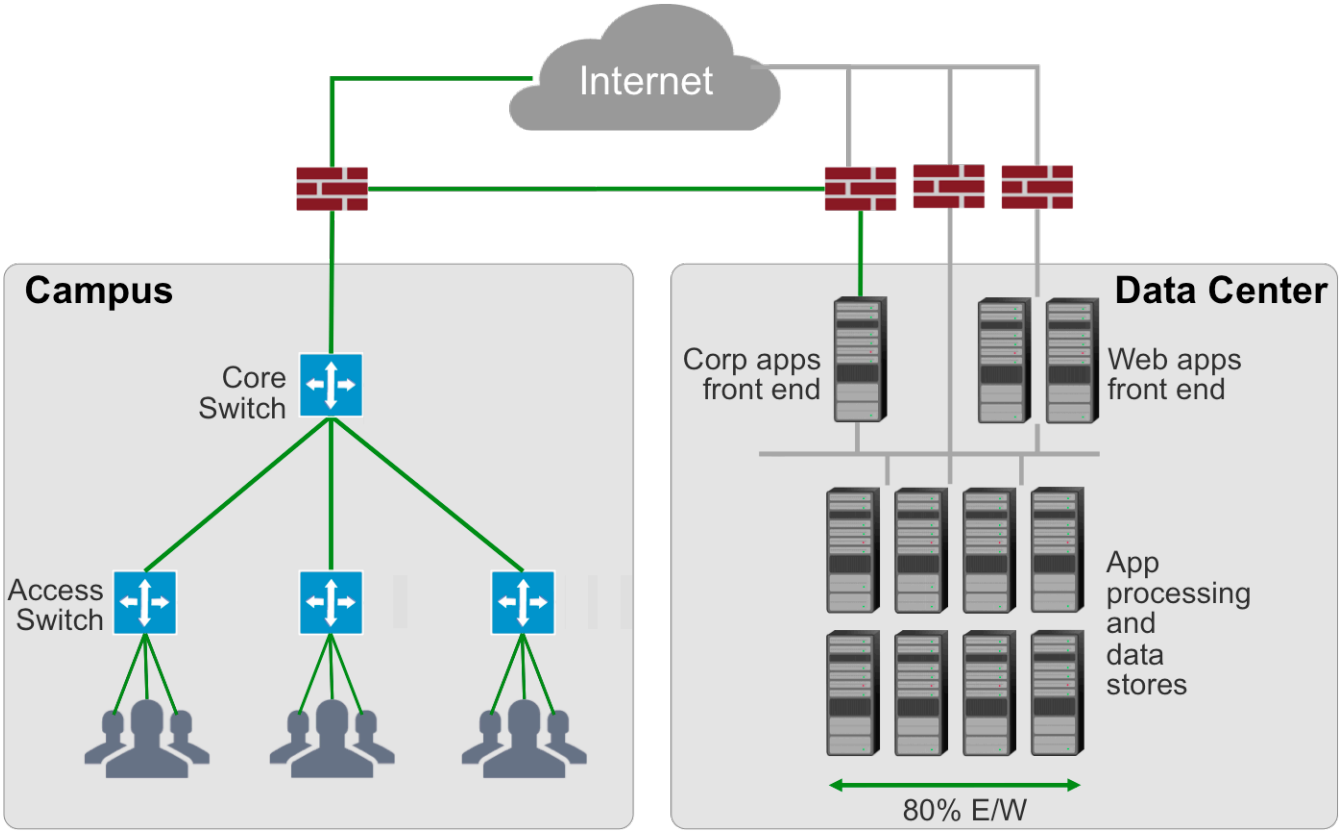

Detection of cyber attacks requires campus and data center visibility

VECTRA
Security that thinks®

Email info@vectranetworks.com **Phone** +1 408-326-2020
www.vectranetworks.com