

Securing Modern Data Centers and Clouds

Table of Contents

Executive Overview	1
Ingredients of security inside data centers	3
Data Center Security Gap: Detection and Response	4
GuardiCore delivers process-level visibility, breach detection, and response that work inside today's dynamic data centers	5
Summary	7

Executive Overview

Attackers require only one success to gain entrance to the entire data center; a breach is typically "game over." Due to the complex and dynamic nature of modern data centers, along with very high traffic rates, organizations face significant challenges in gaining visibility into application communications and putting proper controls in place to secure east-west (server-to-server) traffic. Such traffic now represents over 80 percent of all traffic inside the data center. This means that data center security architecture must be re-evaluated in order to meaningfully address security, shifting focus to application-layer visibility, granular micro-segmentation, real-time detection and automated response. An effective approach must secure a heterogeneous, dynamic environment with extremely heavy traffic rates.

Complete Protection Requires Comprehensive Adaptive Protection



The concept of Adaptive Security Architecture supersedes the blocking and prevention approach. It offers an entire lifecycle of protective capabilities, including detection and response, within an integrated architecture. GuardiCore believes detection and response are essential to security effectiveness.

Gartner, *Designing an Adaptive Security Architecture for Protection From Advanced Attacks*, February 2014



Aligned with the Gartner model of Adaptive Security Architecture,¹ the GuardiCore Centra™ Security Platform helps address this interior data center security challenge by providing a unique combination of process-level visibility, threat deception, semantic-based analysis and automated response. The Centra Security Platform detects, investigates and mitigates data center threats in real-time, reducing exposure, risk, and cost. Its distributed architecture offers full coverage of all traffic inside data centers and scales to very large network sizes and traffic rates, with low impact on hypervisor/ server performance.

Over the last decade, the modern data center has changed dramatically, raising a host of new security challenges. For one, the perimeter, which looks at less than 20 percent of the traffic, is not providing adequate security. Often split across multiple sites and infrastructures, the perimeter is only as strong as its weakest and most vulnerable point. The battle against intrusion is lopsided in favor of the hackers, who can make hundreds of thousands of break-in attempts, while an organization must defend against every one. A single mistake or slip-up is sufficient to endanger the whole data center.

Another challenge is that the aggregate traffic within a data center has swollen to multi-terabit levels. At these traffic rates, popular security technologies such as intrusion detection and prevention systems (IDP) and next-generation firewalls, which are based on deep packet inspection (DPI) and signature-based detection methods, consume too much compute time and become impractical for widespread use.

Hybrid cloud and multi-vendor environments

Organizations often locate their assets in hybrid private and public clouds, and in multiple network and compute environments. The same data center may have servers and workloads both on-premises and in a public cloud. On-premises locations may mix multiple infrastructure types, including legacy “bare metal” servers, VMware ESXi, VMware NSX, OpenStack, Cisco ACI switches, containers, and so on. Securing each of these assets is extremely challenging, and an effective and efficient security solution should cover all of them with a single platform. Legacy security tools don’t offer sufficient levels of agility, flexibility, and visibility for a highly diversified and active data center.

Attacker sophistication is growing

While security approaches are being challenged by rapid data center evolution, the sophistication and menace of advanced persistent threats (APTs) have taken a leap. The adversaries are no longer hacker “script kids.” Lured by huge financial payoffs, organized crime is making major investments in attack tools and teams. Hacking itself has become a professional, well-funded industry.

APTs—a global security crisis

Advanced persistent threats (APTs) have become a professional industry—a source of occupation and global commerce. Criminals have commoditized their threats, attack tools, and even teams as offerings to customer hackers: a supply chain of purchasable kits, fully prepared to invade and corrupt a data center. In the Carbanak case attackers used the Carberp backdoor, whose source code was for sale for \$50k last year, and is now available for free. Headlines reveal that attackers eventually succeed. Their sophistication is growing as they bypass signature-based and perimeter-based security solutions, exploit virtualization and cloud services, and thwart traditional security solutions.

“Rarely do you land where you need to be”

- Rob Joyce, Chief of the NSA's Tailored Access Operations (TAO) elite division Jan. 2016

Ingredients of security inside data centers

The limitations of prevention

Since the early 1990s, most organizations have relied on preventive perimeter defenses, which are designed to keep threats out of the corporate network. These include, but are not limited to, next-generation firewalls, IDP, and sandboxing. While perimeter security remains relevant, it has been outpaced too often by hackers' ingenuity and aggressiveness.

Internal data center segmentation tools, including VLAN separation and end point firewalls, can limit unapproved communication between servers. Micro-segmentation, a state-of-the-art technique involving advanced distributed firewalls, increases the ability to enforce policies on the communication inside data centers.

Separation needs to be dynamic and agile, and must work inside both VLANS and hypervisors. Implementing micro-segmentation within a data center is no simple task, especially in "brownfield" deployments (upgrades or additions to an existing networks that use some legacy components). With these, the IT and security teams must first discover and analyze existing intra-application connections, and make sure that newly installed separation policies will not break existing applications.

Additional preventive practices include patch control, which limits server vulnerability and exposure to certain attack vectors. Unfortunately, patching is not effective against re-use of credentials or zero-day attacks.

Endpoint security, application control, and encryption all play their respective roles, but statistics prove that they fall short of preventing breaches and their consequent damage and loss.

Costly consequences: the epidemic of data breaches

The direction, cost, and frequency of cyber attacks are shearing upward to record levels. The Ponemon Institute reported in 2015 that the total average direct cost of a data breach is now \$3.8 million, up from \$3.5 million a year ago, with the cost of each lost or stolen record up 6 percent.² Direct costs include hiring experts to fix the breach, investigating the cause, setting up hotlines for customers, and offering credit monitoring to victims. Indirect costs can go even higher in terms of lost business and goodwill as current and potential customers depart following a breach.

Data Center Security Gap: Detection and Response

Due to the issues already discussed in this paper, some attacks will inevitably breach traditional blocking and prevention mechanisms, putting more emphasis on the ability to rapidly detect and respond to a breach when it does occur. According to Mandiant,³ the average time to detect a breach is over 200 days, and in 67 percent of these cases, attacks are actually discovered externally.

To properly secure data centers, it's important to note that attackers operate differently once inside the data center as opposed to the perimeter: a server does not infect another server by sending an email with a document to open or a link to follow. Rather, lateral movement between machines inside a data center is more likely to utilize a password harvested at an earlier stage of the attack. Use of zero-day vulnerability is also more common inside data centers than at the perimeter. A security solution within the data center must be designed to detect and respond to relevant attack vectors.

Looking for signatures is not only too compute-intensive for data center traffic, as we mentioned in Part 1. It is also an insufficient defense, as advanced threats often use an unsigned "zero day" or polymorphic code that can change its signature while retaining its original behavior.

Statistical anomaly detection, a useful breach detection tool inside corporate networks, is also extremely challenging for use inside data centers. This is because of the dynamic data center environment, where traffic patterns are constantly changing; virtual machines (VMs) or containers and virtual networks literally "pop up" and disappear in a few seconds or minutes.

End-point detection is of limited value, as it is located on the same trust zone as the attacker, who would typically turn off endpoint security and delete or temper its log files.

The necessities of an effective security strategy

Real time breach detection is essential—and it must have an extremely low "false positive" rate, because a typical security team has very limited time and resources to monitor and investigate all potential security incidents. The detection function must cover all VM-to-VM traffic and scale to massive east-west traffic rates.

An effective security strategy must be able to instantly understand the full nature of the attack: its mode of spreading, its footprint, and where it has already spread. This includes automated analysis that can quickly assist security teams to confirm and prioritize the incidents that require rapid response. Once prioritized, automated response mechanisms, such as the ability to automatically contain the threat by isolating the compromised system, are essential for security teams to respond more quickly and limit the damage. Response should include mitigating the spread of the attack in real time and remediating infected hosts (or simply re-imaging them). Automated response is also extremely valuable in highly dynamic environments, due to the large number of virtual machines constantly appearing and vanishing.

How Badly is Security Broken?

229 days to detection

67% of attacks discovered externally

Mandiant 2014 Malware Report

Checklist for effective detection


- ✓ Detection is made in real time
- ✓ A low rate of "false positives" like configuration errors
- ✓ Coverage of all VM-to-VM traffic
- ✓ Scalable with east-west traffic rates
- ✓ Knowledge of attack method, footprint, and spread
- ✓ Generation of actionable report

GuardiCore delivers process-level visibility, breach detection, and response that work inside today's dynamic data centers

When an organization wishes to implement micro-segmentation, its team now confronts thousands of virtual machines running applications written by many people, some of whom have already left the company. So which applications should be allowed to talk to which others? If a rule blocks two machines from talking, will it take a toll on legitimate communications?

GuardiCore Reveal™, a key component of the GuardiCore Centra™ Security Platform, provides process-level visibility into applications and workloads combined with granular policy definition, giving IT and security teams the ability to discover, visualize, control, and monitor activity inside the data center.

GuardiCore—Security Inside Data Centers



Internal breach detection identifies and informs

The GuardiCore Centra Security Platform leverages state-of-the-art cloud agility and programmability to detect and respond to attacks at an early stage, as they begin lateral movement. It reacts to such “hints” as policy violations and suspicious activity between process-level communications, so it knows which connections to investigate deeper. The platform also seeks out malicious behavior such as backdoor installation, brute-force attempts, and log file manipulation. It confirms active breaches for fast prioritization and creates a footprint, which is then used to scope the impact by automatically identifying compromised systems across the data center. Major ingredients of this solution include:

- **Process-level visibility** provides complete monitoring and visualization of all applications and workloads down to the process level, delivering deep visibility into communications and flows inside the data center
- **Micro-segmentation policy** allows IT and security teams to define granular security policies between processes and monitors those policies for variations and suspicious activity. Variations from defined policies are presented in a comprehensive visual map, and logged as real-time security incidents for further investigation.
- **Distributed threat deception** interrogates, records, and monitors active and ongoing attacker sessions, looking for malicious behavior and gaining deep insights on attacker methods and spread.
- **Real-time forensics**, enabling total visibility and understanding of criminal behavior and security attacks.
- **Automated response**, allowing for real-time attack isolation and remediation of infected files and servers, thus stopping an attack at initial stage, before it has caused any damage.

Ocean's Eleven and the Changing Landscape of Cybercrime

CNN News dubbed a notorious \$1 Billion attack in 2015 the “Ocean’s Eleven of cyber strikes.” Starting with spear-phishing, attackers gained control of a bank employee’s computer and from there, installed a sophisticated modification of the Carberp backdoor, whose source code was then on sale for \$50k (and is now available free).

The next phase included lateral movement through the victim’s network via discovery and exploit of internal servers using Windows and Linux tools, including “legitimate” remote administration utility Ammy and a Secure Shell (SSH) backdoor. Intelligence gathering included collected video, audio, and keyboard tracking from victim computers. This enabled attackers to penetrate money-processing servers, financial accounts and ATM control; to transfer funds, and dispense ATM cash to money mules. The careful attack resembles state-sponsored cyber espionage and demonstrates how differently attackers and vectors behave once inside the data center.

Coping with the new data center environment

Distributed per hypervisor or server, GuardiCore Centra fully covers all traffic inside the data center, scaling to massive east-west traffic rates and network sizes with extremely low impact on hypervisor/server performance.

Centra is built to cope with the current elasticity of virtual infrastructures and their short-lived security policies. The solution scales to address burgeoning server workloads and sprawling virtual machines, keeping pace with aggregate east-west traffic rates of multiple terabits per second. Centralized analysis and detection are based on payload data, as well as on metadata, reflecting actual behavior rather than signatures alone.

Integrated into bare metal, VMWare, OpenStack, CloudStack, and containerized private clouds, as well as AWS and Azure public clouds, Centra provides a wide coverage of data center infrastructure.

Believable threat deception: automated, transparent decoy with low false positives

Utilizing a dynamic decoy tactic on suspect activity, GuardiCore Centra automates creation of decoy ports distributed across the host network. Innovative, high-interaction deception uses real machines, real services, and real IPs and ports—decoying connection attempts to machines which are actually part of the data center network.

This use of real machines rather than emulations reduces an attacker's ability to identify them as a decoy. The process is designed for high interaction with an attacker, without disrupting active sessions. Use of threat deception to detect an attacker based on actual malicious behavior leads to extremely low false positive rates. Centra's deception capability enables resource-constrained security teams to quickly prioritize, investigate, and follow up only on confirmed breaches, instead of having to chase down and investigate thousands of security events to determine their scope. Centra's threat deception features include:

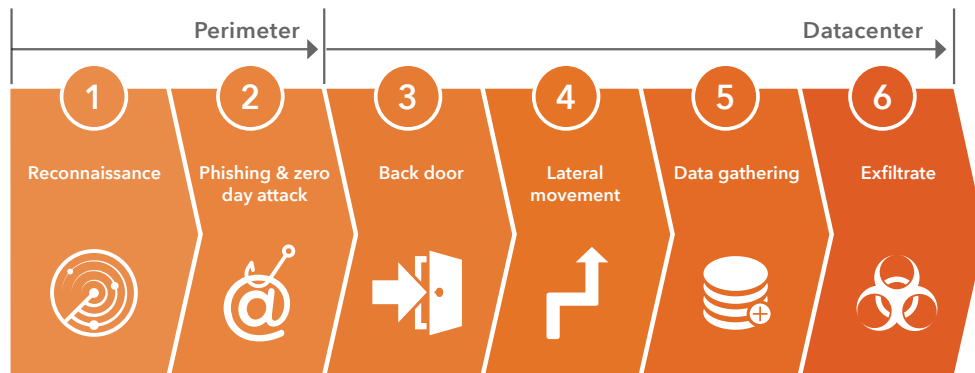
- **Transparent redirection** in order to investigate access attempts on filtered/firewalled/closed ports on existing machines; decoying IPs are seen in the network and targeted by attackers.
- **Dynamic deception engines** that are enabled whenever the attacker is trying to connect and gets blocked. Unlike legacy static honeypots, which use emulation and can be easily avoided by attackers, GuardiCore Centra uses real machines.
- **Active attack session investigation** looks for signs of malicious behavior such as backdoor installation, download, and execution of tools; creation of new users and escalation of their privileges; and tampering with log files.
- **Analysis, attack semantics, and understanding of behavioral characteristics**, including the attacker's specific footprint, and deep forensics, such as user credentials, exploits, and updated tools. GuardiCore's Automated Semantic Analysis™ capability evaluates an incident to identify the detailed semantics of the attack and deliver an actionable incident report.

“Sure, if you can stop something at the perimeter, by all means, do so. But inevitably, attacks will get through. They will spread laterally within our infrastructure. And we need [to be] better able to detect and respond when these types of incidents occur.”

- Neil McDonald, Gartner

Lockheed Martin: Breaking the Attacker's Kill Chain

The ability to detect attacks earlier in the kill chain is instrumental to saving money and mitigating damage. This landmark defeat of a sophisticated hacking attempt (possibly a state-sponsored APT)⁴ involved a tenacious cyber attack on the gargantuan network of Lockheed Martin—a supplier of military and aerospace technology to defense organizations, most prominently the U.S. Pentagon. Following the breach, Lockheed Martin created the methodology of the various steps involved in compromising a network, with the goal of stopping attackers in their tracks, the earlier the better.



This interpretation of the Lockheed-Martin Kill Chain shows six opportunities to stop a hacker who has infiltrated a data center network. The concept behind these stages is that in order to interrupt a hack and mitigate damage, an organization must “think like the hacker” and respond according to the phase of the kill chain in which the attack was detected. The 6th level is “exfiltration”—in which a hacker makes off with stolen information.

Mitigation and Remediation

Upon detection, the GuardiCore Centra Security Platform automatically generates a detailed attack footprint, which is then used to scan the data center and identify compromised systems. Upon identification, it then quarantines and blocks the attack from further spreading, as well as cleaning infected servers.

Summary

As sophisticated attacks become increasingly capable of invading and moving laterally within the data center, the perimeter is no longer a sufficient barrier. Attackers will certainly get through, and security must transform to move closer to the workloads it protects. A distributed strategy, offering continuous visibility and protection within the fabric of the data center is optimal for securing information. Detection is key to defense, capable of looking for the “needle in the haystack” amidst massive, dynamic traffic volumes. The GuardiCore Centra Security Platform provides security visibility, enables effective segmentation and micro-segmentation, and detects breaches once they’ve occurred—but before they can cause damage—identifying infected servers, isolating them, and rendering them harmless while mitigating the breach.

About GuardiCore

GuardiCore is a leader in Internal Data Center Security and Breach Detection. Developed by the top cyber security experts in their field, GuardiCore is changing the way organizations are fighting cyber attacks in their data centers.

More information is available at www.guardicore.com

1. “Designing an Adaptive Security Architecture for Protection From Advanced Attacks,” Gartner, February 12, 2014
2. “Cost of Data Breach Grows as does Frequency of Attacks,” Ponemon Institute, May 27, 2015
3. “Fewer Companies Able to Detect a Cyber Breach,” Matthew Heller, CFO.com, February 24, 2015
4. “Cyber Kill Chain,” Peter Beardmore, RSA Blog, October 2, 2015