Datacenter Revolution and Security

Protecting Modern Infrastructure while Enabling Agile and Efficient IT

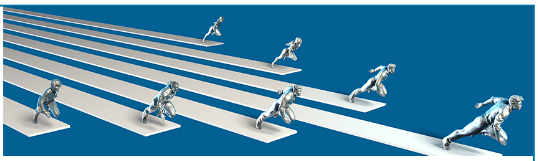

Gartner

Market Guide for Cloud Workload Protection Platforms

Neil MacDonald

26 March 2018

Server workloads in hybrid data centers spanning private and public clouds require a protection strategy different from end-user-facing devices. Security and risk management leaders should evaluate and deploy offerings specifically designed for cloud workload protection.

Key Findings

- Left uncontrolled, cloud environments inevitably spin into unmanageable complexity and have unique security needs that legacy security protection solutions do not address.
- As enterprises implement hybrid data centers, with workloads running on-premises and in multiple infrastructure-as-a-service providers, consistent security becomes difficult.
- The increasing adoption of containers complicates workload protection strategies.
- Numerous cloud workload protection platform vendors are emerging to address these unique requirements, including many smaller startups, which is confusing buyers.
- Signature-based approaches provide little or no value for server workload protection.

Recommendations

Security and risk management leaders tasked with controlling the security and compliance risks inherent in public cloud environments should:

- Use CWPP offerings – not desktop solutions – to protect the scale and dynamism of cloud workloads through native integration, programmability and orchestration.
- Require vendors to protect workloads across physical and virtual machines, containers and multiple public cloud IaaS, all from a single management framework and console.
- Require vendors to natively integrate with VMware, Amazon Web Services and Microsoft Azure, and Google Cloud Platform APIs, as well as tags for policy management.
- Use application control and whitelisting as the primary CWPP protection strategies and use traditional antivirus scanning only when the server hosts a file-sharing repository.
- Require vendors to API-enable security protection functions to be automated and integrated into DevSecOps-style workflows for scanning prior to deployment.

Strategic Planning Assumption

By 2022, 60% of server workloads will use application control in lieu of antivirus, which is an increase from 30% at YE17.

Market Definition

The market for cloud workload protection platforms (CWPPs) is defined by offerings specifically designed for server workload-centric security protection and are typically agent-based for deep workload visibility and attack prevention capabilities (see Note 1).

Left uncontrolled, cloud environments inevitably spin into unmanageable complexity. This makes security difficult or impossible; however, because the environments are software defined, organizations that want to manage this complexity can do so surprisingly easily, if they use the correct CWPP offering.

Market Description

CWPP offerings address the unique requirements of server workload protection in modern, hybrid data center architectures that span on-premises, physical and virtual machines (VMs), and multiple public cloud infrastructure as a service (IaaS) environments. In addition, support for protecting container-based application architectures is becoming a mandatory requirement. Vendors competing in this market offer one or more of the following capabilities for hybrid cloud workload protection:

- Core capabilities:
  - Workload configuration and vulnerability management
  - Network segmentation, firewalling and traffic visibility
  - System integrity measurement, attestation and monitoring
  - Application control
  - Supplemental memory and exploit protection
- Extended capabilities:
  - IaaS data at rest encryption and encryption key management
  - Workload behavior monitoring – essentially endpoint detection and response (EDR) for servers – also referred to as host-based intrusion detection system (HIDS)
  - Host intrusion prevention system (HIPS) and vulnerability shielding
  - Deception capabilities
  - Anti-malware scanning
- Capabilities that augment/verify foundational operational controls:
  - Vulnerability and configuration assessment from outside the workload
  - Multifactor authentication for administrators and basic privileged account management
  - Log management and monitoring

Market Direction

The market for endpoint protection has bifurcated between protecting end-user-facing endpoints (such as desktops and laptops) and protecting server workloads. The CWPP market addresses the protection needs of workloads in modern "hybrid" data centers that run in a mix of physical machines, VMs, containers, and private cloud infrastructure and almost always more than one public cloud IaaS. Leading CWPP offerings provide information security leaders with visibility and control across all of these environments with a "single pane of glass" – a consistent way to manage policy and monitor for risk.

We call this market for hybrid data center cloud workload protection "CWPP." To directly address the unique requirements of cloud workload protection, several of the traditional end-user-facing endpoint protection platform vendors have developed specific CWPP offerings. In addition, many new point solution vendors targeting CWPP have emerged. CWPP is of significant interest to Gartner clients, representing 15% of the 1,489 inquiries associated with the Cloud Security Agenda item during the past 12 months. Gartner has not yet formally sized the CWPP market, but we estimate it to be between $550 million and $600 million at YE17, and it is growing in double digits.

Several key trends are affecting the growth and development of the CWPP market:

- **Server workloads have fundamentally different protection requirements, especially in public clouds.** Most server workloads are restricted to a well-defined set of activities. In VM environments, this is typically one application per VM. In container-type environments, this can be down to a single process or application service. Thus, it is more effective and achievable to apply a default deny application control (also referred to as "whitelisting") model to server workloads than it is on end-user-facing endpoints.

- **Cloud-native applications scale elastically.** This requires protection to scale up and down on demand, with usage-based licensing models that support this. Simply running agents designed for on-premises servers and hoping these will work in IaaS is not sufficient.
- **Public cloud IaaS changes protection requirements.** Encryption of data at rest should be considered a mandatory best practice for public-cloud-based servers, although it's rarely required in on-premises data centers.
- **Most organizations have a stated intention to standardize on at least two IaaS providers.** This requires CWPP offerings that, in addition to supporting private cloud infrastructure, support heterogeneous IaaS cloud environments. They provide policy management consistency and help to close enterprise knowledge gaps.
- **Advanced attacks bypass traditional perimeter and signature-based protection.** Typically, these attacks are financially motivated and specifically target server and application workloads as a way to get to sensitive data or transactions. Advanced attacks have driven several key changes in server workload protection:
  - **Protection models that don't rely on signatures.** The primary protection strategy for CWPP (including container-based implementations) will be based on application control. This involves restricting what applications and supporting code (such as libraries) can run to a predefined set based on policy. Thus, all other code, malicious or not, is blocked by policy.
  - **The need for network traffic isolation, segmentation and visibility.** Advanced attacks will gain a foothold on one system, then spread laterally (east-west) within data centers and cloud deployments. The ability to segment east-west traffic more granularly is another key requirement. To help organizations understand application flows, visibility and visualization of these flows is a critical CWPP use case.
  - **The need for supplemental behavioral monitoring.** If an attacker has bypassed all the preventative controls and compromised a workload, how would you know? Many CWPP offerings provide behavioral monitoring, baselining and anomaly detection for defense in depth against targeted attacks.
- **There is a need for CWPP management to be automated.** In many cases, cloud server workload instantiation will be driven by templates and scripts, requiring security protection vendors to open up their protection capabilities via APIs for automated policy definition and provisioning. DevOps operating models need to incorporate security protection as well, delivering DevSecOps. With the highly automated provisioning requirements of cloud workloads, developers can't slow down to rely on a security professional to go to a CWPP console and set policy to activate a workload.
- **Impact of Spectre/Meltdown**. The early January 2018 disclosures of Spectre/Meltdown have created a focus on the security of public cloud-based workloads. The major public cloud providers have patched their firmware and hypervisor platforms; however, enterprises must still patch their VMs and container host OSs. Performance will be affected (typically by 5% to 20%), so some organizations are forgoing patches for specific workloads. If patches are not applied, network segmentation and application control become critical compensating controls.
- **Application developers are embracing containers.** Containers slipstream the delivery of new services from development into production quickly, often bypassing established security and compliance practices. As a best practice, containers should be scanned for known vulnerability and configuration issues, before they are released into production and protected at runtime. Some of the leading CWPP vendors don't yet have container support. To fill this gap, several new vendors in the CWPP market are designing solely to protect container-based development and deployments.
- **Serverless computing is gaining traction.** Public cloud IaaS offers a wide variety of platform as a service (PaaS) solutions, a subset of which is referred to as "function as a service" (FaaS) or "serverless." These architectures complicate security protection strategies, because there's no OS or container to instrument. In most cases, these services are used in conjunction with VM- and container-based architectures, so a traditional CWPP provides partial protection. Some CWPP vendors are already experimenting with extending approaches into serverless security and it will become a differentiator during the next 12 to 24 months. A deeper discussion of serverless protection strategies will be addressed in research later in 2018.
- **There is a mindset shift toward immutable infrastructure.** This is an operational model in which no configuration changes, patches or software updates are allowed on production systems. Patches and

updates are applied to the base ("golden") images and layers, then the production workloads are built fresh from these images and replaced, rather than serviced. With immutable infrastructure, CWPP protection will shift to a focus on application control and container lockdown at runtime, with a stronger emphasis on scanning in development, before workloads are deployed into production.

- **The legal and regulatory environment is changing, especially cloud.** Many server workload protection requirements are influenced or are direct requirements for compliance with legal and regulatory frameworks. A good example is the requirements for the protection of Payment Card Industry (PCI)-related workloads – specifically, file integrity monitoring, HIDS, patch management, anti-malware scanning or whitelisting (see Note 2), and network isolation.

  Likewise, the requirements for the European Union's (EU's) General Data Protection Regulation (GDPR) has reignited concerns around data residency. This is driving interest in data-at-rest encryption in public cloud IaaS, with customer-controlled keys architected for geofencing in a way that the cloud provider has no access to the keys. CWPP offerings help organizations demonstrate that they are aware of the activities happening in their IaaS environments. These activities are consistent with defined policies and are compliant with the expectations of regulators and other stakeholders.

  Collectively, these trends are creating requirements that are significantly different from traditional, end-user-facing endpoints and traditional physical servers. Information security leaders and architects must understand that new approaches for cloud workload protection are needed. Simply using a solution designed to protect end-user desktops or using an agent-based solution designed for dedicated physical servers in legacy data centers won't work.

Market Analysis

A large number of vendors offer CWPP solutions that vary widely in their capabilities. We recommend that organizations apply a risk-based security approach when developing their server workload protection strategies. Some workloads will host less-sensitive data and require fewer controls. Others with extremely sensitive data are likely to use more controls. Others may be protected behind network-based controls, such as firewalling and intrusion protection systems (IPSs), and not require these capabilities from the CWPP offering.

With this in mind, we have updated our hierarchy of workload protection needs (see Figure 1) to help enterprises prioritize their security investments and to help evaluate vendors with capabilities in this market.

Figure 1. Cloud Workload Protection Controls Hierarchy

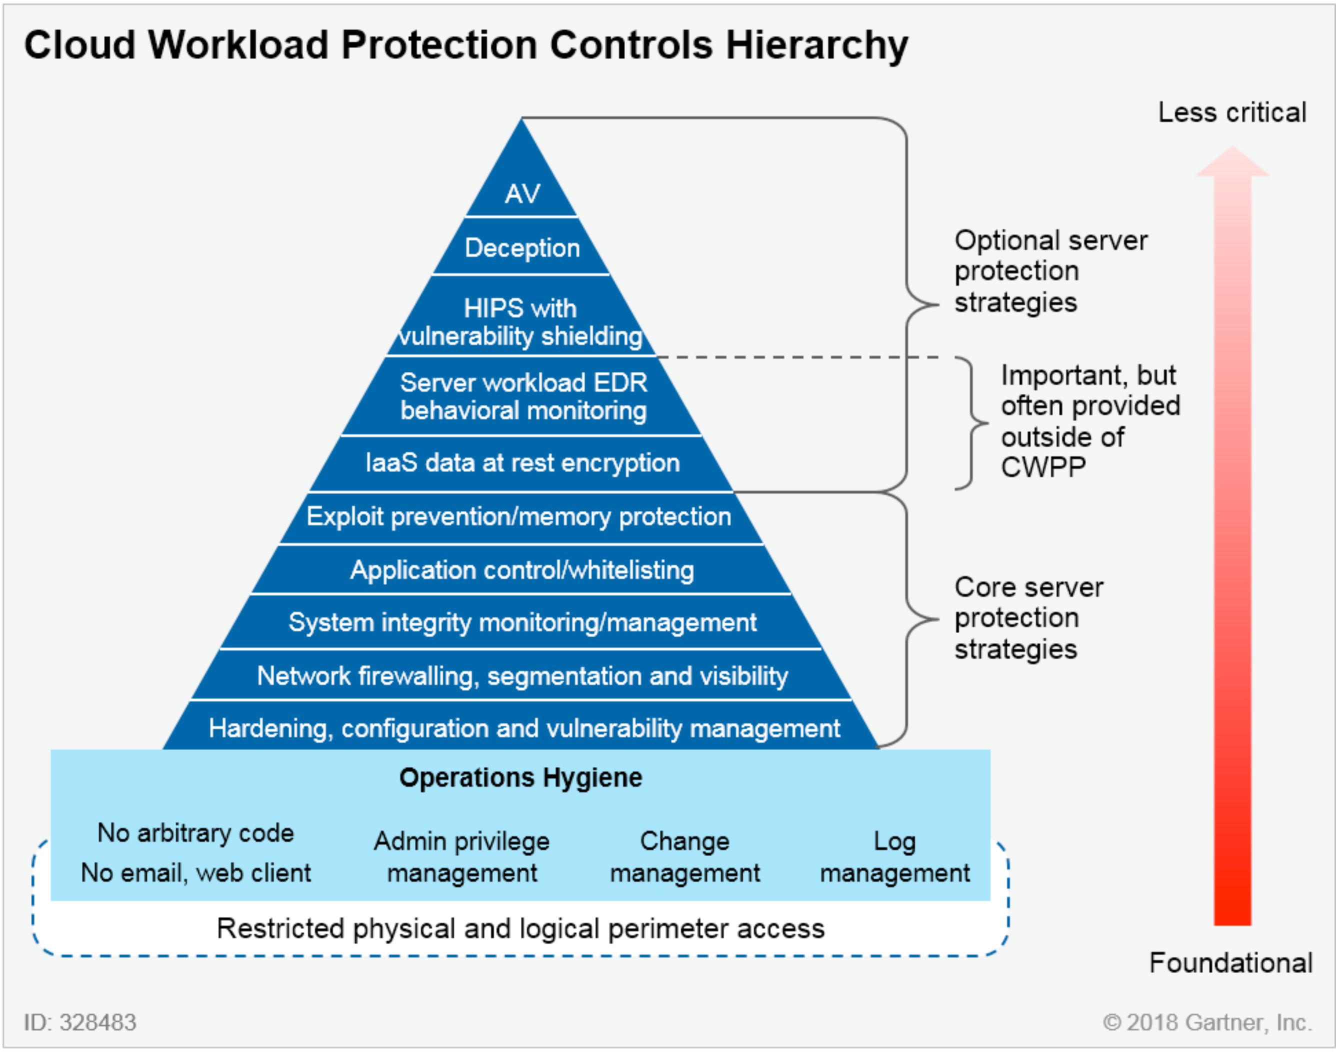
Cloud Workload Protection Controls Hierarchy
AV
Deception
HIPS with vulnerability shielding
Server workload EDR behavioral monitoring
IaaS data at rest encryption
Exploit prevention/memory protection
Application control/whitelisting
System integrity monitoring/management
Network firewalling, segmentation and visibility
Hardening, configuration and vulnerability management
Operations Hygiene
No arbitrary code
No email, web client
Admin privilege management
Change management
Log management
Restricted physical and logical perimeter access
Optional server protection strategies
Important, but often provided outside of CWPP
Core server protection strategies
Less critical
Foundational
ID: 328483
© 2018 Gartner, Inc.

Source: Gartner (March 2018)

Figure 1 illustrates our recommended prioritization of security controls for hybrid cloud server workload protection. Capabilities toward the bottom of the pyramid are more critical (foundational), whereas those toward the top are less important. However, depending on the specific risk profile of the server and the legal/regulatory requirements of the workload and the geography, enterprises may weight their evaluations differently. Some of the capabilities shown may be supplied by the OS provider, cloud IaaS provider or another tool in IT operations (e.g., configuration and patch management). Thus, they may not be heavily weighted for some enterprises. Finally, servers and VMs hosting virtual desktop infrastructure (VDI) are a different use case, and are likely to use a more-traditional, end-user endpoint protection strategy (see Note 3).

Start with solid operational hygiene. At the bottom of Figure 1 is a square box of foundational operational capabilities. Solid server security starts with good operational hygiene. For many organizations, the operational processes and technical solutions for delivering these capabilities are already in place and should be extended to cloud-based workloads, including:

- **Restricted access to the server.** Server workloads should have restricted access —physically and virtually – as to what can reach the workload.
- **Restricted ability for arbitrary code to be placed onto the server.** Local browsers and email clients should be removed or disabled. On physical servers, USB ports, Wi-Fi and similar ways to introduce or remove content other than the network should be removed or disabled.
- **Tight controls around administrative access to the server workloads.** Multifactor authentication or other forms of strong authentication beyond simple usernames and passwords should be mandatory. In addition, strict controls and processes concerning the issuing of administrative credentials should be put in place, using privileged account management (PAM) systems.
- **Well-defined change management processes.** Ideally, these changes would be controlled and managed in conjunction with a PAM system. If runtime changes are allowed at all, changes to workload images should follow a defined change management control process linked to the trouble-ticketing system.

- **Log management and monitoring.** The server workload OS and applications logs should be gathered into a log management system or a security information and event management (SIEM) system. The PAM logs should be managed as well. In virtualized and cloud environments, logs of the activities of cloud administrators should also be managed – for example, using Amazon Web Services (AWS) CloudTrail logs.

Above this foundational operational hygiene level of controls, the following controls should be considered mandatory for the protection of server workloads:

- Configuration and vulnerability management (ideally, scanning the workload before it is released into production)
- Network firewalling, segmentation and traffic visibility
- System integrity monitoring/management
- Application control (whitelisting)
- Exploit prevention and memory protection

Beyond the capabilities listed above, there are additional ways that server workloads may be protected in a deep, layered defense, in-depth strategy. The need for additional protection will be based on multiple factors. These include compliance requirements, the sensitivity of the workload protected, and the presence of such mitigating controls as a network firewall or network IPS. This also involves whether the server can be patched in a timely manner, as well as the risk tolerance of the application, product or service owner.

CWPP Details

The following is a more-detailed description of what we consider to be the key capabilities of solutions that compete in this market. Starting from the bottom layer of the pyramid in Figure 1, the core components of a server protection platform are:

- **Hardening, configuration and vulnerability management.** Unnecessary components, such as Telnet, FTP and other services, should be removed. Images should be hardened using industry-standard guidelines as the starting point. This responsibility may be managed by IT operations. However, information security is responsible for ensuring that systems are hardened and configured according to the organization's guidelines, and systems are kept patched and up-to-date in a timely manner, according to the organization's policies. Standard configuration baselines are available from organizations such as the Center for Internet Security. This group has established baselines for AWS ("Tag: CIS AWS Foundations Benchmark" (https://aws.amazon.com/blogs/security/tag/cis-aws-foundations-benchmark/)), Azure ("CIS Microsoft Azure Foundations Benchmark v1.0.0 Now Available" (https://www.cisecurity.org/cis-microsoft-azure-foundations-benchmark-v1-0-0-now-available/)) and environments, such as Docker ("CIS Docker Benchmarks" (https://www.cisecurity.org/benchmark/docker/)).

  In many cases, this functionality will be achieved using an external scanning tool or service – for example, Cavirin, Qualys, Tenable.io (Nessus) or Rapid7. However, some of the CWPP solutions in this Market Guide can also assess the system configuration, compliance and vulnerability status from the "inside out," using their agents to provide this visibility. In these cases, CWPPs should provide specific policy recommendations for the workload hardening, based on the workload's contents. Another hardening approach is referred to as moving target defense (https://www.dhs.gov/science-and-technology/csd-mtd) – randomizing the OS kernel, libraries and applications so that each system differs in its memory layout to prevent attacks.
- **Workload segmentation, traffic visibility and optional network traffic encryption.** A foundation of solid workload security is isolation and segmentation of its ability to communicate with external resources. Some of the workload protection solutions provide their own firewalling capabilities, whereas others manage the built-in firewalls of Microsoft Windows and Linux. Some will manage the built-in segmentation of AWS Security Groups and Azure Network Security Groups. The solution should support the growing requirement for "microsegmentation" (more-granular segmentation) of east/west traffic in data centers.

  In addition, several of the solutions provide visibility and monitoring of the communication flows. Visualization tools enable operations and security administrators to understand flow patterns, set policies and monitor for deviations. Finally, several vendors offer optional encryption of the network traffic (typically, point-to-point IPsec transport mode security associations) among workloads for the protection of data in motion, and provide cryptographic network isolation among workloads.

- **System integrity monitoring/management.** Capabilities here span two areas:
  - **Preboot** – The ability to measure the basic input/output system (BIOS), firmware, hypervisor, VMs and container system images before they are loaded, which is typically achieved using trust measurements rooted in hardware for physical systems. In the public cloud, this will be limited to measuring the integrity of the system images and containers before mount.
  - **Postboot** – The real-time monitoring of the integrity of critical system files after the workloads are booted. Like stand-alone antivirus, the value of file integrity monitoring (FIM) alone is minimal. However, it may be required by auditors, because FIM is a requirement of multiple regulations. Advanced solutions also monitor the integrity of the Windows registry, startup folders, drivers, bootloader and other critical system areas.
- **Application control (whitelisting).** Most workloads in on-premises VMs and in public cloud IaaS run a single application. This is almost always the case with containers hosting microservices-based applications. The use of whitelisting to control what executables are run on a server provides an extremely powerful security protection strategy. All malware that manifests itself as a file to be executed is blocked by default. Many CWPP solutions provide built-in application control capabilities, or dedicated point solutions offer them.

  Alternatively, the built-in application control capabilities of the OS might be used, such as software restriction policies, AppLocker and Defender Device Guard with Windows, or Security-Enhanced Linux (SELinux) or AppArmor with Linux, or AppDefense with VMware. Some of the application control vendors can further constrain the runtime behavior of whitelisted applications, using more-granular policy enforcement.
- **Exploit prevention and memory protection.** Application control solutions are fallible and must be combined with exploit prevention and memory protection capabilities, either from the OS – for example, ASLR and seccomp (see "seccomp" (https://en.wikipedia.org/wiki/Seccomp)) – or from the CWPP vendor. We consider this a mandatory capability to protect from the scenario in which a vulnerability in a whitelisted application is attacked. The injected code runs entirely from memory, and doesn't manifest itself as a separately executed and controllable file (referred to as "fileless malware"). In addition, exploit prevention and memory protection solutions can provide broad protection against attacks, without the overhead of traditional, signature-based antivirus solutions. They can also be used as mitigating controls when patches are not available. Moving target defense solutions also provide exploit prevention.

Additional CWPP layers include:

- **IaaS data-at-rest protection.** Encryption of data at rest should be a standard best practice for workloads running in public cloud IaaS. With the use of Intel's AES-NI for cryptographic operation acceleration, the impact on performance is minimal. In addition, many enterprises are making this a standard requirement in their on-premises data centers. We have not made this a core requirement for CWPP selection, because many OSs now provide full drive encryption for free and support a "headless" mode specifically for server protection scenarios.

  Amazon also provides free full-volume encryption in AWS, as well as free solutions for RDS and S3. Microsoft provides a similar capability with Azure Disk Encryption. With any encryption, there is a need for secure storage and management of encryption keys, as well as a need to support customer-managed keys. More-advanced solutions support features such as key management, cross-cloud encryption and automatic key rotation.
- **Server EDR for behavioral monitoring.** This layer should also be mandatory; however, this can be largely achieved via monitoring from outside the workload. Server EDR goes beyond the system integrity monitoring (a basic form of EDR). Server EDR monitoring looks at behaviors such as network communications, processes launched, files opened and log entries for behavior patterns that indicate malicious activity, including within containers. Another technique is to establish patterns of expected behaviors from whitelisted applications and to look for deviations in behavior.

  Several of the end-user EDR point solution vendors specifically target server workload protection use cases. These capabilities are focused on detection and response, rather than prevention of attacks. Some organizations will achieve this with network-based monitoring, rather than host-based agents. Thus, we haven't made this a core requirement of CWPP. Another common use case for server EDR will be to quickly scan all systems for the presence of a specific file by name or hash in the event of an outbreak. This is similar to signature-based antivirus scanning, but is used in detection/response scenarios.

- **Host IPS including vulnerability-facing HIPS.** Here, in addition to traditional network IPS protection against known attacks, the CWPP vendor deeply inspects the incoming network traffic stream for attacks against known vulnerabilities and prevents them. This layer may be redundant, with network IPSs protecting the data center; however, those may not protect from inter-VM or intercontainer-based attacks. HIPS becomes a valuable defense in depth control to shield from attacks on a zero-day vulnerability, until the patch can be applied or the VM/container is rebuilt.

  HIPS is used by some organizations to reduce the frequency of server patching. HIPS is also useful for protecting cloud workloads where network IPS may be difficult, expensive and difficult to scale dynamically, especially when dealing with encrypted traffic. HIPS may also be critical for protecting server workloads that are difficult to patch or that are no longer supported with patches by the vendor (such as Windows Server 2003, which fell out of support in 2015).
- **Deception.** Deception has an important role to play on servers, but is typically provided by dedicated deception vendors, rather than CWPP providers. This emerging security protection capability creates fake vulnerabilities, systems, shares, cookies, etc., on the server (sometimes referred to as "honey data" or "honey tokens"). If an attacker tries to access or use these fake resources, it is a strong indicator that an attack is in progress, because a legitimate workload should not see or try to access these resources. Some of the deception technologies for network, application, endpoint and data are entirely agent-based, residing on the server workload. Thus, they fall under the scope of this CWPP research.
- **Signature-based antivirus.** Signature-based antivirus and anti-malware scanning provides little to no value on well-managed server workloads. Use an application control whitelisting model as the primary control for server workload protection. In some cases, signature-based file scanning makes sense – for example, if the server workload is serving as a general-purpose file repository, such as a file share, a Network File System (NFS) server, an FTP server or a SharePoint server. In these cases, the file repository should be scanned, but this can be performed externally. (The same is true with object stores in public cloud IaaS, which should also be scanned.)

  Another exception requiring antivirus would be where regulatory requirements specify the use of antivirus and it is not negotiable with the auditor. Here, basic file system scanning to meet compliance requirements using a minimal open-source software (OSS) engine, such as ClamAV, is a possible strategy. Alternatively, use your incumbent endpoint antivirus solution (configured to minimize the impact on server performance by disabling real-time scanning and reducing the frequency of on-demand scans). This has the advantage of being managed under the same policy management system as other workloads.

The above CWPP capabilities typically run within a workload and are the focus of this research. However, as information security architects develop a comprehensive protection strategy for server workloads, security protection services external to the workload are also likely to be needed, and are likely to involve different vendors and capabilities. An overall cloud service protection strategy will occur in multiple layers, as shown in Figure 2.

Figure 2. Cloud Workload Protection Outside the Workload

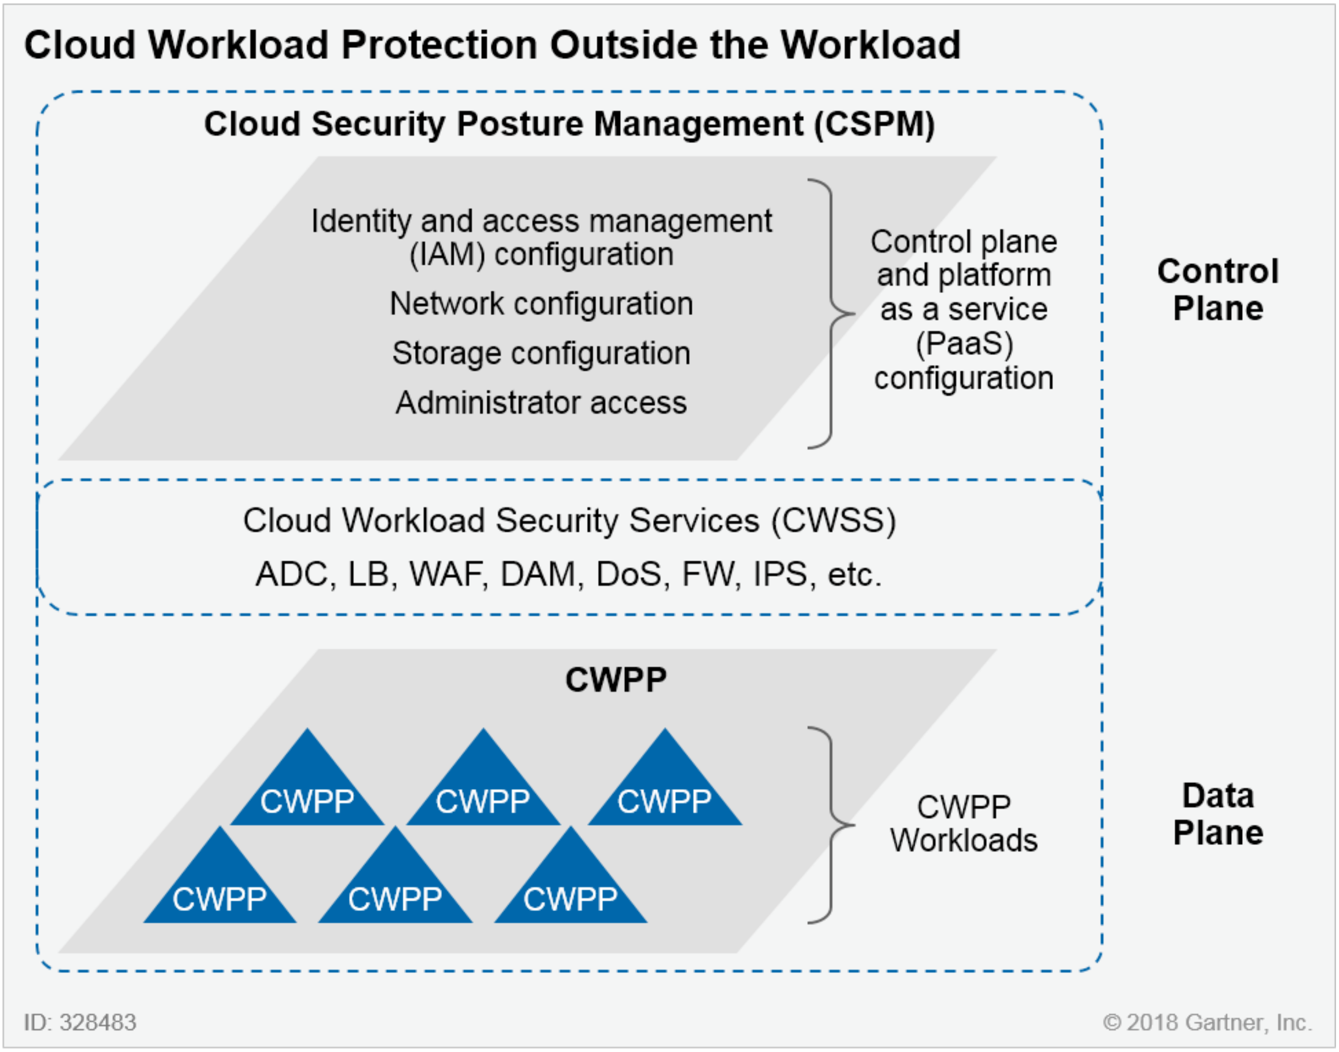
Cloud Workload Protection Outside the Workload
Cloud Security Posture Management (CSPM)
Identity and access management (IAM) configuration
Network configuration
Storage configuration
Administrator access
Control plane and platform as a service (PaaS) configuration
Control Plane
Cloud Workload Security Services (CWSS)
ADC, LB, WAF, DAM, DoS, FW, IPS, etc.
CWPP
CWPP
CWPP
CWPP
CWPP
CWPP
CWPP
CWPP Workloads
Data Plane
ID: 328483
© 2018 Gartner, Inc.

Source: Gartner (March 2018)

CWPP protected workloads are shown in triangles labeled "CWPP" in the data plane in Figure 2. Surrounding these workloads is a set of cloud workload security services (CWSS) that are external to the workload and should also be considered (outside the focus of this research). These optional capabilities provide application-specific protection, such as web application firewalls (WAFs), database activity monitoring, load balancing, denial of service (DoS) protection, network-based firewalling and network-based IPS. The use of these may obviate the need for the specific capabilities from the CWPP provider.

Above this in Figure 2, there is a set of surrounding control plane services used to provision/deprovision, configure and manage the workload. For example, identity and access management (IAM) services, network connectivity/configuration, storage configuration and PaaS services. In larger cloud environments, correct configuration of the control plane has become extremely complex, leaving the organization's information and workloads at risk. As an example of this, consider the frequent occurrence of sensitive data left sitting in AWS S3 buckets exposed directly to the public internet without protection.

To assess and manage the security posture of the cloud control plane, a market is emerging for cloud security posture management (CSPM), previously called CISPA (see Note 4). Several of the CWPP vendors in this Market Guide have begun offering CSPM capabilities to assess and manage the configuration risk of cloud services. The better offerings provide this across multiple public cloud providers for consistent policy enforcement (for example, alerting or blocking when network groups in any IaaS are directly exposed to the public internet). For large, cloud-based workload deployments, CSPM capabilities should be considered mandatory.

CWPP Architectural Considerations

When evaluating CWPP solutions, there are several key architectural considerations that vary among the solution providers.

Support for hybrid cloud environments. One of the most-critical considerations is that the solution work in hybrid cloud environments that span on-premises workloads, VMs, containers and deployments in public cloud IaaS from multiple cloud providers. For enterprises that still have physical servers, support for these systems may be a requirement.

Server OSs supported. Most vendors support Windows and Linux. If Windows is supported, clarify which versions and whether both 32- and 64-bit versions are supported. Few vendors support HP-UX, IBM AIX or Oracle Solaris. Some specialize in supporting out-of-support server OSs, such as Windows 2000 Server and Windows Server 2003. If Linux is supported, look for specific support of your enterprise distributions, as well as 32- and 64-bit support. Determine whether the product is at feature parity with Windows. Because Linux use dominates cloud IaaS in AWS and GCP, your CWPP provider should offer support for common Linux distributions used in public cloud IaaS. For example, CentOS, Debian, Ubuntu, Red Hat, SUSE and, in AWS, Amazon Linux with documented historical timely updates for support as updates to Linux are released, especially kernel updates.

Container support. CWPP vendors need to be able to provide visibility into containers and to distinguish and apply policies, based on individual containers, including network segmentation. This is an emerging critical requirement for organizations using containers to support microservices-style architectures and rapid DevSecOps workflows. Three primary methods are being used to protect container-based workloads at runtime. A more-traditional architecture, with agents running in the host OS, may be used if the agent is enlightened/aware of containers. Alternatively, a "privileged" container can be used as a peer to the other containers – a security container – then use the container management system to provision these one per physical host.

Another approach is to "inject" or layer the security controls into each container as they are constructed before release into production. Linux container support is the primary driver here using the standard Docker container format. Windows containers have shipped, but adoption is slow, and few CWPP vendors support them.

Full API enablement. In highly automated cloud environments, security protection needs to be automatically and programmatically applied to workloads. Rather than requiring expensive and slow manual configuration via "human middleware" to configure security policy via consoles, security policy is applied automatically via APIs using the scripts, recipes and templates common in these highly automated development environments. All functionality available in the console should be available via APIs, and, ideally, the vendor's console is built entirely on its own APIs.

Explicit SDL integration. As enterprises shift to more-rapid DevSecOps-style development, security scanning needs to be integrated directly into the continuous integration/continuous delivery (CI/CD) toolchain. As new workloads are created via tools such as Chef, Puppet and Ansible, or when using cloud management platforms, such as OpenStack, the security policy – such as vulnerability and configuration scanning – can be applied automatically. It can be scanned via APIs, with security controls automatically provisioned and configured.

Impact on runtime performance. Depending on the capabilities the CWPP delivers, there may be a measurable impact on the system footprint and performance. For example, deep-packet-inspection-based HIPS can be resource-intensive. Encryption should use hardware acceleration capabilities. Signature-based, anti-malware scanning creates a measurable impact on performance, when real-time scanning is kept activated, and when crawling and scanning the file system.

"Agentless" protection. In VMware environments, multiple providers have linked into its vSphere hypervisor APIs for agentless, anti-malware scanning. One vendor, Trend Micro, supports agentless file integrity monitoring. With VMware's NSX, agentless IPS is possible and several vendors support this deployment option, including several CWPP vendors. With VMware's AppDefense (https://www.vmware.com/products/appdefense.html),[1] (https://www.gartner.com/document/3869864?ref=solrAll&refval=200827736&qid=5835bf6f2911c243ff72730b4fd8f437#dv_1_introduction_to) it has provided capabilities for workload whitelisting and behavioral monitoring and has partnered with Carbon Black for server EDR on top of AppDefense telemetry. With container-based architectures, several vendors avoid traditional agents and use a privileged container model or layered security control insertion to enforce security policy. Finally, Bracket Computing uses an innovative VM "wrappering" approach that protects individual workloads without agents.

Native integration and support for leading virtualization and cloud providers. For effective protection in public and private cloud-based environments, the CWPP should understand and integrate with native tagging capabilities of the platform, so that policies can be applied based on these tags. Furthermore, integration with the APIs of the cloud provider can signal the console when new workloads have been created, potentially without security protection installed. Finally, understanding the native segmentation of the cloud provider, such as network and security groups, will help define segmentation strategies.

Management console capabilities. Most enterprises prefer web-based management consoles that don't depend on a specific OS to access the console. The vendor's policy framework should support delegation of administrative capabilities with full, role-based access control supporting administrators with different responsibilities. The console should support the ability of the administrator to apply logical tags to workloads to apply policies for similar workloads – for example, applying a "PCI" tag to all PCI workloads for policy enforcement. In addition, the console should support the native tagging of the cloud provider where the workloads are running (AWS, Azure, Google Cloud Platform [GCP] and VMware). The console should be able to be run on-premises or in the public cloud as a virtual appliance based on customer preference.

Console as a service. Increasingly, enterprises don't want to install or manage a console at all. They prefer to consume the management console as a managed service, provided by the CWPP vendor. Delivering a cloud-based, scalable, multitenant console as a service is a significant architectural shift for CWPP vendors, and not every vendor offers this option. This is useful for smaller enterprises that don't want the hassle, complexity and cost of setting up their own management servers and for the CWPP vendor's channel partners to offer managed CWPP services to their customers.

Compliance reporting. For organizations with specific regulatory requirements, the ability to provide specific compliance reports reduces the workload when auditors ask for evidence of compliance (for example, PCI, GDPR and Health Insurance Portability and Accountability Act [HIPAA] compliance reporting). In addition, reports against configuration and hardening best practices, such as the CIS guidelines, are needed.

Ability to securely bootstrap. Systems that are rapidly provisioned with security agents embedded may not be able to know in advance the policies that will need to be applied at runtime. The agents should be able to be provisioned using templates and upon boot, and to securely reach out, download and apply the appropriate policy, based on the context of the workload (e.g., the location of the workload) or based on its tagging.

Machine learning. Machine learning will apply at nearly every layer shown in Figure 1. For example, modeling network traffic patterns and identifying anomalous behavior is extremely useful for baselining and understanding initial traffic patterns and grouping them for the application of microsegmentation policies. Likewise, as enterprises move toward a default deny approach for application execution (application control), manual management of these policies and rules won't scale for cloud workloads.

The CWPP vendor should use learning and observation (ideally starting in development) to build and maintain the whitelist over time and be adaptable by relearning updates and changes without requiring rule changes. Server EDR also benefits from machine learning by grouping patterns of similar workloads from a behavioral level. Even traditional anti-malware scanning benefits from machine learning. However, it does not rely solely on signatures to determine whether code is malicious prior to execution. CWPP vendors that offer anti-malware scanning should apply machine learning prior to execution, in addition to traditional signatures.

Pricing model flexibility. The ideal solution enables the enterprise to choose the mix of licensing models that makes the most sense. Most vendors set prices using a subscription model of per VM, per year, whether on-premises or in the public cloud. Others offer pricing per CPU socket. For highly elastic workloads, a pricing model based on actual usage in VM or container hours or minutes, or other usage-based metrics, may be the better choice in IaaS.

Auditing and logging. All administrative activities and events in the console should be logged, and these logs should be exportable to leading SIEM systems.

Threat intelligence and community intelligence. The vendor's lab research capabilities should provide global threat intelligence to inform security operators of changing attack patterns and trends, and, ideally, feed directly into its protection solution. The vendor's customer community should enable participants to share visibility and intelligence information to better protect from threats.

Representative Vendors

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

The following vendors offer solutions designed to satisfy one or more of the requirements noted in the previous section:

- Alcide
- Amazon
- Bitdefender
- Bracket Computing
- Capsule8
- Carbon Black
- Cisco
- CloudAware
- CloudPassage
- Cloud Raxak
- Dome9
- Edgewise
- GuardiCore
- HyTrust
- Illumio
- Kaspersky Labs (see Note 5)
- Lacework
- McAfee
- Microsoft
- Polyverse
- Qingteng (China only)
- Security Code
- Sophos
- Symantec
- Threat Stack
- Trend Micro
- Tripwire
- Virsec
- VMware

Container-focused CWPP providers:

- Aqua Security
- Ericsson (Apcera)
- Aporeto
- Twistlock
- Layered Insight
- StackRox

Market Recommendations

The rapid adoption of private and public cloud computing models, containers and DevSecOps is fundamentally reshaping the security requirements for protecting private and public server workloads. The time has come for enterprises to start protecting cloud-based workloads, using a protection strategy that is different from end-user-facing desktops and laptops. Cloud server workload protection strategies should be based on a foundation of solid operational hygiene, including proper administrative control, patching discipline and configuration management. With the widespread adoption of VMs and containers, server workloads tend to be allocated to a specific application or service. This leads to the imperative to adopt a core workload protection strategy anchored in reducing the surface area for attack and preventing the execution of unknown code using application control, combined with exploit prevention and memory protection techniques.

For servers in which sensitive information is handled or stored, other types of protection beyond the core set of capabilities (such as EDR-type behavioral monitoring) add defense in depth. Furthermore, if the server OS is out of support and can't be patched (e.g., Windows Server 2003), then vulnerability-facing IPS (network or host-based) should be a key part of the protection strategy, in addition to the core application control strategy.

Signature-based, anti-malware scanning should be deactivated or scaled back (and, in many cases, removed) from all server workloads that don't serve as general-purpose, file-sharing repositories in favor of a whitelisting-centric approach using application control. If regulatory requirements specify antivirus scanning and are not negotiable, use less-frequent, scheduled file-based scanning without the performance overhead of continuous, memory-based scanning. Alternatively, keep the agent installed, but only activated on-demand as a response tool for security operations staff, if a scan is needed.

Source: Gartner Research Note G00328483, Neil MacDonald, 26 March 2018

Return to Home (index.html)

Note 1. Representative Vendor Selection

Representative vendors provide generally available, workload-centric protection offerings, which are typically agent-based. Designed to protect workloads in hybrid data centers that span private and public clouds, they may include support for container-based workloads, as well as legacy physical servers. Representative vendors offer one or more of the CWPP capabilities shown in Figure 1.

Note 2. PCI and Application Control

The PCI Data Security Standard (DSS) requirement is for an anti-malware security control. QSA should and will accept application control directly to meet the PCI DSS requirement.

"Cb Defense Meets PCI DSS. Attestation Report by Coalfire" Carbon Black

Note 3. VDI

VDI is a special use case in which end-user-facing endpoint sessions are hosted on a server. These should be secured using a more-traditional endpoint security approach, in which these sessions are kept strongly isolated from the rest of the data center network. Signature-based anti-malware scanning should be considered mandatory. However, because these VDI sessions are hosted on servers using virtualization platforms, agentless, anti-malware scanning solutions are often favored to reduce resource contention.

Note 4. Example CSPM Providers

In prior Gartner research, CSPM was referred to as cloud infrastructure security posture assessment (CISPA). We have updated this term slightly – changing "assessment" to "management, reflecting the ability of these providers to take action on policy violations. We have also removed the word "infrastructure," because these capabilities are useful for the PaaS (and SaaS) layers as well:

- Alert Logic
- BMC
- Cavirin
- Cloud Conformity

- CloudAware
- CloudCheckr
- Cloudnosys
- Cloudvisory
- DivvyCloud
- Dome9
- Evident.io (to be acquired by Palo Alto Networks[2])
- RedLock
- Saviynt
- Turbot

In addition, leading cloud access security broker (CASB) providers are also adding comprehensive CSPM capabilities by integrating with IaaS APIs in the same way they provide visibility and control to SaaS via APIs:

- Bitglass
- McAfee (Skyhigh Networks)
- Netskope
- Oracle
- Palo Alto Networks
- Symantec

Note 5. Kaspersky Labs and U.S. Government Dispute

In early September 2017, the U.S. government ordered all federal agencies to remove Kaspersky Lab's software from their systems. This action occurred after several media reports, citing unnamed intelligence sources, claimed that Kaspersky's software was being used by the Russian government to access sensitive information. Although the U.S. government has not given any official explanation for the ban, Kaspersky Lab vehemently refutes the unsubstantiated claims and is seeking an appeal in U.S. federal court on the ban (Department of Homeland Security Binding Operational Directive 17-01). Gartner clients, especially those who work closely with U.S. federal agencies, should continue to monitor this situation for updates. Kaspersky has commenced legal action against the U.S. government and has denied all allegations against it.

(http://www.gartner.com)

About Gartner (/technology/about.jsp) | Careers (/technology/careers/) | Newsroom (/it/products/newsroom/) | Policies (/technology/about/policies/guidelines_ov.jsp) | Site Index (/technology/site-index.jsp) | IT Glossary (/technology/it-glossary) | Contact Gartner (/technology/contact/contact_gartner.jsp)