

VM-SERIES ON GOOGLE CLOUD PLATFORM

Big data and context-rich application initiatives on Google® Cloud Platform are transforming data centers into hybrid clouds, yet the risks of data loss and business disruption remain. Embedding the VM-Series into your GCP™ application development lifecycle prevents data loss and business disruption, allowing your adoption to move at the speed of the cloud.

VM-Series Virtualized Next-Generation Firewall on GCP

- Complements the native Google Cloud firewall with application enablement policies that prevent threats and data loss.
- Enables security to be transparently embedded into the application development process through automation and centralized management.
- Delivers managed scale and cloud-centric high availability through integration with Google Cloud load balancing.

Introduction

As Google Cloud Platform becomes a more significant deployment platform for your business-critical applications, protecting the increased public cloud footprint from threats, data loss and associated business disruption remains challenging. The VM-Series on GCP solves these challenges, enabling you to:

- Protect your GCP workloads through unmatched application visibility, control and advanced threat prevention.
- Prevent threats from moving laterally between workloads and stop data exfiltration.
- Embed security into your highly available and scalable application development framework.

The VM-Series protects your GCP workloads with next-generation security features that deliver superior visibility, control and threat prevention at the application level when compared to other cloud-oriented security solutions. Automation features and centralized management allow you to embed security into your application development process, ensuring security can keep pace with the speed of the cloud.

Google Cloud Firewall or VM-Series Next-Generation Firewall?

Organizations are migrating their enterprise applications onto GCP for many reasons, including business agility and a desire to reduce data center footprints. In nearly all cases, the GCP deployment is connected to the corporate network, making the GCP resources network accessible by users – and possibly attackers. Security best practices dictate that your public cloud security posture should mimic your data center security approach: understand your threat exposure through application visibility, use policies to reduce the attack surface area, then prevent threats and data exfiltration within the allowed traffic.

The question then becomes: if you were building a new data center today, would you follow a different approach that relied solely on Layer 4 security when it has historically been woefully inadequate? It's well-known that attackers will use SSL, tunnel across TCP/80 or use non-standard ports to bypass port-based controls in order to compromise your deployment and exfiltrate data. Yet customers regularly forego application-centric, prevention-based approaches, believing native, Layer 4 security services like GCP firewall are sufficient.

GCP firewall rules performs port-based filtering to control access to the GCP resources deployed. They are a required feature in that they must be enabled for the cloud deployment to be operational. They also:

- Follow a positive control security model, using port-based policies to allow traffic and deny all else. GCP firewall rules cannot be used to explicitly deny traffic on GCP.
- Allow all outbound traffic by default. More granular policies can be defined to further reduce outbound traffic flows, but only by whitelisting IPs.
- Enable you to add or remove rules at any time, meaning there is no traditional policy commit process.

The VM-Series complements GCP firewall port-based controls by reducing your attack surface through application enablement, preventing threats and stopping data exfiltration.

VM-Series on GCP

Palo Alto Networks® VM-Series on GCP allows you to embrace a prevention-based approach to protecting your applications and data on GCP. Automation and centralized management features enable you to embed next-generation security into your GCP application development lifecycle, helping you protect your data and limit business disruption.

Complete Visibility Improves Security Decisions

Understanding which applications, including those that may be SSL-encrypted, are traversing your GCP deployment – where they are coming from and going to, and the user's identity – are just a few of the data points the VM-Series provides to help you make better-informed security policy decisions.

Segmentation and Application Whitelisting Aid Data Security and Compliance

Using application whitelisting to enforce a positive security model reduces the attack surface by allowing specific applications and denying all else. You can align application usage to business needs, control application functions (e.g., allow SharePoint® documents for all but limit SharePoint administration access to the IT group), and stop threats from accessing and moving laterally across your GCP deployment. Whitelisting policies also allow you to segment applications communicating with each other across different subnets and between VPCs for regulatory compliance.

User-Based Policies Improve Security Posture

Integration with a range of on-premise user repositories, such as Microsoft® Active Directory®, LDAP and Microsoft Exchange, lets you grant access to critical applications and data based on user credentials and respective need. For example, your dev group can have full access to the dev VPC, while only IT admins have RDP/SSH access to the production VPC. When deployed in conjunction with GlobalProtect™ network security for endpoints, the VM-Series on GCP can extend your corporate security policies to mobile devices and users, regardless of their location.

Applications and Data Are Protected Against Known and Unknown Threats

Attacks, like many applications, can use any port, rendering traditional prevention mechanisms ineffective. Enabling Threat Prevention and WildFire® cloud-based threat analysis service as segmentation policy elements can prevent exploits, malware, and previously unknown threats from both inbound and lateral movement perspectives.

Multiple Defenses Block Data Exfiltration and Unauthorized File Transfers

Data exfiltration can be prevented using a combination of application enablement and threat prevention features. File transfers can be controlled by looking inside the file (as opposed to only at the file extension) to determine if the transfer action should be allowed. Executable files, found in drive-by downloads or secondary payloads, along with malware command and control as well as associated data theft, can be blocked. Data filtering features can detect and control the flow of confidential data patterns (credit card and Social Security numbers, as well as custom patterns).

Centralized Management for Policy Consistency

Panorama™ network security management enables you to manage your VM-Series firewalls across multiple cloud deployments alongside your physical appliances, ensuring consistent and cohesive policy. Rich, centralized logging and reporting capabilities provide visibility into virtualized applications, users and content.

Panorama comprises Panorama Manager and the Log Collector, allowing you to centrally manage your VM-Series firewalls in a distributed manner. You can also use Panorama in conjunction with Palo Alto Networks Logging Service. For more information, please review the [Panorama datasheet](#).

Automation to Support App Dev Workflows

The VM-Series on GCP includes management and automation features that enable you to embed security into your application development workflow process:

- Bootstrapping can automatically provision a VM-Series with a working configuration, complete with licenses and subscriptions, and then auto-register with Panorama.
- A fully documented XML API, Dynamic Address Groups and External Dynamic Lists allow you to automate VM-Series configuration changes and consume external data to dynamically drive security policy updates.
- HTTP Log Forwarding allows you to drive actions based on observed incidents.

In conjunction with Google Cloud Templates or third-party tools, you can deploy next-generation security at the speed of the cloud.

Automating Deployments With Terraform and Ansible

Organizations using multiple public and private cloud platforms, or that want to embed VM-Series deployments into their application development processes, can deploy and configure the VM-Series using third-party toolsets, such as Terraform® and Ansible®. The combination of these tools and VM-Series automation features enables organizations to deploy and configure heterogeneous environments at scale with great agility, allowing them to embed next-generation security into their application development frameworks. Check out the templates library [here](#).

Health Monitoring With Google Stackdriver

VM-Series firewalls on GCP can send internal metrics to Google Cloud Stackdriver as a means of monitoring the capacity, health status and availability of your VM-Series, along with other resources in your GCP environment. The internal metrics that can be sent to GCP Stackdriver include:

- Session utilization %
- Total active sessions
- Dataplane CPU utilization %
- Dataplane packet buffer utilization %
- SSL proxy utilization %
- GlobalProtect active tunnels
- GlobalProtect tunnel utilization %

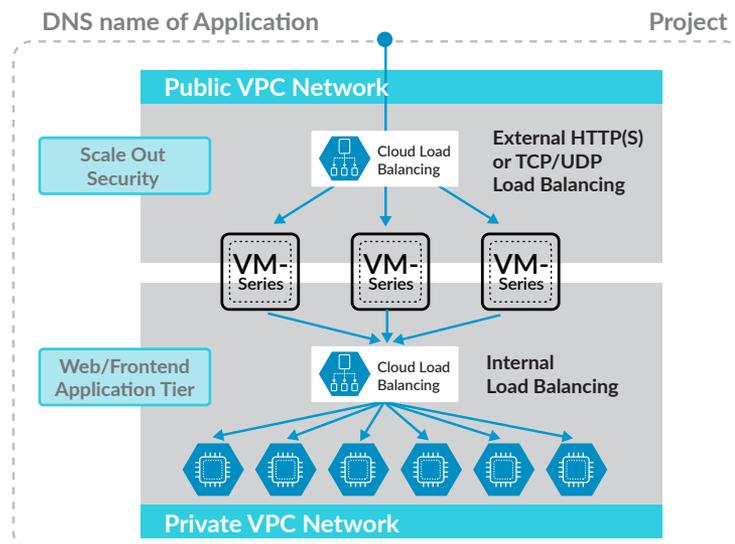


Figure 1: VM-Series scale-out architecture on GCP

VM-Series on GCP Scalability and Availability

Cloud-native designs provide higher capacity and availability by using the native services of the cloud provider, such as load balancing, DNS and built-in networking services. These focus on service reliability rather than session availability, requiring that each component of the architecture, especially applications, maintain its own state information. The VM-Series leverages several of the native services of Google Cloud to enable scale-out of security architectures.

Scalability Options With GCP Load Balancing

The VM-Series can be deployed in conjunction with GCP Load Balancing to enable a more scalable and highly available solution for both inbound and outbound traffic:

- **Scale-out with GCP load balancing:** To secure inbound VPC traffic at scale, you can deploy the VM-Series using a “load balancer sandwich” which comprises external GCP HTTP(S) or TCP/UDP load balancing and internal GCP load balancing that distribute traffic across your GCP workloads. Both the VM-Series and the application tiers, deployed behind the GCP load balancing services, can be distributed across multiple GCP availability zones. To achieve higher capacity, application performance and geographic availability, you can also use multiple GCP regions. This scale-out architecture also provides a cloud-native design that guards against a single point of failure.
- **Securing outbound and east-west flows:** For high availability and resilience of outbound traffic – both internet-facing and back to the corporate data center – and east-west traffic between VPCs, you can use GCP routing and forwarding rules to route traffic to a pair of VM-Series firewalls deployed as primary and secondary “next hops” to the same destination. The primary VM-Series firewall will be assigned a higher-priority route, using a lower metric in GCP, ensuring all traffic flows through it. If the primary VM-Series firewall fails, GCP will detect this and redirect traffic to the secondary VM-Series firewall, typically within 30 seconds.

Deploying business-critical applications in GCP dictates the need for a security solution that scales in a managed manner and is resilient. Utilizing cloud services supported and maintained by Google Cloud, combined with the VM-Series, allows you to build secure, cloud-centric architectures.

VM-Series on GCP Use Cases

The VM-Series can be deployed on GCP to address several different use cases.

Hybrid Cloud: Securely Enable App Dev and Test

Securely migrate application development and testing onto GCP through a hybrid deployment that integrates your existing development environment with GCP via a secure connection. This allows your application development and testing teams to get started while maintaining a strong security posture. When deployed on GCP, the VM-Series can act as an IPsec VPN termination point to enable secure communications to and from GCP. You can also layer application control and threat prevention policies atop the IPsec VPN tunnel or GCP Direct Connect as added security elements.

Segmentation Gateway: Separation for Security and Compliance

High-profile breaches have shown that cybercriminals are adept at hiding in plain sight, bypassing perimeter controls and moving at will across networks, physical or virtualized. A GCP virtual private cloud, or VPC, provides an isolation and security boundary for your workloads. The VM-Series can augment that separation through application-level segmentation policies to control traffic between the VPCs and across subnets. With application-level policies, you have greater control over application traffic moving laterally, and you can apply threat prevention policies to block their movement as well. If traffic is flowing between VPCs in different regions across the internet, you can enable encryption for added protection.

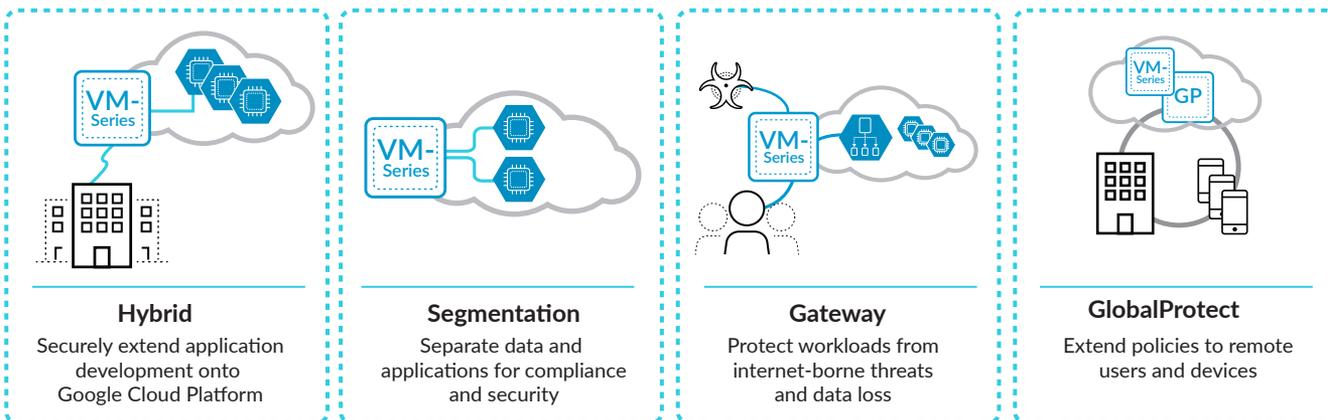


Figure 2: VM-Series on GCP deployment scenarios

Internet Gateway: Protect Production Workloads

As your GCP deployment expands to include public-facing workloads, you can use the VM-Series on GCP as an internet gateway to protect web-facing applications from known and unknown threats. Additionally, you can enable direct access to web-based developer resources, tools and software updates, thereby minimizing the traffic that flows back to corporate and then out to the web.

GlobalProtect: Extend Security to Users and Devices

GlobalProtect will enable you to extend perimeter security to your remote users and mobile devices wherever they are. GlobalProtect establishes a secure connection to protect the user from internet threats and enforces application-based access control policies. Whether the need is for access to the internet, your data center or SaaS applications, users will enjoy the full protection provided by the VM-Series.

Flexible Licensing Options

The VM-Series on GCP supports several licensing options, including consumption-based licensing via the GCP Launcher Marketplace, bring your own license and the VM-Series Enterprise Licensing Agreement, or ELA:

- **Pay-As-You-Go (PAYG) aka consumption-based licensing:** Use your Google Cloud Launcher to purchase and deploy VM-Series bundles directly from the Google Cloud Launcher at an hourly rate, with per-minute metering and billing:
 - **Bundle 1 contents:** VM-300 firewall license, Threat Prevention (inclusive of IPS, AV, malware prevention) subscription and Premium Support (written and spoken English only).
 - **Bundle 2 contents:** VM-300 firewall license, Threat Prevention (inclusive of IPS, AV, malware prevention), WildFire, URL Filtering, GlobalProtect subscriptions, and Premium Support (written and spoken English only).
- **Bring your own license (BYOL):** You can purchase any one of the VM-Series models, along with the associated subscriptions and support, via normal Palo Alto Networks channels. From the Google Cloud Launcher, you would then deploy the VM-Series and apply the authorization code to license the VM-Series.
- **VM-Series ELA:** For large-scale deployments on GCP or across multiple virtualization or cloud environments, the VM-Series ELA allows you to forecast, and purchase upfront, the number of VM-Series firewalls to be deployed over a one- or three-year period. The VM-Series ELA gives you a single license authorization code used for the life of the term, providing predictable security spend and simplifying the licensing process by establishing a single start and end date for all VM-Series licenses and subscriptions. Each VM-Series ELA includes a VM-Series firewall, subscriptions for Threat Prevention, URL Filtering, WildFire, and GlobalProtect Gateway, plus unlimited Panorama virtual machine licenses and support.

Performance and Capacities

Many factors, such as the shared tenancy of a public cloud environment, GCP instance size and number of cores, can impact performance. The performance and capacities listed below have been generated using the indicated GCP instance size, support for DPDK, and the following test conditions:

- VM-Series is deployed on GCP as a firewall between clients and servers in the same availability zone and same region. Throughput is measured based on bidirectional traffic (sent and received) by the VM-Series.
- Firewall throughput and IPsec VPN are measured with App-ID™ and User-ID™ technology features enabled, utilizing 64K HTTP transactions.
- Threat Prevention throughput is measured with App-ID, User-ID, IPS, antivirus and anti-spyware features enabled, utilizing 64K HTTP transactions.
- IPsec VPN performance is tested between two VM-Series instances in the same VPC, same availability zone and region. Performance will depend on GCP instance type and network topology, that is, whether connecting on-premises hardware to VM-Series on GCP; from VM-Series in an GCP VPC to an GCP VPN Gateway in another VPC; or VM-Series to VM-Series between regions.
- Connections per second is measured with OK HTTP transactions.

We recommend additional testing within your environment to ensure your performance and capacity requirements are met. For a complete listing of all VM-Series features and capacities, please visit our [firewall comparison tool page](#).

Model	VM-50/ VM-50 Lite ¹	VM-100/VM-200	VM-300/ VM-1000-HV	VM-500	VM-700
GCP right-sized instance tested (vCPUs, RAM, Interfaces)	N/A	n1-standard-4 ³ (4, 15, 4)	n1-standard-4 (4, 15, 4)	n1-standard-8 (8, 30, 8)	n1-standard-16 (16, 60, 8)
Firewall throughput (App-ID enabled)	N/A	500 Mbps	750 Mbps	2 Gbps	7 Gbps
Threat Prevention throughput	N/A	250 Mbps	500 Mbps	1.5 Gbps	5 Gbps
IPsec VPN throughput	N/A	In Process	In Process	In Process	In Process
GCP max-sized instance tested (vCPUs, RAM, Interfaces) ⁴	N/A	n1-standard-64 (64, 240, 8)	n1-standard-64 (64, 240, 8)	n1-standard-64 (64, 240, 8)	n1-standard-64 (64, 240, 8)
Firewall throughput (App-ID enabled)	N/A	2 Gbps	3.5 Gbps	7 Gbps	7-10 Gbps ²
Threat Prevention throughput	N/A	1 Gbps	1.75 Gbps	3.25 Gbps	8 Gbps
IPsec VPN throughput	N/A	In Process	In Process	In Process	In Process
New sessions per second (All Instance sizes)	N/A	9K	9K	20K	40K
Max sessions (All Instance sizes)	N/A	250K	800K	2M	10M
System Requirements					
Cores supported (min)	N/A	4	4	8	16
Memory (min)	N/A	6.5 GB	9 GB	16 GB	56 GB
Disk drive capacity (min)	N/A	60 GB	60 G	60 GB	60 GB
Licensing options	N/A	BYOL, VM-Series ELA	BYOL, VM-Series ELA or GCP Launcher	BYOL, VM-Series ELA	BYOL, VM-Series ELA

¹ The VM-50 and VM-50 Lite are not supported on GCP.

² Performance can vary depending on GCP allocation of resources to the VM's host system.

³ VM-100/200 can run on the n1-standard-2, but this size only offers 2 NICs, so n1-standard-4 is recommended.

⁴ When using larger sized instances only the minimum required vCPUs are used by VM-Series, but GCP enables higher throughput and allows use of additional network interfaces.



3000 Tannery Way
Santa Clara, CA 95054
Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087
www.paloaltonetworks.com

© 2018 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.
vm-series-google-cloud-platform-ds-070918