VM-Series Sizing Guide for Google Cloud

Shiva V
Sr. TME, Network Security

Why sizing is important

- Sizing provides an fairly accurate picture of how many firewalls would be needed to handle the customer’s traffic.
- This directly impacts the number of credits that must be purchased to license the firewalls.

Questions to ask the customer

- What are the minimum and/or maximum bandwidth requirements from the firewalls?
- Which traffic use-cases does the customer want to secure? Inbound, Outbound, East-West, Remote connectivity?
  - Most cases, there are multiple use-cases to secure.
- Does the customer expect the traffic to scale up and down a lot?
- Do they need session resiliency?

VM-Series on Google Cloud

Performance and Capacity

MODEL	SOFTWARE NGFW CREDITS (2 VCPUS) VM-100	SOFTWARE NGFW CREDITS (4 VCPUS) VM-300	SOFTWARE NGFW CREDITS (8 VCPUS) VM-500	SOFTWARE NGFW CREDITS (16 VCPUS) VM-700	SOFTWARE NGFW CREDITS (32 VCPUS)
GCP instance size tested (recommended)	n2-standard-4	n2-standard-4	n2-standard-8	n2-standard-16	n2-standard-32
Firewall throughput (App-ID enabled)	1.7 Gbps	3.1 Gbps	6.7 Gbps	13.9 Gbps	18.5 Gbps
Threat Prevention throughput	0.9 Gbps	1.6 Gbps	4.4 Gbps	8.2 Gbps	15.1 Gbps
Licensing options	VM-Series Software NGFW Credits, ELA, or BYOL	VM-Series Software NGFW Credits, ELA, PAYGO, or BYOL	VM-Series Software NGFW Credits, ELA, or BYOL	VM-Series Software NGFW Credits, ELA, or BYOL	VM-Series Software NGFW Credits

***The data above is for PAN-OS v11.0. For performance results on specific versions, please visit the TechDocs portal here.*

Choosing the right Machine Type and Family

Consider when choosing the machine type for VM-Series,

- VM-Series needs a minimum of 3 interfaces – one each for **Public**, **Private** and **Mgmt**.
- Number of vCPUs on Compute Instances scale at $\mathbf{2^n}$.
- Minimum number of vCPUs required for VM-Series is **4**.
- The recommended Machine Family for VM-Series is N2-Standard family.
- For specific traffic parameters like Active Connections per second, etc., consider how much memory you allocate to your VM-Series instance. Refer this page for more details.

**Bonus – You can choose a higher machine type (higher vCPU and memory) and yet license a lower number of vCPUs, if required. Refer this page for more details.

Best practices for Sizing for Active-Active design

Consider before choosing the machine type

- Use an instance template with a Managed Instance Group (MIG).
- For resiliency, use the same instance template with MIGs across at least 2 zones.
- Choose the machine type based on the maximum bandwidth required divided by the number of firewalls. Consider dividing that for firewalls across zones for resiliency.
    - For example, if the traffic bandwidth is 10Gbps, you can choose the **n2-standard-8** machine type to deploy 3 (TP enabled) firewalls with MIGs across 2 zones.
- If you are considering autoscaling, remember that autoscaling is a MIG feature, so you can choose to use a different instance template with a lower vCPU machine type for the MIG that you want to configure for autoscaling.

Best practices for Sizing for Active-Passive design

Consider before choosing the machine type

- Two VM Instances (one each for Active and Standby) will be deployed.
- Even though traffic will be passing through only one instance, both firewalls will need to be licensed.
- Make sure to determine peak bandwidth requirements before deploying.
- Choose machine type that will support bandwidth requirements from a single VM-Series instance.
    - For example, if the traffic bandwidth is 10Gbps, you would need to choose **n2-standard-32** machine type to deploy 2 (TP enabled) firewalls across 2 zones.

Best practices for Sizing Overall

- Consider splitting the firewalls to physically segment the security posture for Inbound, Outbound, East-West and Remote Connections as required.
  - While this does mean more firewalls to manage, you will have much granular control over the security applied for each use-case.
  - This design also provides the flexibility scale out per use-case, rather than the entire set.
  - If managed by separate Device Groups for each use-case, the firewalls and the security applied are easier to manage through Panorama as well.

Appendix

- Managed Instance Groups on Google Cloud
- Instance Templates on Google Cloud
- Improving resiliency with Cloud Load Balancing
- Autoscaling on Google Cloud
- VM-Series performance and capacity on Google Cloud
- Maximum Limits based on Tier and Memory for VM-Series

Thank you

paloaltonetworks.com