
Implementation Guide - CloudGen Firewall in AWS

<https://campus.barracuda.com/doc/17193/>

Click [here](#) to download the AWS Implementation Guide in PDF form.

Amazon Web Services follows the shared security responsibility model. AWS is responsible for security of the cloud. This includes physical security, servers, networking hardware, and the hypervisor. The customer, on the other hand, is responsible for everything running in the cloud, such as securing and managing the operating system, network configuration, data, and connections to the cloud.

The Barracuda CloudGen Firewall is a next generation firewall built to integrate seamlessly with the AWS cloud platform. The flexibility of the CloudGen Firewall allows cloud architects to easily select a reference architecture by the intended use case and size of the workload. A CloudFormation template is supplied with each reference architecture, making it easy to deploy or integrate with your current cloud resources. The CloudGen Firewall adds network layer security controls, visibility, and connectivity to your cloud network. Depending on the use case, CloudGen Firewall deployment is selected to satisfy the correct balance among the following criteria:

- Support for required firewall features such as firewalling or VPN
- High availability
- Scalability
- Cost optimization
- Failover or recovery times


Barracuda CloudGen Firewall Common Use Cases

- **Edge Firewall**
- **Secure Remote Access**
- **Office to Cloud / Hybrid Cloud**
- **Segmentation**

Edge Firewall

Common use cases:

- Network security enforcement with firewall and IPS.
- Default (outbound) gateway for cloud resources in the same VPC.




The CloudGen Firewall secures access to the AWS cloud resources from the Internet by enforcing granular firewall access policies and scanning incoming traffic for malware and exploits. The next generation firewall features replace or extend the native AWS security groups and NACLs by:

- Protecting against network-based attacks and exploits with the built-in IPS
- Virus scanning and Advanced Threat Protection (ATP) (BYOL only)
- Geolocation-based access control
- Traffic Shaping (QoS) to protect business-critical traffic.

CloudGen Firewall High Availability Cluster with Route Shifting

- **High Availability** - Yes
- **Failover / Recovery time** - Seconds to minutes, depending on the AWS API
- **Auto Scaling** - No
- **Default Gateway for instances in the VPC** - Yes

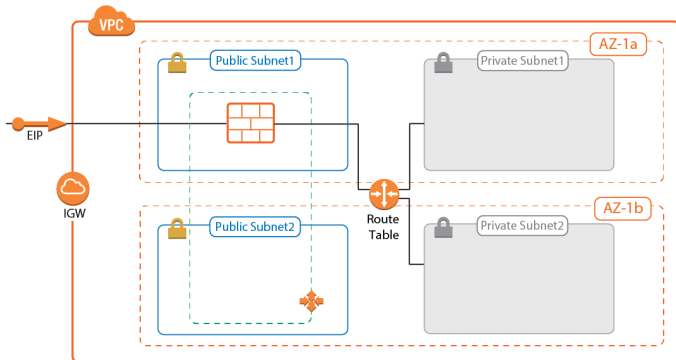


The CloudGen Firewall High Availability Cluster supports all firewalling and the default gateway features required to act as an edge firewall. The firewalls are in an active-passive cluster that syncs session information and configurations. All outgoing traffic from the private subnets is routed over the active firewall. In the event of a failover, the passive firewall takes over and connects to the cloud fabric to rewrite all routes to use the now-active firewall in the High Availability Cluster as the target. Routes added after deployment that use the firewall as the gateway are automatically detected and, in the case of a failover, are also rewritten.

For more information, see [AWS Reference Architecture - CloudGen Firewall HA Cluster with Route Shifting](#).

CloudGen Firewall Cold Standby Cluster

- **High Availability** - No
- **Failover / Recovery time** - Multiple minutes
- **Auto Scaling** - No
- **Default Gateway for instances in the VPC** - Yes, with manual changes required for new routes.

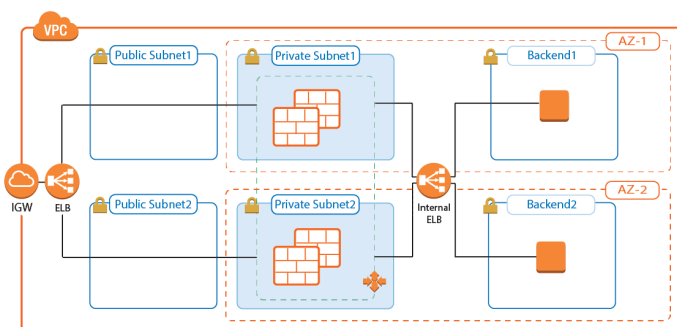


The Cold Standby Cluster is a cost-effective solution that offers the full range of next generation firewall features. In case the firewall instance becomes unresponsive, the firewall instance is automatically replaced. Routes for the private subnets are rewritten, but must be adjusted manually in the CloudFormation template to match your architecture. Route tables are not monitored automatically; additional routes or changes to existing routes must be completed by first updating the template and then updating the CloudFormation stack.

For more information, see [AWS Reference Architecture - CloudGen Firewall Cold Standby Cluster](#).

CloudGen Firewall Auto Scaling Cluster

- **High Availability** - Yes
- **Failover / Recovery time** - Instant
- **Auto Scaling**- Yes
- **Default Gateway for instances in the VPC** - No, source NAT is required for inbound traffic.



Sizing the firewall for highly dynamic traffic can be difficult. You can easily incur unnecessary costs for

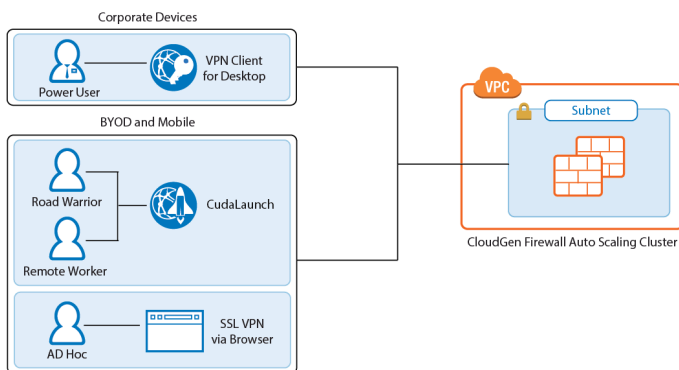
instances that are too large, or else you can run the risk of creating bottlenecks in your architectures if the firewall cannot keep up with current demand. The CloudGen Firewall Auto Scaling Cluster scales automatically to match your workload. One or more Elastic Load Balancers distribute traffic over the firewall instances in the Auto Scaling group. Custom Firewall metrics collected by CloudWatch allow custom-tailored scaling policies that match your cloud applications. Since the source IP address must be rewritten on the firewall, the CloudGen Firewall Auto Scaling Cluster cannot be used as a default gateway for outbound traffic for instances in the private networks.

For more information, see [AWS Reference Architecture - CloudGen Firewall Auto Scaling Cluster](#).

Secure Remote Access

Common use cases:

- Remote access for unknown or highly dynamic workloads.
- Remote access for predictable workloads.



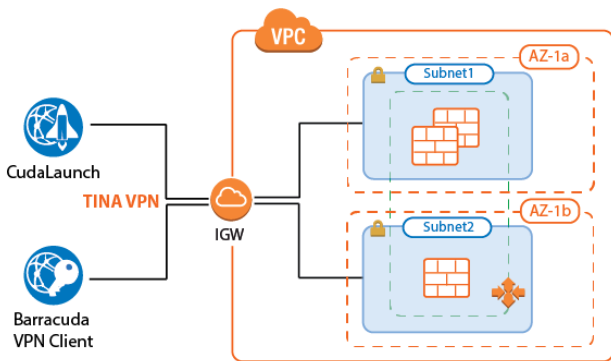
Remote access features offer remote users secure access to their organization's cloud applications and resources from virtually any device. Depending on the workload, full client-to-site VPN or SSL VPN are available, with CudaLaunch offering a richer level of remote access spanning both client-to-site and SSL VPN.

For power users, or users with centrally managed corporate devices, the client-to-site VPN offers transparent access to the corporate network. The Barracuda VPN client uses the TINA VPN protocol, specifically designed for robust VPN connections. VPN clients can be authenticated through client certificates, external and internal authentication schemes, or a combination thereof.

The SSL VPN service provides seamless integration without having to install a client app. CudaLaunch works with the SSL VPN service to provide more advanced SSL VPN features such as SSL tunneling or native app support. The number of simultaneous users using the SSL VPN is limited only by the performance of the AWS instances.

CloudGen Firewall Auto Scaling Cluster


Remote access workloads tend to be cyclical in nature. Remote workers sign in with their VPN clients in the morning, and disconnect at the end of their work day. By using a CloudGen Firewall Auto Scaling Cluster, the number of firewalls is scaled automatically to meet the current demand. The cluster can also be scaled according to a schedule, depending on how predictable the workload is. The firewall instances are automatically deployed into two or more Availability Zones. Custom firewall and VPN metrics collected by AWS CloudWatch allow the admin to configure customized scaling policies. Auto Scaling is limited to the PAYG images of the Barracuda CloudGen Firewall.



For more information, see [AWS Reference Architecture - CloudGen Firewall Auto Scaling Cluster](#).

CloudGen Firewall Cold Standby Cluster

For a small number of remote users with predictable traffic patterns, the Cold Standby Cluster is a very cost-effective remote access solution. The single firewall running is automatically replaced within minutes after a failure. The configuration is stored on an S3 bucket and can optionally be fetched from a CloudGen Control Center. Using a Control Center allows for the use of BYOL pool licenses for the instance. For single firewalls, the PAYG image is used. Cold Standby Clusters must be scaled up manually to meet increased demand.




For more information, see [AWS Reference Architecture - CloudGen Firewall Cold Standby Cluster](#).

Office to Cloud / Hybrid Cloud

Common use cases:

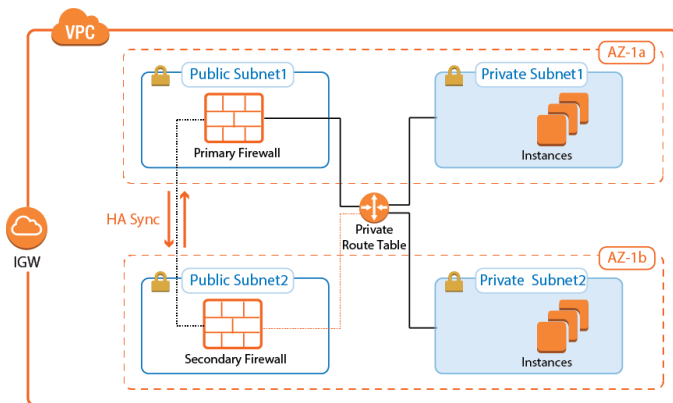
- Hybrid cloud using site-to-site VPN.
- Default (outbound) gateway for cloud resources.
- Secure traffic on the Direct Connect MPLS line.



Create site-to-site VPN connections to transparently connect your on-premises networks with your applications and services hosted in the cloud. For VPN tunnels using the proprietary TINA VPN protocol, SD-WAN allows you to split a VPN tunnel into up to 24 VPN transports, each using a different WAN connection to the firewall in the cloud. For the user, this happens completely transparently. In addition, SD-WAN also allows you to route traffic dynamically based on bandwidth or latency requirements. Offloading traffic to cheaper connections allows you to use smaller bandwidth Direct Connect connections, or to increase the quality for business-critical or latency-sensitive information.

CloudGen Firewall High Availability Cluster

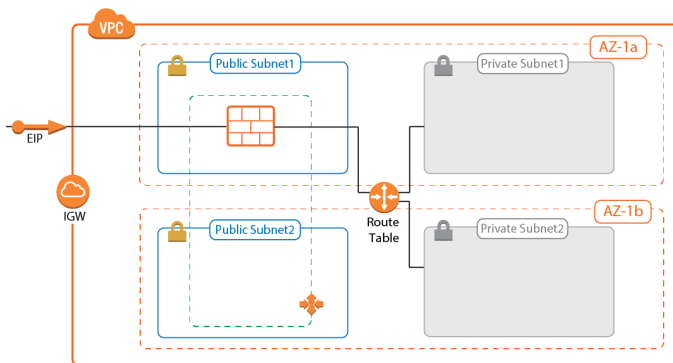
The CloudGen Firewall High Availability Cluster supports both TINA and IPsec IKEv1 and IKEv2 site-to-site VPN tunnels. For IPsec tunnels, Route 53 must be used for incoming traffic since the Elastic Load Balancer does not support UDP. Optionally, a CloudGen Control Center can be used to retrieve and manage the firewall configuration and to monitor the remote firewalls in one central location. If only TINA VPN tunnels are used, no incoming load balancing is required since TINA VPN tunnels can be configured to use two public IP addresses as the VPN endpoint. CloudGen Firewall High Availability Clusters must be scaled up manually if the workload increases.



For more information, see [AWS Reference Architecture - CloudGen Firewall HA Cluster with Route Shifting](#).

CloudGen Firewall Cold Standby Cluster

The CloudGen Firewall Cold Standby Cluster supports the same VPN features as the High Availability Cluster. The single firewall instance runs in an Auto Scaling group of one with the firewall configuration stored on an S3 bucket. In case the firewall becomes unavailable, it is automatically replaced. By default, only PAYG licenses are supported. However, it is possible to use a Firewall Control Center to manage the firewall. This allows for the use of BYOL pool licenses. The Cold Standby Cluster must be sized to meet peak demand because it does not scale dynamically.



For more information, see [AWS Reference Architecture - CloudGen Firewall Cold Standby Cluster](#).

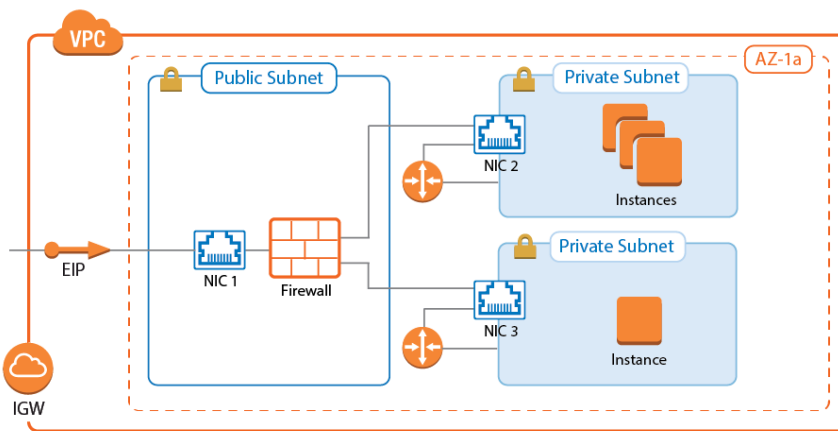
Segmentation

Common use cases:

- Provide network segmentation (INS) in the cloud.

Segmentation Firewall for Single AZ VPCs

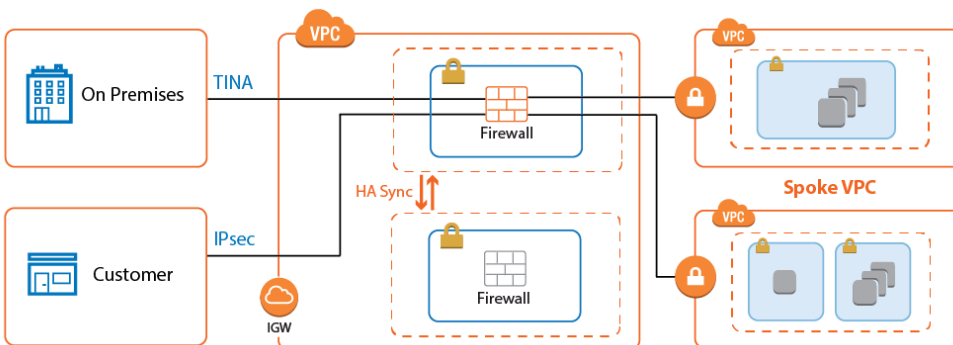
Traditional network security approaches rely heavily on network segmentation to secure the network with internal firewalls that allow only defined traffic between the different services and networks. When these on-premises applications are migrated to the cloud, the internal firewall is replaced by a CloudGen Firewall with multiple network interfaces. This allows the application to be moved to the cloud without a costly and time-consuming revamp of the architecture. Firewall access rules and the next generation firewall capabilities provide fine-grained security policies and real-time traffic visibility. Since the Elastic Network Interfaces attached to the firewall instance must be in the same Availability Zone, this solution is limited to single AZ applications.



For more information, see [AWS Reference Architecture - Segmentation Firewall for Single AZ VPCs](#).

Transit VPC

For cloud-native applications to take full advantage of the AWS cloud platform, each application is hosted in a dedicated VPC. This allows the application to be the logical context for segmentation. To organize and secure these highly dynamic VPCs, connect them in a hub and spoke architecture, with a firewall cluster in the central Transit VPC. The Transit VPC architecture is very flexible: it can be combined with High Availability Clusters, Cold Standby Clusters, or Auto Scaling Clusters, depending on the workload and predominant use case.



For more information, see [AWS Reference Architecture - Transit VPC using CloudGen Firewall](#).

Figures

1. aws_ips.png
2. multi_AZ_routeshifting_ha_0.png
3. cold_standby.png
4. aws_remote_access_autoscaling_group.png
5. remote_access_overview.png
6. aws_autoscale_cluster_c2s.png
7. cold_standby_01.png
8. aws_direct_connect.png
9. route_shifting_ha_5.png
10. cold_standby.png
11. segmentation.png
12. vpc_vpn_spoke_01.png

© Barracuda Networks Inc., 2026 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.