Site-to-site VPN

Introduction

Purpose

Create a secure connection from enterprise data center to the cloud infrastructure. This allows users to access the guest VMs by establishing a VPN connection to the VR of the account form a device in the datacenter of the enterprise. This eliminates the need to establish VPN connections to individual VMsState.

References

- Site-to-Site VPN
- Inter-VLAN Routing

Document History

Glossary

Feature Specifications

Requirements

- **Requirements**

- - Create VPN connection: Admin is allowed to establish a VPN connection from VR to a device in the datacenter of an enterprise
  - Destroy VPN connection: Admin should be able to delete a VPN connection between VR and external device
  - Reset VPN connection: Allow user to reset a VPN connection (in the event it's not functioning correctly)
  - Usage: CS should generate usage records for the VPN in terms of # of Bytes
  - VPN Logs: Ensure VPN configuration process is logged appropriately. These logs should aid troubleshooting when VPN does not function correctly.
- **Non-requirements**
  - VPN connection between two virtual routers within or across Zones within CloudStack
  - Support for end-points other than Cisco or Juniper mentioned above
- **Test guidelines**
  - Functional test
- **Configuration**

- - No global configuration.
- **System requirement:**
  - The remote VPN end-points must be:
    - Cisco ISR with IOS 12.4 or later
    - Juniper J-Series routers with JunOS 9.5 or later
    - Even though specific Cisco and Juniper devices are listed above the expectation is that any Cisco, Juniper device running the supported OSs will be able to establish VPN connections
- **Performance & scalability implications when feature is used from small scale to large scale**
  - One point VR may got heavy work for the whole network(VPC level) VPN connection
- **User permission**
  - Normal users would have privilege to create an site-to-site VPN.

Use cases

Here is the steps to create an site-to-site VPN:

1. Create a VPC(refer to Inter-VLAN Routing)
2. Create a customer VPN gateway.
3. Create a VPN gateway for the VPC.
4. Create VPN connection from CS VPN gateway to customer VPN gateway. After this step, VPN connection is established.
5. Destroy VPN connection.

User can also reset VPN connection, which would disconnect then connect again.

Architecture and Design description

- The software we used to support site-to-site VPN is OpenSwan.
- Use preshared key(PSK).
- The VPN protocol would be IPsec.
    - SSL is easier to penetrate firewall, but not interoperable standard.
- Support Phase 1(ISAKMP) and phase 2(ESP) encryption/hash:
    - AES128, AES192, AES256, 3DES
    - MD5, SHA1
    - Diffie-Hellman: Group 2, Group 5.
- Tables:
    - s2s_vpn_connection table
    - s2s_vpn_gateways
    - s2s_customer_gateways.
- Remote subnets are needed, and traffic target to remote subnets would be sent to VPN. No routing-based VPN supported.
- VPN connection monitor:
    - For every router.check.interval(30s by default), mgmt server would check the status of VPN connection. And if it's down, it would send out alert.
    - Mgmt server would only check the connections which states are "Connected" or "Disconnected". It would only change the state of VPN connection to one of these two states. It would not check connection in "Pending" or "Error" state.
    - One VPC router would only be checked if there are any implemented networks inside the VPC

Notes about VPN Connection:

1. This VPN connection would established a connection between VPC's CIDR specified subnet with all the subnets specified by customer gateway.
2. All the traffic between VPC's CIDR specified subnet and customer gateways' subnets would considered as VPN traffic.
3. If DPD is enabled, then dpddelay=30, dpdtimeout=120, dpdaction=restart is set. (Check ipsec.conf about these parameters).
4. If the VR which is handling VPN connection is restarted, the previous VPN connections would be established again automatically after router boot up(except the connection in Error state).

Web services APIs

- CreateVpnCustomerGateway: This API would give CS the information about the customer's gateway.
    - Parameters:
        - gateway: The remote gateway ip address
        - cidrlist: the guest CIDR list of remote subnets.
            - Divided by ",".
        - ipsecpsk: the preshared key of remote gateway
        - ikepolicy: phase 1 policy. Combination of (aes128|aes192|aes256|3des)-(sha1|md5);(modp1024|modp1536)
            - e.g. "aes-sha1", "3des-sha1;modp1536",
        - esppolicy: phase 2 policy. Same format as ikepolicy. If group is specified, then PFS is enabled.
        - ikelifetime: phase 1 lifetime in seconds. Default to 86400 seconds(1 day).
        - esplifetime: phase 2 lifetime in seconds. Default to 3600 seconds(1 hour).
        - dpd: if DPD(Dead Peer Detection) is enabled
        - account: Account name of owner.
        - domainid: Domain ID of owner.
    - Return:
        - Site2SiteVpnCustomerGatewayResponse.
    - Note
        - This API wouldn't validate the status of remote gateway.
        - Only one entry for one remote gateway IP.
- CreateVpnGateway: This API would create a CS VPN gateway. It would assign a public ip for the VPN gateway of VPC. The CIDR of VPC would become the source subnet of VPN connection.
    - Parameters:
        - vpcId: The ID of vpc which the gateway belong to.
    - Return:
        - Site2SiteVpnGatewayResponse.
    - Note:
        - Only one VPN gateway for one VPC.
        - CS would choose address for the ip automatically, current it only choose source nat.
- CreateVpnConnection: This API would established the connection between CS's VPN gateway and customer's VPN gateway.
    - Parameter:
        - s2scustomergatewayid
        - s2svpngatewayid
    - Return:
        - Site2SiteVpnConnectionResponse.
    - Note
        - Only one connection between two gateways.
        - Only one connection each gateway can have.
        - The status would kept in vpn gateway table.
        - Status of VPN connection:
            - Pending: Trying to make connection to customer gateway.
                - ResetVpnConnection won't work with Pending state.
            - Connected: The VPN connection is established.
            - Disconnected: The VPN connection is lost, mainly due to router reboot or detected by CS connection status update thread.
            - Error: Fail to make VPN connection to customer gateway, mainly due to wrong configuration.
- DeleteVpnConnection:
    - Parameter:

        - VPN connection ID
    - Return:
        - Success or not.
- ResetVpnConnection:
    - Parameter:
        - VPN connection ID
        - VPN should only in
    - Return:
        - VPN connection ID, or failed.
- DeleteVpnCustomerGateway:
    - Parameter:
        - Vpn Customer Gateway ID.
    - Return
- DeleteVpnGateway:
    - Parameter:
        - Vpn Gateway ID
    - Return
- UpdateVpnCustomerGateway:
    - Parameter:
        - Vpn CustomerGateway ID.
        - Other the same as CreateVpnCustomerGateway.
    - Return:
        - Site2SiteVpnCustomerGatewayResponse.
- ListVpnCustomerGateways: (to be extended)
    - Parameter:
        - GatewayID
    - Return
- ListVpnGateways: (to be extended)
    - Parameter:
        - Vpn gateway ID.
    - Return
- ListVpnConnections: (to be extended)
    - Parameter
        - Connection ID
    - Return

UI flow

Current limitation(updated by Aug 31st, 2012):

1. Don't support isolate network for site-to-site VPN. Only support to use with VPC.
2. No scalability test.
3. Source NAT ip would be reused as site 2 site VPN IP.
4. We don't allow different VPC connect to the same Customer gateway at the same time, even for different accounts, this would implicit following:
    a. Once one user is create a customer gateway ip as a.b.c.d, any other user cannot create the customer gateway with the same ip.
    b. If one user already have one VPC connected to the customer gateway, it's not allowed to have another VPC belong to the same or different user connect to the same customer gateway.

Configuration reference for CISCO router:

```
crypto isakmp policy 1
encr 3des
hash md5
authentication pre-share
group 5
crypto isakmp key 123123 address 10.223.165.25 <-----this is VPN gateway ip
!
!
crypto ipsec transform-set proposal4 esp-3des esp-md5-hmac
!
crypto map cmap 2 ipsec-isakmp
set peer 10.223.165.25
set transform-set proposal4
match address 165

ip nat inside source list 185 interface GigabitEthernet0 overload
```

```
interface GigabitEthernet0
ip address 10.223.85.18 255.255.255.0
ip access-group 102 in
ip access-group 102 out
ip nat outside
ip virtual-reassembly
duplex auto
speed auto
fair-queue
crypto map cmap

access-list 102 permit ip any any

access-list 165 permit ip 172.16.10.0 0.0.0.255 10.10.0.0 0.0.255.255        <- We allow two subnets here
access-list 165 permit ip 192.168.10.0 0.0.0.255 10.10.0.0 0.0.255.255

access-list 185 deny ip 172.16.10.0 0.0.0.255 10.10.0.0 0.0.255.255        <- you need to disable NAT for certain subnet
access-list 185 permit ip 172.16.10.0 0.0.0.255 any
```

Configuration Reference for Juniper SRX Router

SRX-S2S-VPN.pdf

Comments

- Is there any provision in UI/API for enabling/disabling vpn for the accounts?
  [Sheng] No, because it should only be enabled for Admin
    - UDPATE: There is still no provision in UI/API for enable/disable VPN feature for accounts. But we would support it for normal users.
- What are the VPN protocols supported by VR and supported algorithms as well?
  [Sheng] Would update that in the FS. Protocol is IPsec. Use pre-shared key for authentication. Support encryption including aes, des, 3des; hash including md5 and sha1.
- Do we support both site-to-site vpn and remote access vpn on the VR at the same time?
  [Sheng] No.
- How do we specify the datacenter device (another end of site-to-site vpn) details in CS ?
  [Sheng] Would update that in the FS.

- Does site-to-site vpn supported only on SourceNat IP or any aquired IP as well?
  [Sheng] It's supported on all acquired (public) IP.

- Could you explain the requirement :

- Add Route: As part of the configuration, users should be able to specify a list of routes. Traffic to these routes should be directed to the VPN tunnel interface
  [Sheng] Don't need to do it anymore IIUC. VPN software would deal with that route. I would update the FS.

- Could you give the details on usage tables to look for usage records for the VPN?
  [Sheng] Haven't got it done. Would update the spec later.

- Could you specify the log file details to look for vpn logs?
  [Sheng]Currently they're at /var/log/auth.log and /var/log/daemon.log in the VR.

- We have Juniper SRX with JUNOS10.4 . Can we use it as one end point of the vpn?
  [Sheng]Should be fine. I have only tested with Cisco router now. Would test Juniper SRX soon.
- What are the hypervisors we are supporting?
  [Sheng] All major ones.

- Will S2S vpn supported in upgrade environment?
  [Sheng] I don't think we support upgrade to VPC?

- Can a customerVpnGateway participate in more than one s2s vpn ? Or one customerVpnGateway can have connections with multiple CS Vpn Gateways?
  [Sheng] Currently they are all 1:1.

- If a multitier architecture has 3 guest networks and all are attached to vpcVR, what are the networks included in vpn ?
  [Sheng] All. They should in the same CIDR covered by VPC.

- I assume that the customerVpnGateway configuration details are supplied from the customer and we just give CS this information using createVpnCustomerGateway API will do this.
  We don't configure the customerVpnGateway. Please let me know if my understanding is wrong.
  [Sheng] You're right.

- vpcVR has souceNat enabled by default. In general s2s Vpn will disable the source nat on the vpn gateway. Vpn gateway will encapsulate the IP packet with vpn ips.
  Could you explain the behavior in our s2s vpn implementation?
  [Sheng] NAT is not needed if you communicate between subnets, and it won't affect the normal traffic to the outside. What's the issue here?

Appendix

Appendix A:

Appendix B:

UI Flow .