Link Balancing and Failover

https://campus.barracuda.com/doc/13304031/

On the Barracuda NextGen Firewall X-Series, you can configure inbound link balancing, outbound link balancing, and outbound link failover. Link balancing is also sometimes called 'link aggregation'.

Outbound Link Balancing and Failover

To achieve outbound link load balancing, create a connection object that balances the traffic among multiple links. Then use this connection object in the firewall rules that direct outgoing traffic. The connection object specifies what happens if multiple links are configured. Options include:

- If one interface becomes unavailable, the traffic fails over to the next available link in the sequence.
- Use a set of interfaces in weighted round robin fashion. You can specify the weights for each interface in the connection object.
- Randomly choose one of a list of interfaces.

For more information about configuring connection objects, see How to Configure Outbound Loadbalancing and Failover.

Inbound Link Balancing and Failover Using DNS

You can use DNS to balance inbound traffic among multiple links. Associate your domain name (or names) with multiple IP addresses, each of which represents an external interface. When the DNS request for the domain name is resolved, all of these IP addresses are included in the answer. The DNS server can vary the order of the IP addresses, and the client uses the first entry in the list to access your site. You can add multiple DNS entries with the same IP address to send more queries to the preferred WAN interface. Configure the X-Series Firewall as the authoritative DNS resolver for the domain name.

For more information, see How to Configure Authoritative DNS.

Inbound Failover and Load Balancing Using DNAT Access Rules

You can use load balancing and failover in a DNAT access rule to distribute incoming traffic to multiple internal servers. Add additional IP addresses to the network object referred to in the rule, or enter

them in the **Redirect** list of the rule. Depending on the configuration, all traffic is initially sent to the first IP address and, if this address is no longer reachable, to the second, and so forth (fallback mode), or distributed to all IP addresses depending on the mode set in the rule: round robin or cycle.

For more information, see Example - Configuring a DNAT Access Rule.