OSF : Open Service Framework

A high-speed flow steering and load balancing framework

Anand Gorti
IBM

Vijoy Pandey
Blade Network Technologies

DC-CAVES
September 14th, 2009

Security Appliance Deployment Without OSF

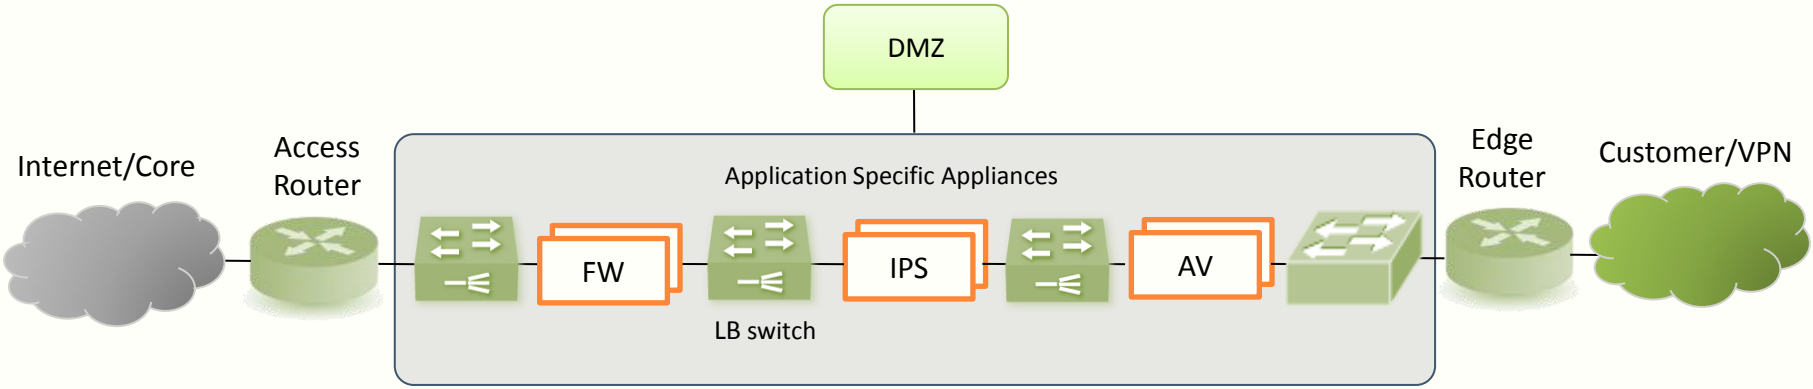

- Deployment for security applications require sequential processing of traffic through a series of application groups.
- Require ability to filter and load balance traffic for each specific application group
  - Packets pass through each server
  - Lower performance – degrades @ every step
  - High server overhead to perform tasks such as LB; poor load balancing
  - No flexibility in services offered

Security Appliance Deployment With OSF

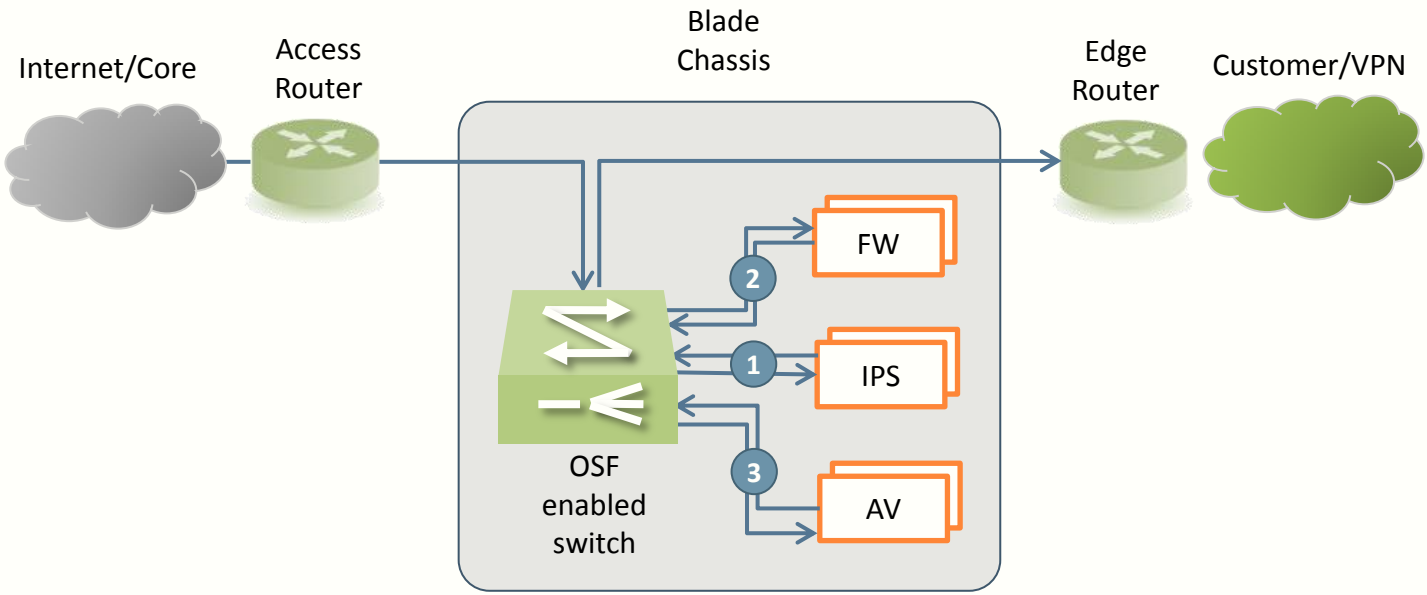

Open Service Framework

- Provide infrastructure to intelligently steer traffic between application groups
- Deployed as a transparent Ethernet bridge (“bump in the wire”) at the Service Provider edge network.
  - Does not alter layer 3 network design
  - Requires no IP address in the data-plane
- Provide session-less load balancing at line rate performance
- Consolidates various security applications on standard servers or specialized network processing hardware to be deployed in a blade server environment

OSF : Sequencing

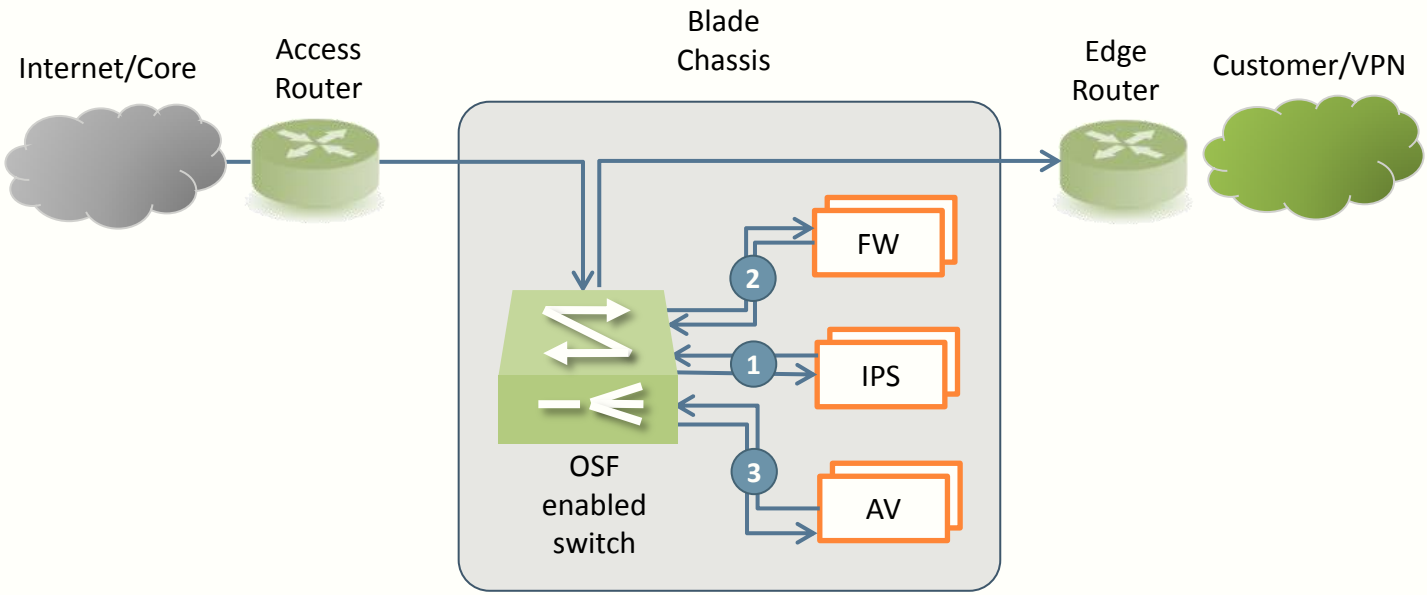

- Simplest incarnation, traffic can be mapped as

  <ingress port : app group>

  or

  <ingress port : L2 switch>

- Sequence is defined as one or more of the above mappings

  <sequence> = <ingress port : app group>, …, <ingress port : L2 switch>

- Allows for service chaining

OSF : Sequencing & Load Balancing

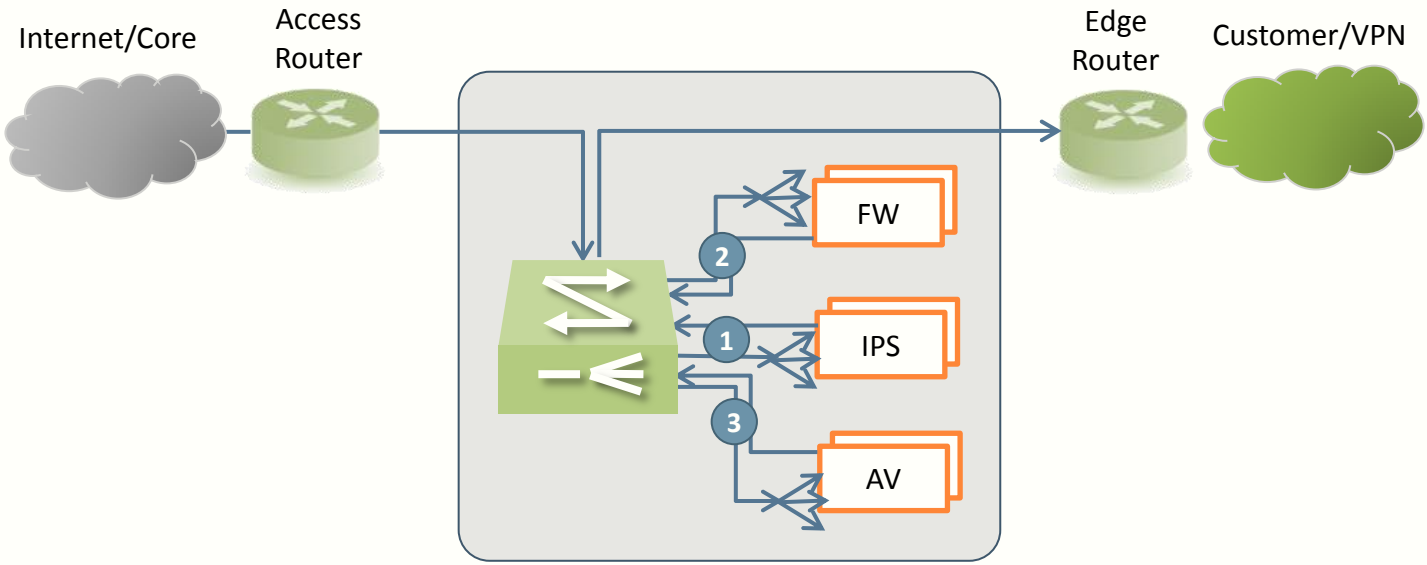

- Server persistence maintained via Layer 2-4 attributes
- N+1 redundancy on application groups
- Flow stickiness maintained on server failure - Only traffic from the failed blade will be redistributed to the remaining active servers in the group.

OSF : Sequencing & LB per flow

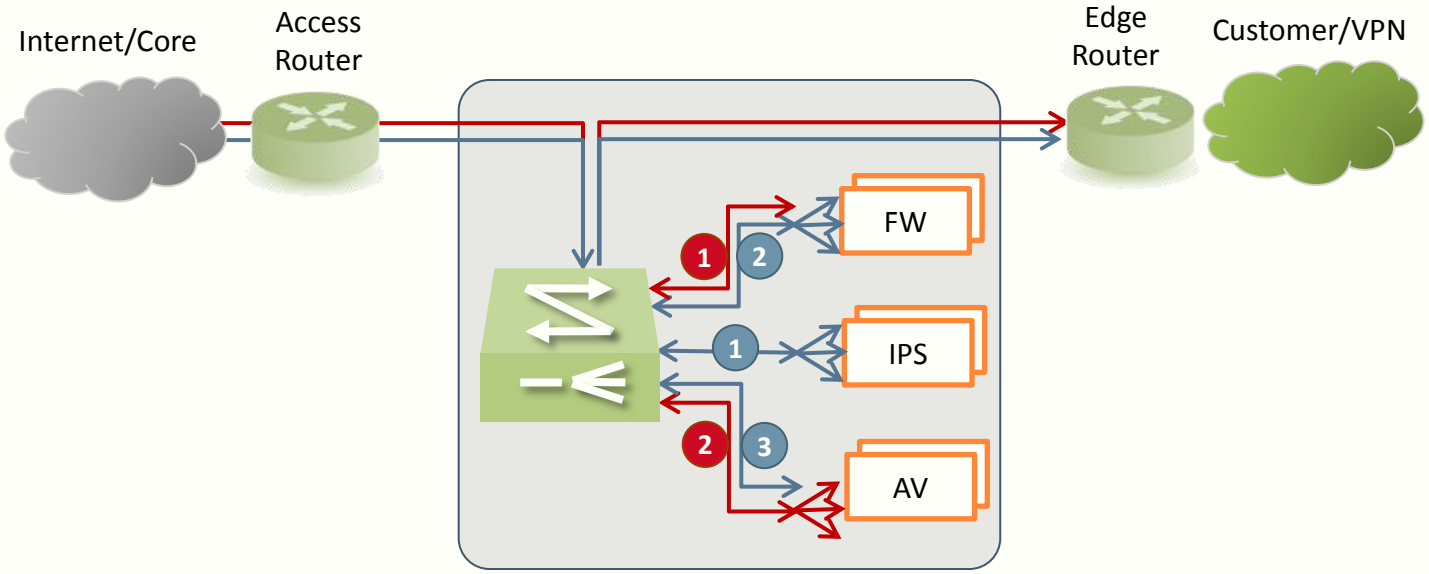

- User defines flows into system based on :

MAC SA, DA, Ether Type	VLAN ID	IP SRC, DST, Proto	TCP S-Port, D-Port	• •	<sequence>

- Flows are mapped to service sequences to LB groups
  - Allows for service chaining and forking **per flow**
  - Allows for service customizations and tiers
- Load balance at every service group

OSF : Service-aware Health Monitoring

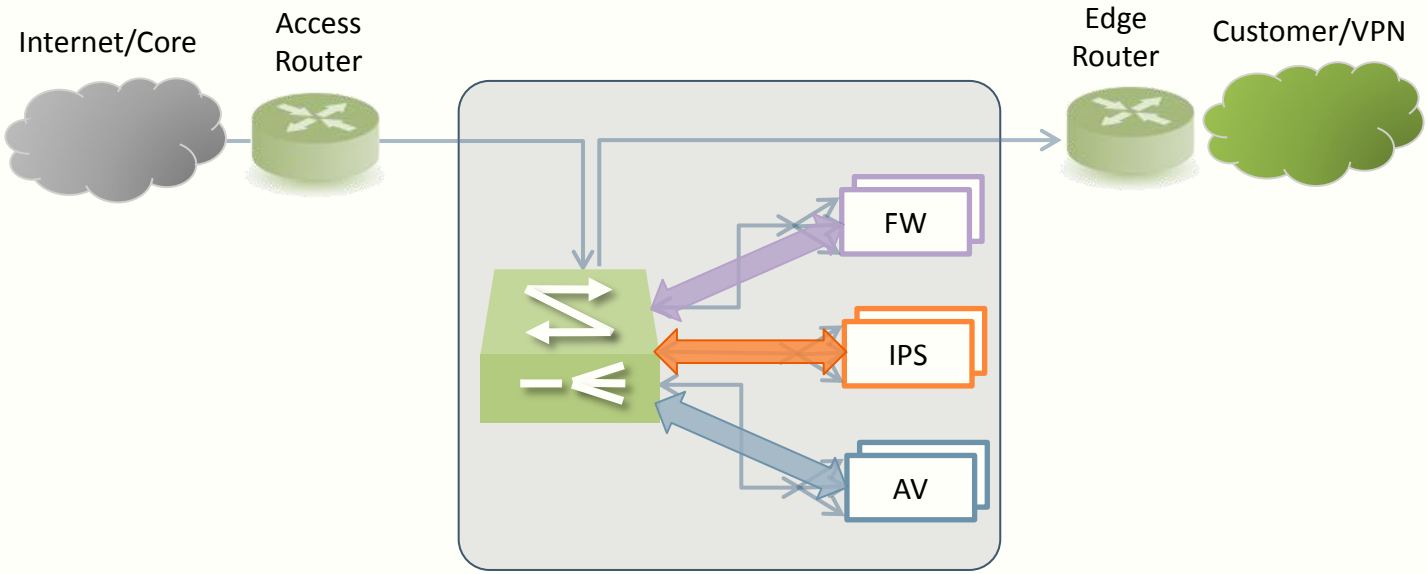

- Link State (Default)
- ICMP Ping or ARP
- TCP Health Check
- Application Specific Health Checks
  - HTTP
  - HTTPS/SSL
  - SMTP
  - Session Initiated Protocol (SIP)
  - UDP-based DNS
- Scriptable for state-full health monitoring

OSF : High Availability Sandwich

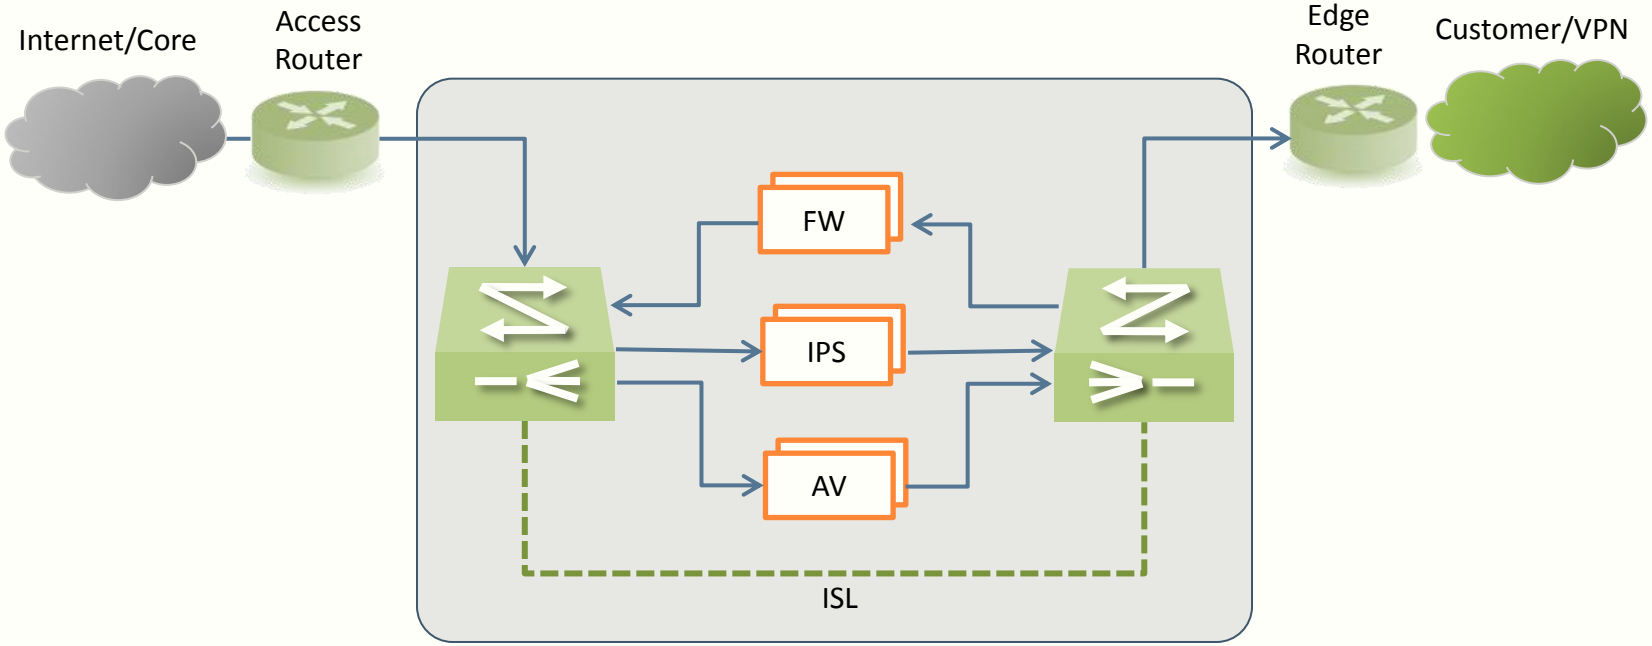

OSF : High Availability Sandwich

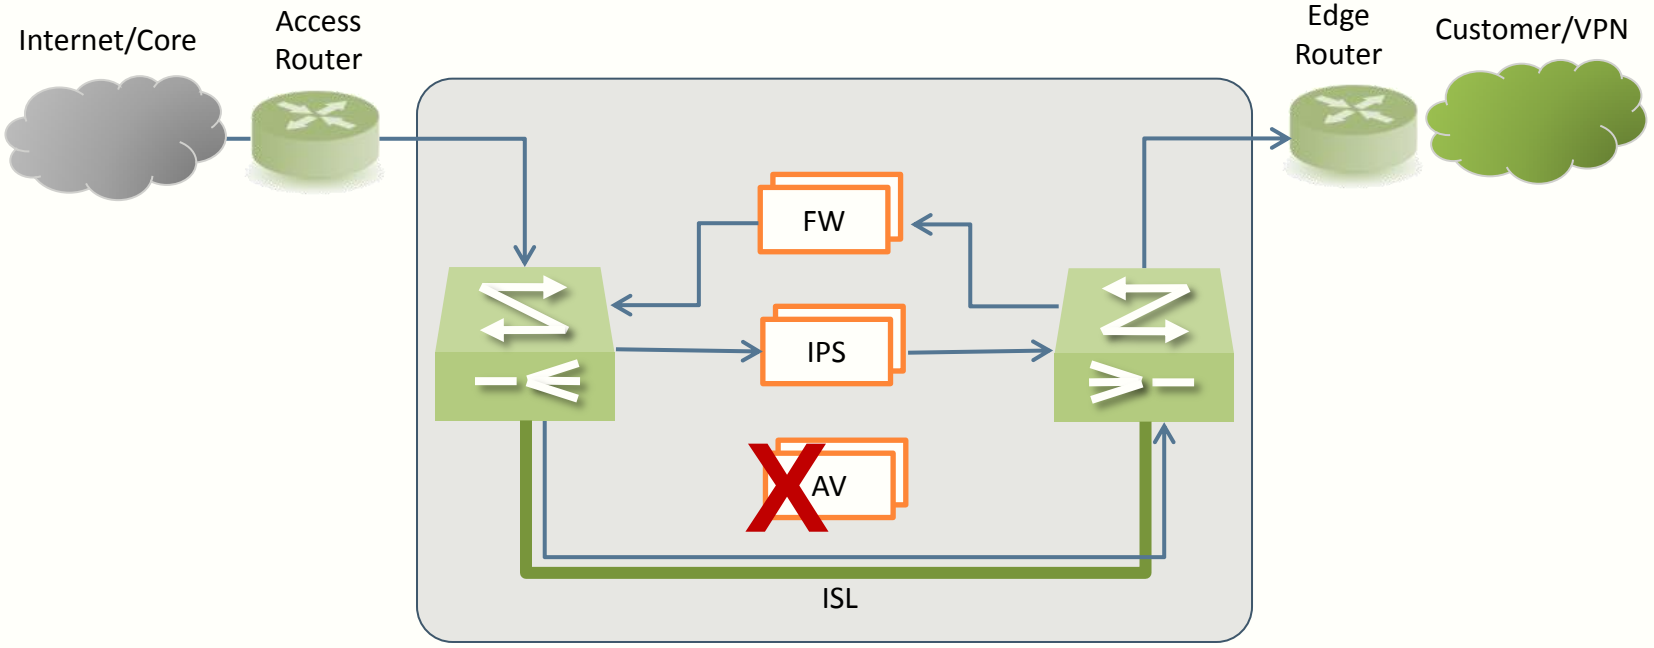

- Sandwich configurations do not require changes to application software
  - Most applications work under the assumption of a dual-home design
- ISL is available on some commercially available blade chassis by default
- Upon service failure [all blades in app group down, physically or service-wise]
  - Option of dropping packets destined to that application (e.g., IPS)
  - Option of bypassing the application

OSF : Architecture

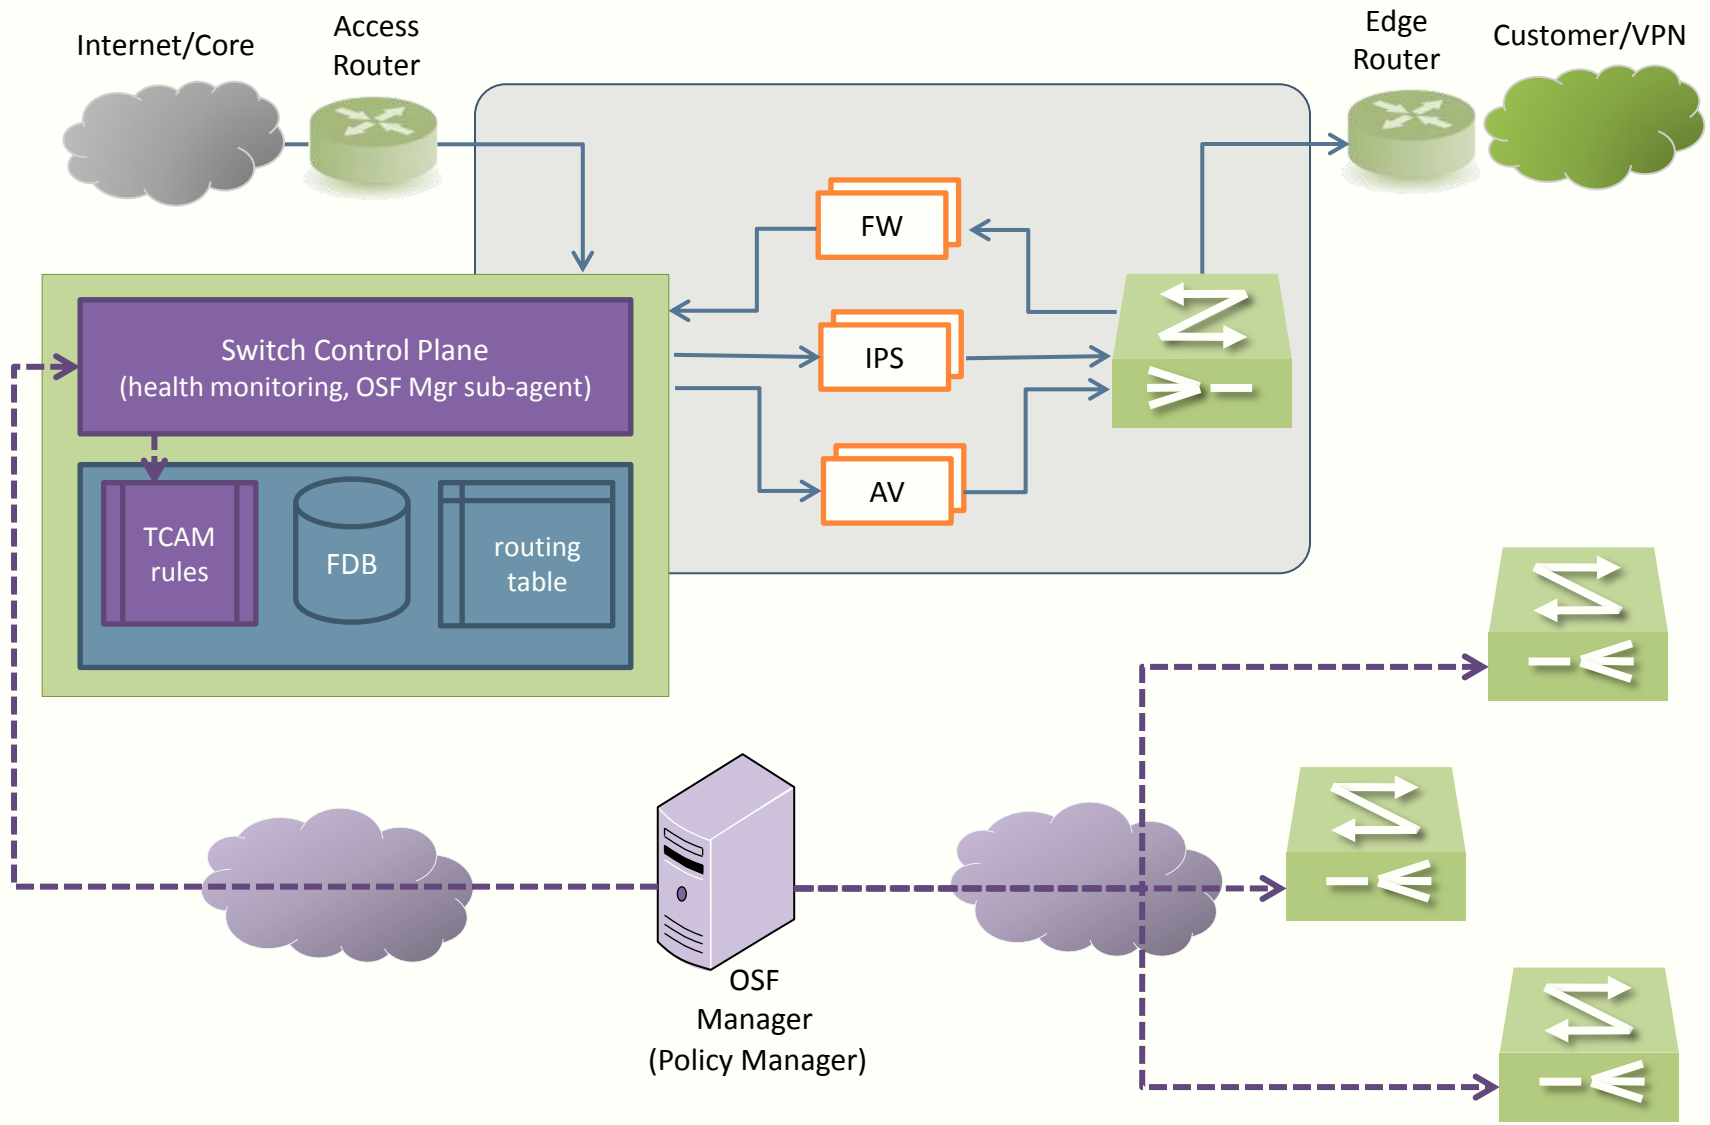

OSF Next Steps : Virtual Appliances - I

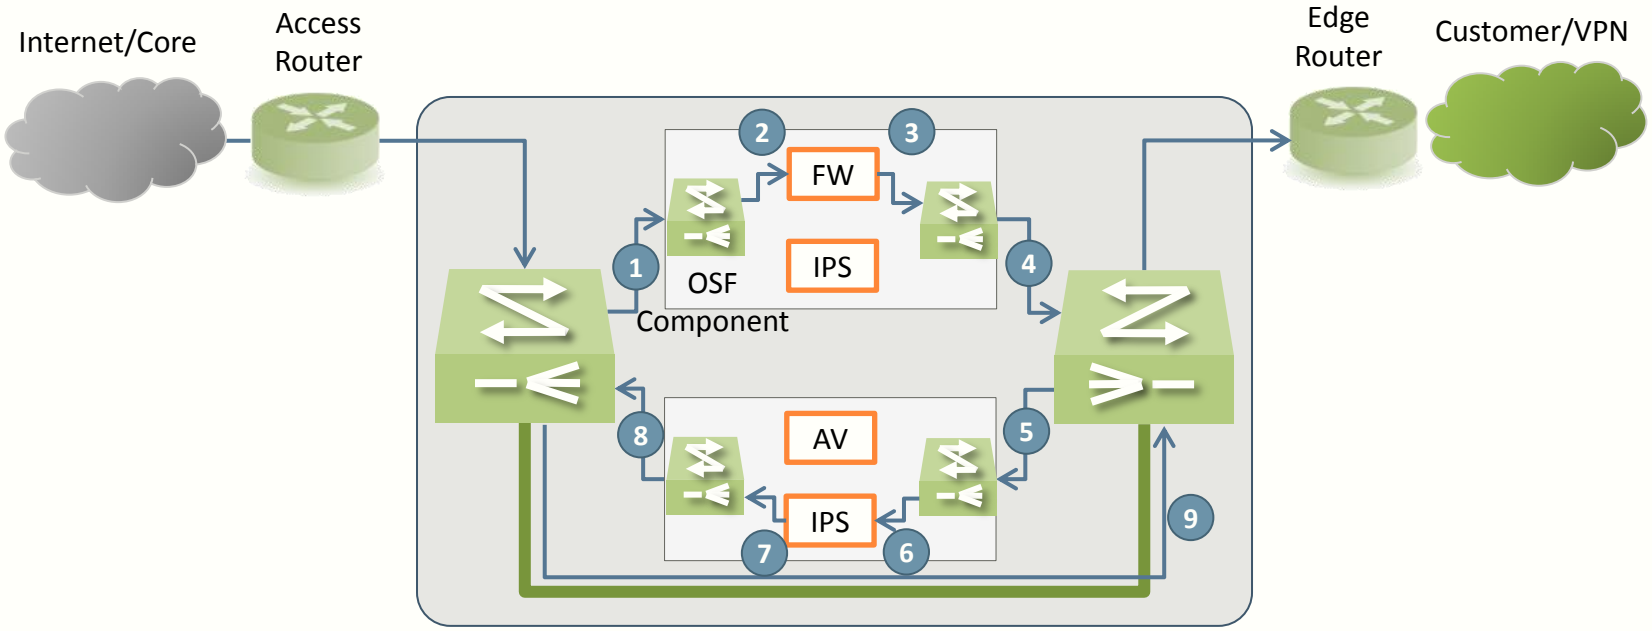

- Sequencing and load balancing needed between virtual appliances
- With VEB, modeled as a nested OSF domain
  - More intelligence needs to reside in the OSF Manager

OSF Next Steps : Virtual Appliances - II

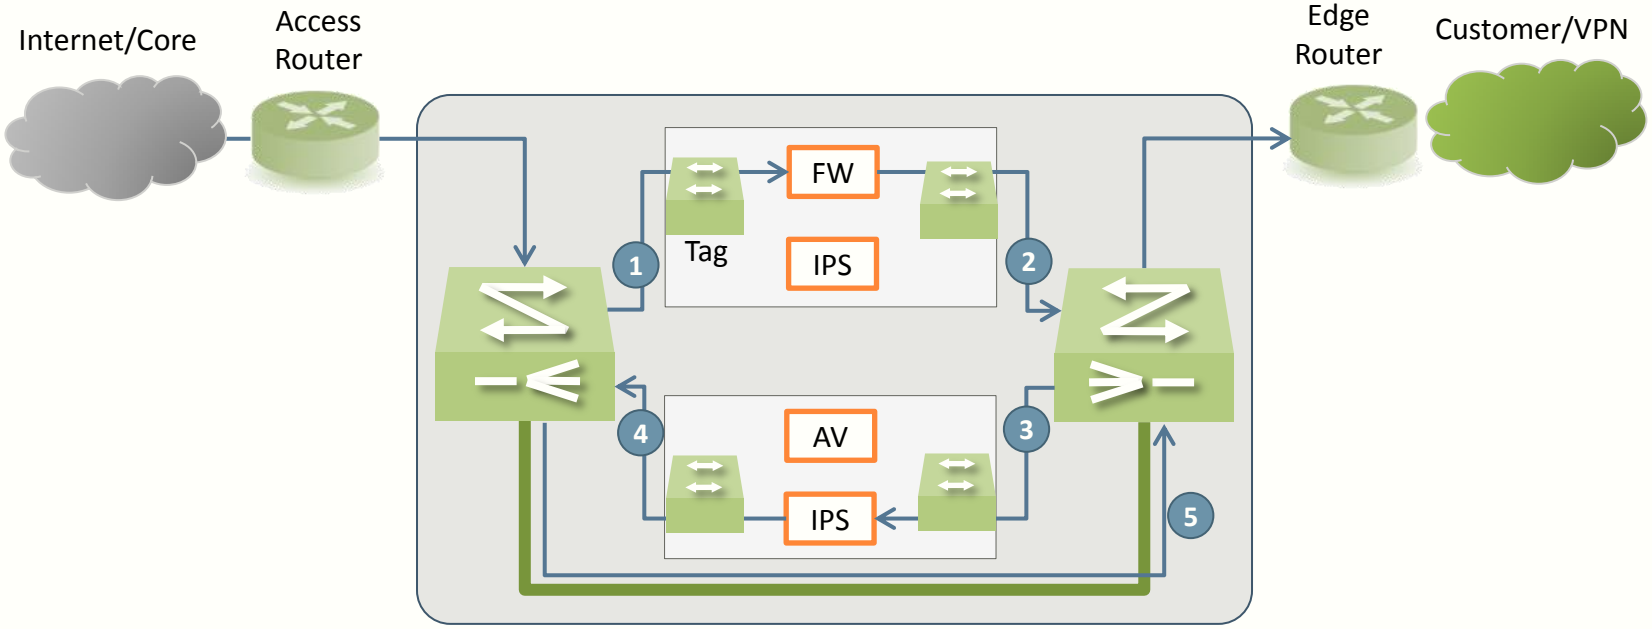

- Sequencing and load balancing needed between virtual appliances
- With a tag approach [S-Tag, VN-Tag]
  - OSF Manager only grows in scale to handle virtual ports
  - Bladed switch performs sequencing and load balancing on virtual ports

OSF : Open Service Framework

Thank You