Cisco live!

What You Make Possible

TOMORROW
starts here.

CISCO

Advances in Routing

BRKRST-3370

TOMORROW starts here.

Cisco Open Network Environment Platform Kit (onePK)

The Open Network Environment

What

- Open Network Environment – Complementing the Intelligent Network

    Preserve what is working:
    Resiliency, Scale and Security, Comprehensive feature-set

    Evolve for Emerging Requirements:
    Operational Simplicity, Programmability, Application-awareness

- The Open Network Environment integrates with existing infrastructure

    Software Defined Network concepts are a component of the Open Network Environment

    The OpenFlow protocol can be used to link agents and controllers, and as such is component of SDN as well

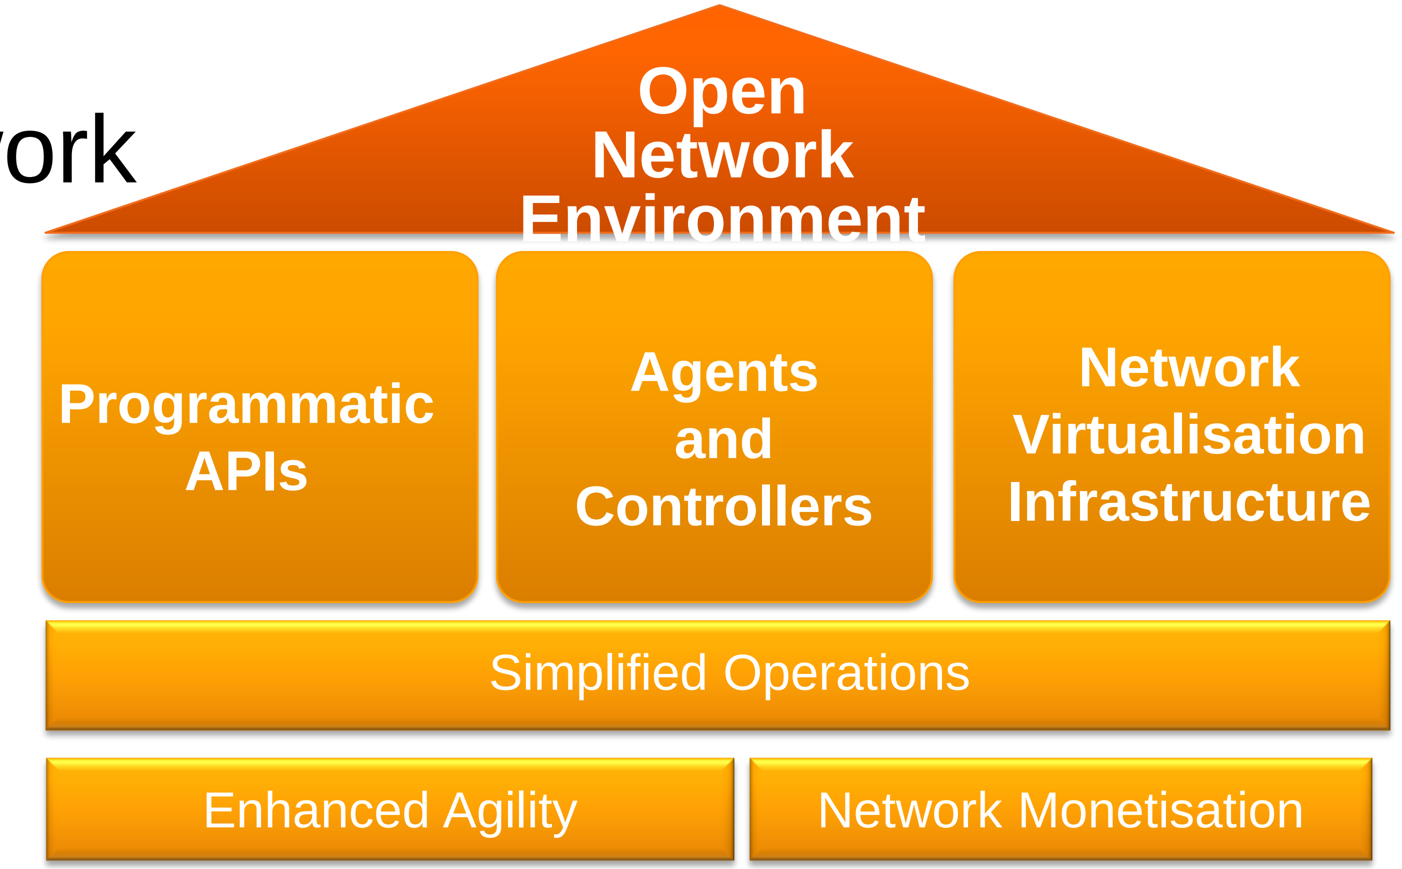

Evolving How We Interact With The Network Operating System

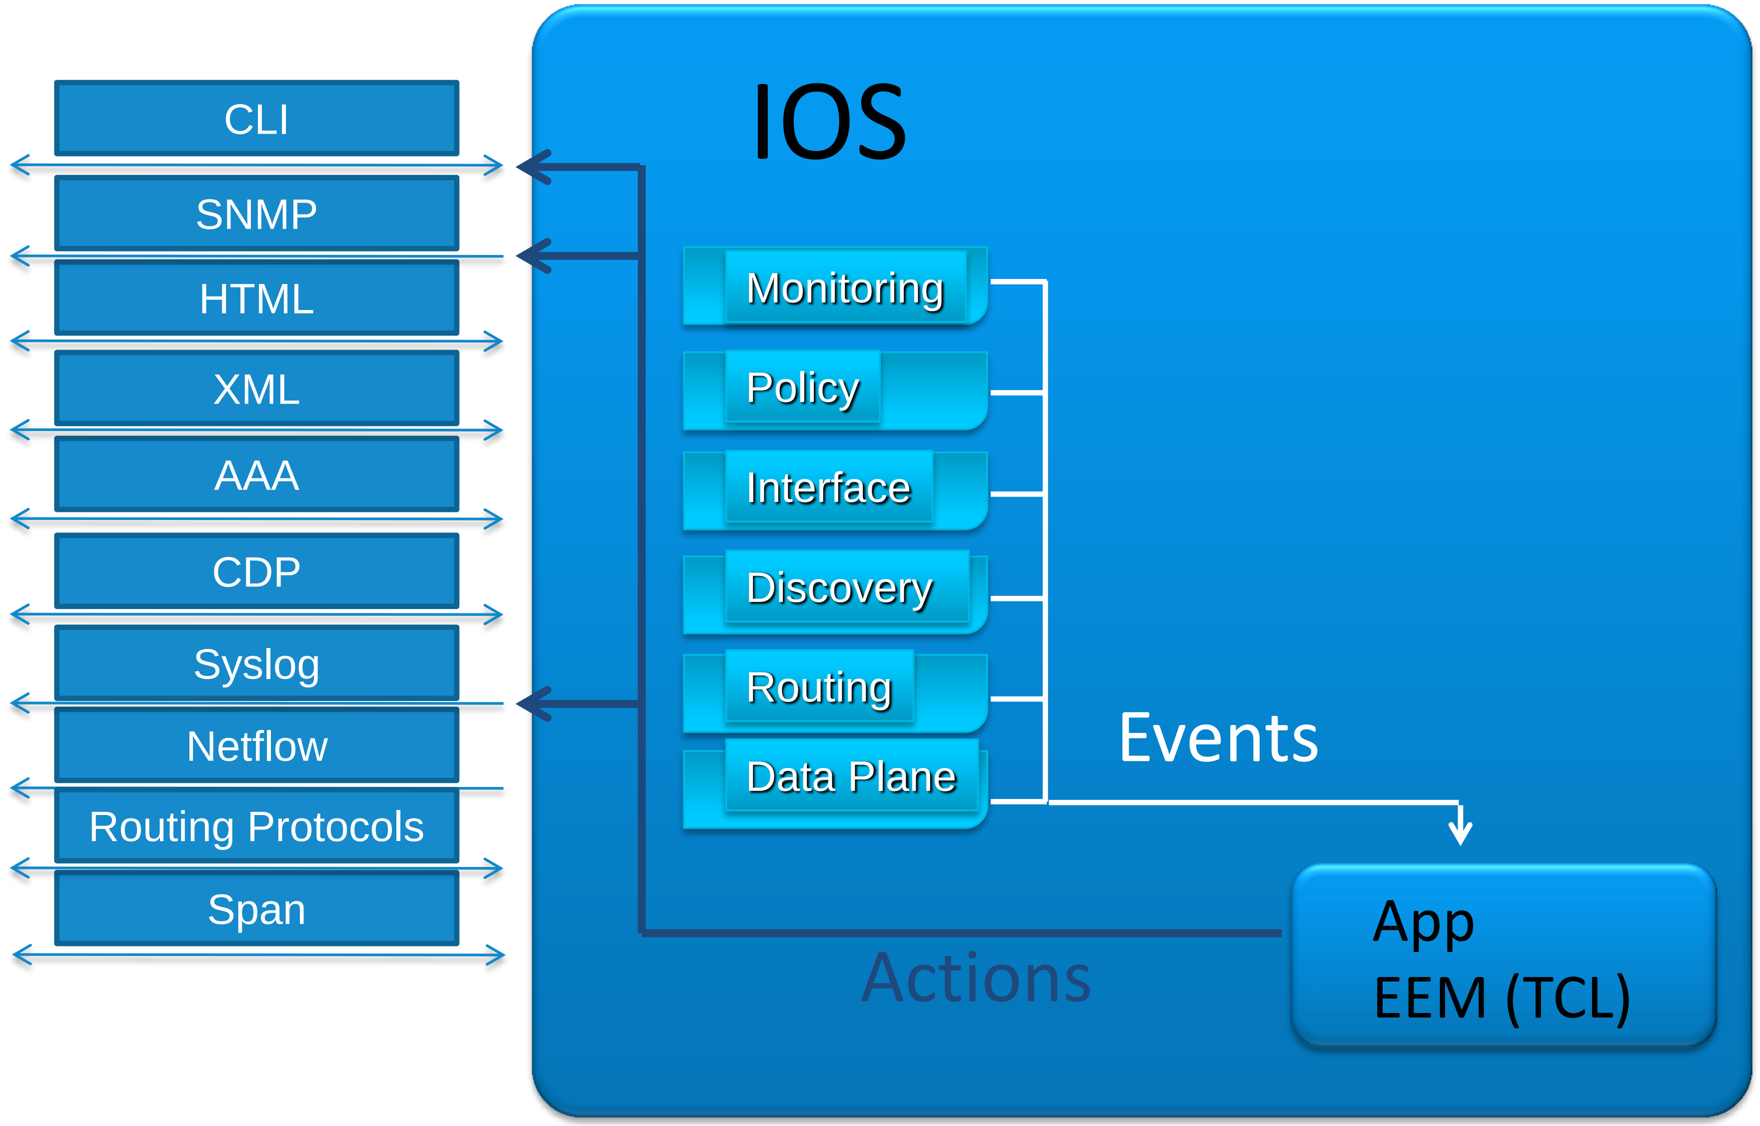

Introducing One Platform Kit (onePK)

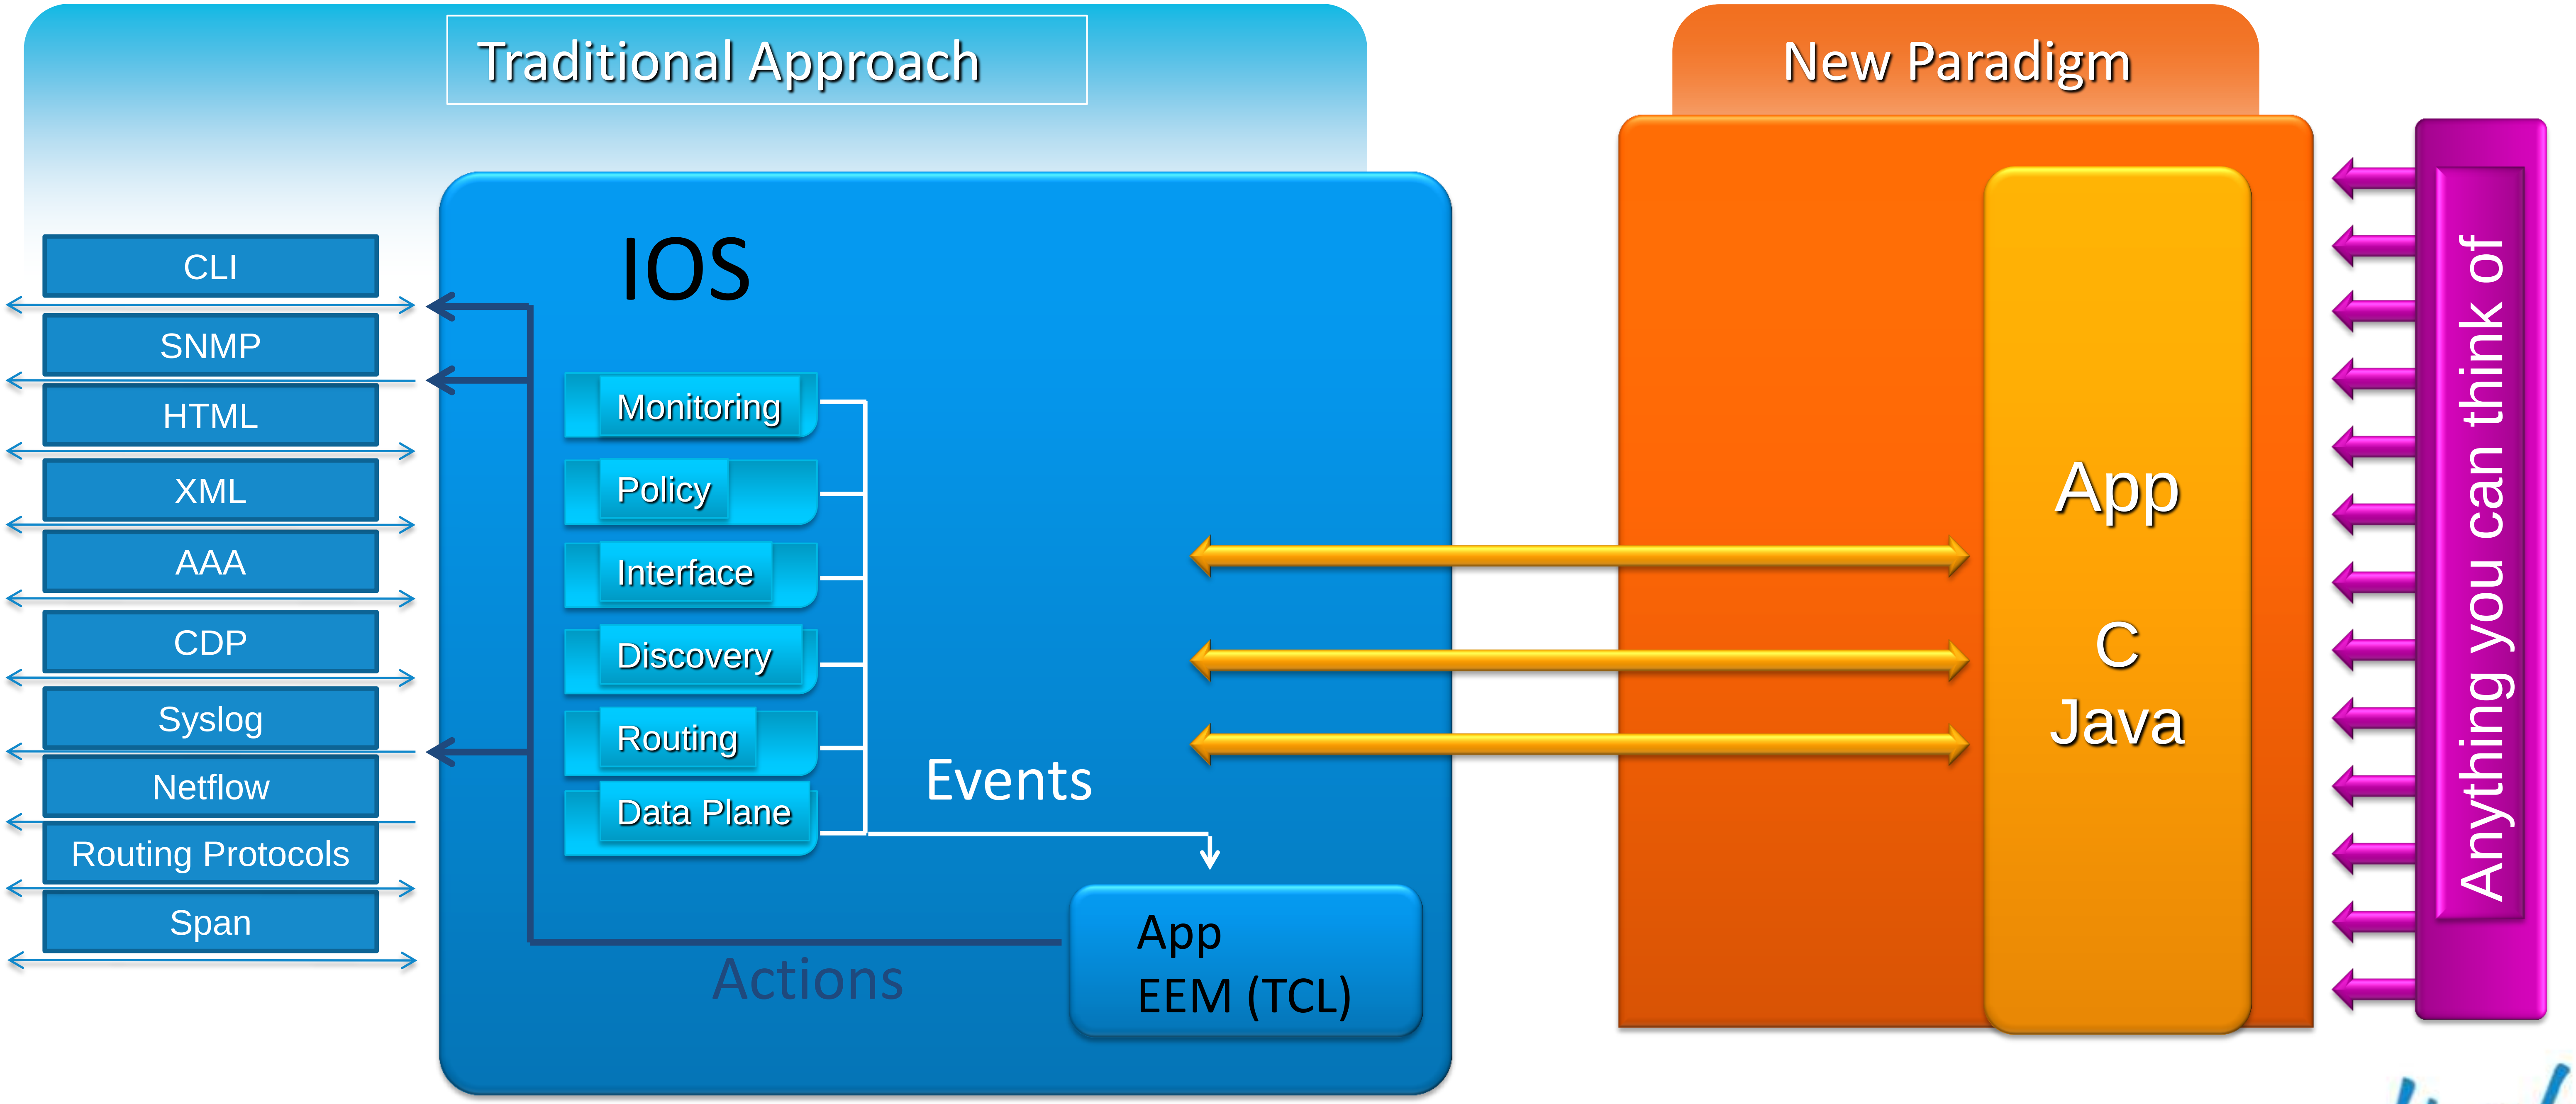

Introducing One Platform Kit (OnePK)

Flexible development environment to:

Innovate

Extend

Automate

Customise

Enhance

Modify

OnePK Architecture

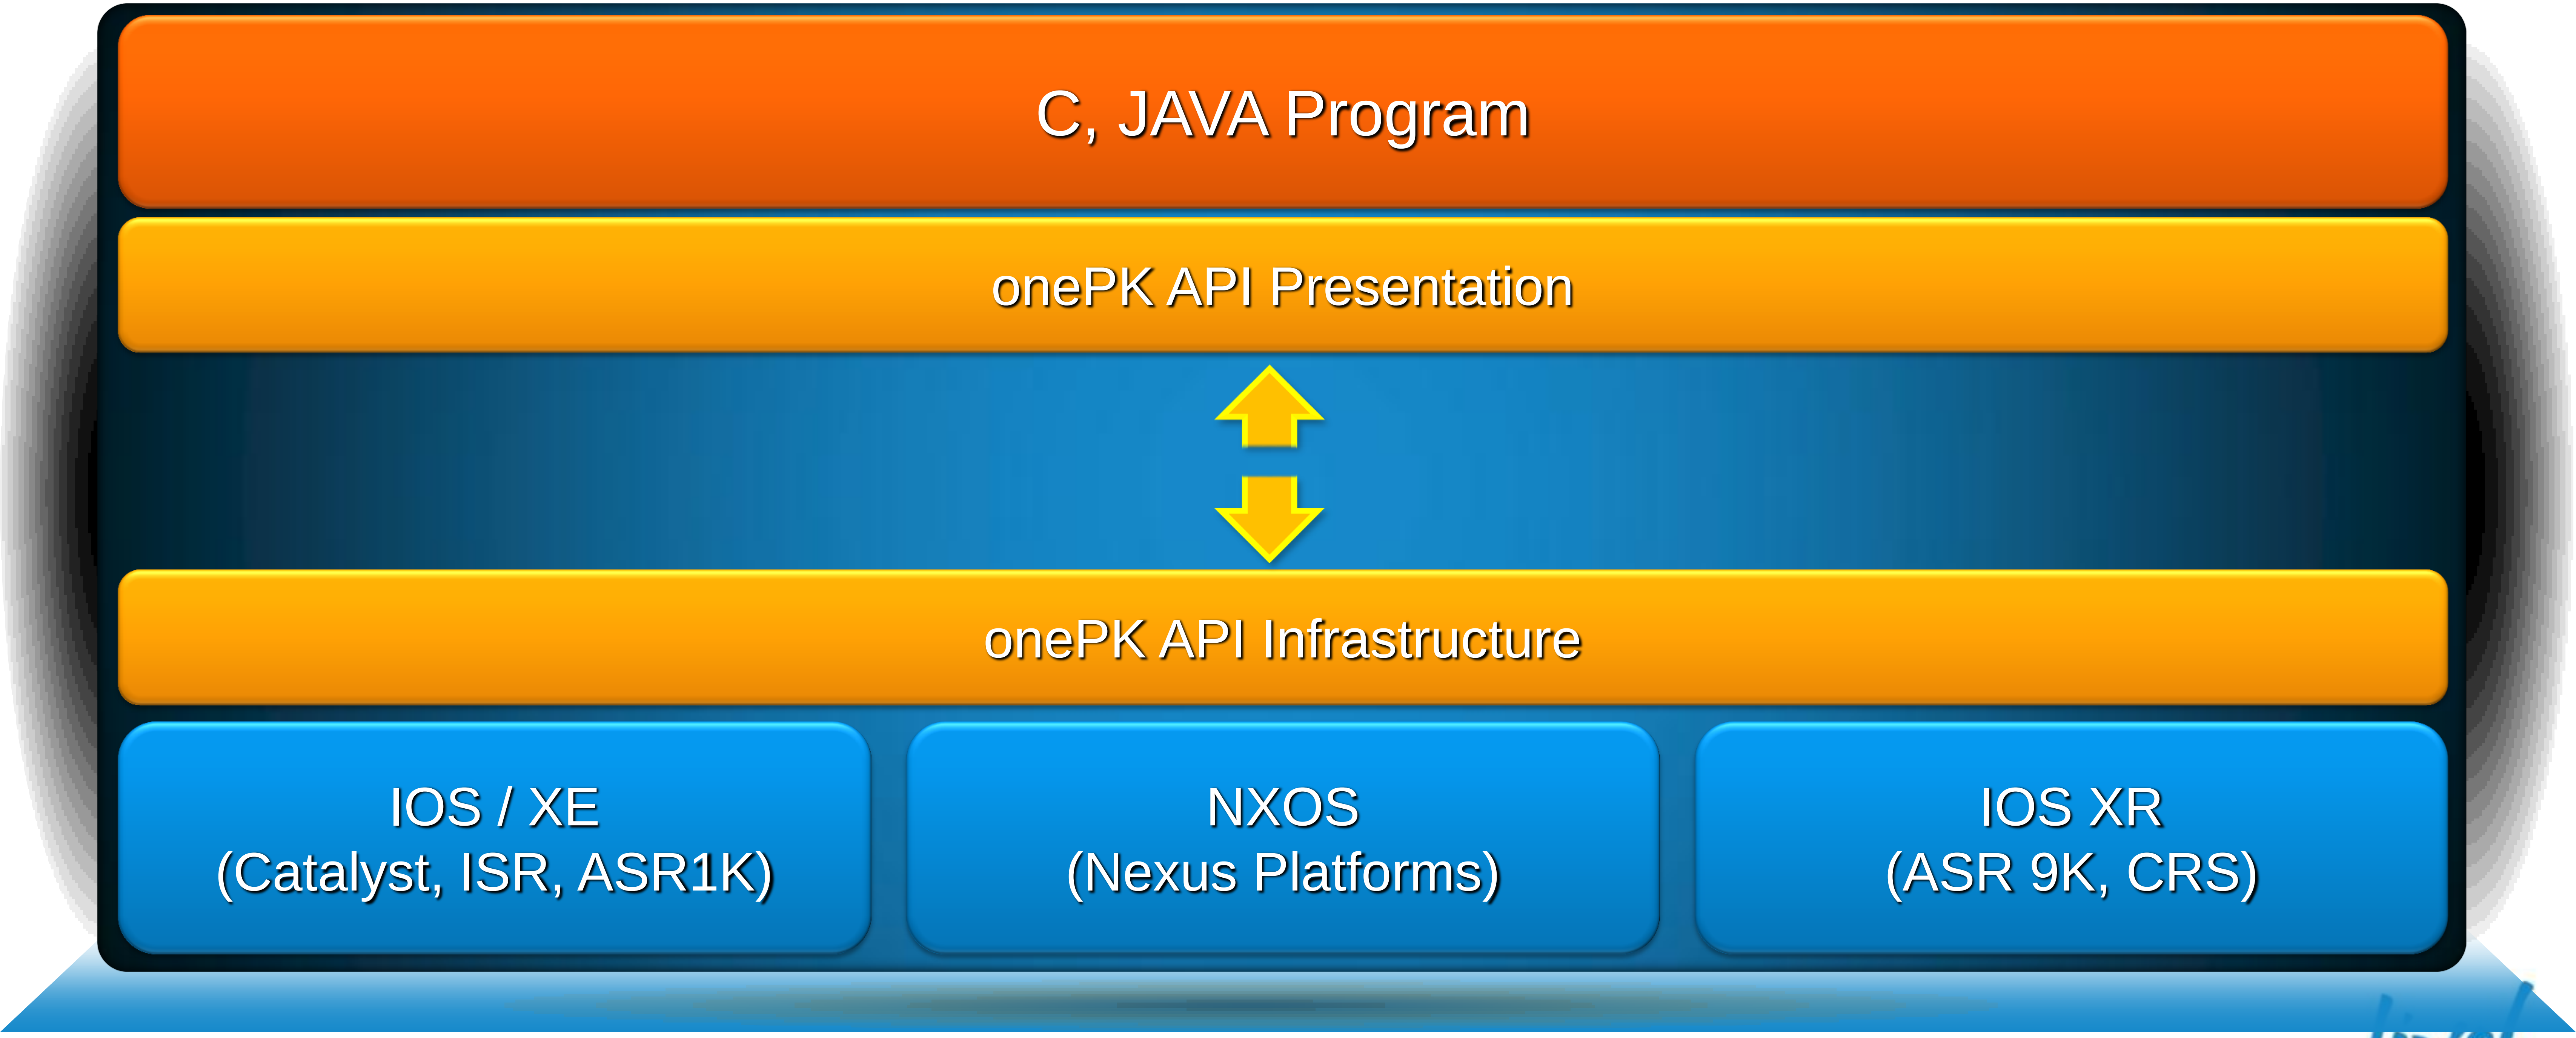

OnePK Application Hosting Options

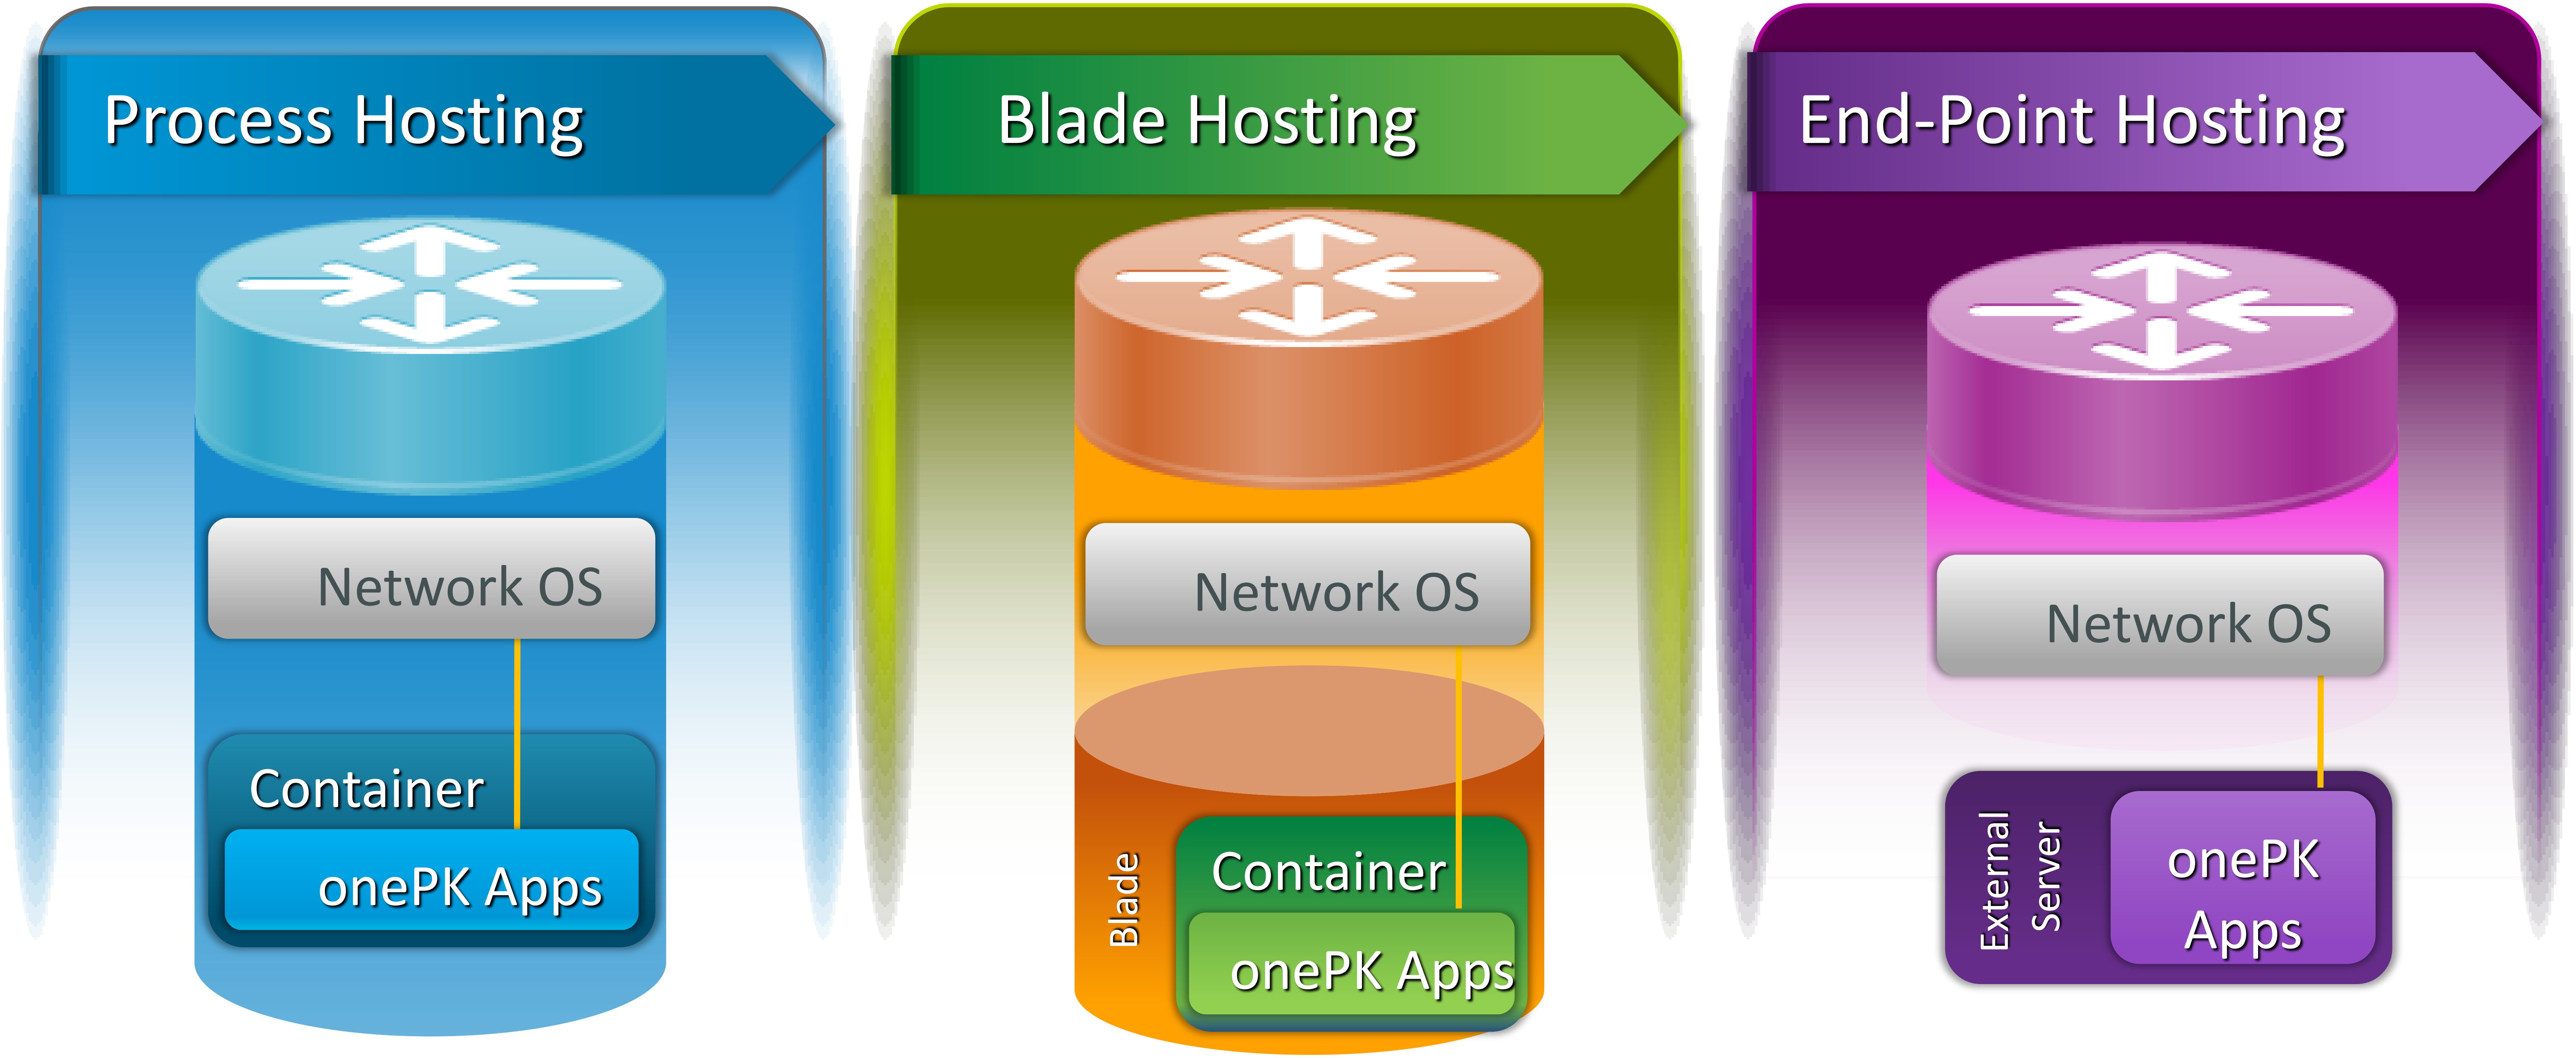

Write Once, Run Anywhere

OnePK

Why – SDN drivers

"A way to optimise link enhanced, applica

Simplified Operations

open solution for customised flow forwarding control in and between Data Centers"

"An open solution for VM mobility in the Data-Center"

"A pl

"Develop solutions at software speeds: I don't want to work with my network vendor or go through lengthy standardisation."

"A way to reduce the CAPEX of my network and leverage commodity switches"

"A means to get assured quality of experience for my cloud service offerings"

"A solution to build a very large scale layer-2 network"

"A means to do traffic engineering without MPLS"

"A solution to build virtual topologies with optimum multicast forwarding behavior"

Diverse Drivers

Common Concepts

Different E

"A means to scale my fixed/mo gateways and optimise their placement"

Enhanced Agility

adcast TV del placement an ection"

New Business Opportunities

ld my own ion solution"

"A way to scale my firewalls and load balancers"

"A way to distribute policy/intent, e. for DDoS prevention, in the network"

"A way to configure my entire network as a whole rather than individual devices"

"A solution to get a global view of the network – topology and state"

OnePK Service Sets

How

Base Service Set	Description
Data Path	Provides packet delivery service to application: Copy, Punt, Inject
Policy	Provides filtering (NBAR, ACL), classification (Class-maps, Policy-maps), actions (Marking, Policing, Queuing, Copy, Punt) and applying policies to interfaces on network elements
Routing	Read RIB routes, add/remove routes, receive RIB notifications
Element	Get element properties, CPU/memory statistics, network interfaces, element and interface events
Discovery	L3 topology and local service discovery
Utility	Syslog events notification, Path tracing capabilities (ingress/egress and interface stats, next-hop info, etc.)
Developer	Debug capability, CLI extension which allows application to extend/integrate application’s CLIs with network element

OnePK
Getting Properties and Statistics

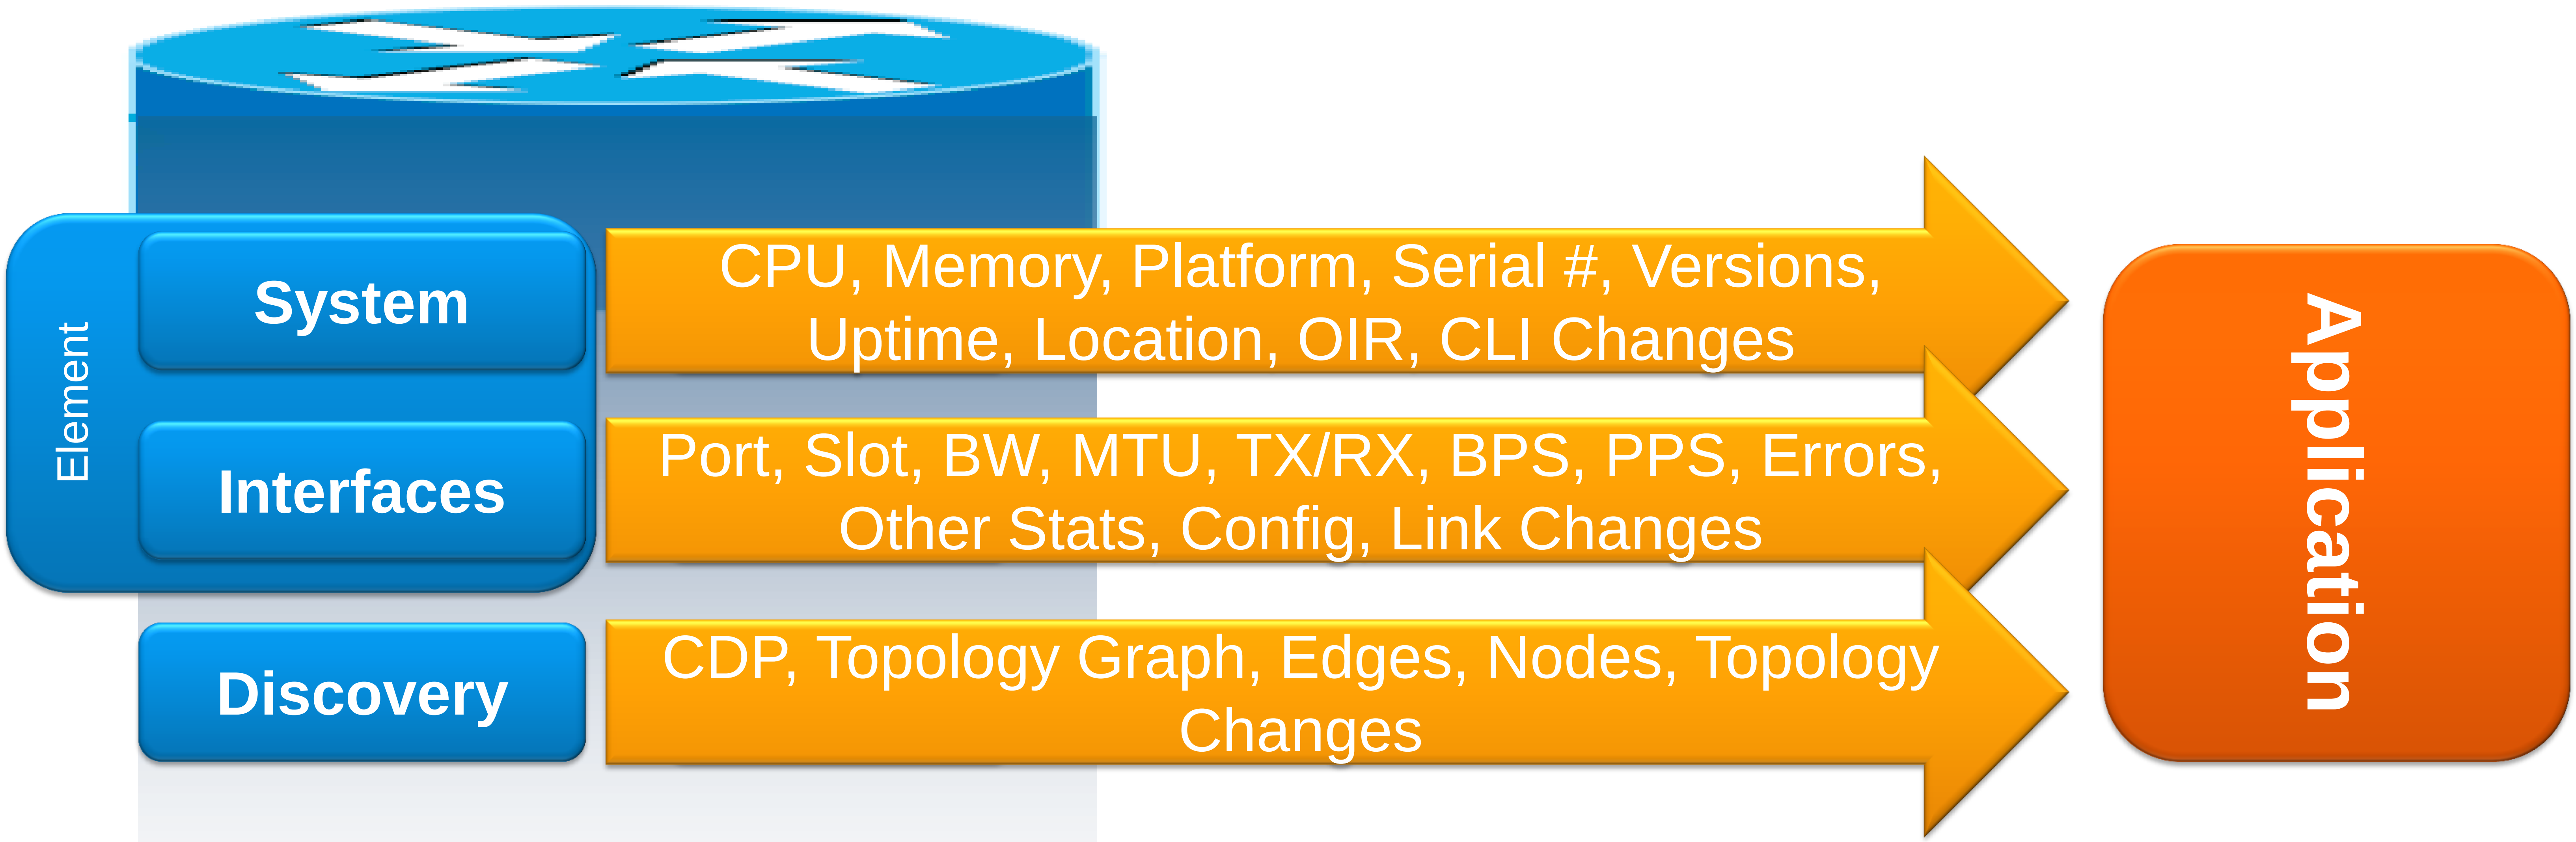

OnePK

Setting Properties and Statistics

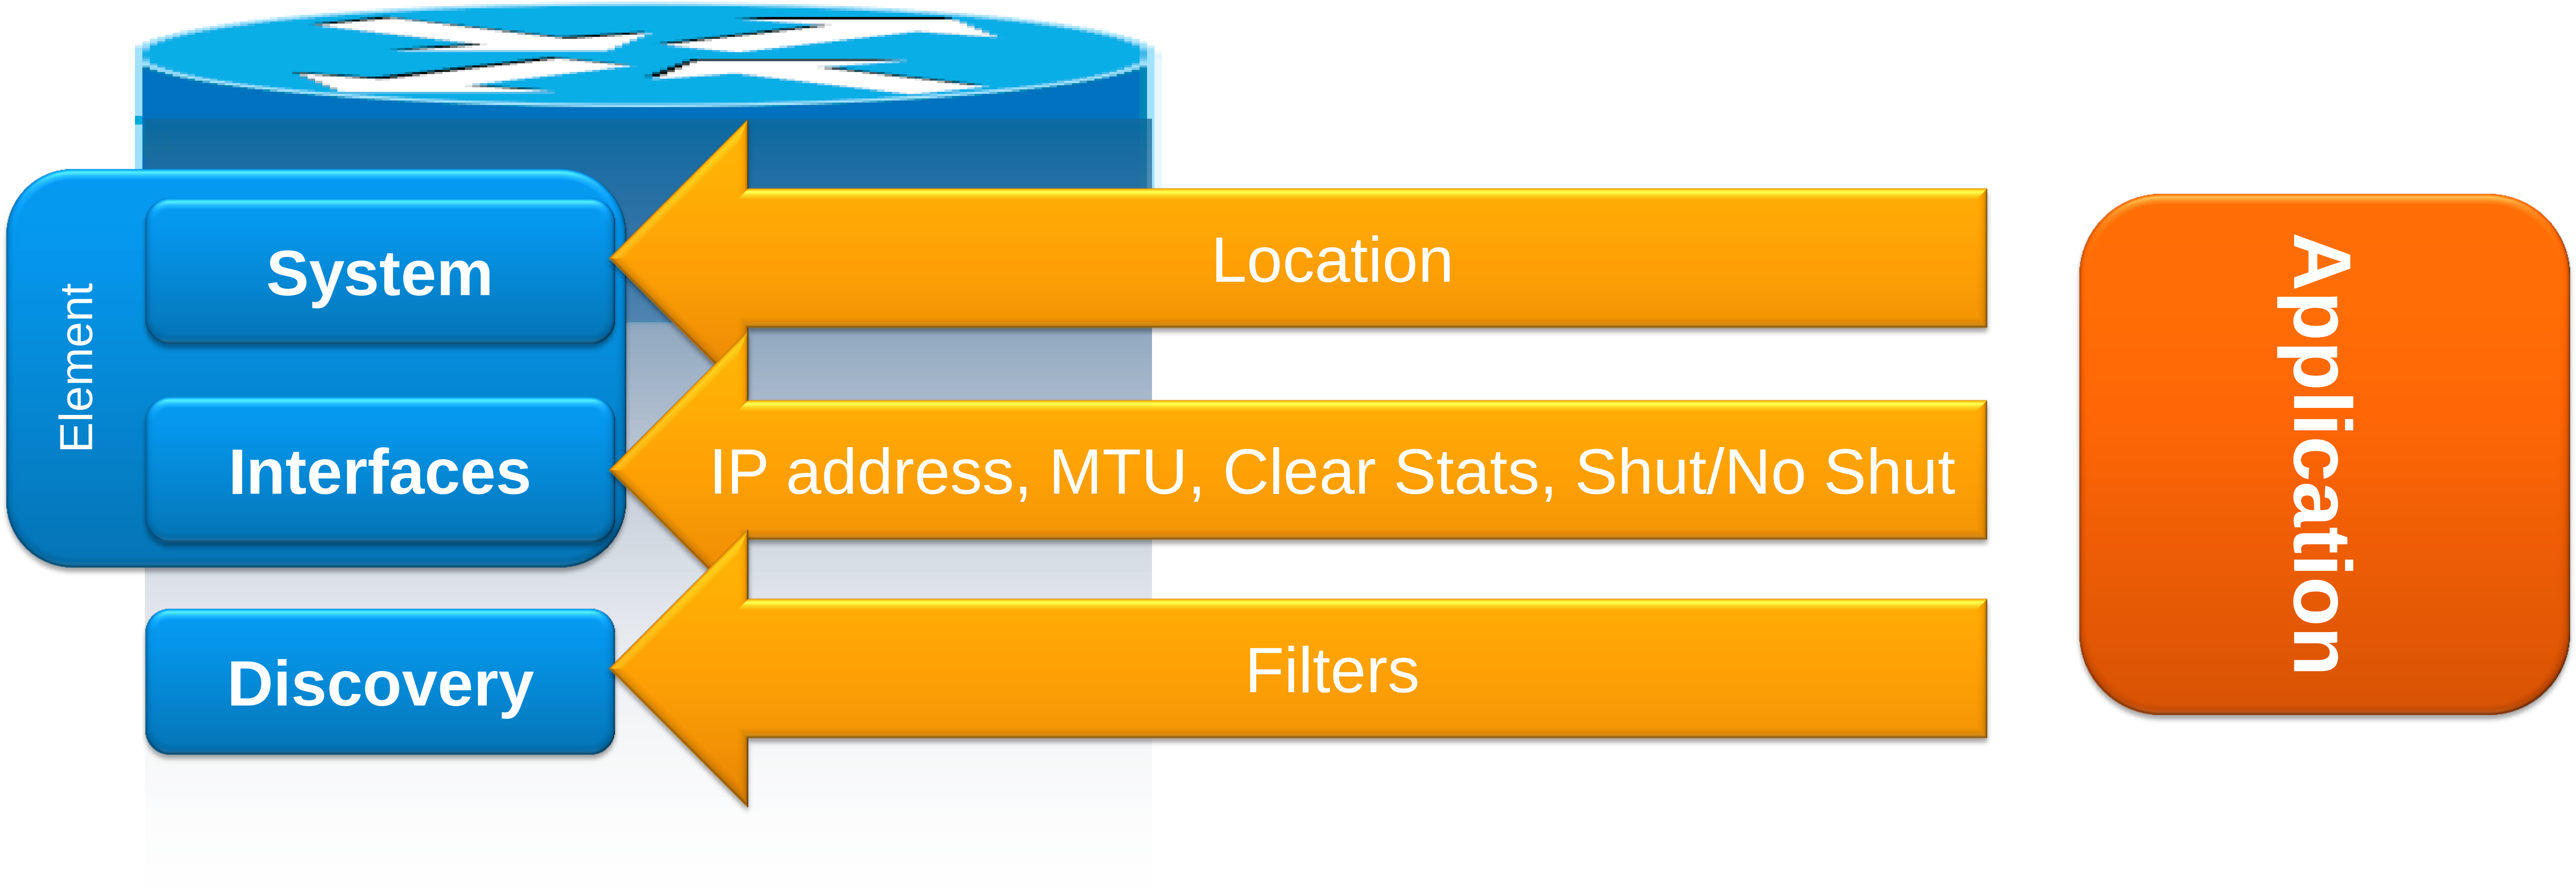

OnePK

How (C)

```
[scadora@localhost task103]$ bin/task103
Successful connection to network element

Element Info:
NetworkElement [ 172.20.165.44 ]
        Product ID      : ASR1001
        Processor       : 1RU
        Serial No       : SSI16050CJ5
        sysName         : ASR1K
        sysUpTime       : 546414
        sysDescr        : Cisco IOS Software, IOS-XE Software (X86_64_LINUX_IOSD-UN
IVERSAL-M), Experimental Version 15.3(20120510:014633) [mcp_dev-BLD-BLD_MCP_DEV_
LATEST_20120510_002552-ios 157]
Copyright (c) 1986-2012 by Cisco Systems, Inc.
Compiled Wed 09-May-12 21:44 by mcpre
```

```
        }
}
```

Example – Properties and Statistics

MTU Management

- Problem: Misconfigurations cause network outages, degrade performance, impact SLAs.

1. Network begins with mismatched parameters on either side of link (e.g. MTU)
2. Application checks parameters on either side and identifies mismatches (red lines)
3. Application sets parameters to match (lines turn green)
4. Application registers for events related to parameters change.
5. Users logs into console and manually changes parameter. Topology indicates change.

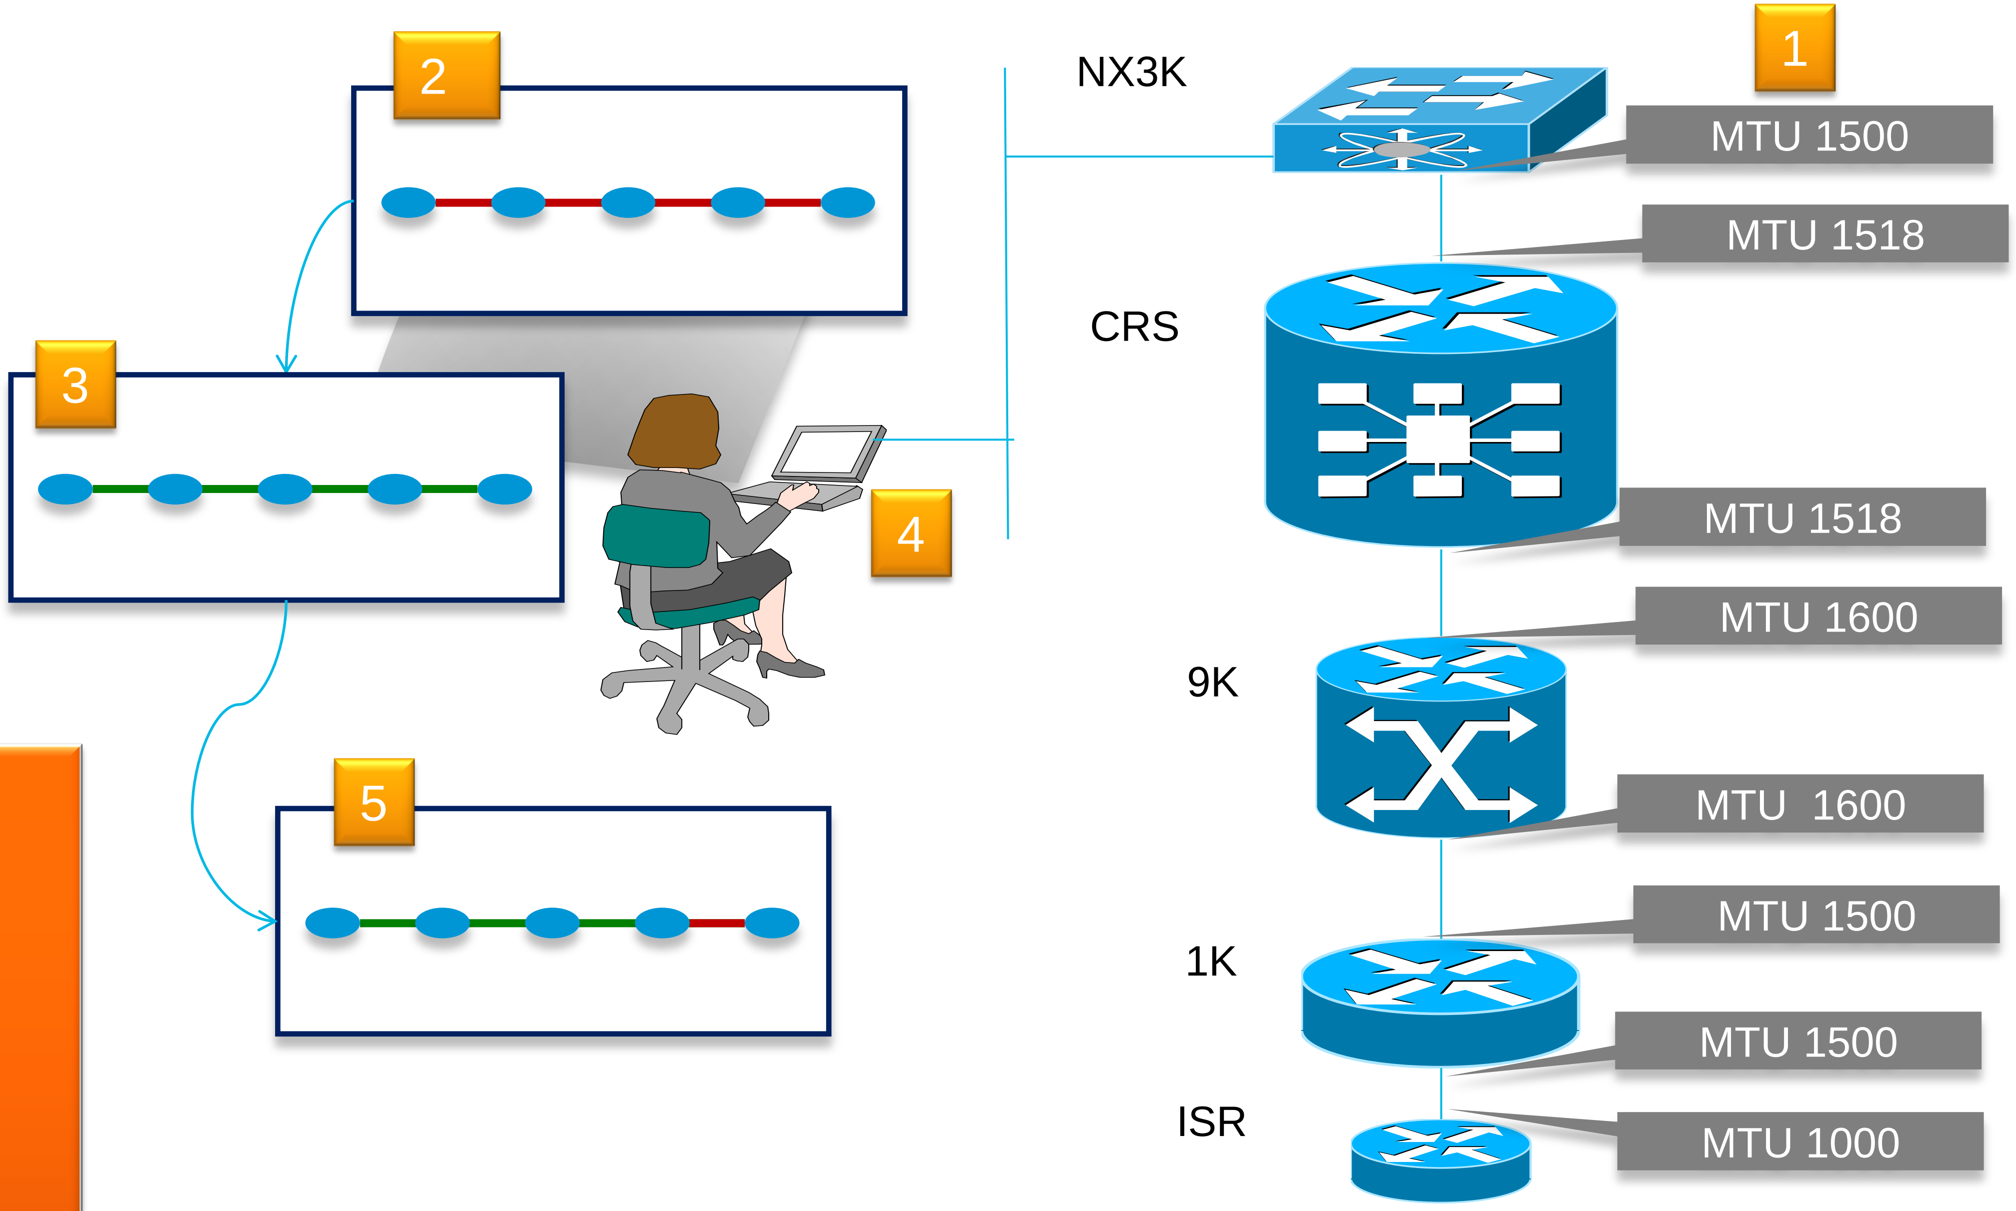

Example – Properties and Statistics

MTU Management Application Output

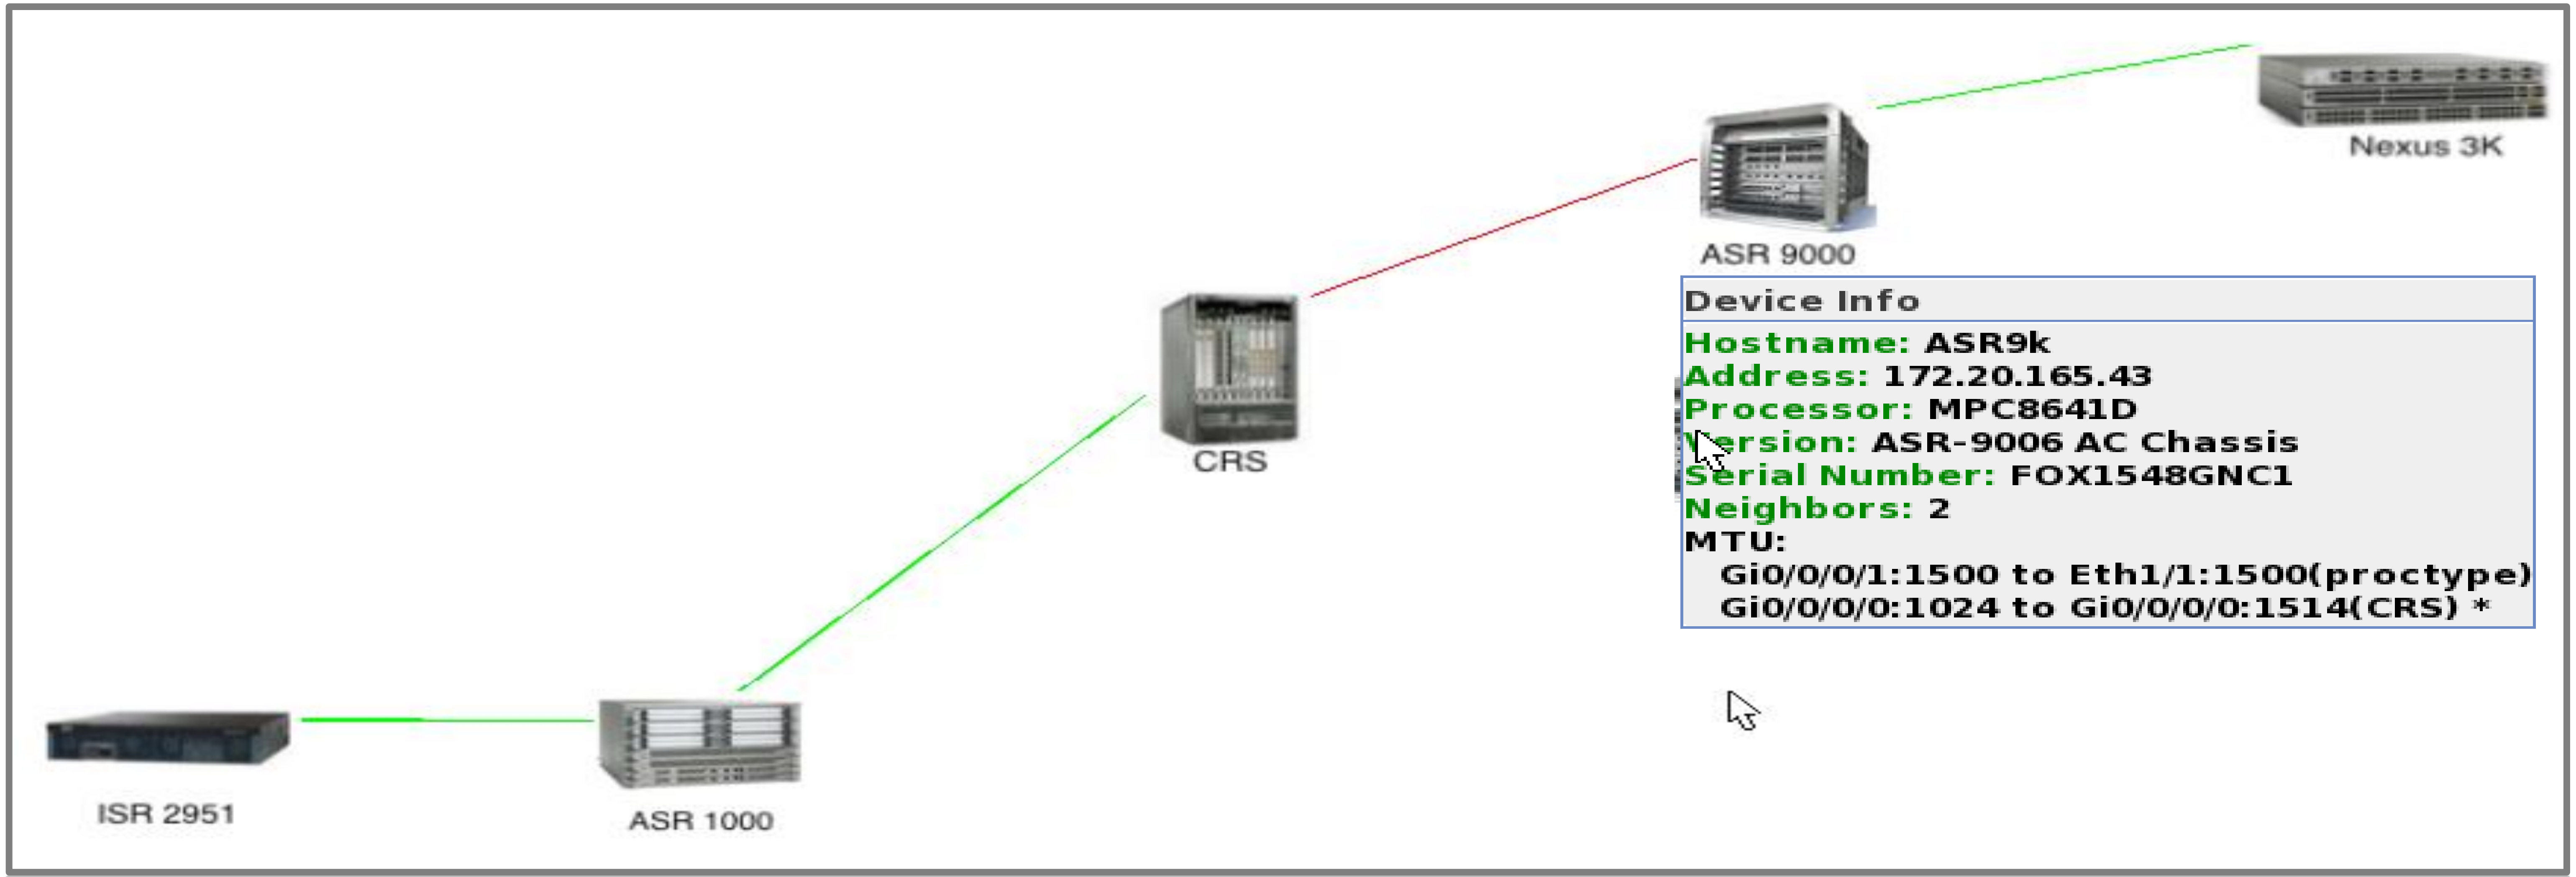

OnePK

Getting Policies and Routes

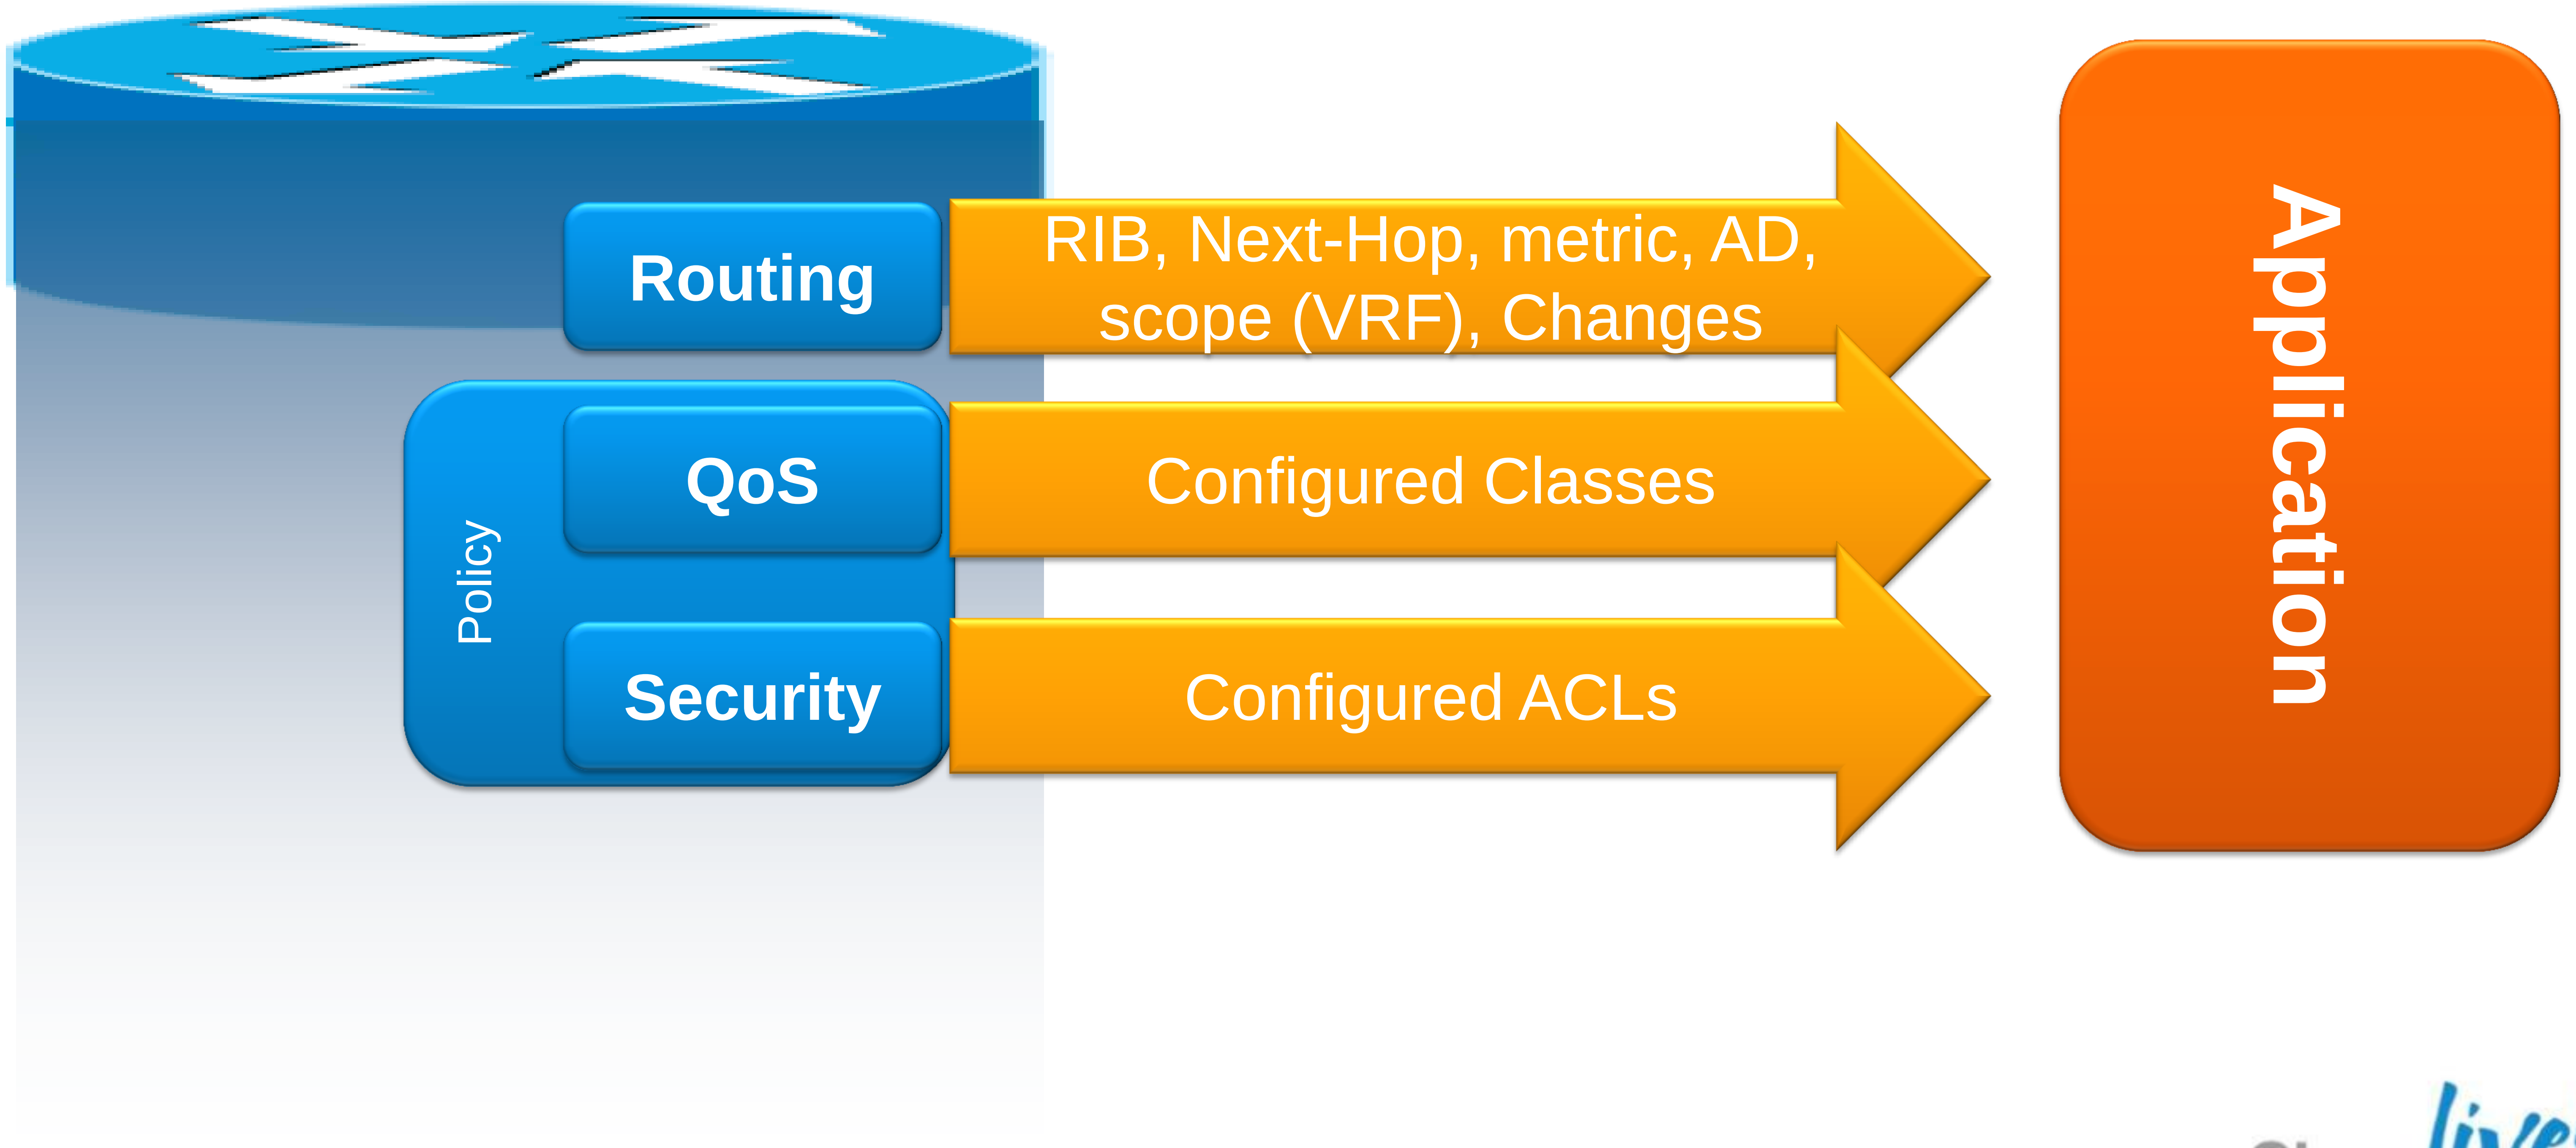

OnePK

Setting Policies and Routes

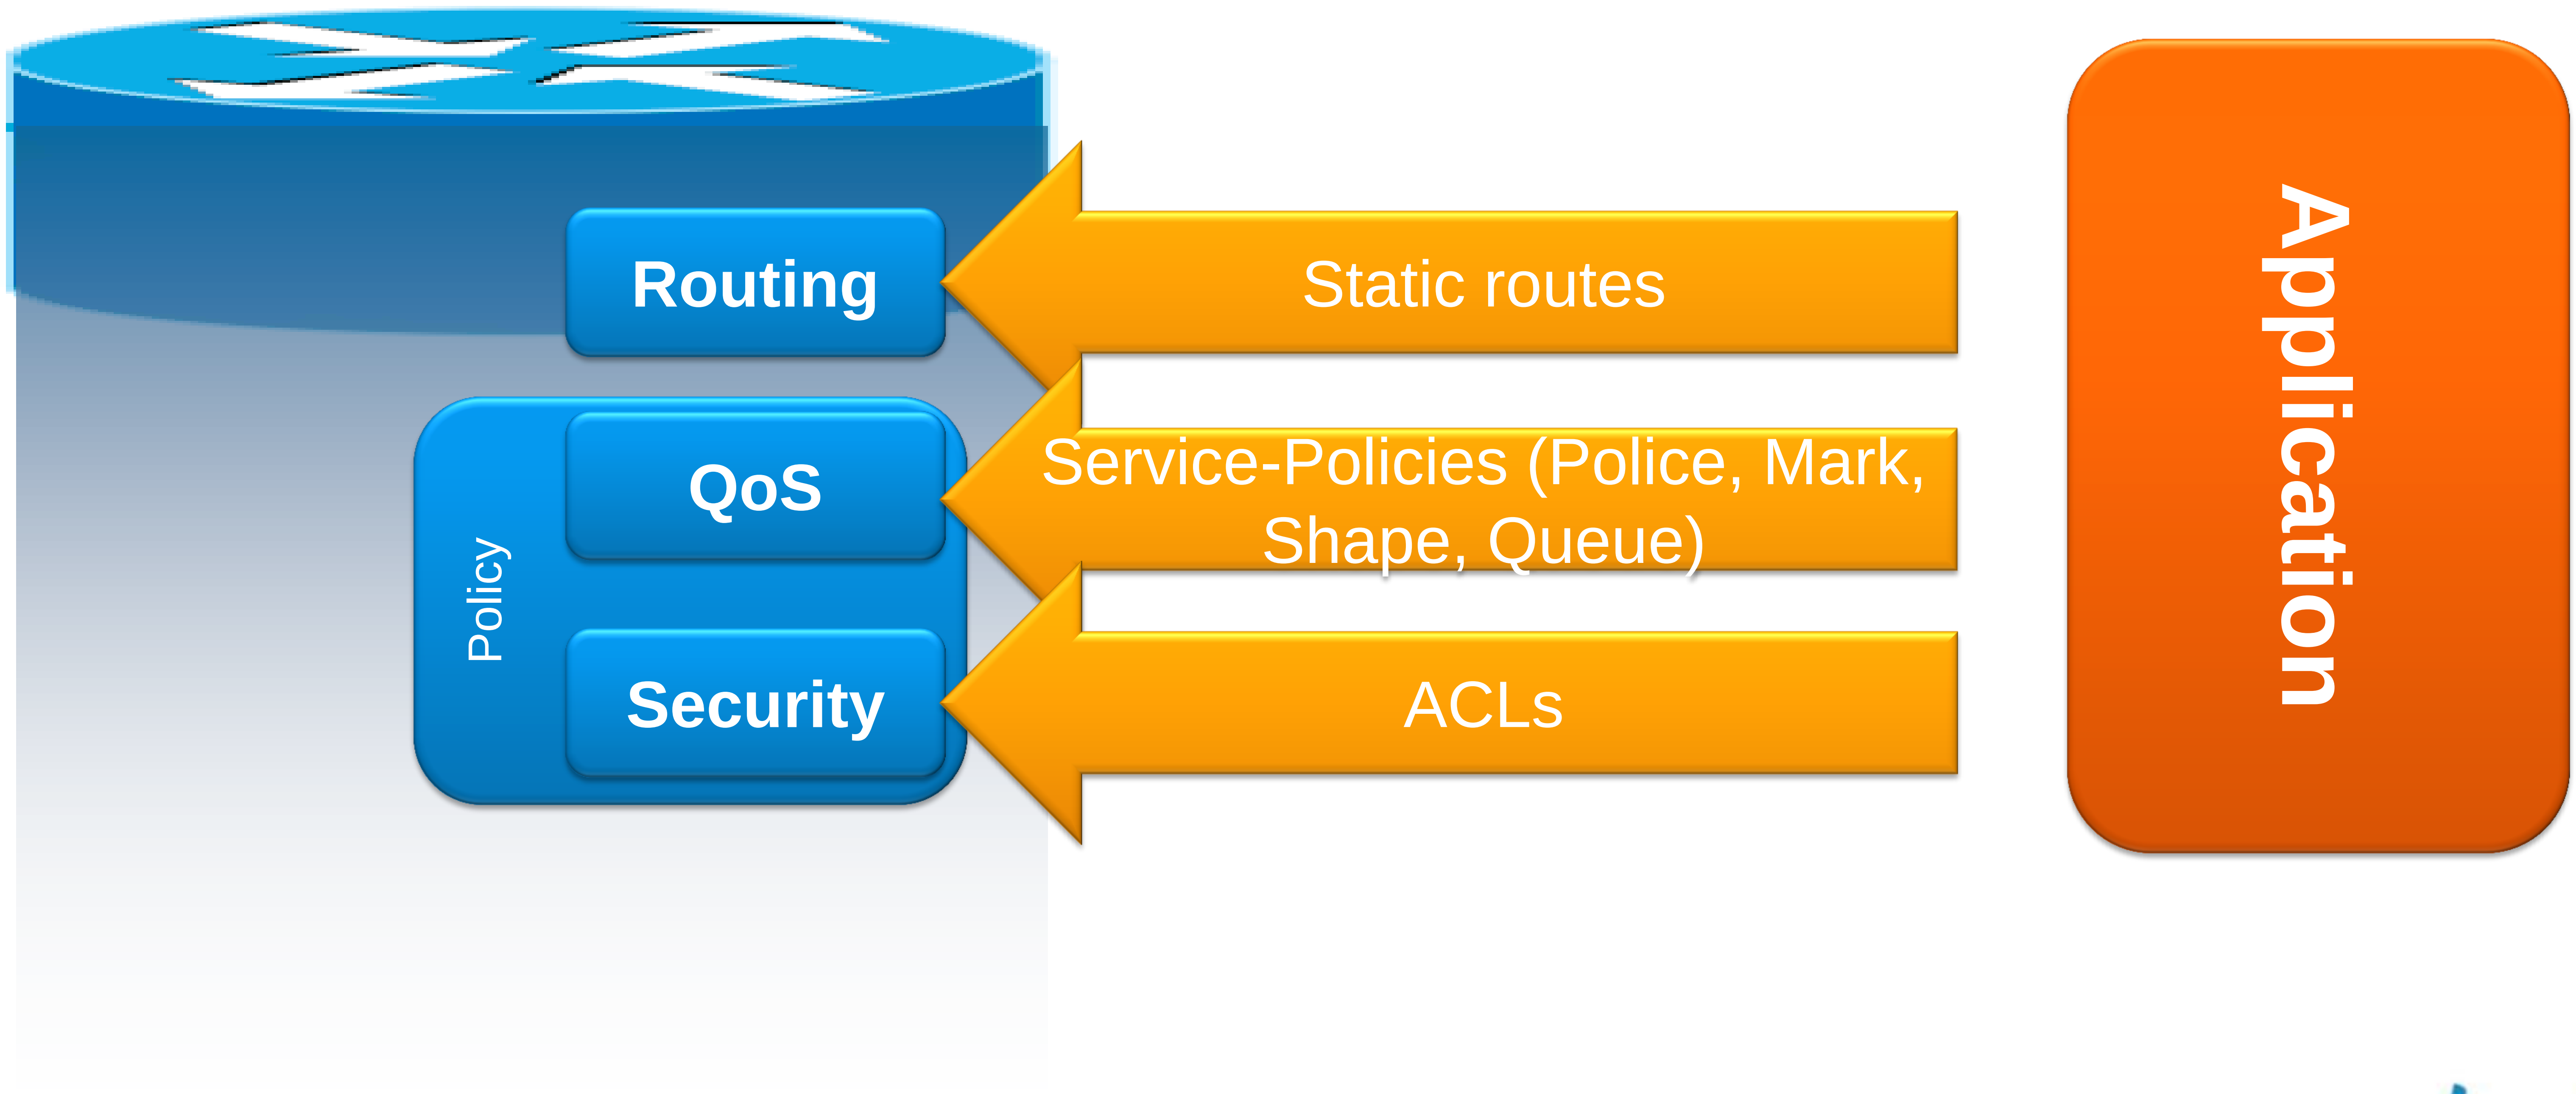

OnePK

How (java)

```
L3UnicastScope scope = new L3UnicastScope("", AFIType.IPV4, SAFIType.UNICAST, "");
NetworkPrefix prefix = new NetworkPrefix(InetAddress.getByName("0.0.0.0"), 0);
L3UnicastRIBFilter ribFilter = new L3UnicastRIBFilter(OwnerType.NONE, "NONE", prefix);
L3UnicastRouteRange range = new L3UnicastRouteRange(prefix, RouteRange.RangeType.EQUAL_OR_LARGER, 100);
List<TopoNode> mynodes = TopoNode.getAllNodes();
for(TopoNode thisnode : mynodes) {
    Routing routing = Routing.getInstance(thisnode.ne);
    RIB rib = routing.getRib();
    List<Route> routeList = rib.getRouteList(scope, ribFilter, range);
    for (Route route : routeList) {
```

Get Routes

```
L3UnicastRoute aRoute = new L3UnicastRoute(prefix, nextHopList);
aRoute.setAdminDistance(1);
RouteOperation op = new L3UnicastRouteOperation(RouteOperationType.ADD, aRoute);
List<RouteOperation> opList = new ArrayList<RouteOperation>();
opList.add(op);
AppRouteTable art = routing.getAppRouteTable();
art.updateRoutes(scope, opList);
```

Set Routes

Example – Policies and Routes

Custom Routing Algorithm

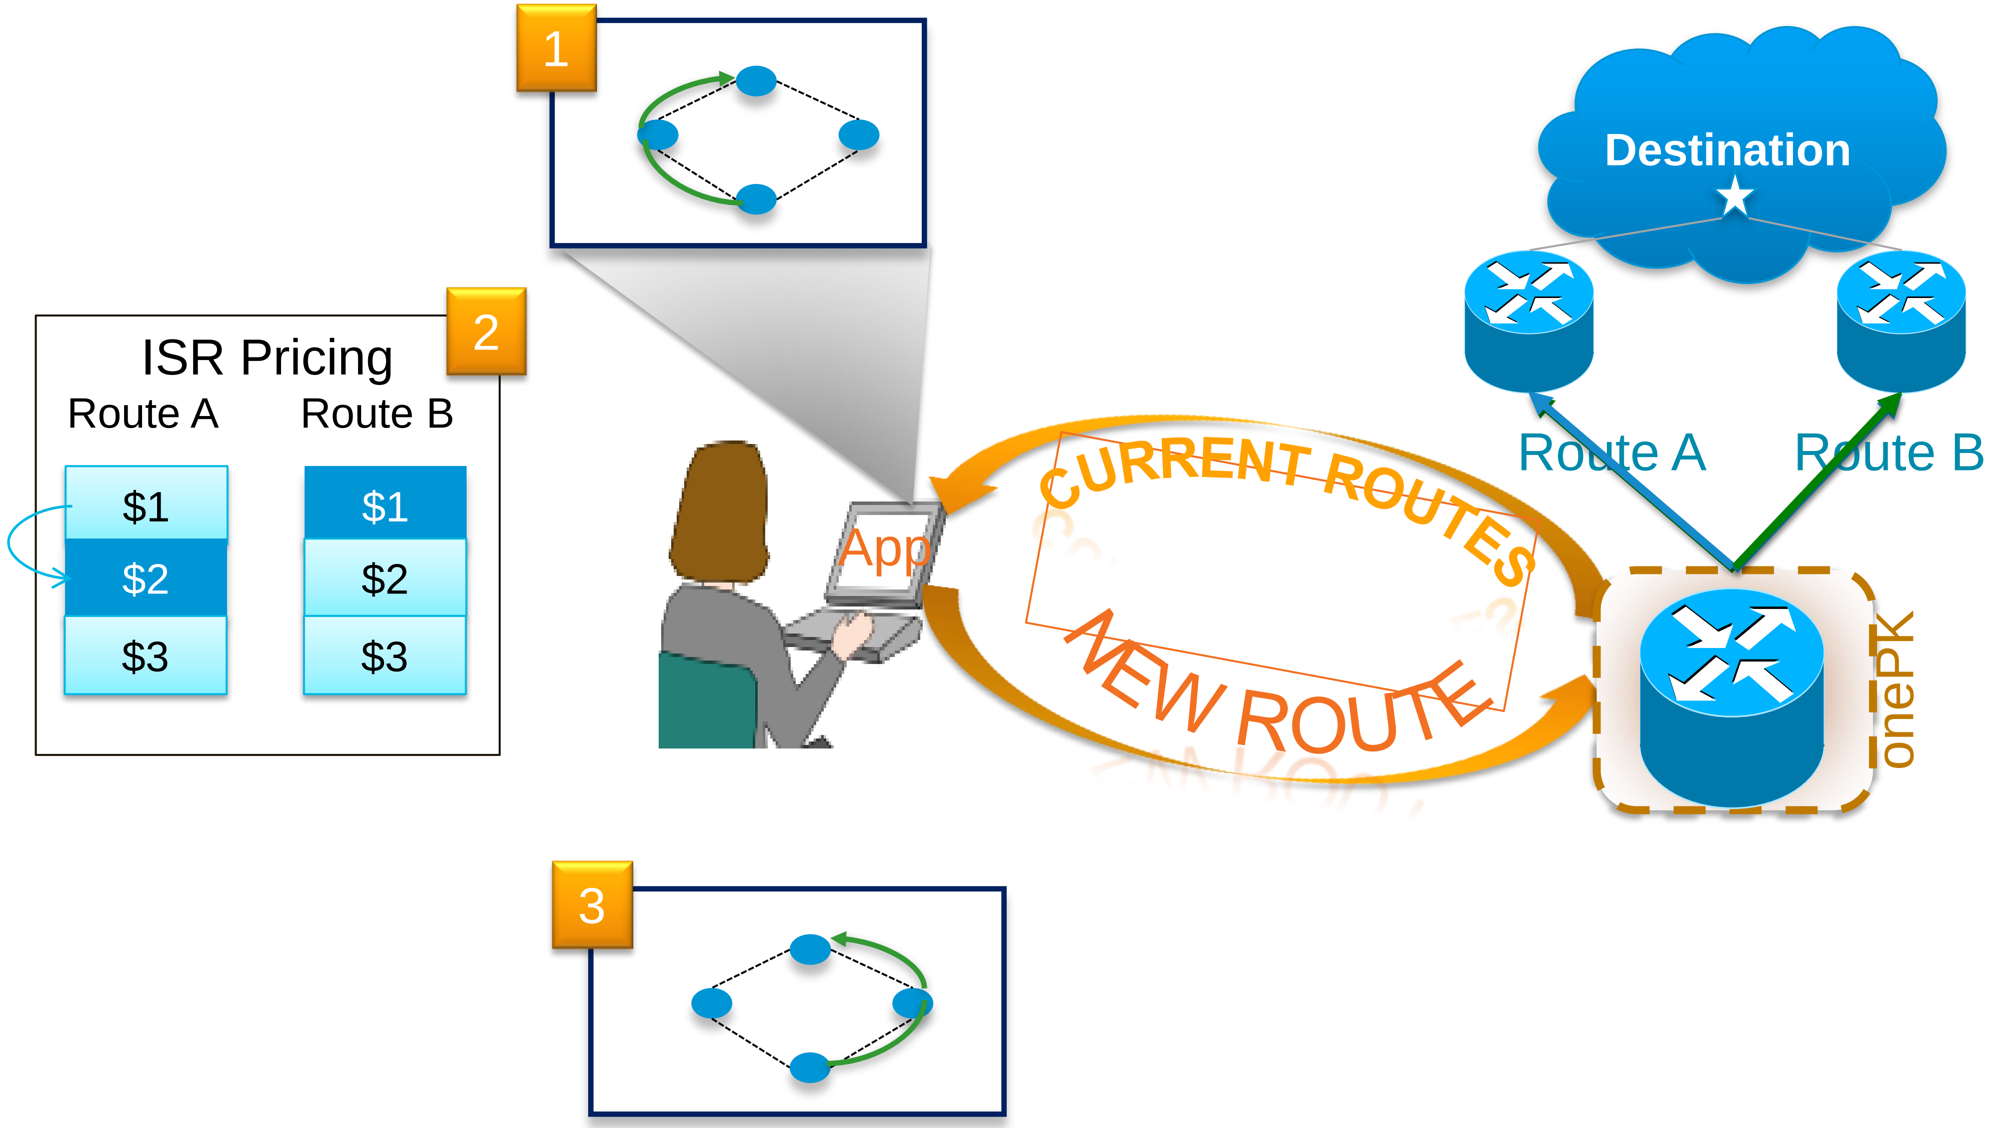

Unique Data Forwarding Algorithm Highly Optimised for the Network Operator's Application

Example – Policies and Routes

Custom Routing Algorithm Output - Default EIGRP Routing

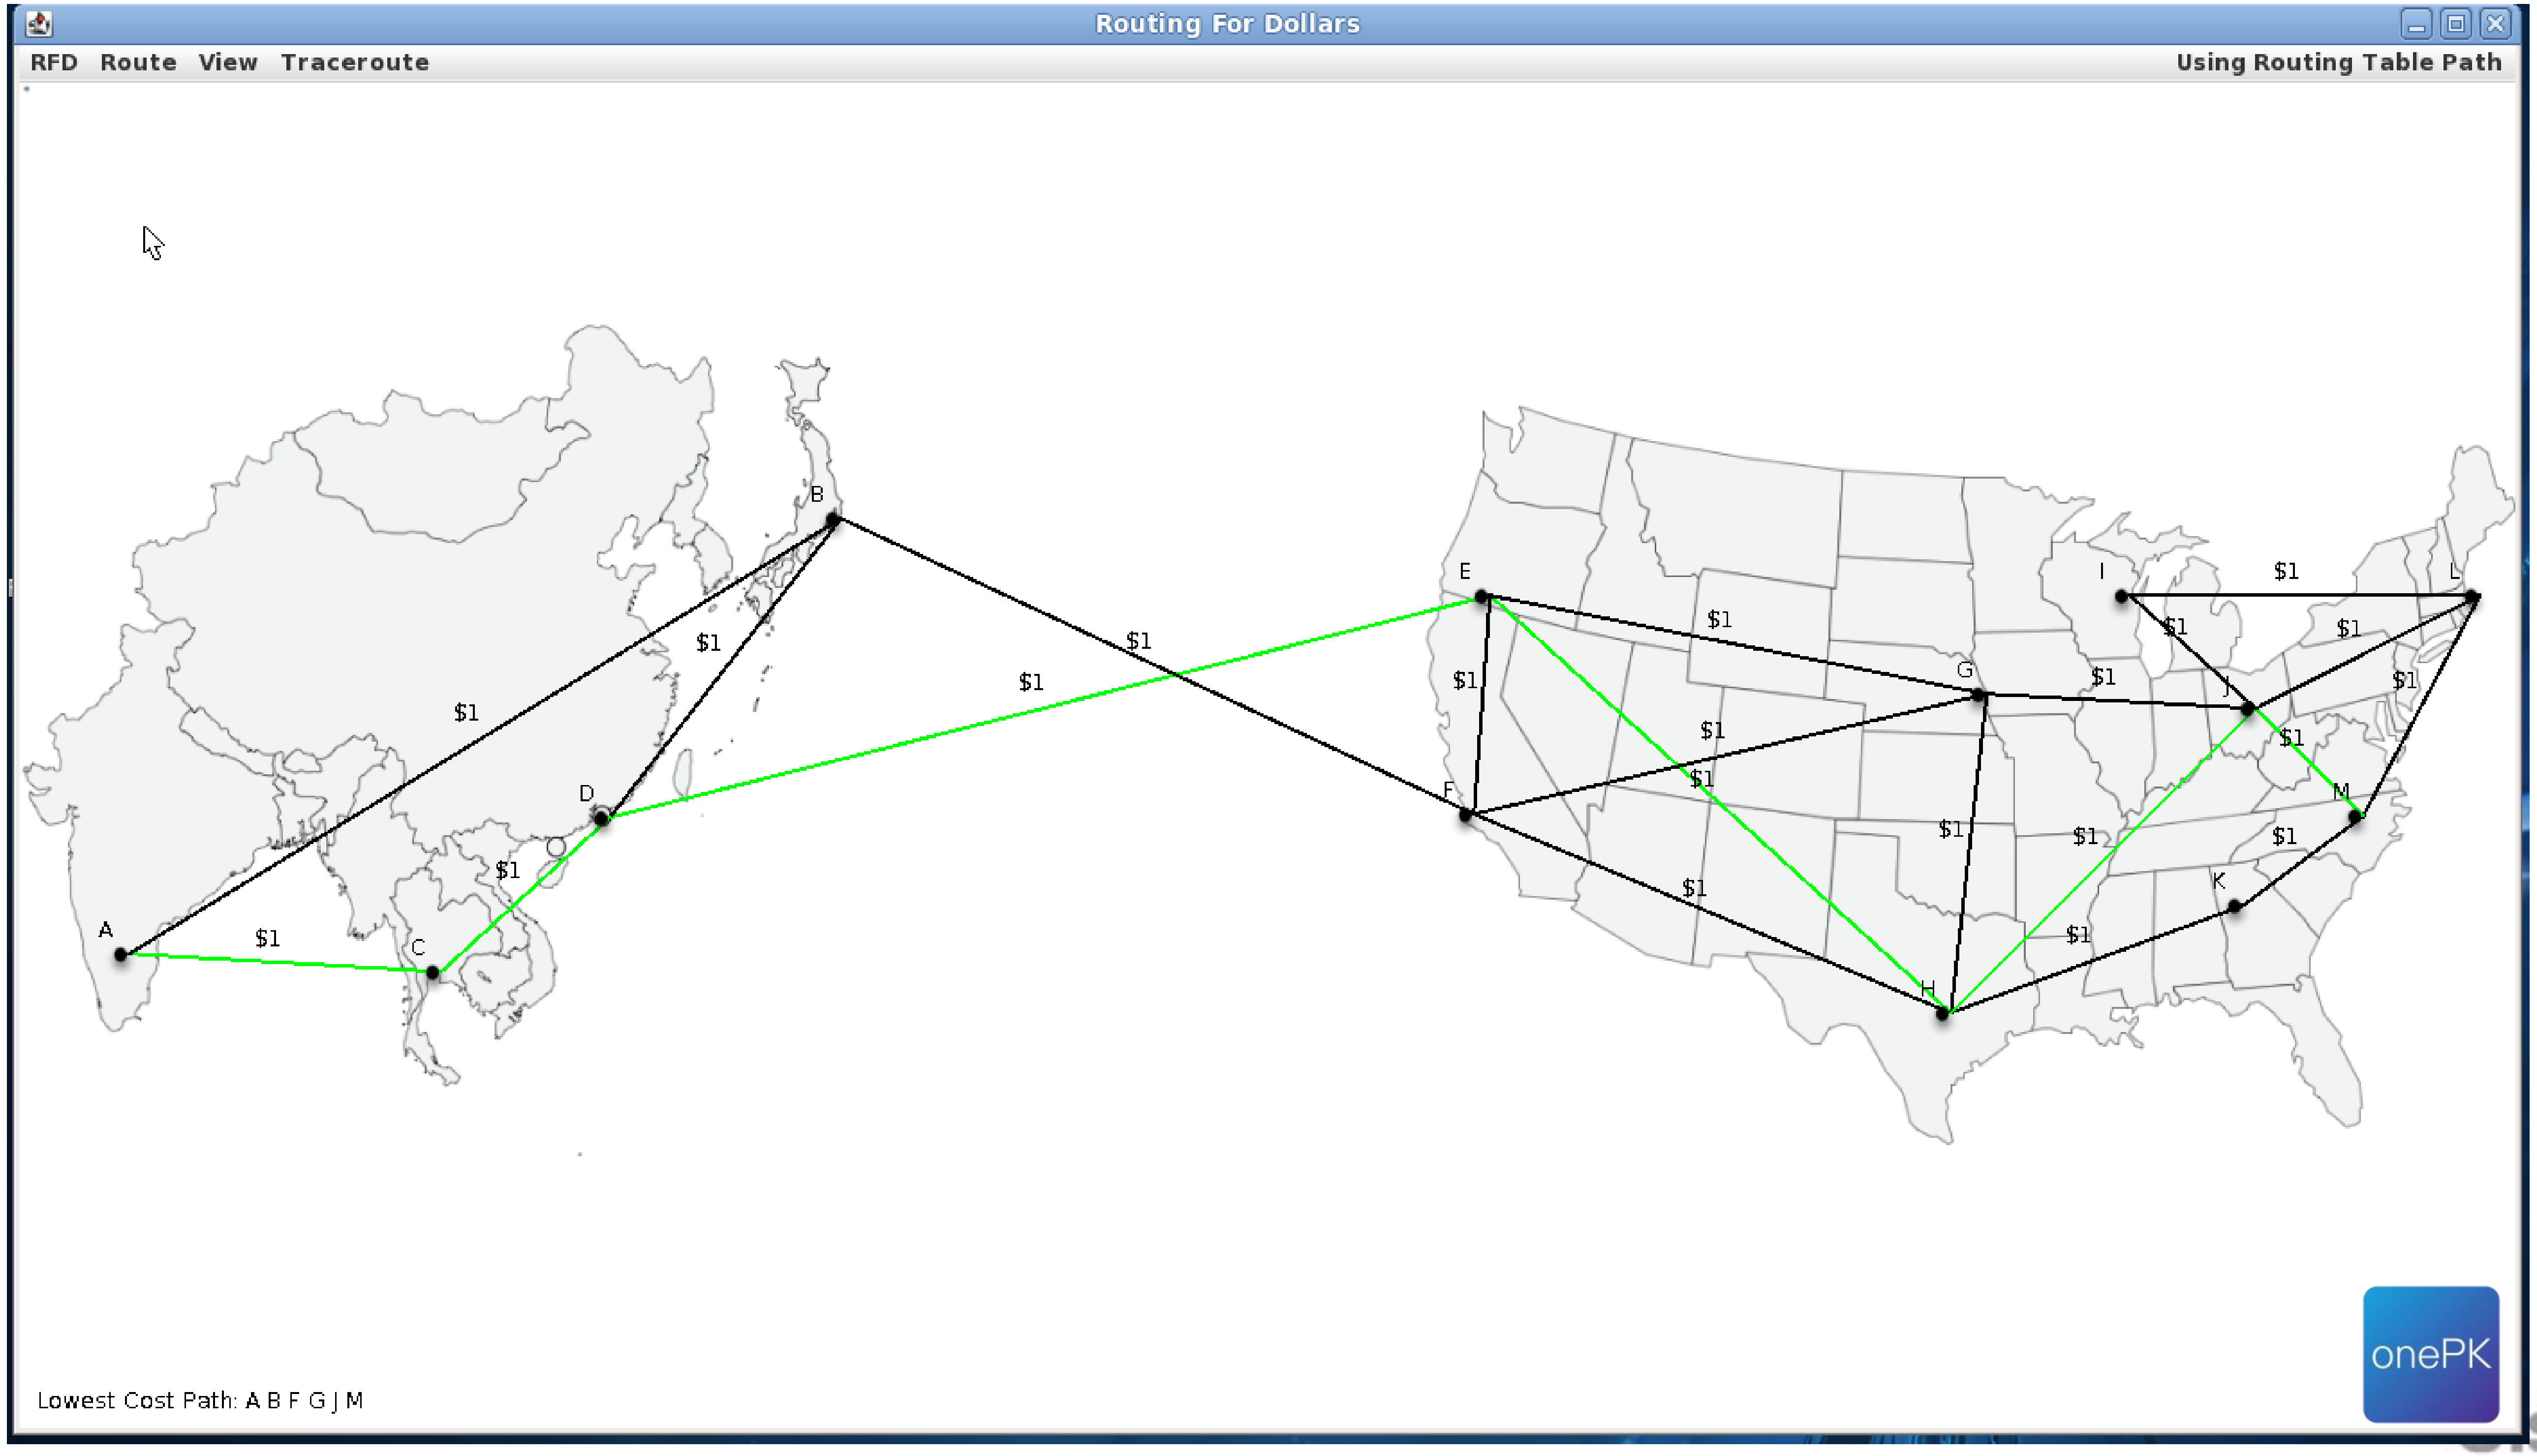

Example – Policies and Routes

Custom Routing Algorithm Output - OnePK Application Routes Applied

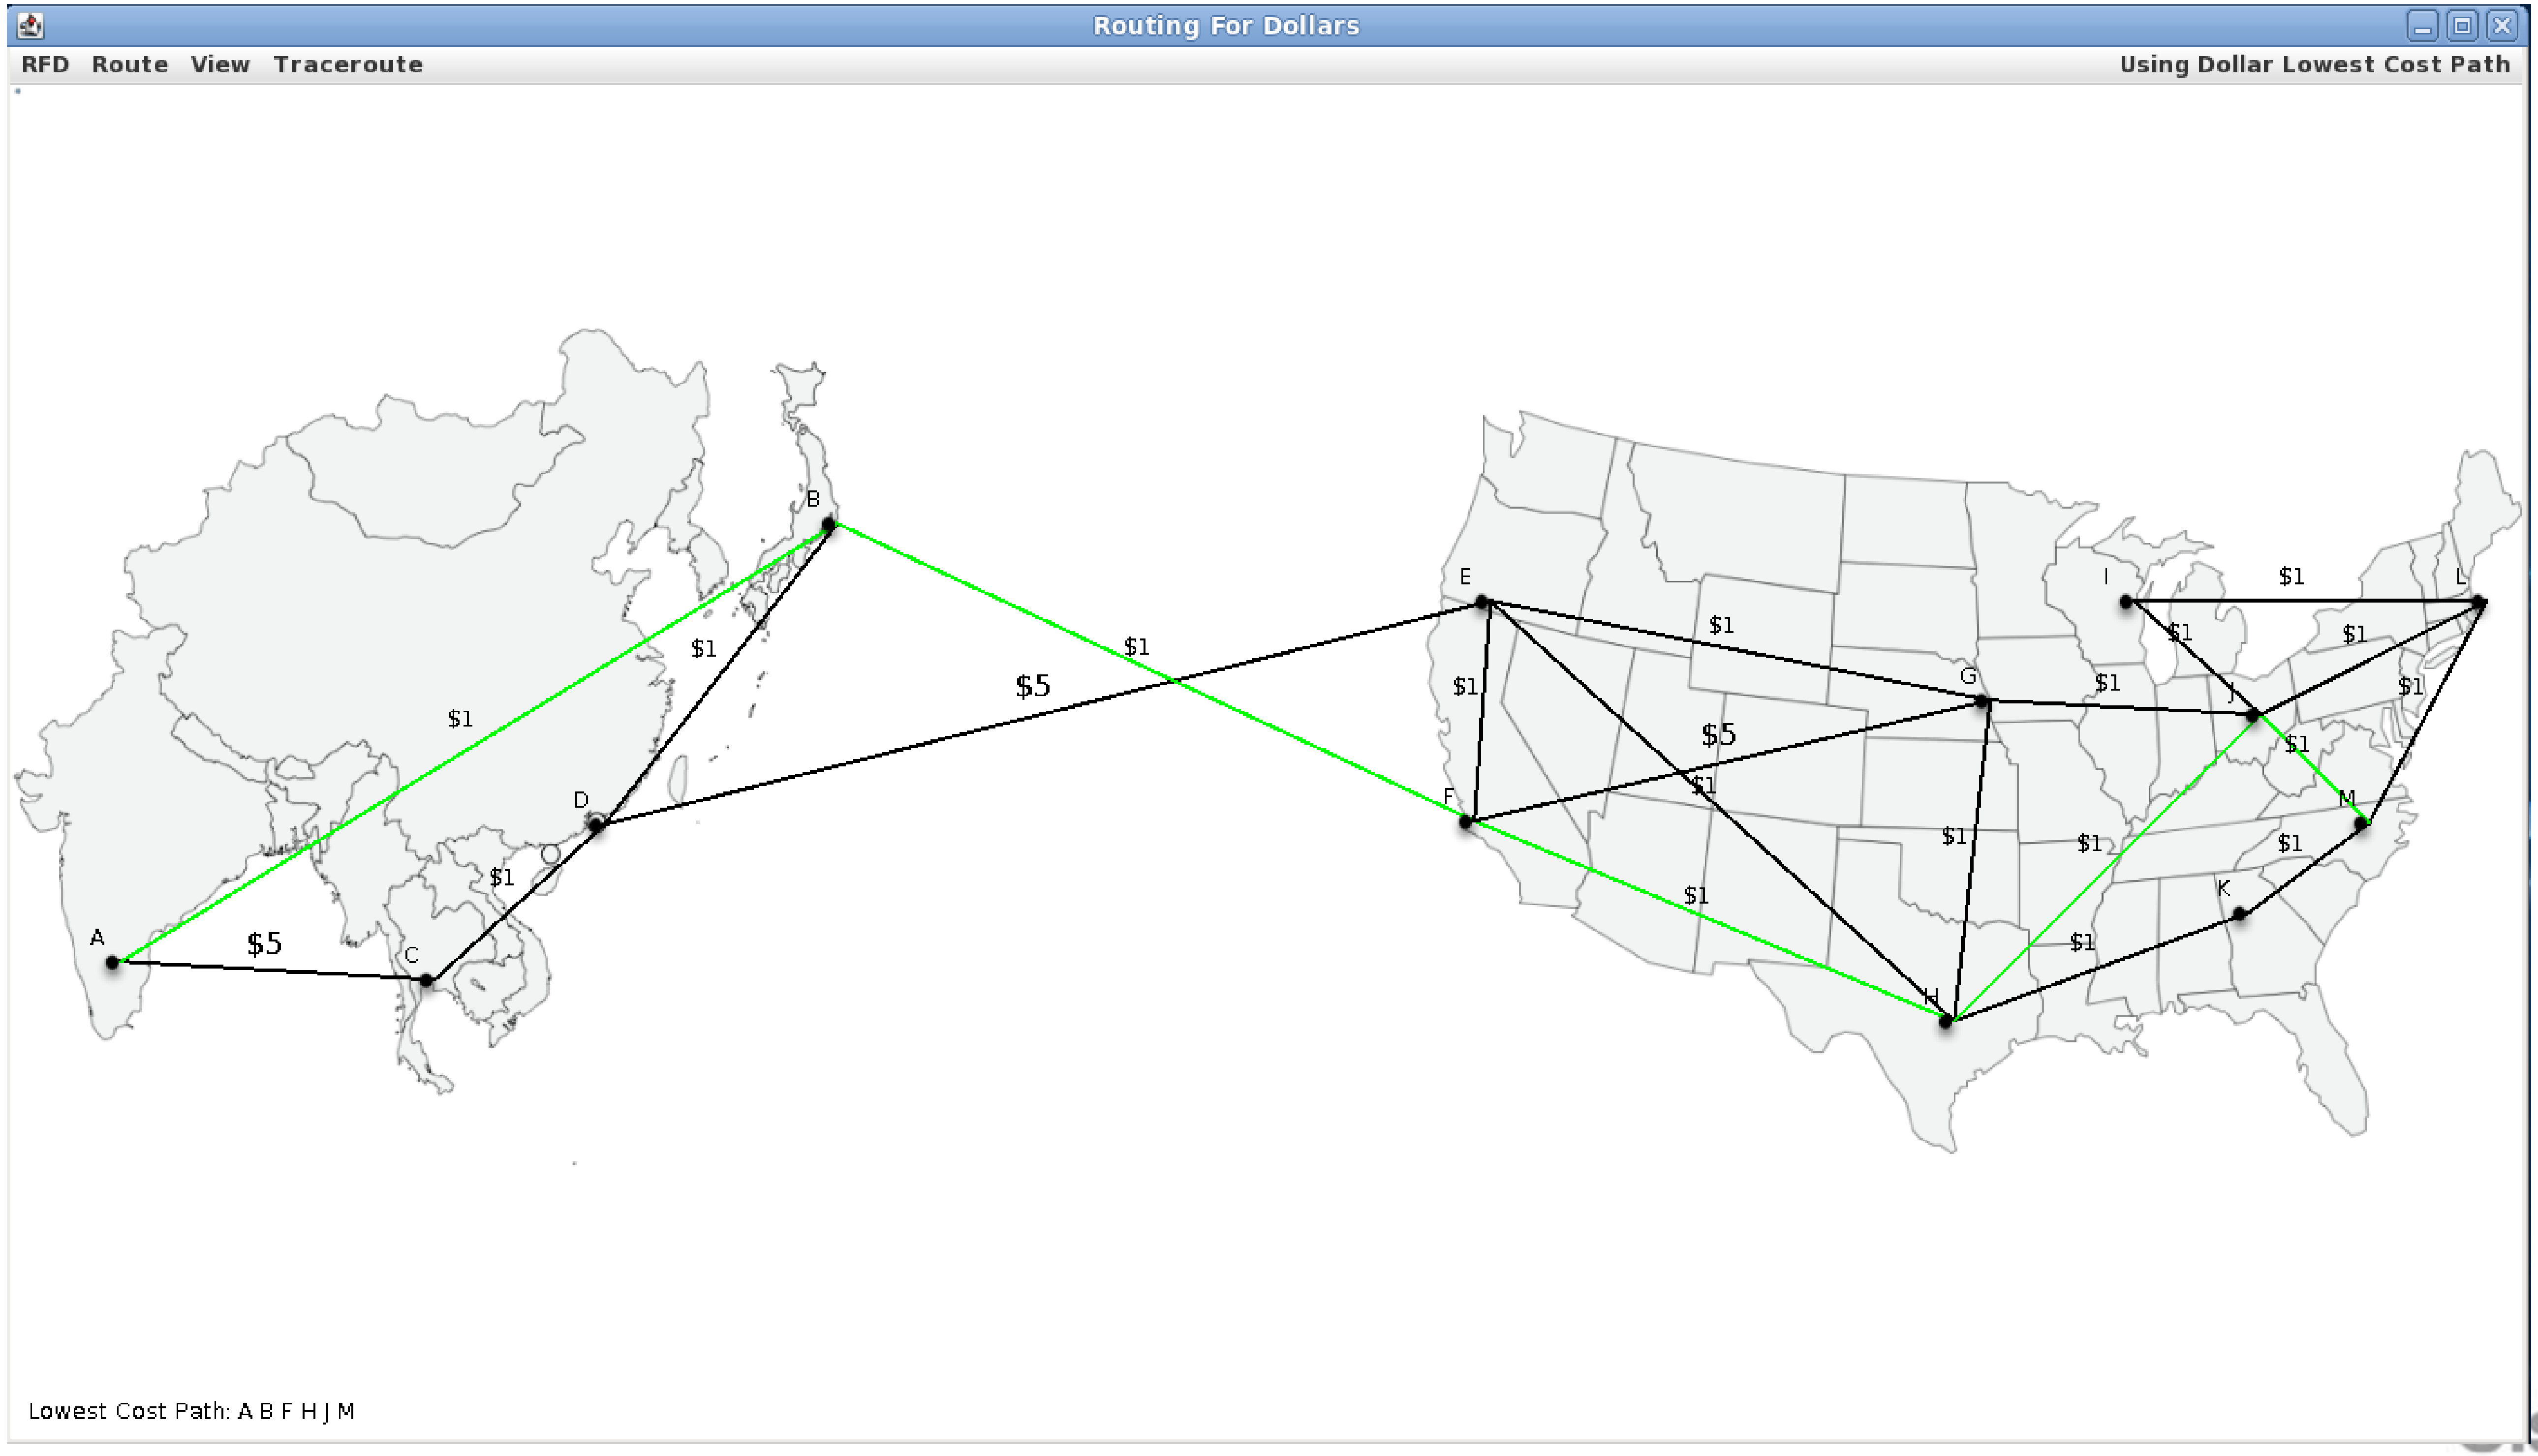

Example – Policies and Routes

Custom Routing Algorithm Output - OnePK Application Routes Applied

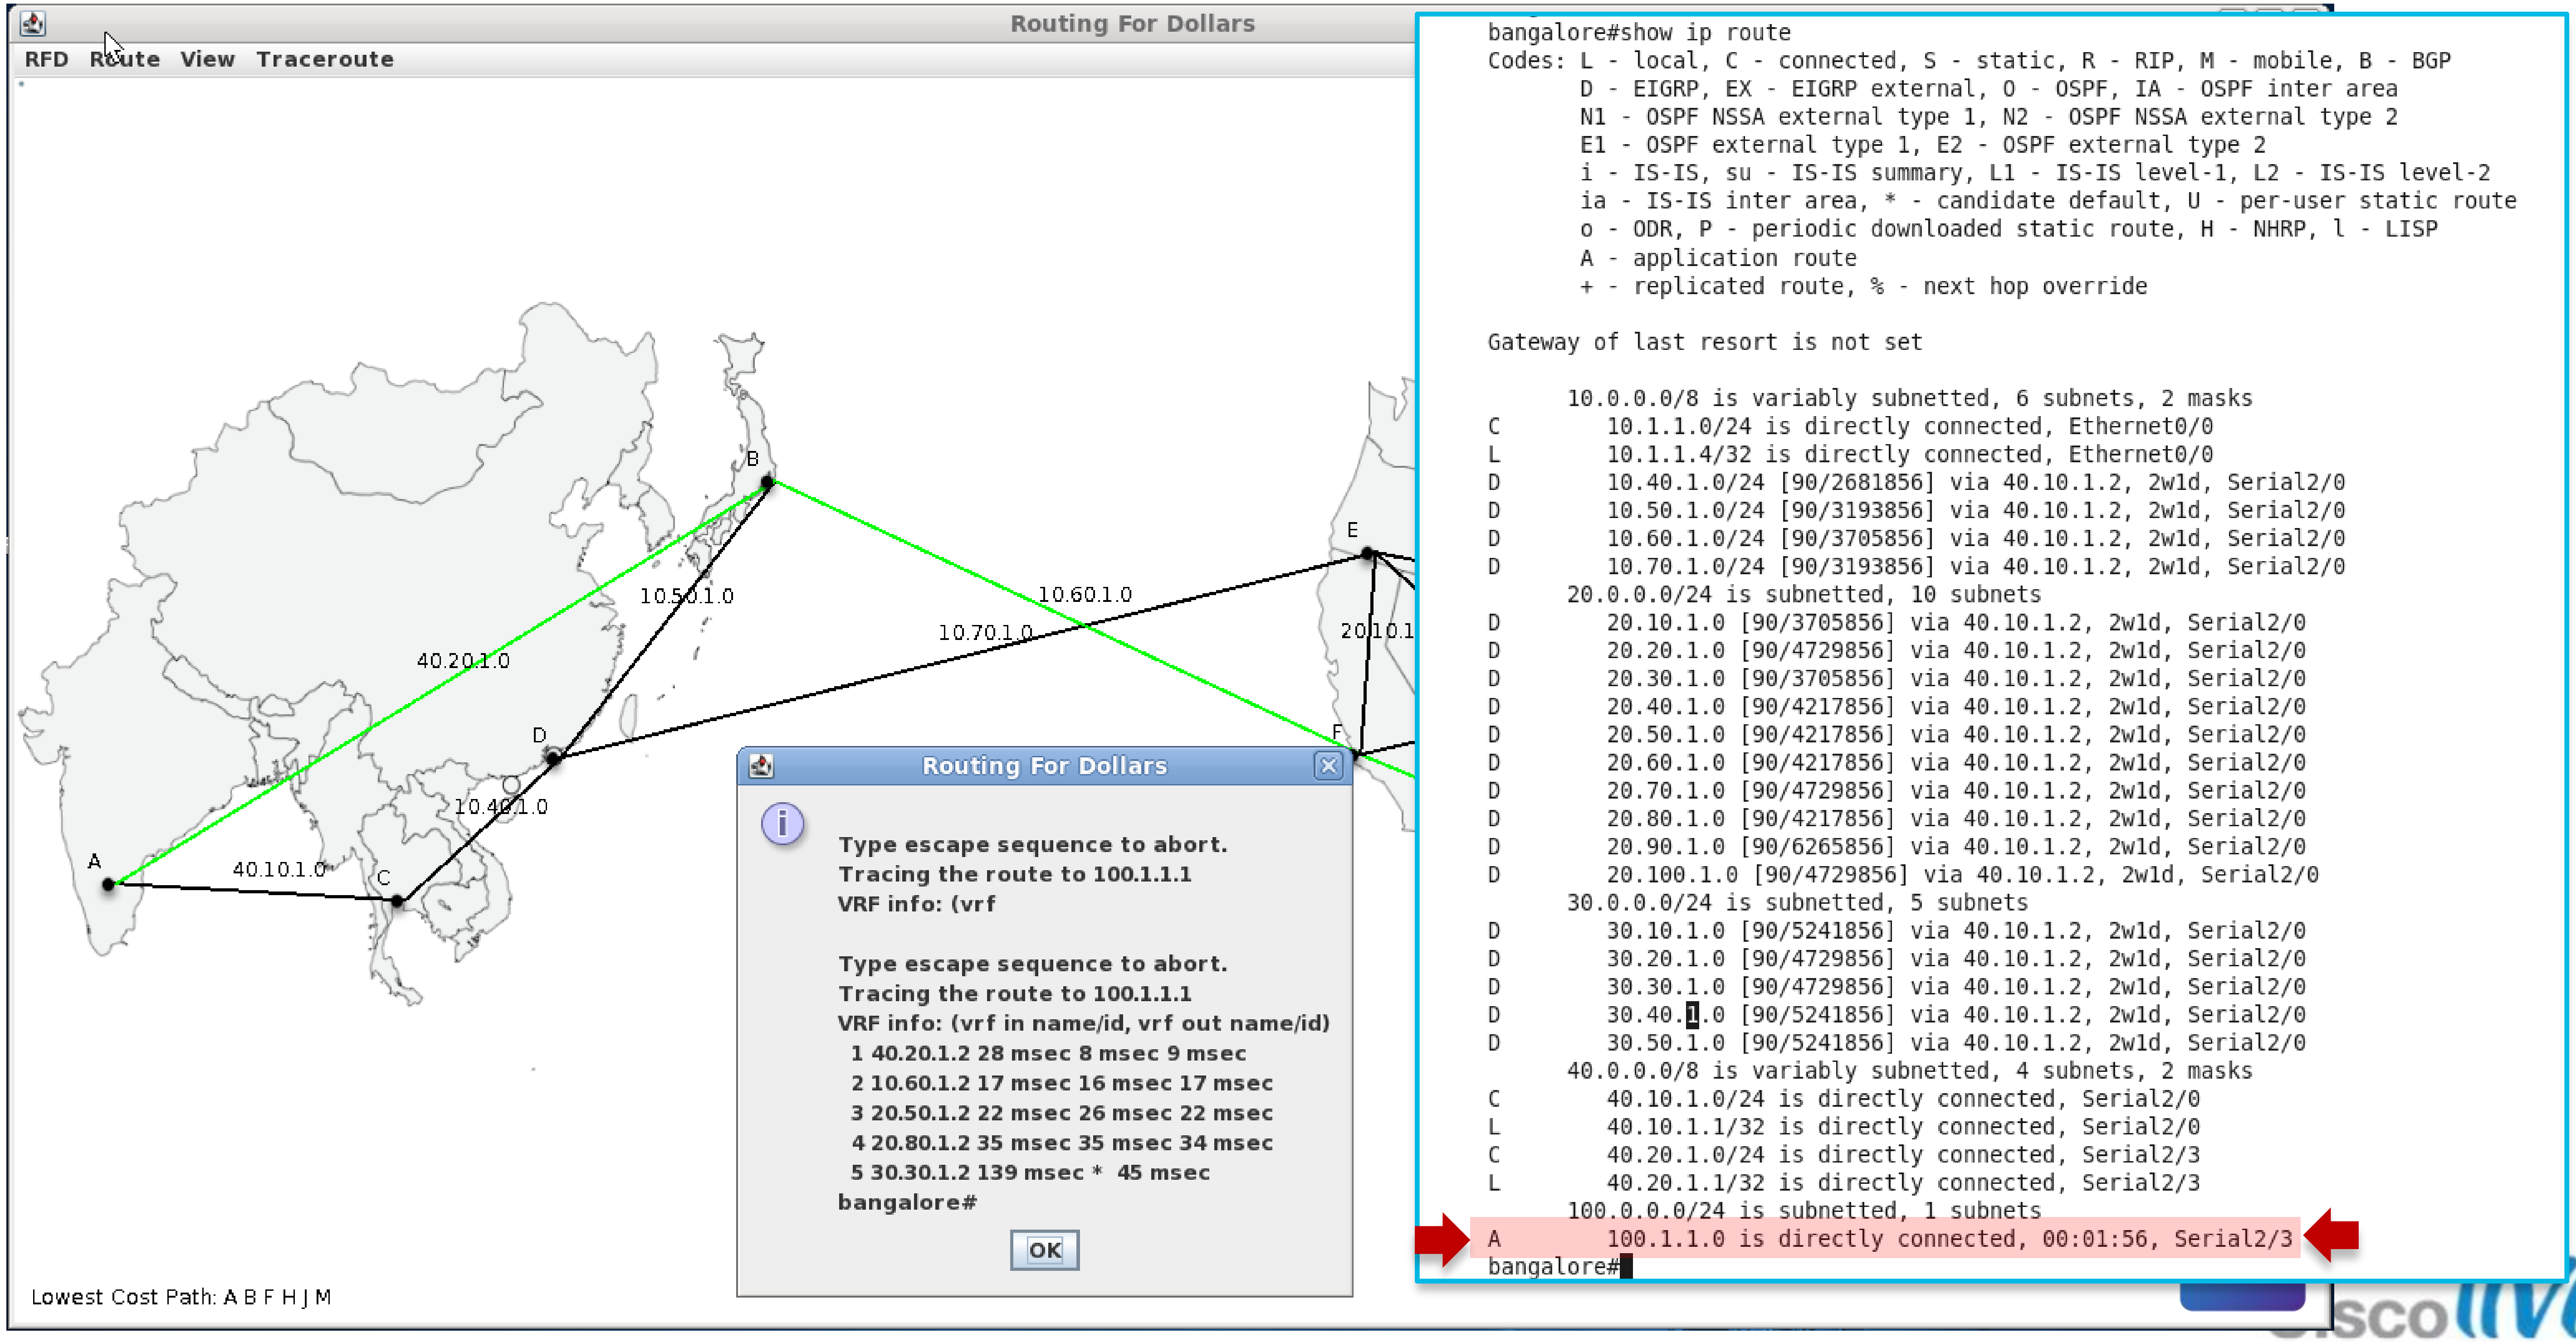

OnePK
Getting Packets

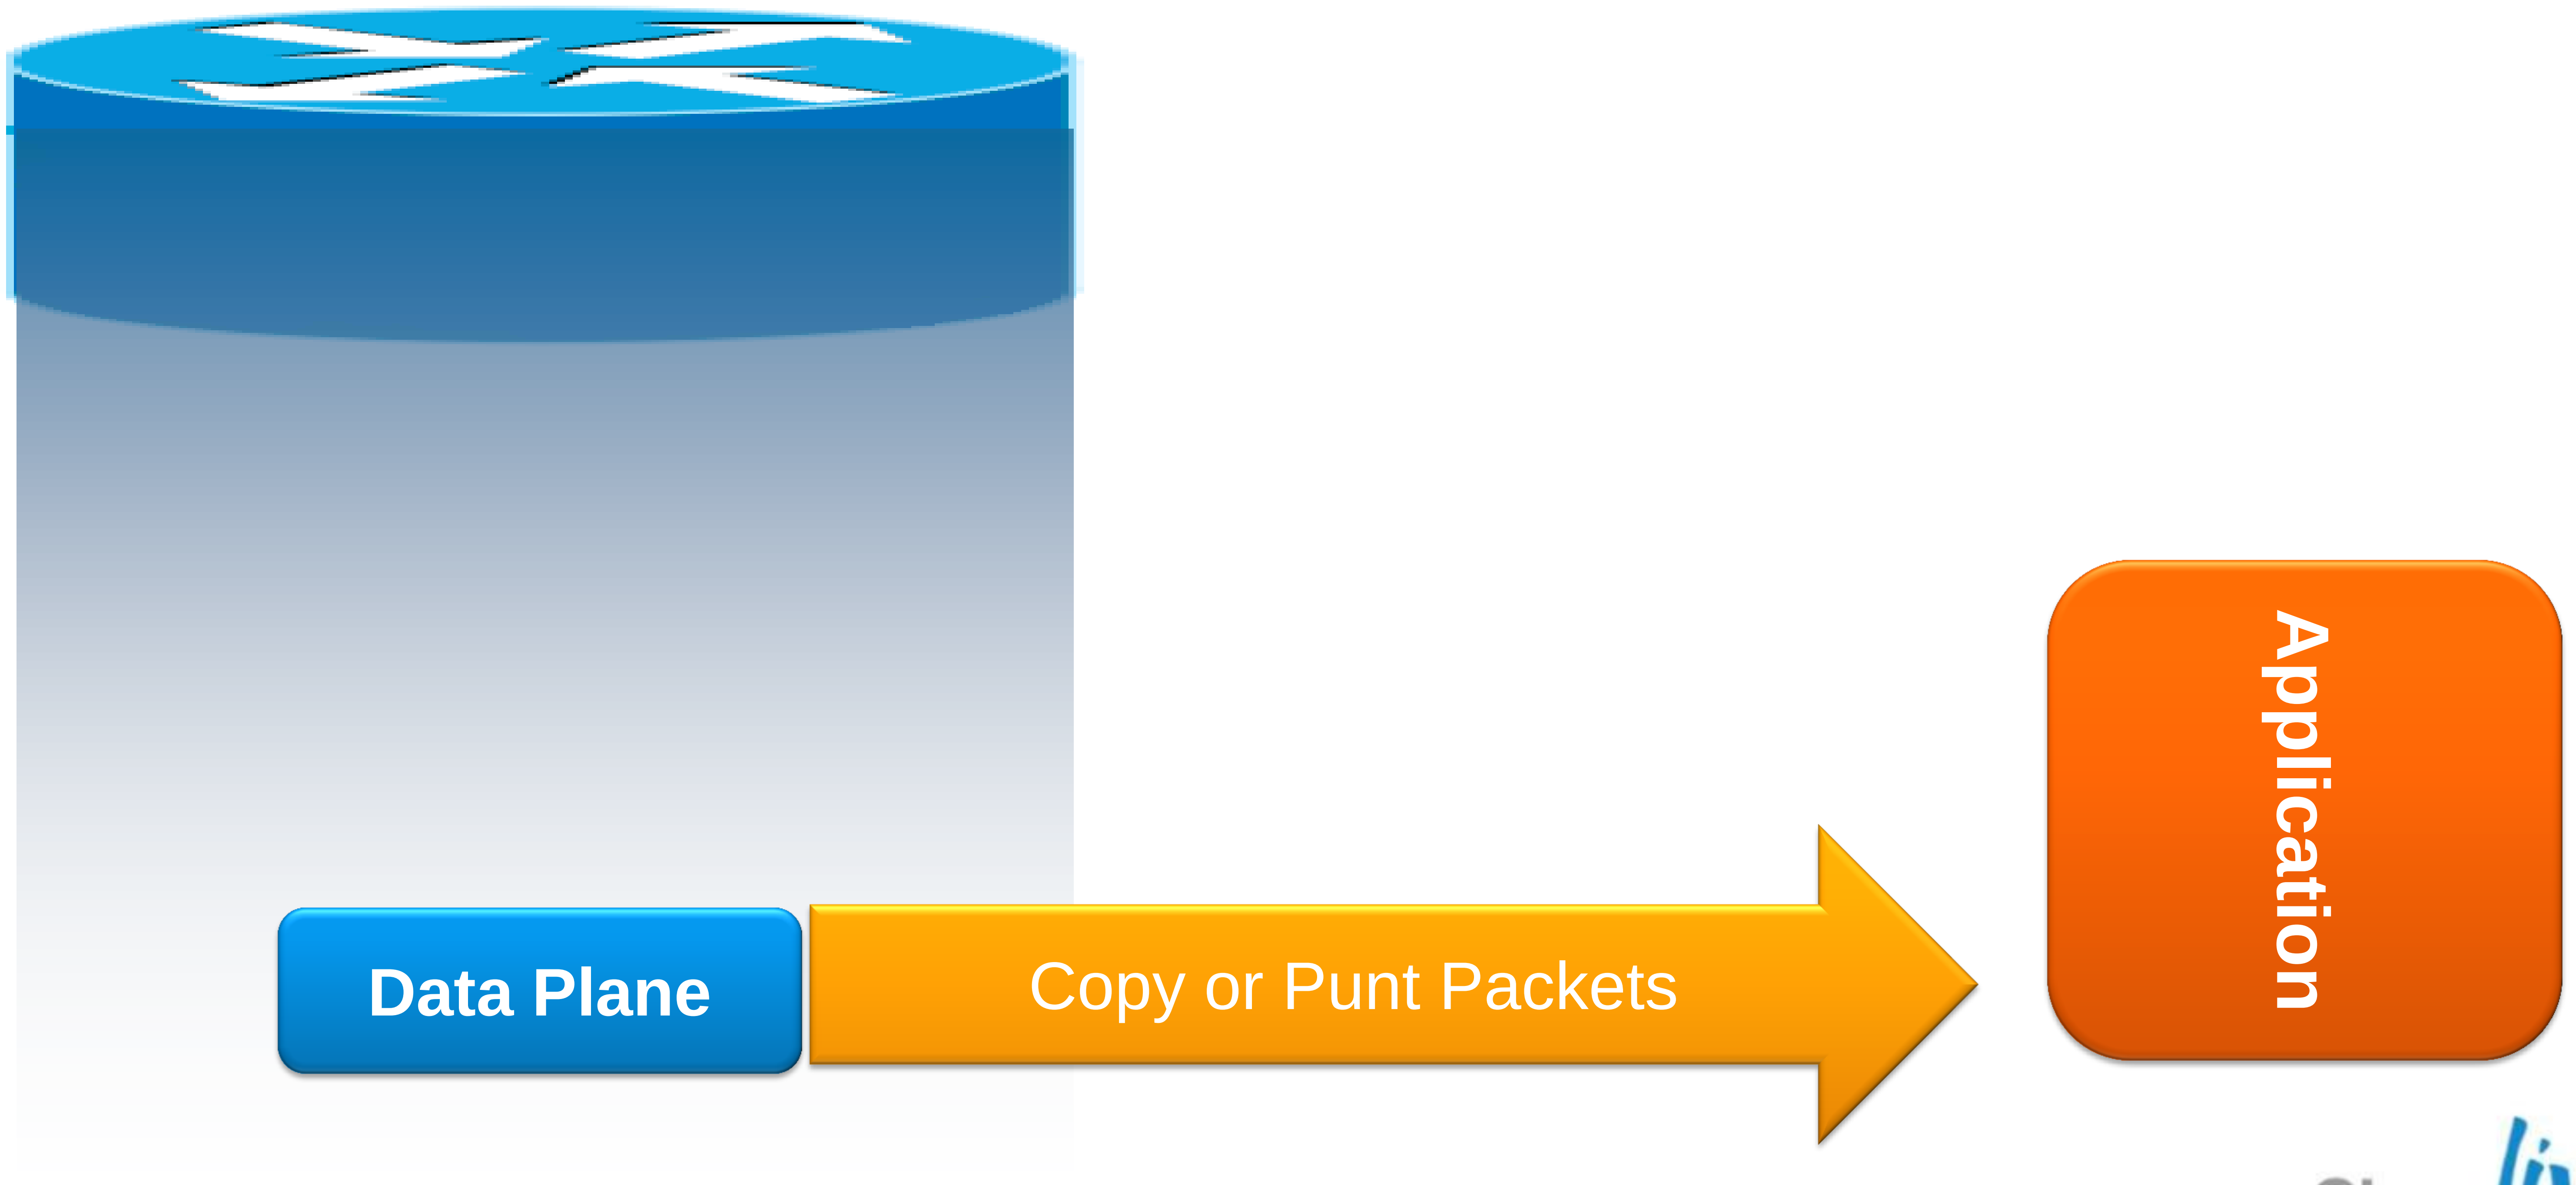

OnePK

Injecting Packets

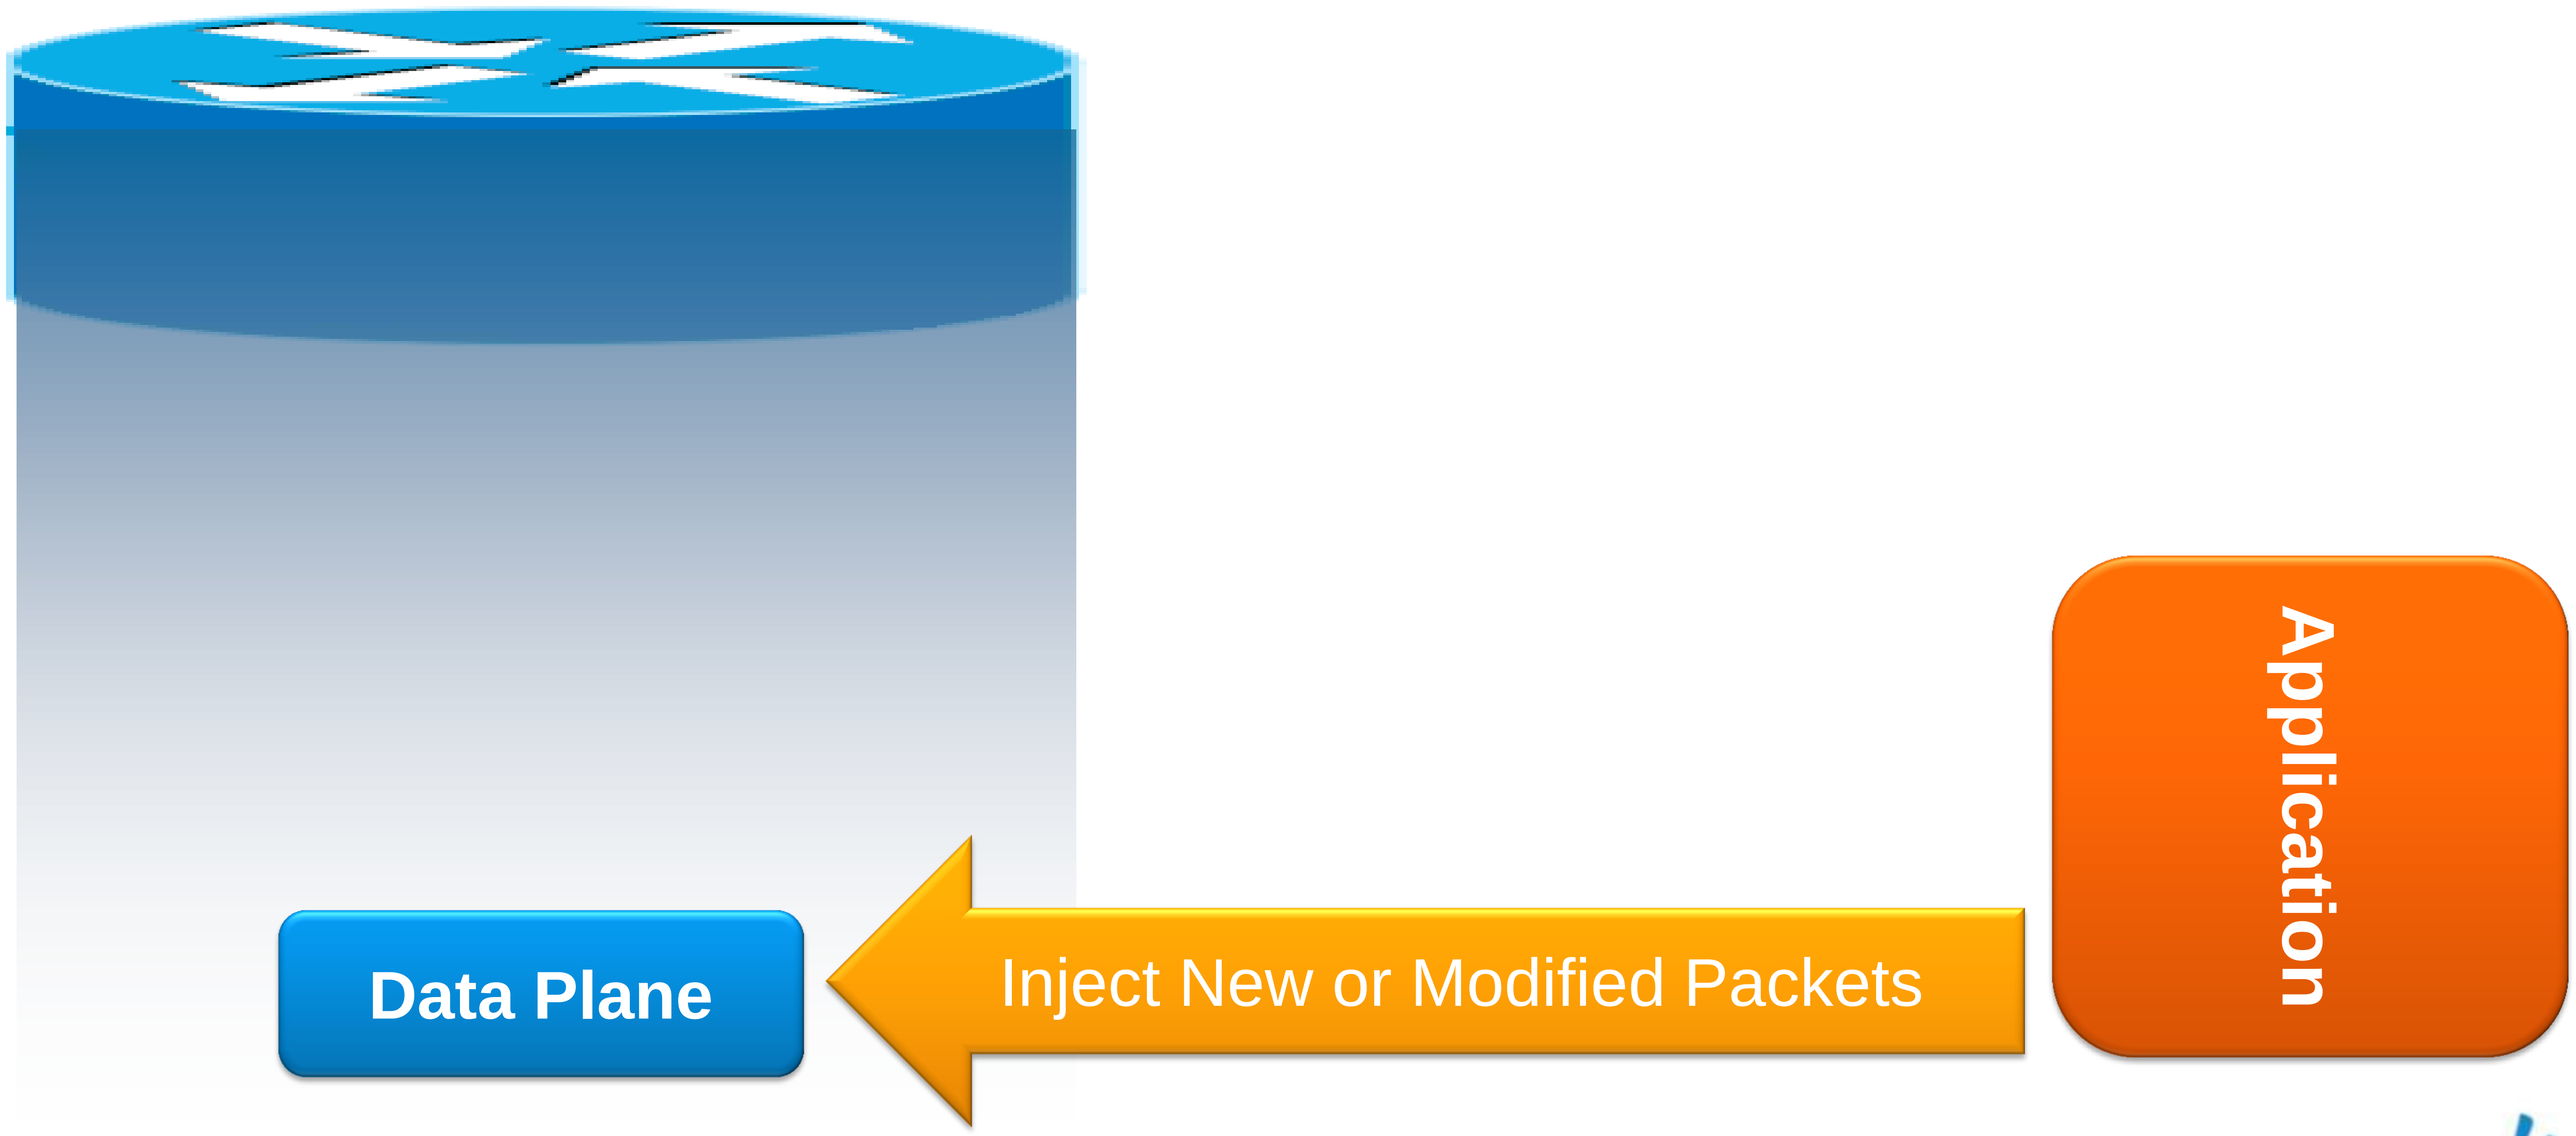

OnePK Demo

OnePK
Security

- Security Five Ways

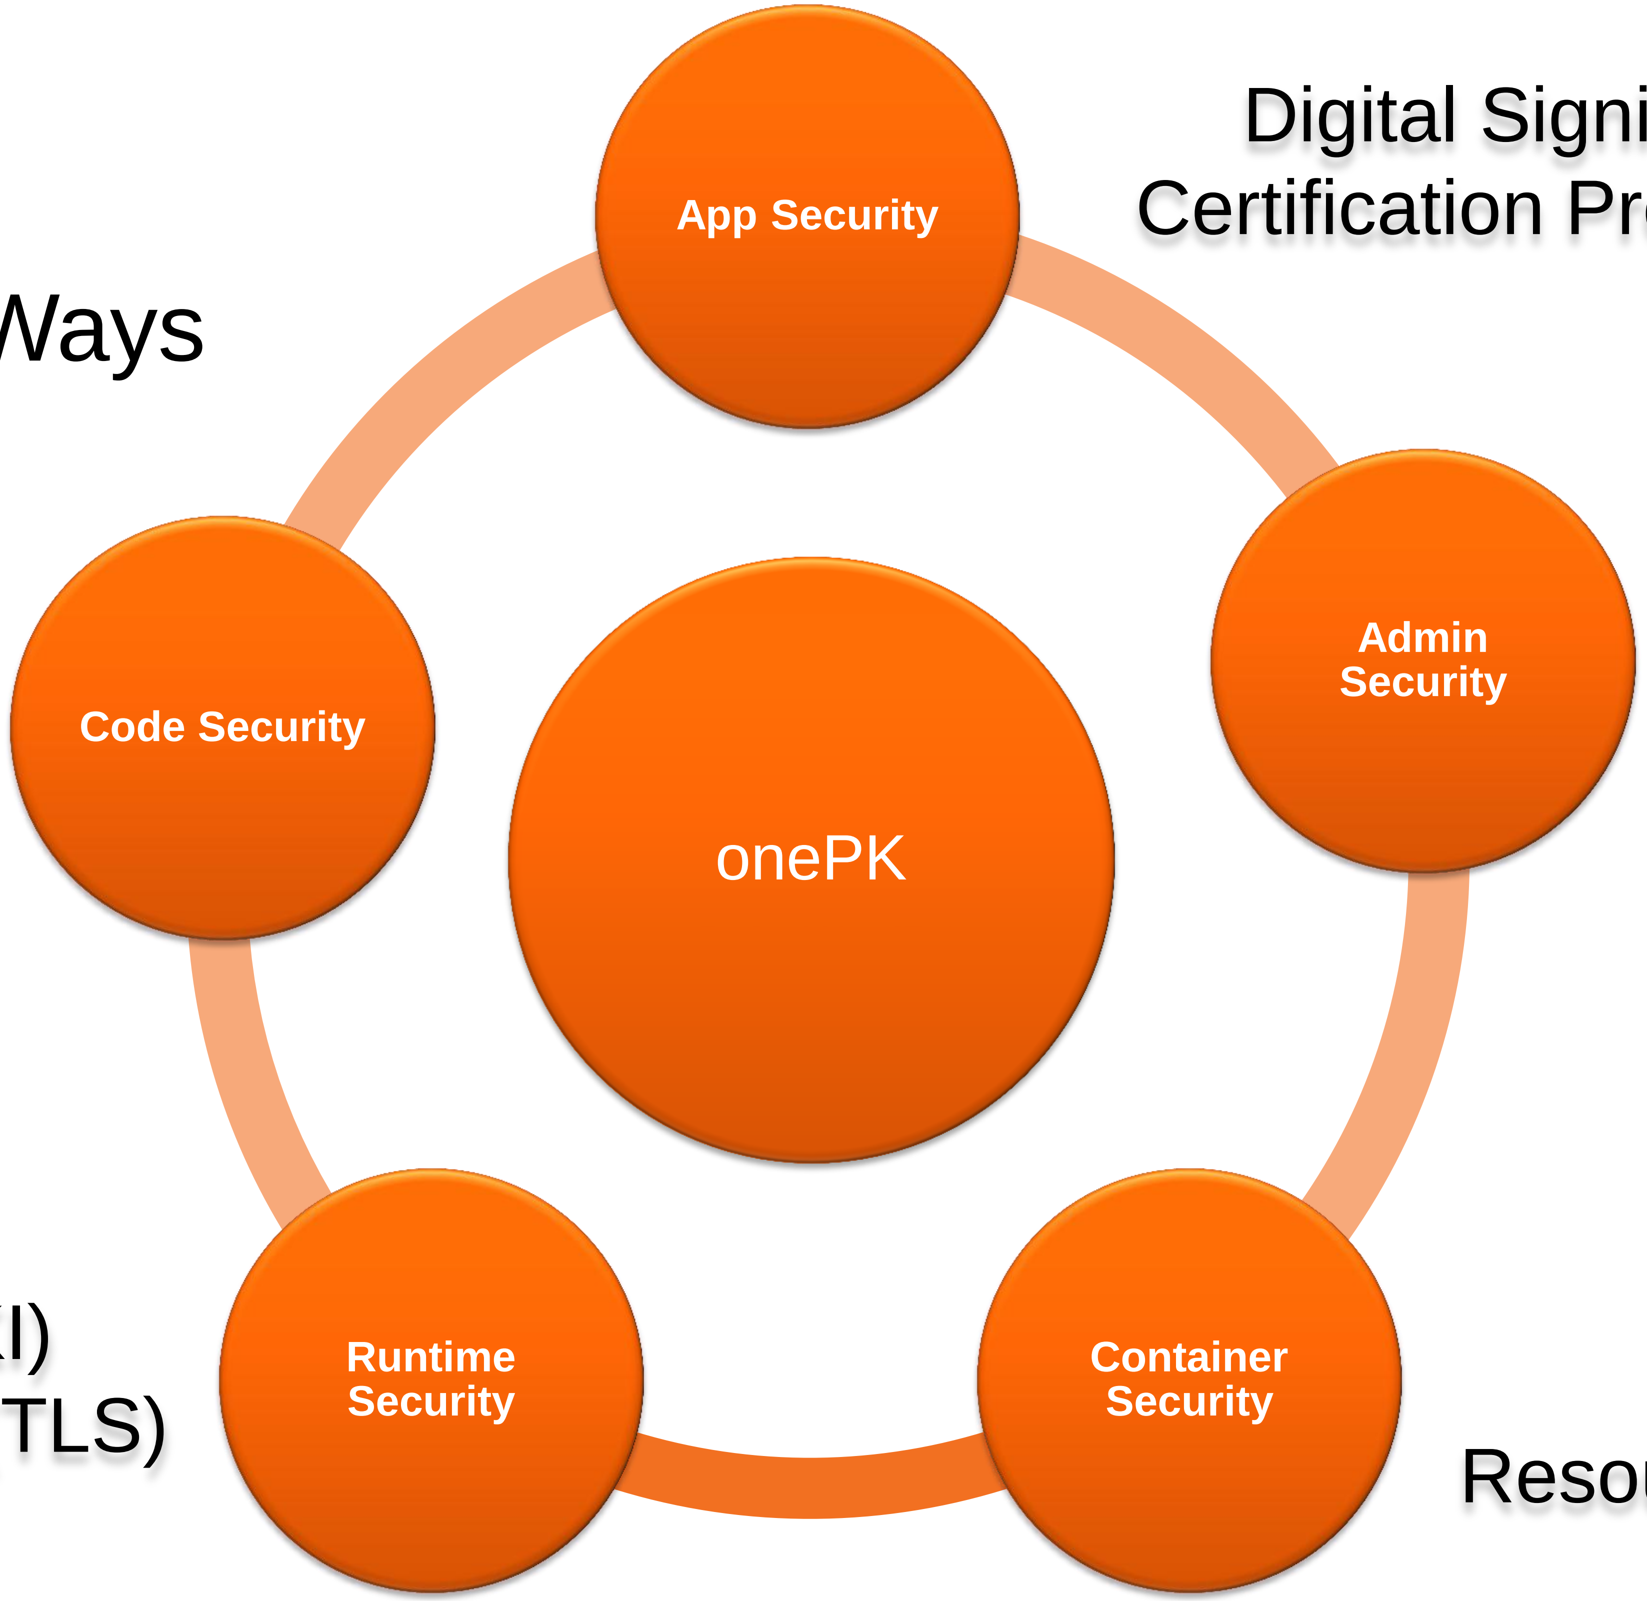

Digital Signing
Certification Process

CLI Control
Resource Allocation

Isolation
Resource Consumption

AAA (PKI)
Encryption (TLS)

Code Isolation
Strong Typing

Key Takeaways

OnePK Platform

- Build, Automate, Improve
- Speed and Faster Adaptability
- Extend
- Revenue and Cost Savings
- Simplicity, Integration, and the power of choice

A Platform For Innovation

Thinking Caps On….

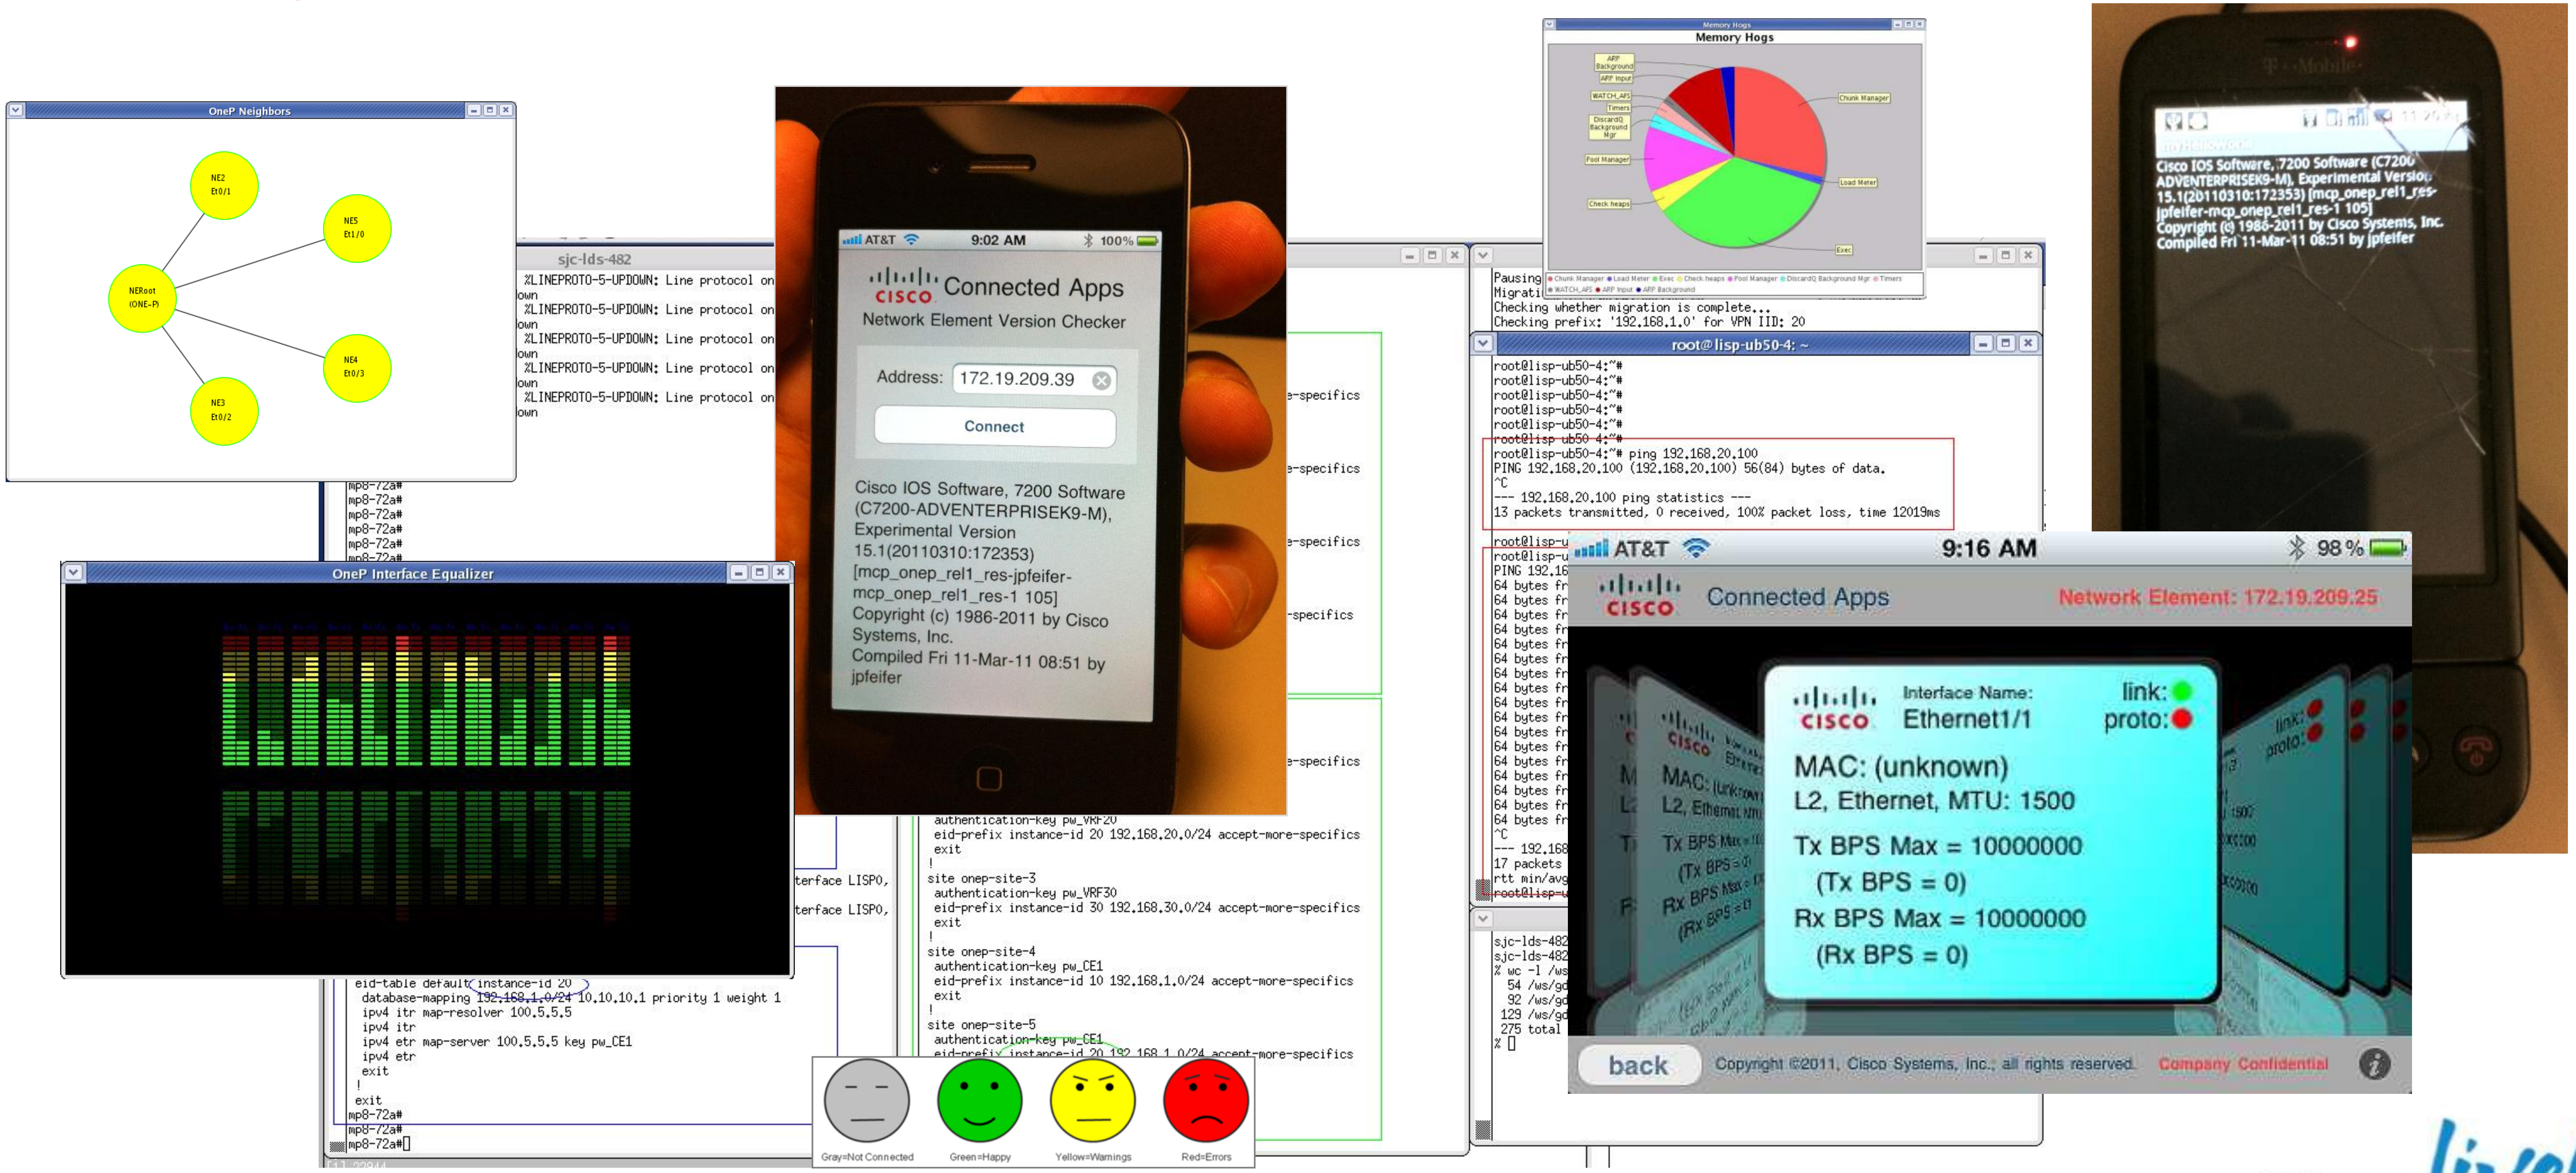

More Information

- Main OnePK home page

  http://www.cisco.com/go/onepk

FlexVPN

FlexVPN

What

- Internet Key Exchange Version (IKEv2), a next-generation key management protocol based on RFC 4306, is an enhancement of the IKE Protocol.
- IKEv2 is used for performing mutual authentication and establishing and maintaining security associations (SAs).
- FlexVPN is Cisco's implementation of the IKEv2 standard featuring a unified paradigm and CLI that combines site to site, remote access, hub and spoke topologies and partial meshes (spoke to spoke direct).

FlexVPN vs EasyVPN, DMVPN, and Crypto Maps

Why

```
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp client configur
 key cisco123
 pool dvti
 acl 100
crypto isakmp profile dvti
   match identity group cisco
   client authentication list
   isakmp authorization list
   client configuration addre
   virtual-template 1
crypto ipsec transform-set dv
crypto ipsec profile dvti
 set transform-set dvti
 set isakmp-profile dvti
interface Virtual-Template1 t
 ip unnumbered Ethernet0/0
 tunnel mode ipsec ipv4
 tunnel protection ipsec prof
ip local pool dvti 192.168.2.
ip route 0.0.0.0 0.0.0.0 10.0
access-list 100 permit ip 192
```

```
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto ipsec transform-set vpn-ts-set es
 mode transport
crypto ipsec profile vpnprofile
 set transform-set vpn-ts-set
interface Tunnel0
 ip address 10.0.0.254 255.255.255.0
 ip nhrp map multicast dynamic
 ip nhrp network-id 1
 tunnel source Serial1/0
 tunnel mode gre multipoint
 tunnel protection ipsec profile vpnprof

ip route 192.168.0.0 255.255.0.0 Null0ro
bgp log-neighbor-changes
redistribute static
 neighbor DMVPN peer-group
 bgp listen range 10.0.0.0/24 peer-group
 neighbor DMVPN remote-as 1
 no auto-summary
```

```
crypto isakmp policy 1
 encr 3des
 authentication pre-share
 group 2
crypto isakmp client configuration group cisco
 key pr3sh@r3dk3y
 pool vpnpool
 acl 110
crypto ipsec transform-set vpn-ts-set esp-3des esp-sha-hmac
crypto dynamic-map dynamicmap 10
 set transform-set vpn-ts-set
 reverse-route
crypto map client-vpn-map client authentication list userauthen
crypto map client-vpn-map isakmp authorization list groupauthor
crypto map client-vpn-map client configuration address initiate
crypto map client-vpn-map client configuration address respond
crypto map client-vpn-map 10 ipsec-isakmp dynamic dynamicmap
interface FastEthernet0/0
 ip address 83.137.194.62 255.255.255.240
 crypto map client-vpn-map
ip local pool vpnpool 10.10.1.1 10.10.1.254
access-list 110 permit ip 192.168.1.0 0.0.0.255 10.10.1.0 0.0.0.255
```

VPN Technology Selection

Why

- 3rd party and legacy support
- Hub & Spoke
- AAA Manageability
- Failover time
- Spoke – Spoke Direct
- Solution vs Components
- IPv4/IPv6 dual stack
- Failure detection method
- Design complexity
- Route Injection
- Dual DMVPN
- Dynamic Routing
- Crypto Map or Tunnels
- Feature order
- Multi-Hub Homing
- Per peer ACL's
- Multicast
- Scalability
- Multi-ISP Homing
- QoS support
- High Availability

FlexVPN Unifies

Unified Overlay VPNs

- One VPN to learn and deploy

VPN	Interop	Dynamic Routing	IPsec Routing	Spoke-spoke direct (shortcut)	Remote Access	Simple Failover	Source Failover	Config push	Per-peer config	Per-Peer QoS	Full AAA Management
Easy VPN	No	No	Yes	No	Yes	Yes	No	Yes	Yes	Yes	Yes
DMVPN	No	Yes	No	Yes	No	partial	No	No	No	group	No
Crypto Map	Yes	No	Yes	No	Yes	poor	No	No	No	No	No
Flex VPN	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Cisco live!

FlexVPN

How

All parameters tunable "per-peer" via AAA

```
crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local fqdn R1.cisco.com
 authentication local rsa-sig
 authentication remote eap
 pki trustpoint TP sign
 aaa authentication eap default
 aaa authorization user eap
 virtual-template 1
interface Virtual-Template1 type tunnel
 ip unnumbered loopback0
 tunnel protection ipsec profile default
 ip nhrp network-id 1
 tunnel mode ipsec ipv4
```

IKEv2 Parameters

Hub & Spoke

Remote Access

Spoke-Spoke
shortcut switching

Dual Stack v4/v6

Legacy
crypto map peer

IKEv2

Key Comparisons with IKEv1

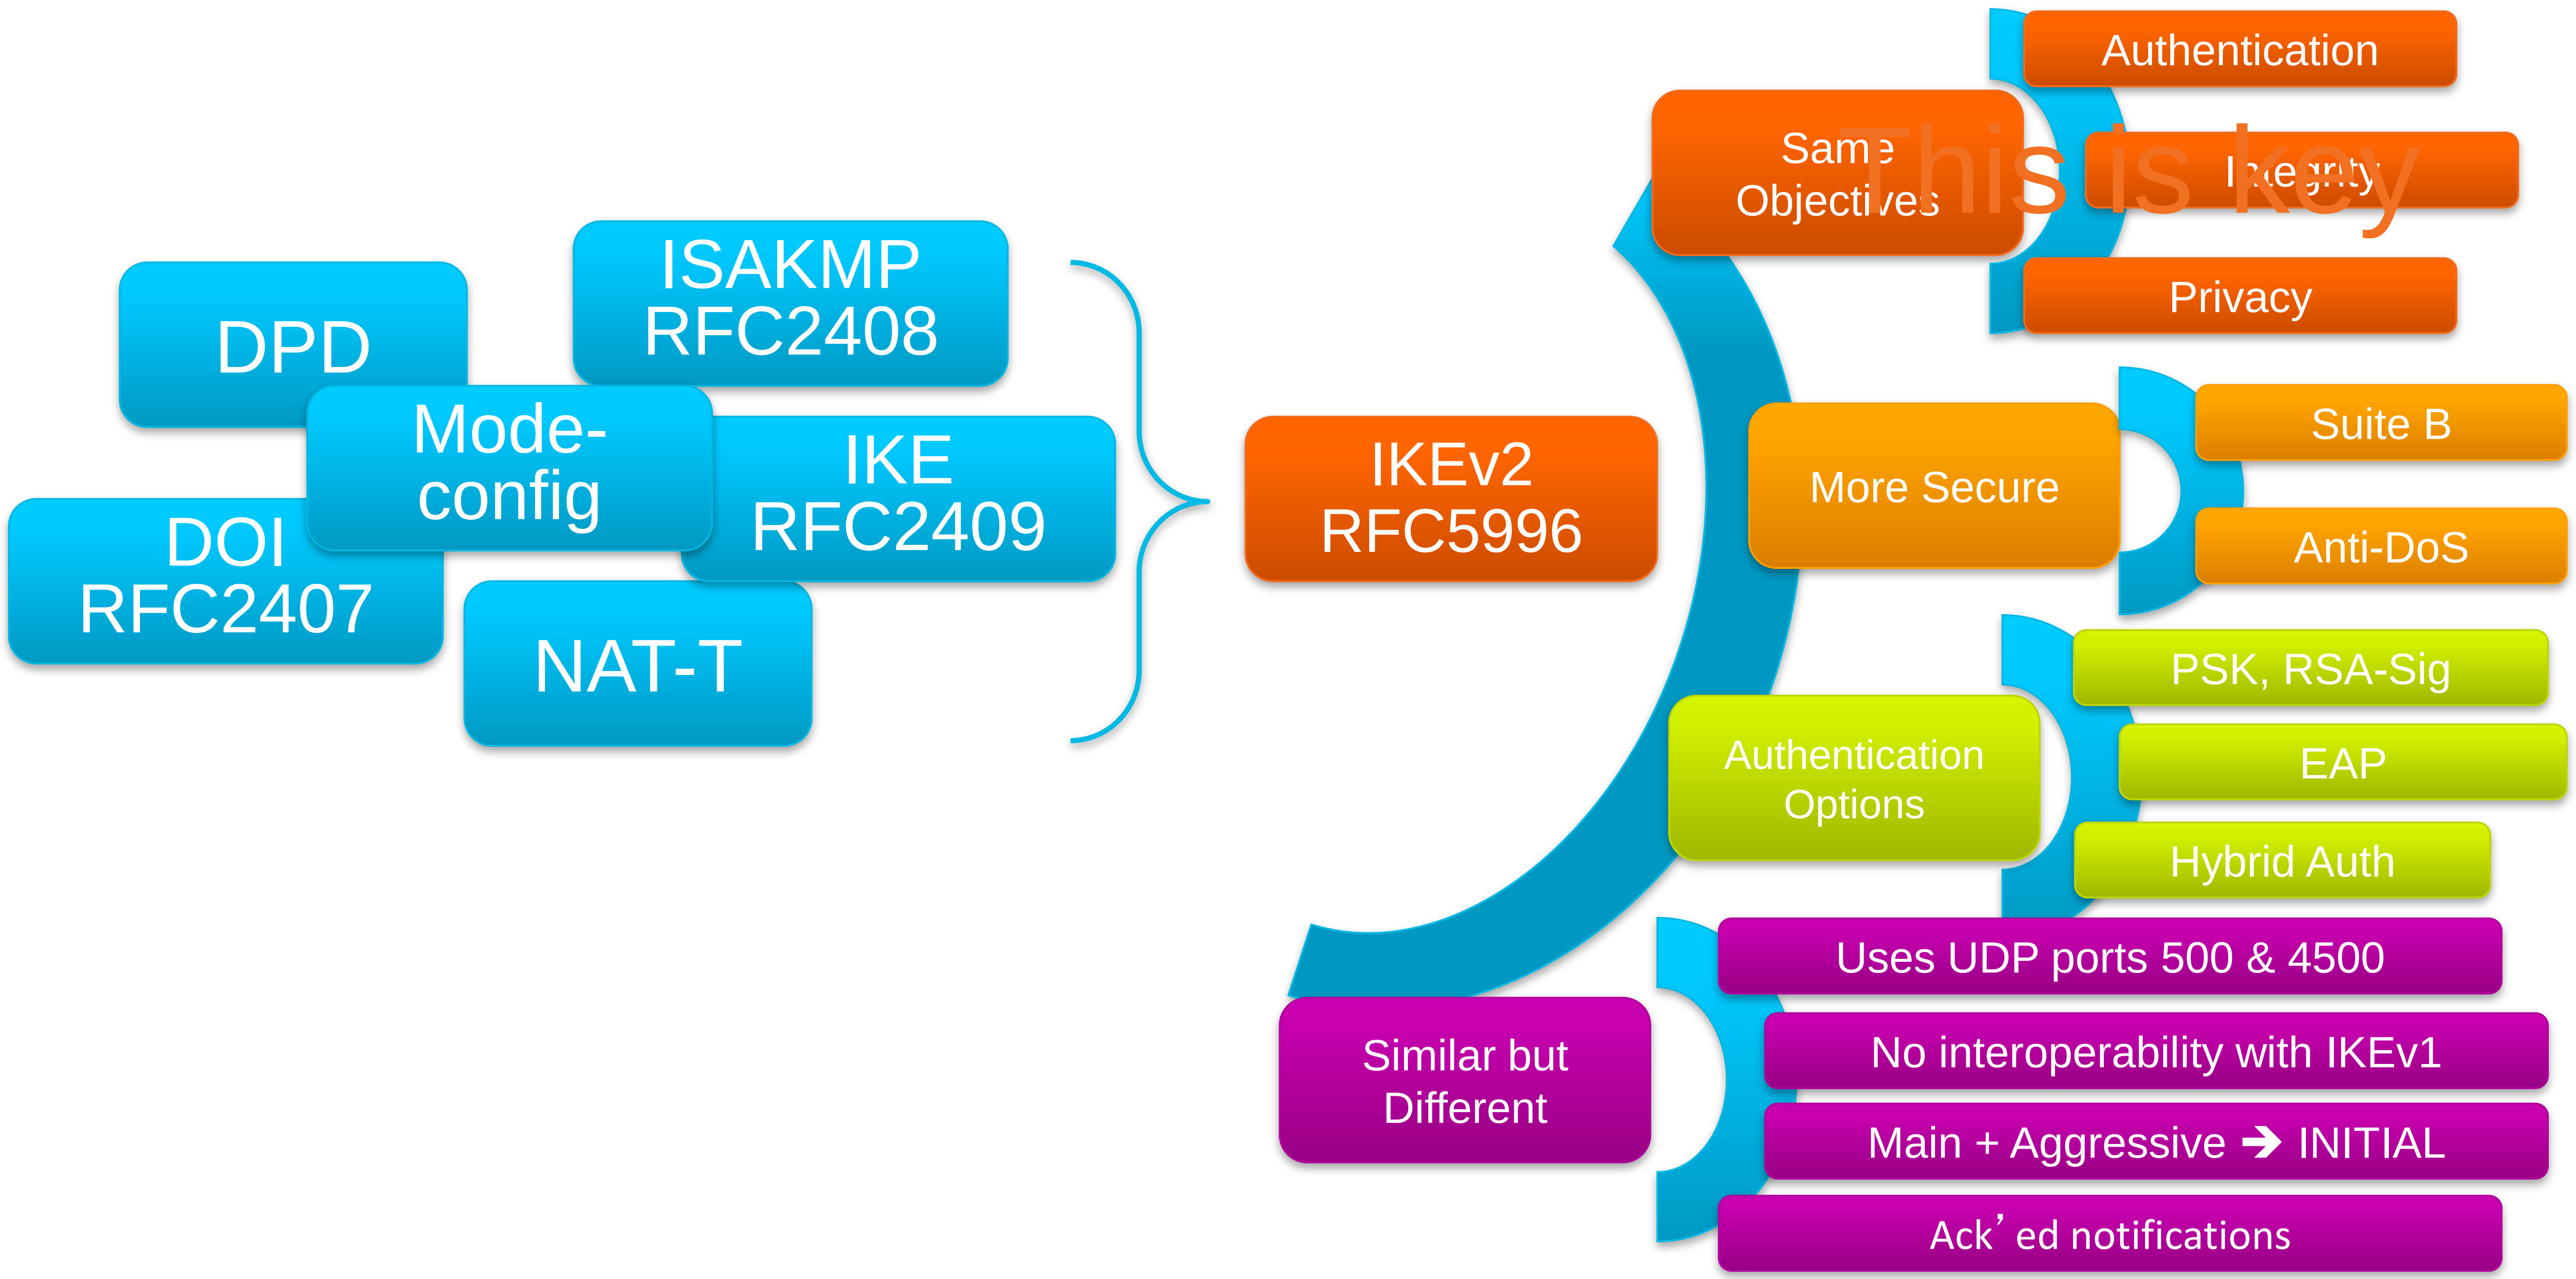

Extensible Authentication Protocol
(EAP)

- No X-AUTH in IKEv2; EAP instead
- EAP – authentication framework that provides common functions for various methods:
  - Tunnelling - EAP-TLS, EAP/PSK, EAP-PEAP…
  - Non-tunnelling – EAP-MSCHAPv2, EAP-GTC, EAP-MD5,…
- Implemented as additional IKE_AUTH exchanges
- Only used to authenticate the initiator to responder
- Responder MUST use Certificate
- Can severely increase number of messages (12-16)

Non-Tunnelling Recommended

Smart Defaults

Intelligent, reconfigurable defaults

- Pre-existing constructs:

  crypto ikev2 proposal

  AES-CBC 256, 196,128 , 3DES / SHA-512,384,256, SHA-1, MD5 / group 5, 2

  crypto ikev2 policy (match any)

  crypto ipsec transform-set (AES-128, 3DES / SHA, MD5)

  crypto ipsec profile default (default transform set, ikev2 profile default)

- Only an IKEv2 profile called "default" needs to be created

```
crypto ikev2 profile default
 match identity remote address 10.0.1.1
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint TP
!
interface Tunnel0
 ip address 192.168.0.1 255.255.255.252
 tunnel protection ipsec profile default
```

Example full config using smart defaults

Reconfigurable Defaults

All defaults can be modified, deactivated and restored

- Default proposals pre-configured
  - for IKEv2
  - for IPsec
- Modifying defaults
- Restoring defaults
- Disabling defaults

```
crypto ikev2 proposal default
  encryption aes-cbc-128
  hash md5

crypto ipsec transform-set default aes-cbc 256 sha-hmac
```

```
default crypto ikev2 proposal

default crypto ipsec transform-set
```

```
no crypto ikev2 proposal default

no crypto ipsec transform-set default
```

FlexVPN Usage

Anything Can Be Done

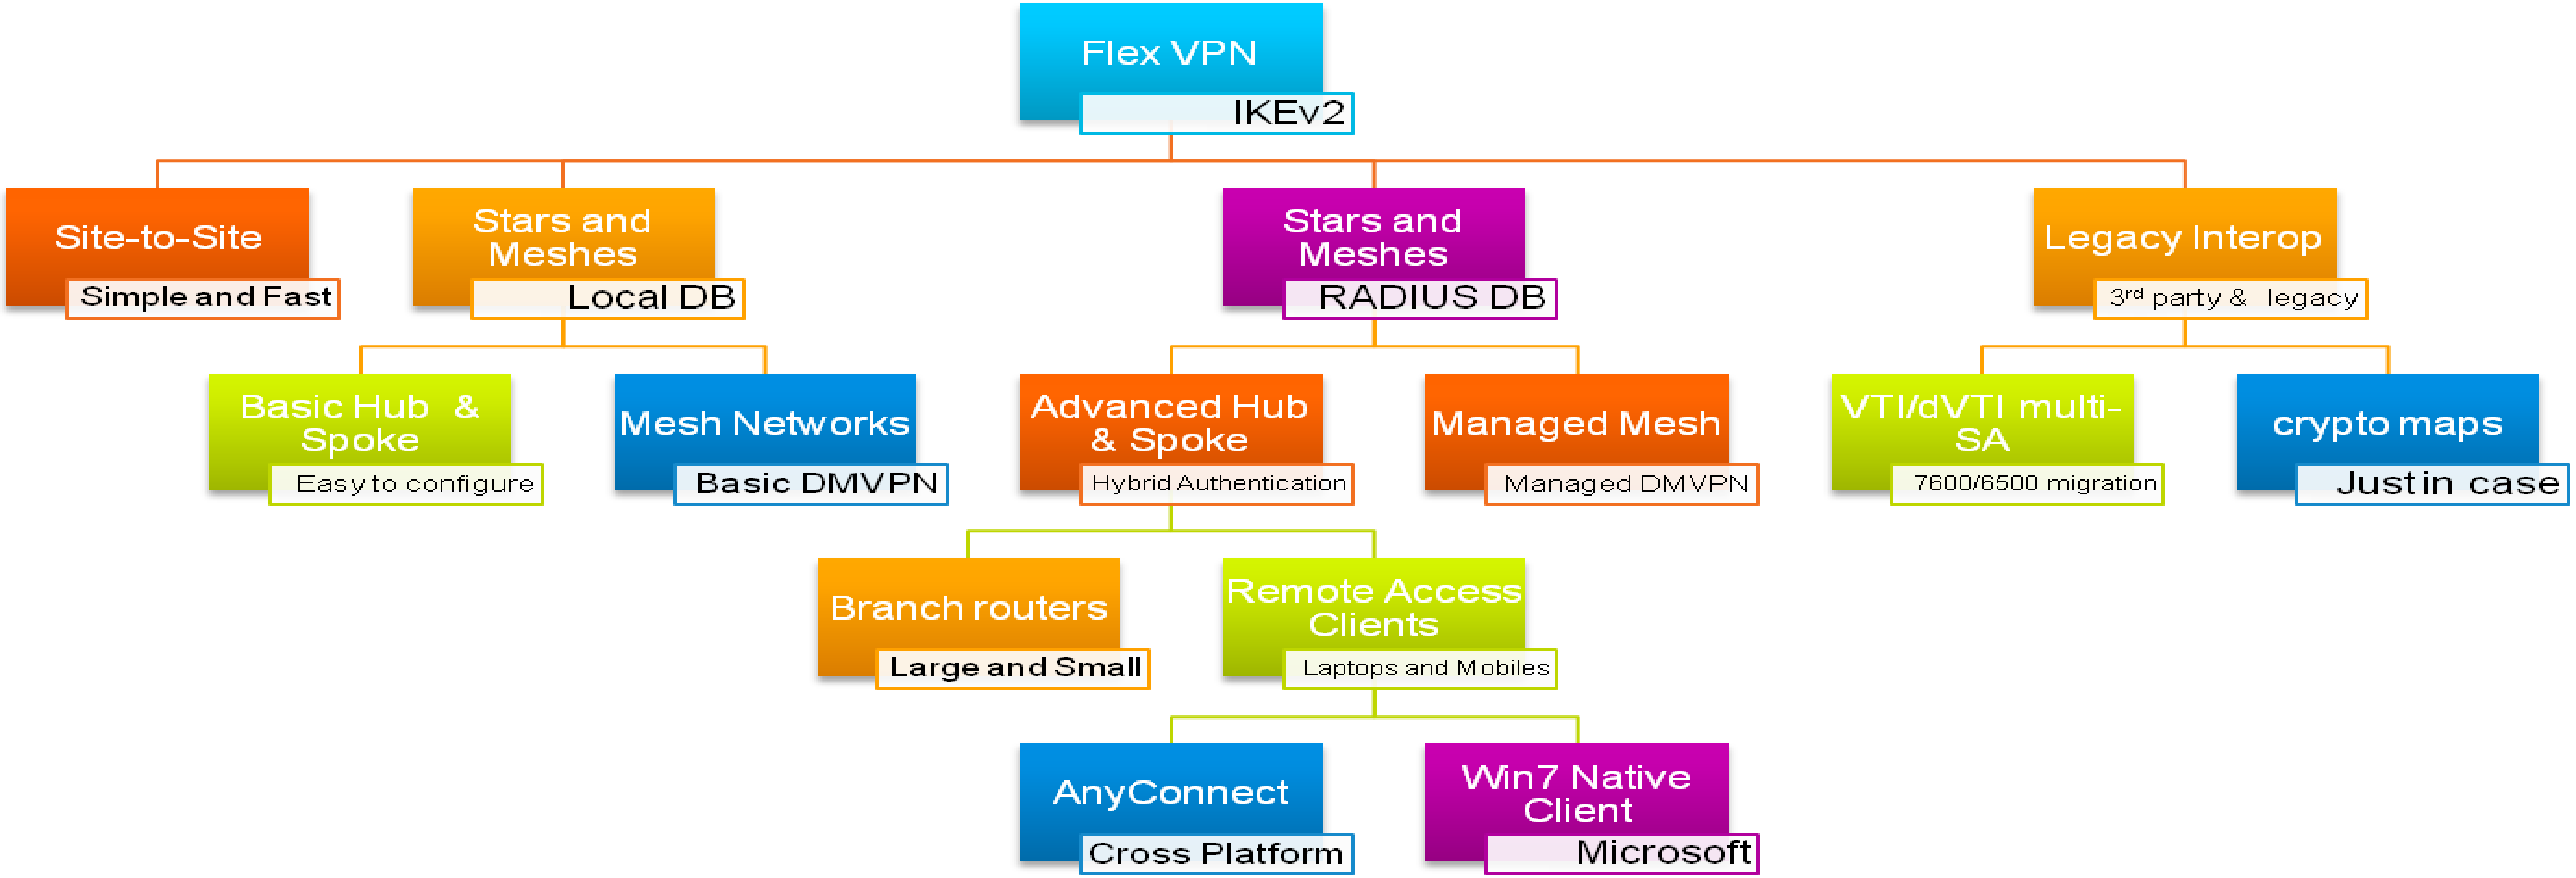

Pick and Choose

(Almost) Never Lose

Tunnelling	Authentication Method	Tunnel Config	Config Source
GRE/IPsec	Certificate	Static	Local config
Pure IPsec	Pre-shared Key	Dynamic	RADIUS
	EAP (initiator)	crypto map	Hybrid

Security policy & routing
IKEv2 “routing”
BGP
Static routes
Reverse-Route Injection
EIGRP or anything else!

Example 1a – Site to Site

IPv4 and Static Routing

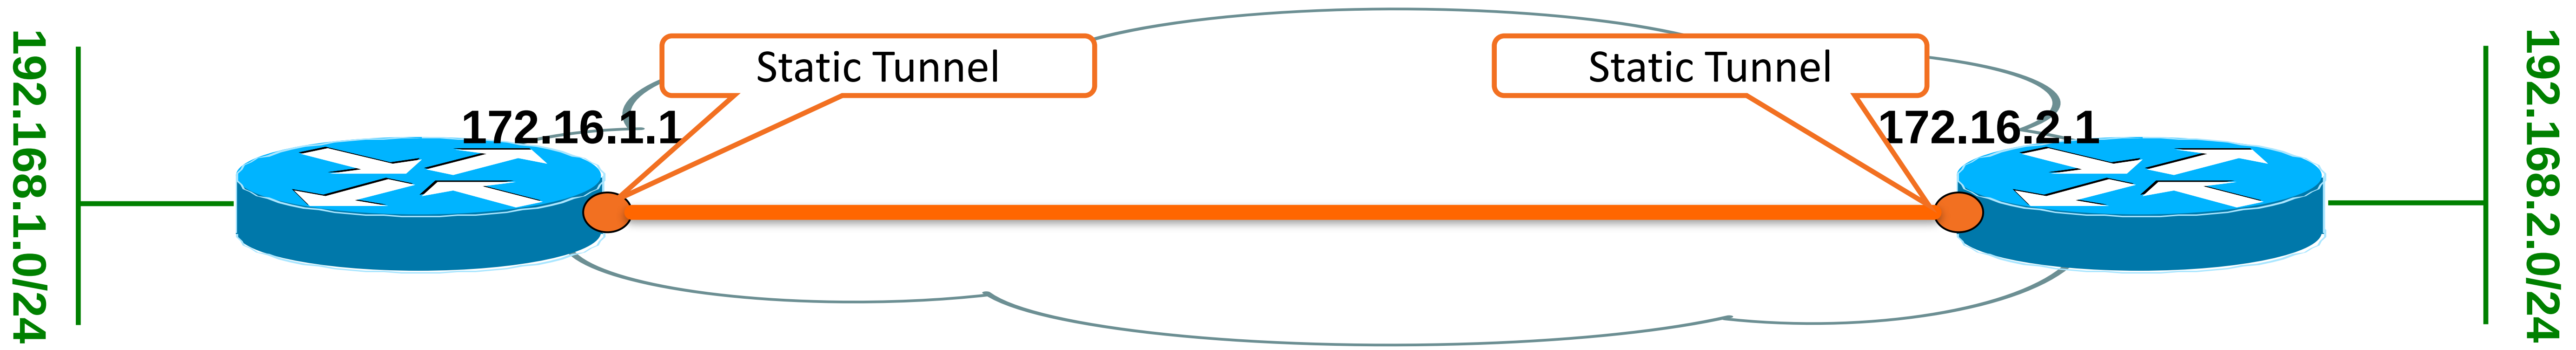


```
crypto ikev2 keyring KR
 peer RightPeer
  address 172.16.2.1
  pre-shared-key local CISCO
  pre-shared key remote OCSIC

crypto ikev2 profile default
 match identity fqdn RouterRight.cisco.com
 identity local fqdn RouterLeft.cisco.com
 authentication local pre-shared
 authentication remote pre-shared
 keyring local KR

interface Tunnel0
 ip address 10.0.0.1 255.255.255.252
 tunnel source FastEthernet0/0
 tunnel destination 172.16.2.1
 tunnel protection ipsec profile default

ip route 192.168.2.0 255.255.255.0 Tunnel0
```

Just a string

Peer address

Could use a routing protocol (IGP/BGP)

```
crypto ikev2 keyring KR
 peer LeftPeer
  address 172.16.1.1
  pre-shared-key local OCSIC
  pre-shared key remote CISCO

crypto ikev2 profile default
 match identity fqdn RouterLeft.cisco.com
 identity local fqdn RouterRight.cisco.com
 authentication local pre-shared
 authentication remote pre-shared
 keyring local KR

interface Tunnel0
 ip address 10.0.0.2 255.255.255.252
 tunnel source FastEthernet0/0
 tunnel destination 172.16.1.1
 tunnel protection ipsec profile default

ip route 192.168.1.0 255.255.255.0 Tunnel0
```

Example 1b – Site to Site

IPv6 and OSPF

For Your Reference

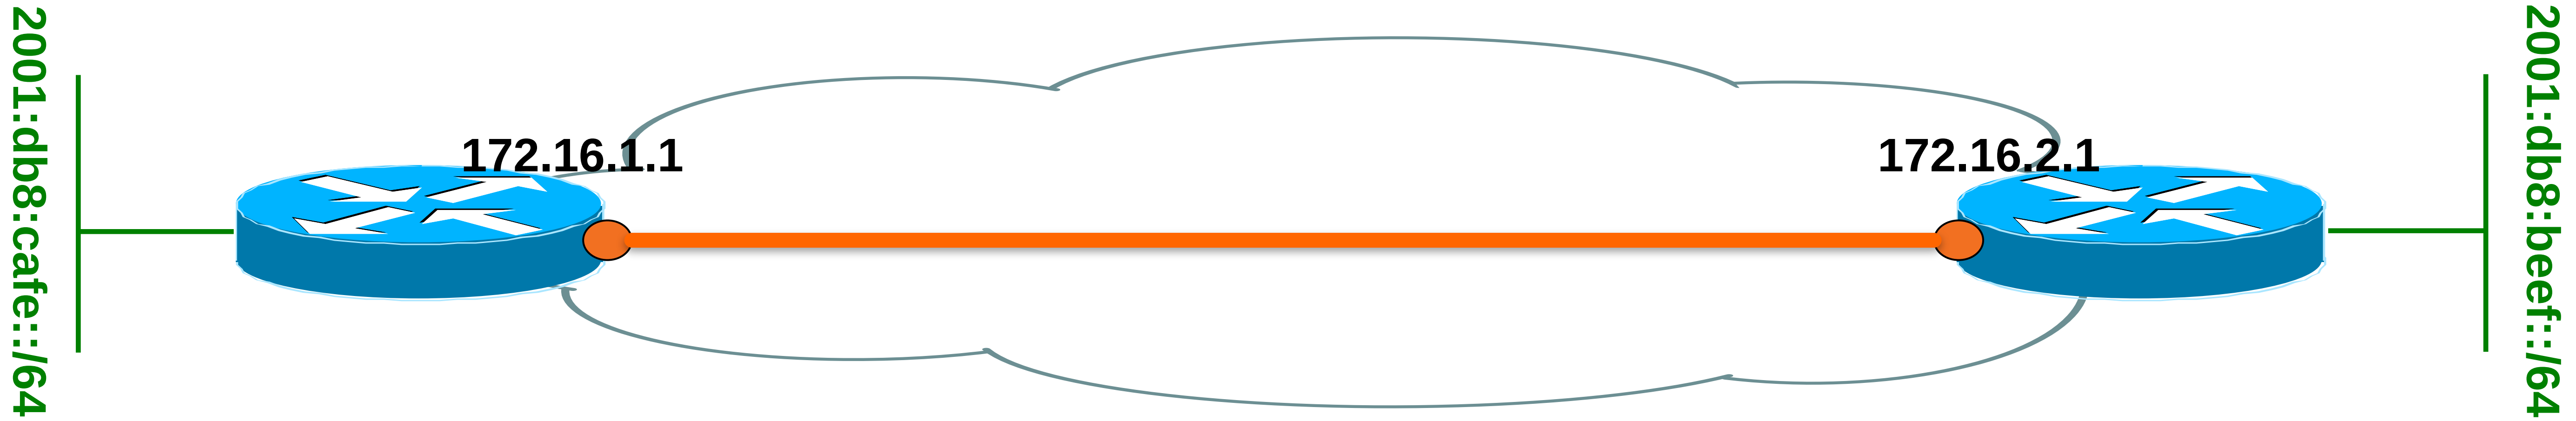


```
…

ipv6 unicast-routing

interface Tunnel0
 ipv6 address FE80::1 link-local
 ipv6 ospf 1 area 0
 tunnel source FastEthernet0/0
 tunnel destination 172.16.2.1
 tunnel protection ipsec profile default

interface E0/0
 ipv6 address 2001:db8:cafe::1/64
 ipv6 ospf 1 area 0
```

```
…

ipv6 unicast-routing

interface Tunnel0
 ipv6 address FE80::2 link-local
 ipv6 ospf 1 area 0
 tunnel source FastEthernet0/0
 tunnel destination 172.16.1.1
 tunnel protection ipsec profile default

interface E0/0
 ipv6 address 2001:db8:beef::1/64
 ipv6 ospf 1 area 0
```

Example 2 – Remote Access

Software Clients Connect to a Hub

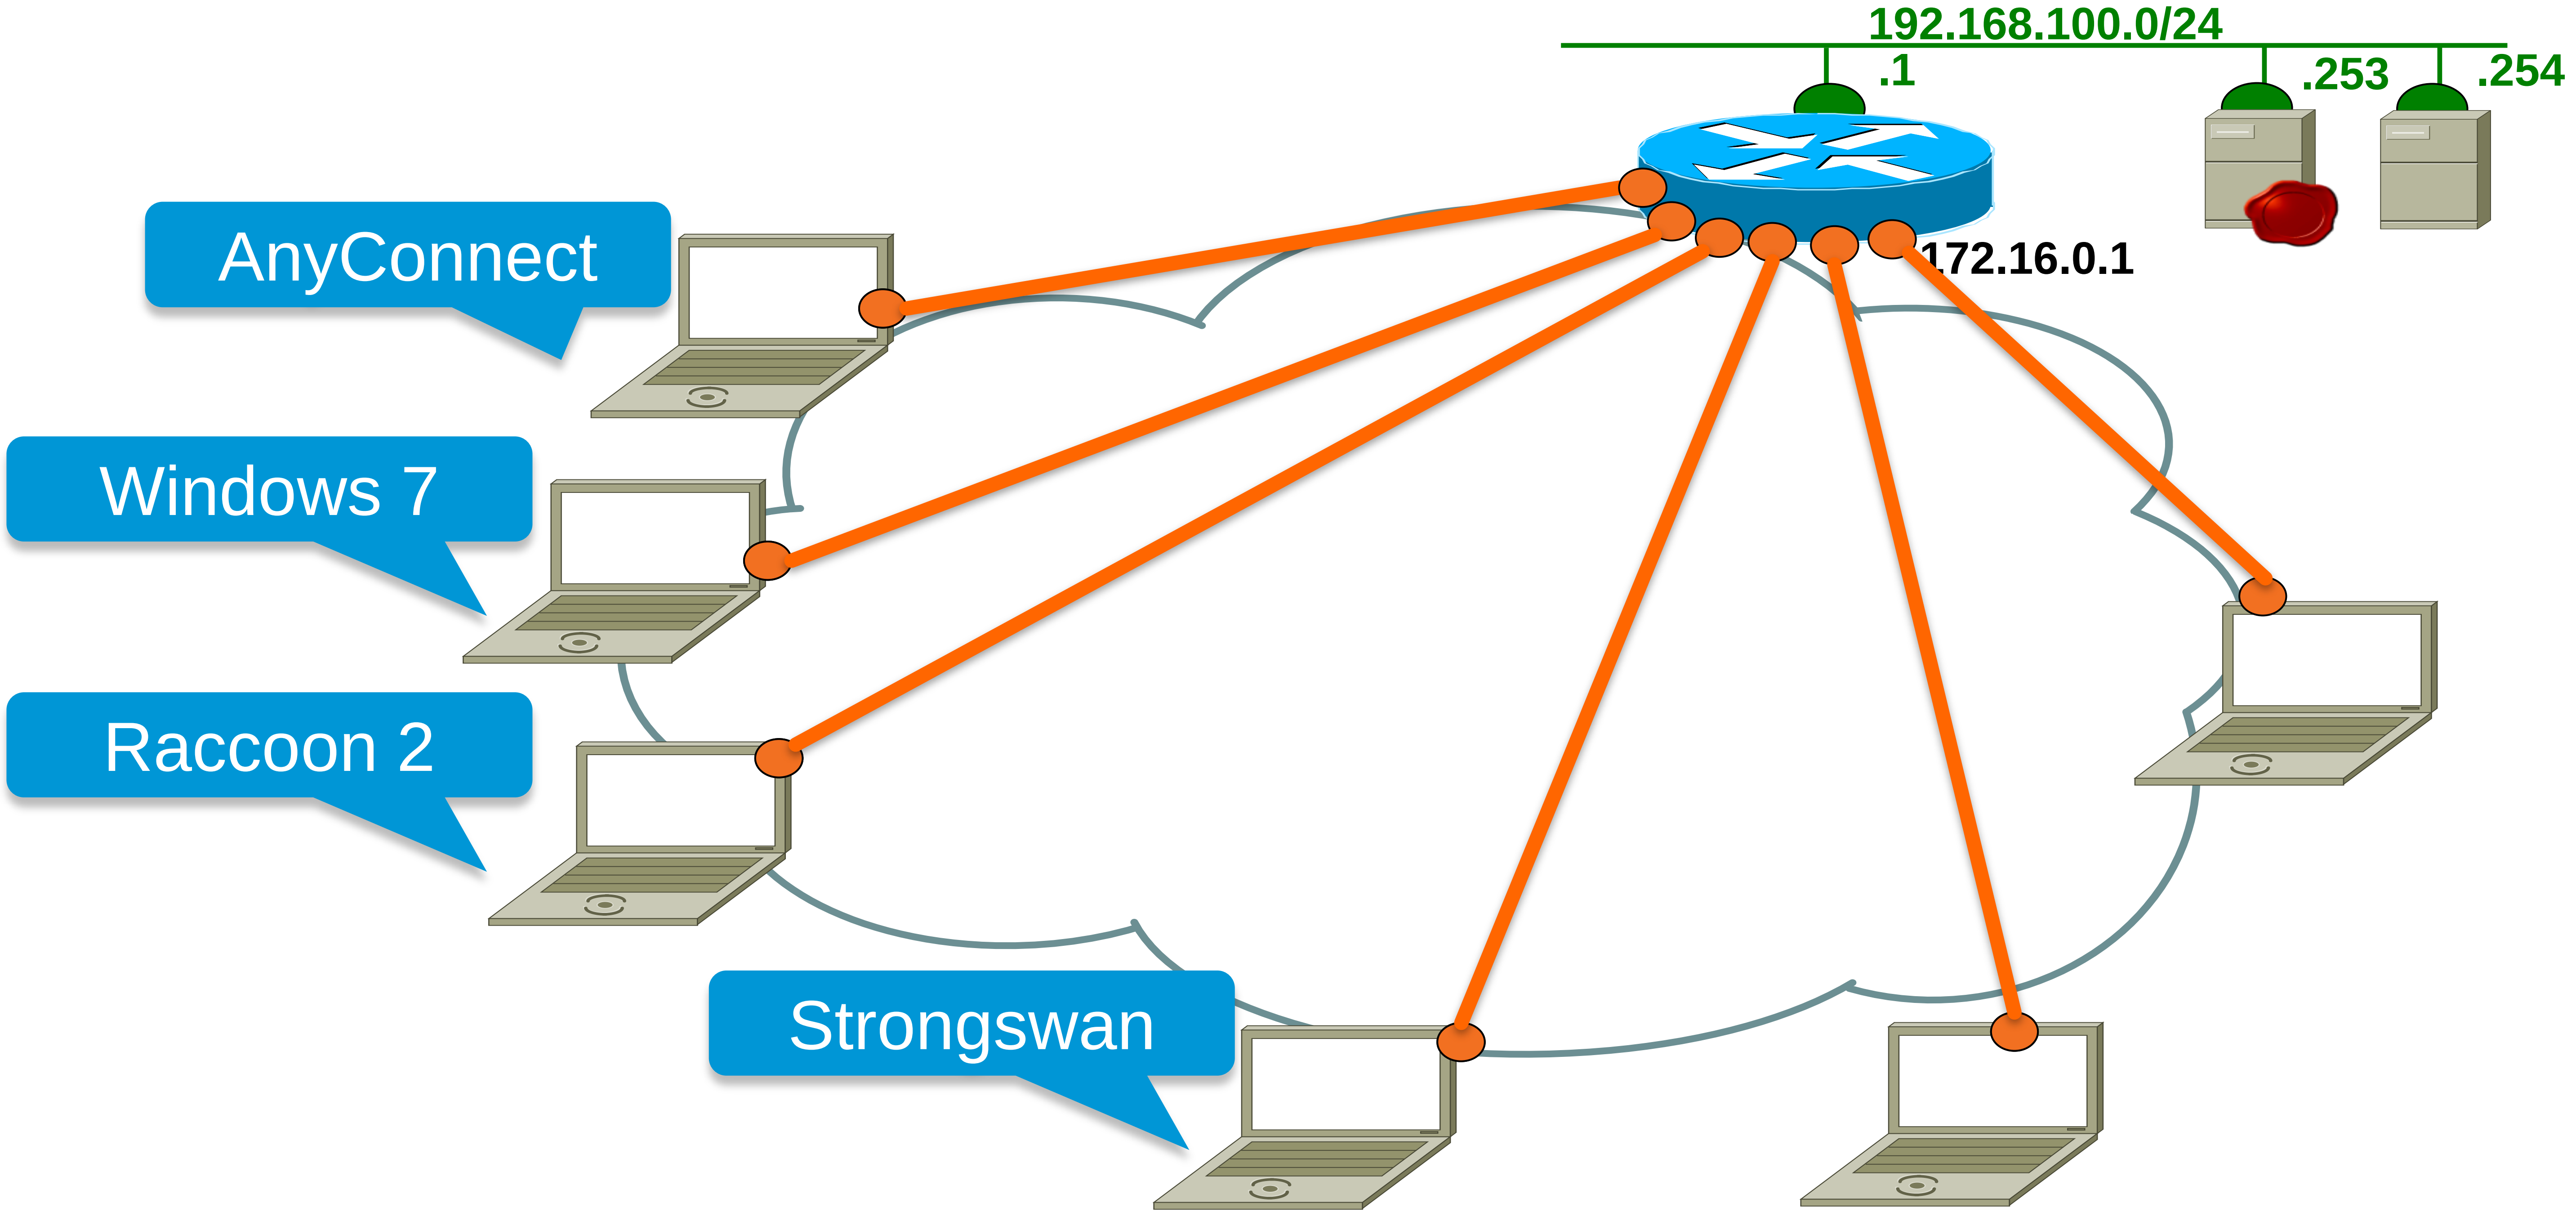

Example 2 – Remote Access

AC Requirements

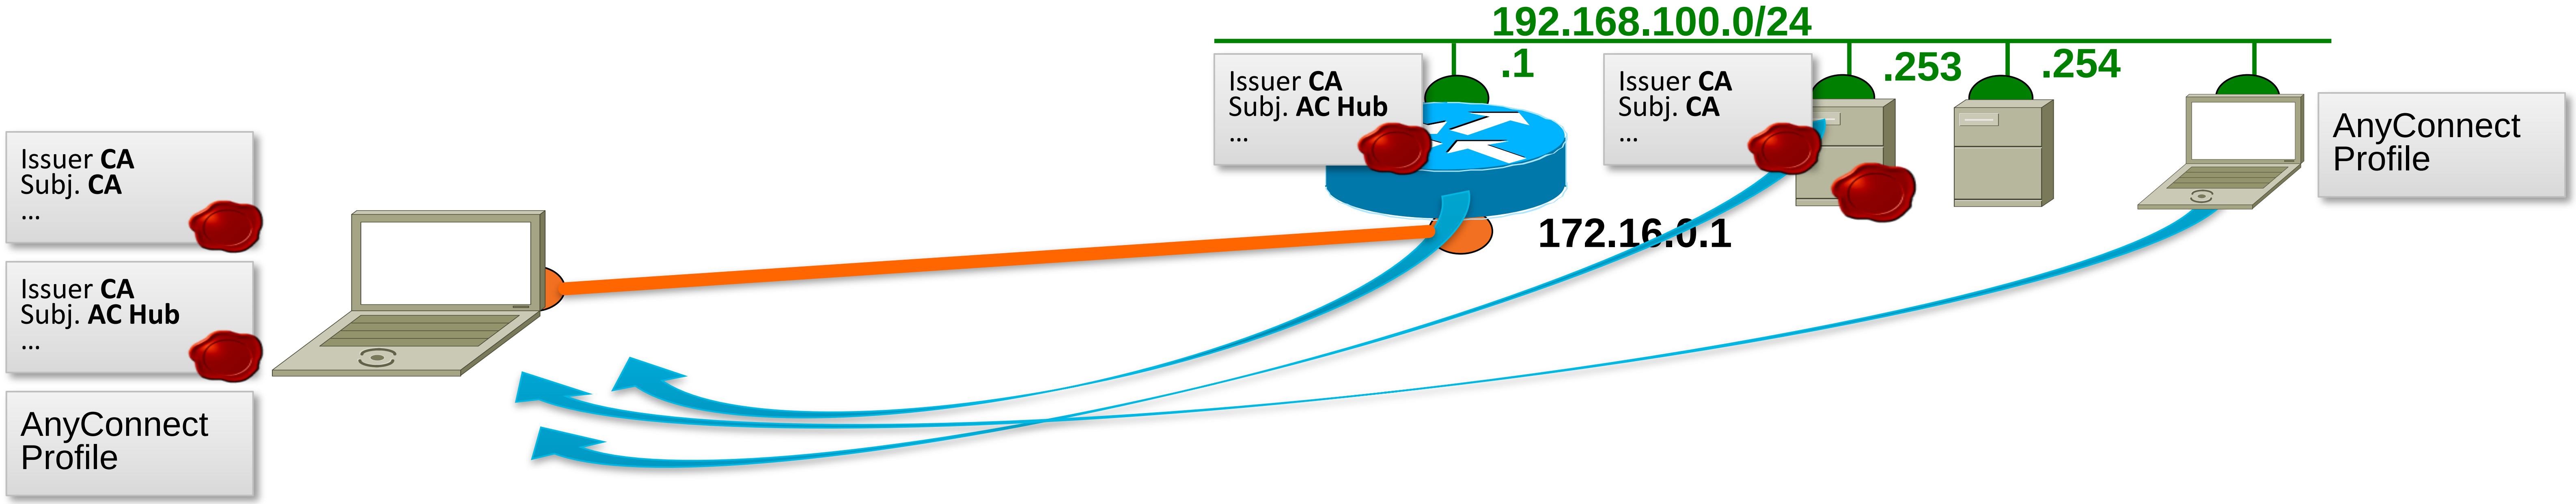

Hub Certificate:

Subject:
CN=**AC Hub**,
OU=TAC, O=Cisco, C=BE
...

AnyConnect Profile

PrimaryProtocol: IPsec
StandardAuthenticationOnly: true
AuthMethodDuringIKENegotiation: EAP-MD5 IKEIdentity: MyAnyConnect
HostName: AC Hub

Deploy Hub Certificate **or** CA Certificate on client

Any certificate store can be used
Local Computer store if tunnel set up before logon

Deploy AnyConnect profile on client
Watch for dependencies

Example 2 – Remote Access

Hub Configuration

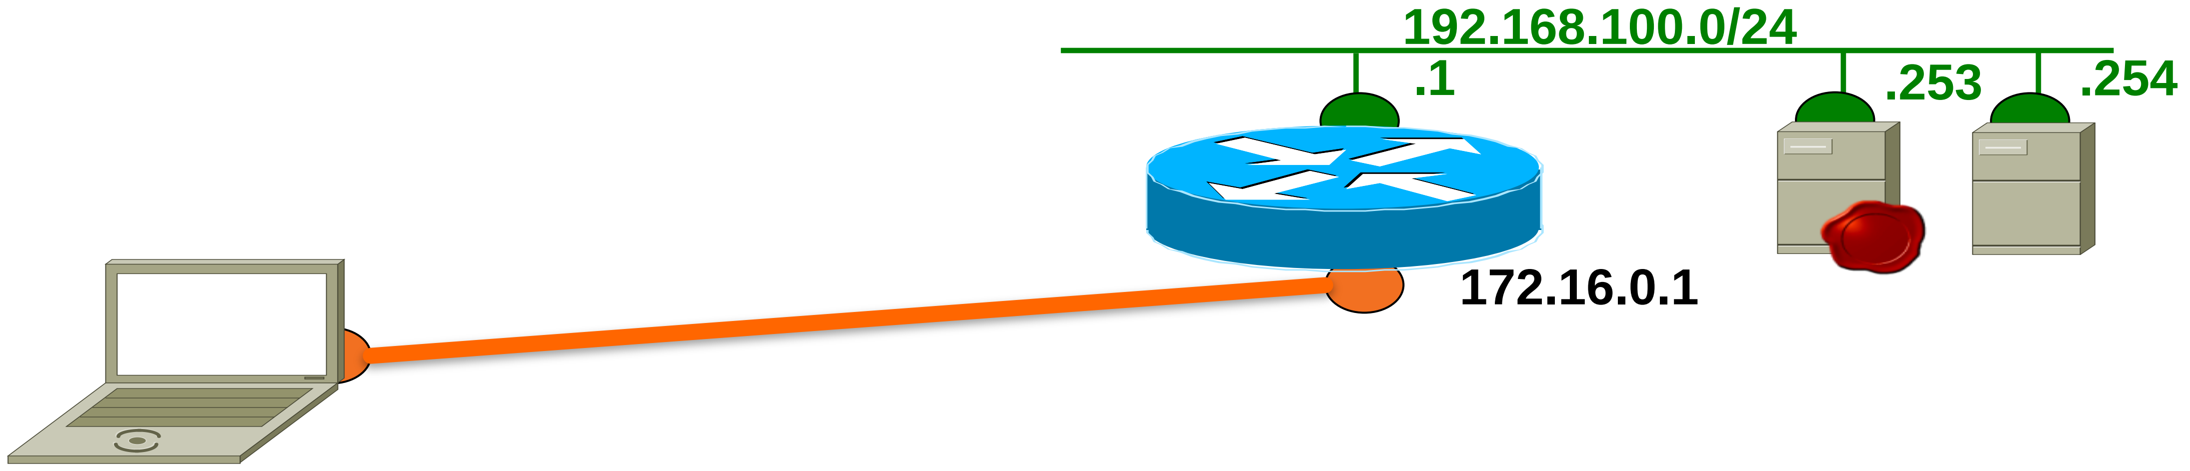


```
crypto ikev2 profile default
 match identity key-id MyAnyConnect
 identity local dn
 authentication local rsa-sig
 authentication remote eap query-identity
 pki trustpoint CA
 aaa authentication user eap AC
 aaa authorization user eap cached
 virtual-template 1

interface virtual-template 1 type tunnel
 ip unnumbered loopback0
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile default
```

Must match client profile configuration

RADIUS server authenticates client (EAP)

Authentication profile reused for authorisation

```
Fred@cisco.com
Cleartext-Password := "MyPass",
Framed-Pool=FlexPool
ipsec:route-accept=any
ipsec:route-set=interface
```

```
aaa authentication login AC group R
ip local pool FlexPool 10.0.0.1 10.0.0.253

aaa group server radius R
 server-private 192.168.100.254
  auth-port 1812 acct-port 1813
  key cisco123
```

Pool managed by Hub

Example 3 – Flex Mesh

Network Diagram

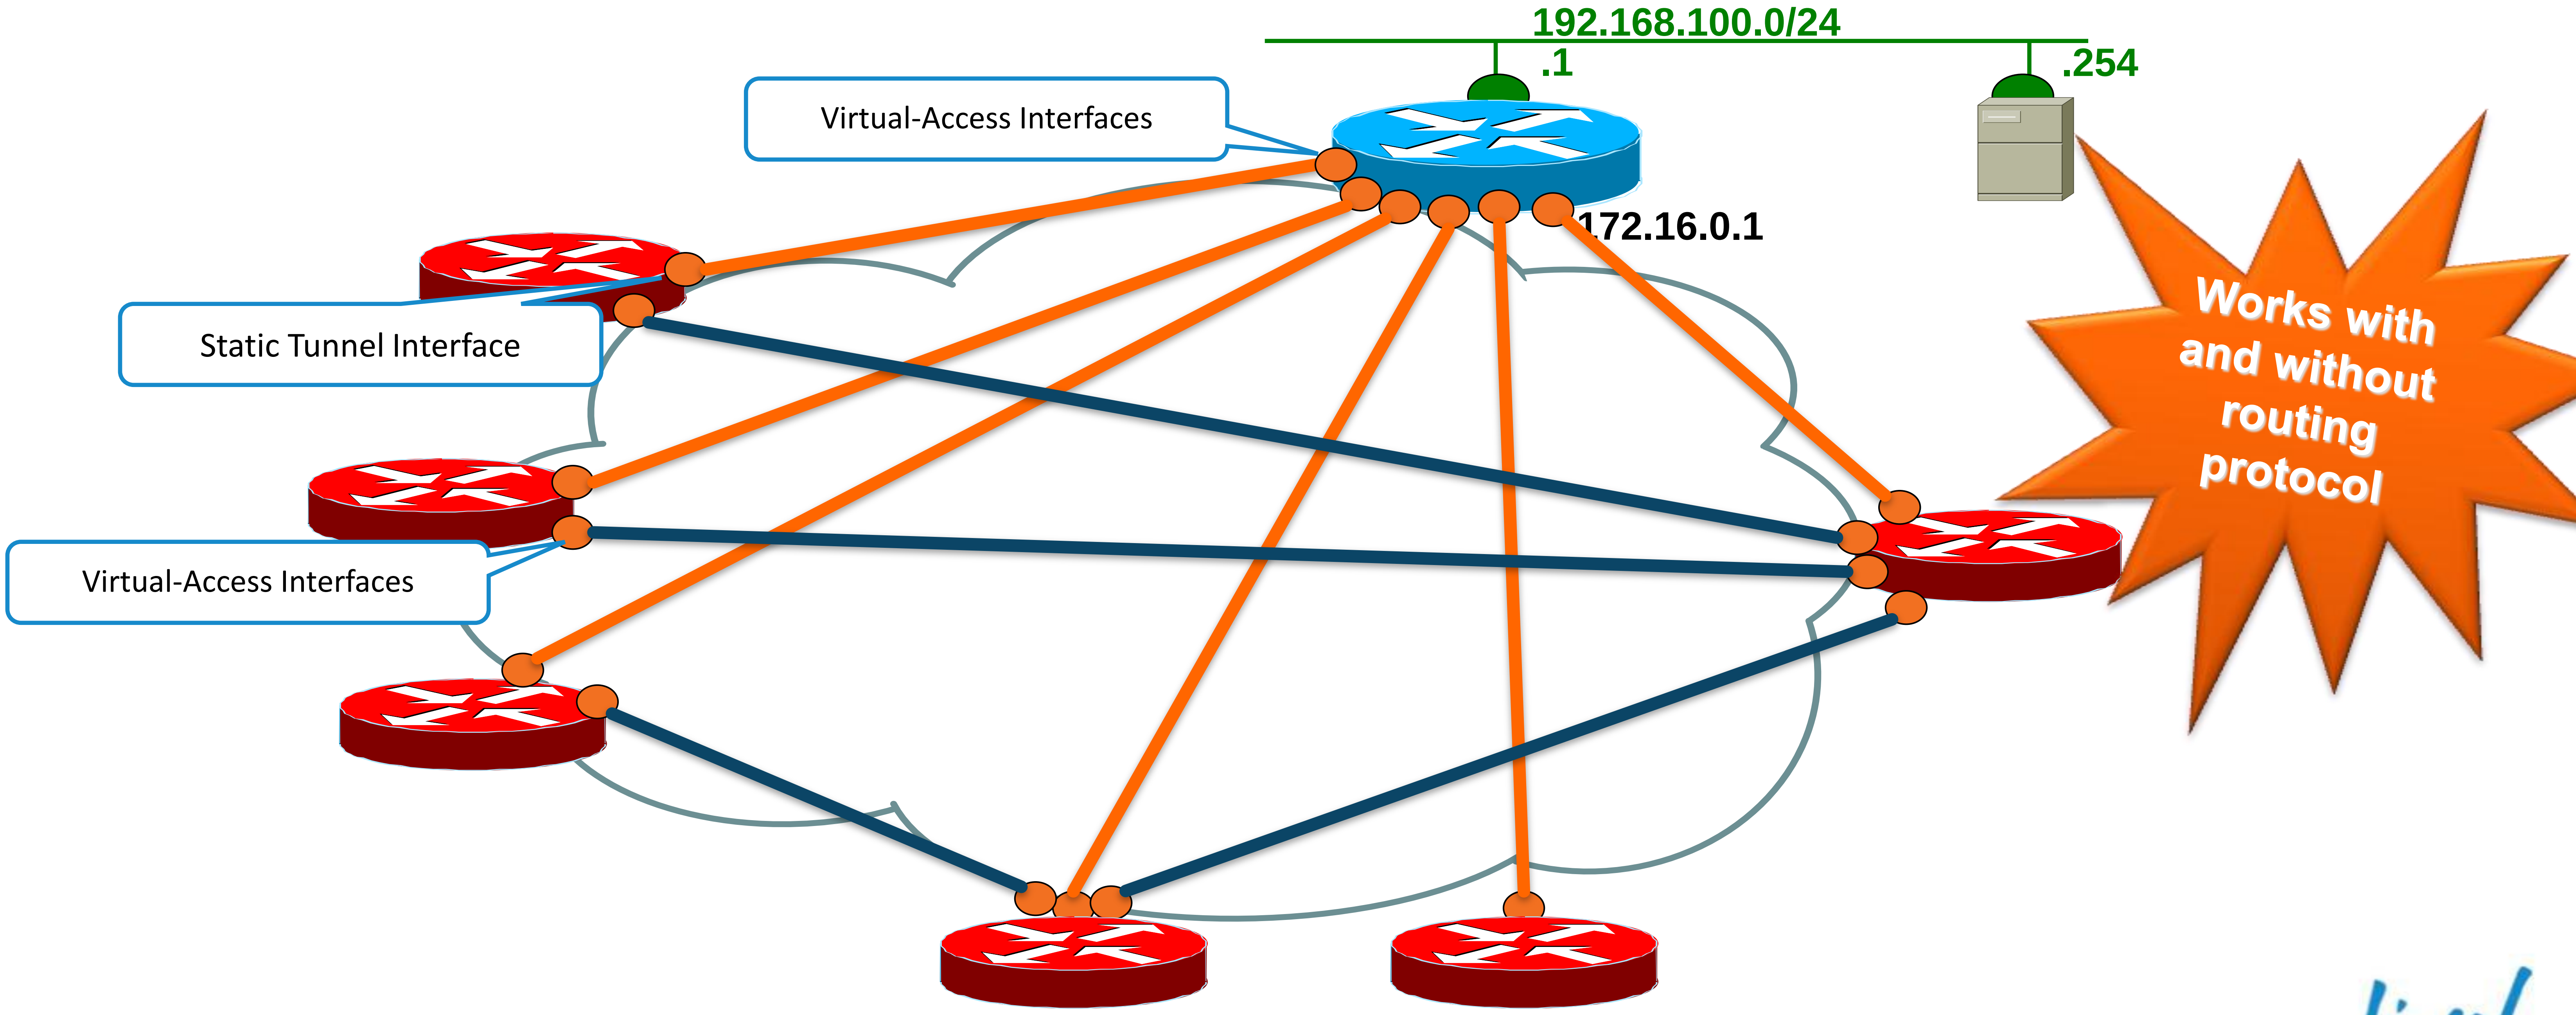

Example 3 – Flex Mesh

Hub Configuration


```
aaa new-model
aaa authorization network default group radius

crypto ikev2 profile default
 match identity remote fqdn domain cisco.com
 identity local dn
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint CA
 aaa authorization user cert list default
 virtual-template 1

ip route 192.168.0.0 255.255.0.0 Null0

router bgp 1
 neighbor Spokes peer-group
 neighbor Spokes remote-as 1
 bgp listen 10.0.0.0/8 peer-group Spokes
 redistribute static
```

Symmetric cert authentication

AAA Authorisation possible!

Routing via BGP

BGP Dynamic peering

```
interface-config=service-policy out PM
framed-ip=10.0.0.1
ipsec:route-set=interface
ipsec:route-accept=any
```

```
interface loopback0
 ip address 10.0.0.254

interface virtual-template1 type tunnel
 ip unnumbered loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
```

Example 3 – Flex Mesh

Spoke Configuration


```
aaa authorization network default local

crypto ikev2 profile default
 match certificate HUBMAP
 identity local fqdn Spoke1.cisco.com
 authentication local rsa-sig
 authentication remote rsa-sig
 pki trustpoint CA
 aaa authorization group cert list default default
 virtual-template 1

interface Tunnel0
 ip address negotiated
 tunnel source FastEthernet0/0
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 1
 tunnel destination 172.16.0.1
 tunnel protection ipsec profile default
```

Tunnel address and routes from config-exchange

Hub assigned address

No NHRP registration except if NAT/PAT needed

Shortcut switching

Each V-Access can host per peer policies (via AAA or v-template)

Shortcut tunnel template

```
interface virtual-template 1 type tunnel
 ip unnumbered tunnel 0
 ip nhrp network-id 1
 tunnel protection ipsec profile default
 ip nhrp shortcut virtual-template 1

router bgp 1
 neighbor Hub peer-group
 neighbor Hub remote-as 1
 neighbor 10.0.0.254 peer-group Hub
 network 192.168.0.0 255.255.255.248
```

Routing via BGP

Spoke specific subnet

Route Exchange Protocol Selection

Branch-Hub		Use case				
IKEv2 Recommended	Simple, large scale	Static (No redistribution IGP→IKE)	Simple branches (< 20 prefixes)	Identity-based route filtering	Lossy networks	High density hubs
BGP Recommended	Simple to complex, large scale	Dynamic (Redistribution IGP → BGP)	Complex branches (> 20 prefixes)	Powerful route filtering – not identity based	Lossy networks	High density hubs up to 350K routes
EIGRP **not recommended at large scale**	Simple to complex	Dynamic (Redistribution IGP → IGP)	Semi-complex branches (> 20 prefixes)	Intermediate route filtering – not identity based	Lossless networks (very rare)	< 5000 prefixes at hub

Hub-Hub	Use case		
BGP Recommended	Large amount of prefixes (up to 1M)	Road to scalability	Powerful route filtering
IGP (EIGRP, OSPF)	< 5000 prefixes total	Perceived simplicity	

More Information

- Main FlexVPN documentation:

  http://www.cisco.com/en/US/docs/ios-xml/ios/sec_conn_ike2vpn/configuration/15-2mt/sec-flex-vpn-15-2mt-book.html

- FlexVPN RADIUS Attributes:

  http://www.cisco.com/en/US/docs/ios-xml/ios/sec_conn_ike2vpn/configuration/15-2mt/sec-apx-flex-rad.html

IPv6 First Hop Security

IPv6 First Hop Security

What

- A range of functions to protect the operation of IPv6 first hop protocols
- Preventing man-in-the-middle layer 2 access attacks
- Preventing layer 2 Denial of Service layer 2 access attacks
- The same best practice equivalents of IPv4 First Hop Security

IPv6 First Hop Security

Why

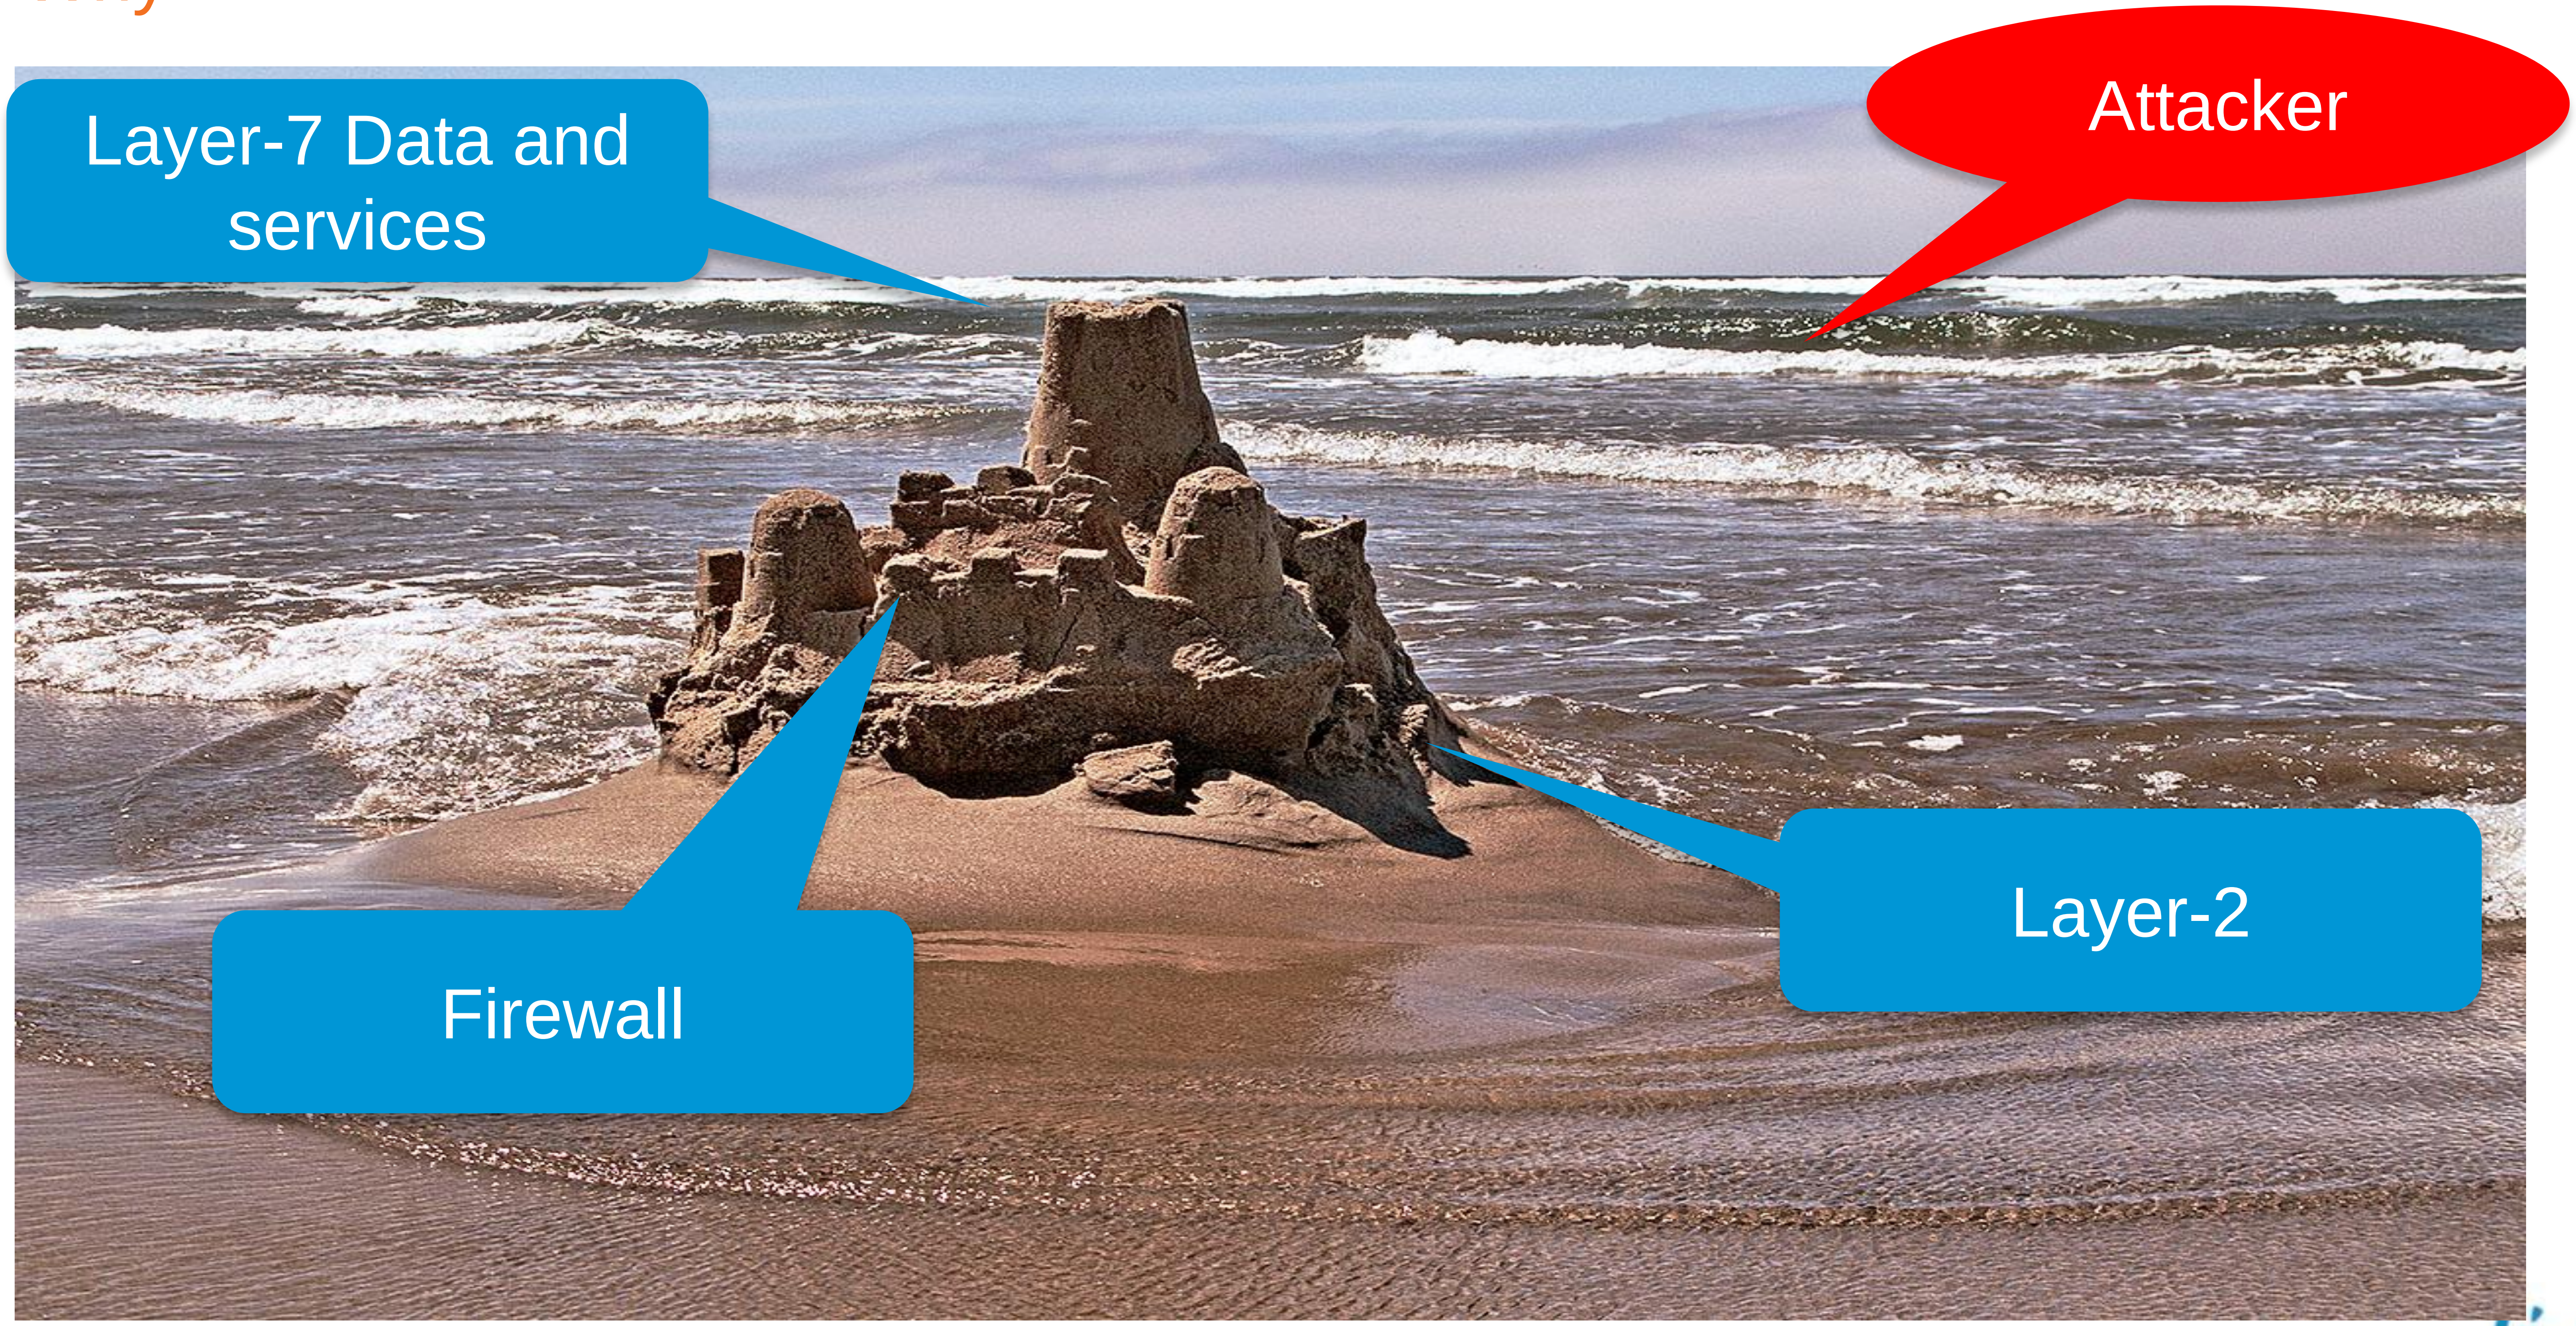

Courtesy of Curt Smith

IPv4 Vulnerabilities and Countermeasures

Catalyst Integrated Security Features (CISF)

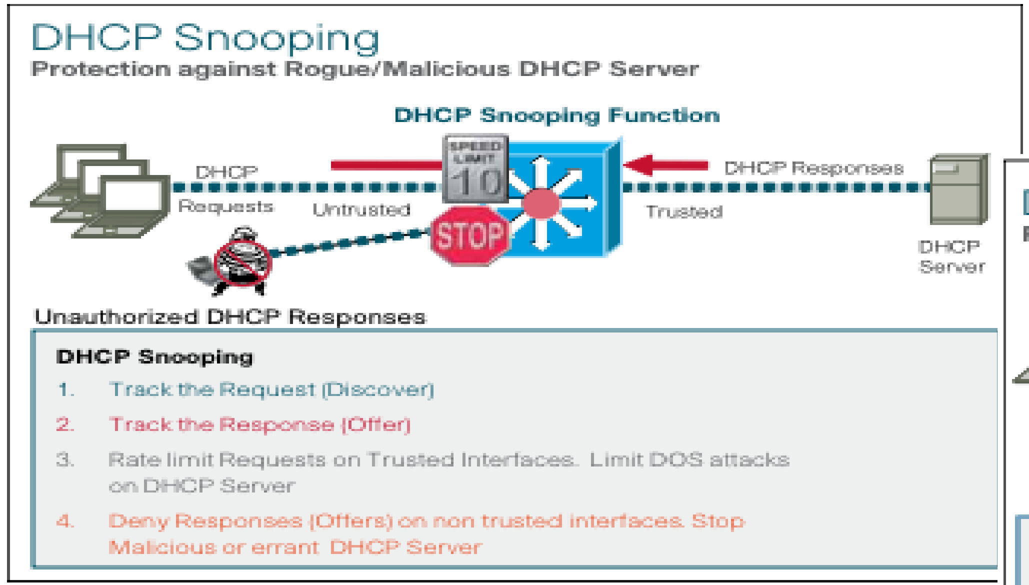

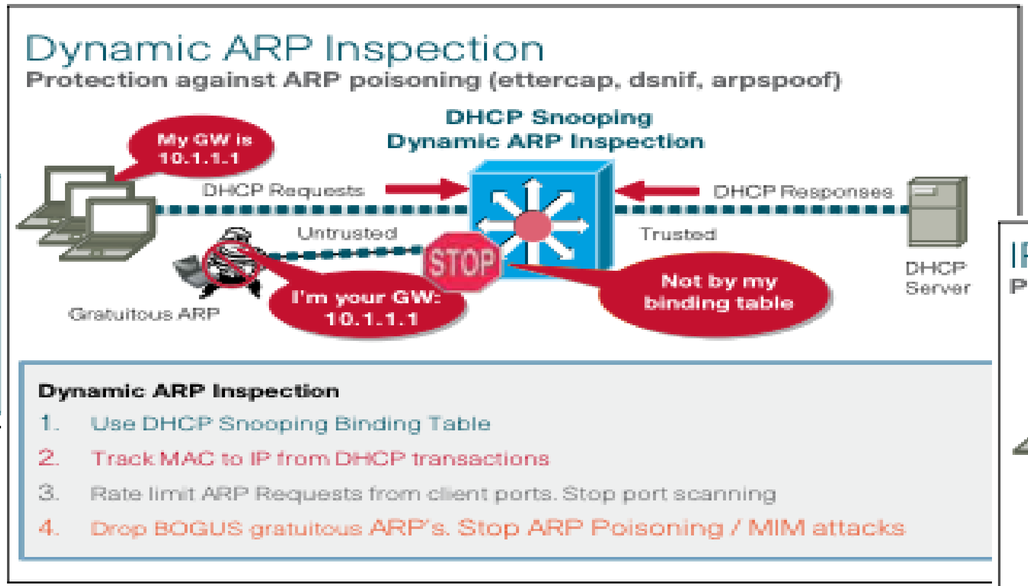

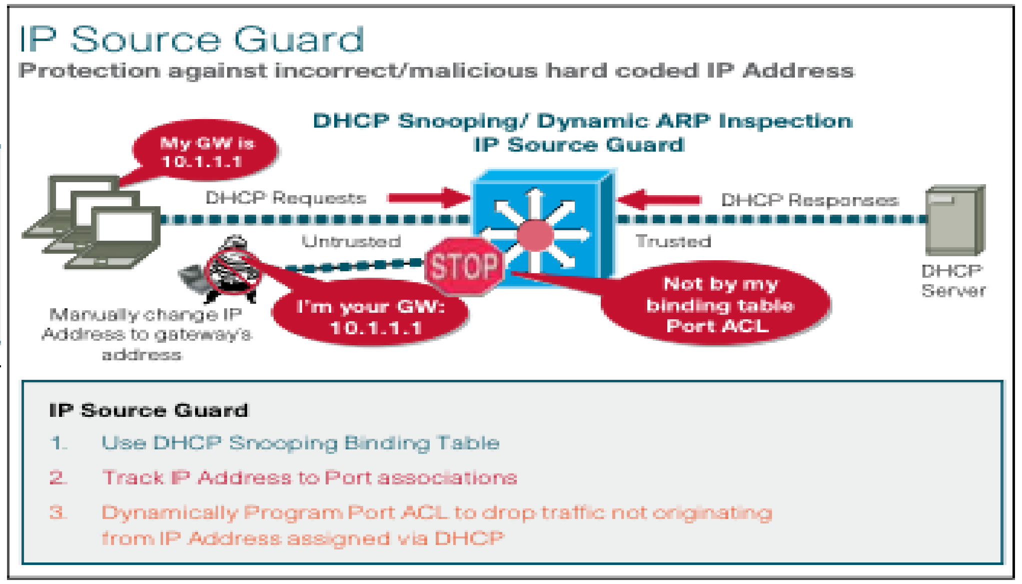

For more info: http://www.cisco.com/web/strategy/docs/gov/turniton_cisf.pdf

IPv6 Link Operations

What are Link Operations?

Operations contained within the link boundaries, necessary for a node to communicate with its neighbors, including the link exit points.

It encompasses:

- Address configuration parameters
- Address initialization
- Address resolution
- Default gateway discovery
- Local network configuration
- Neighbor reachability tracking

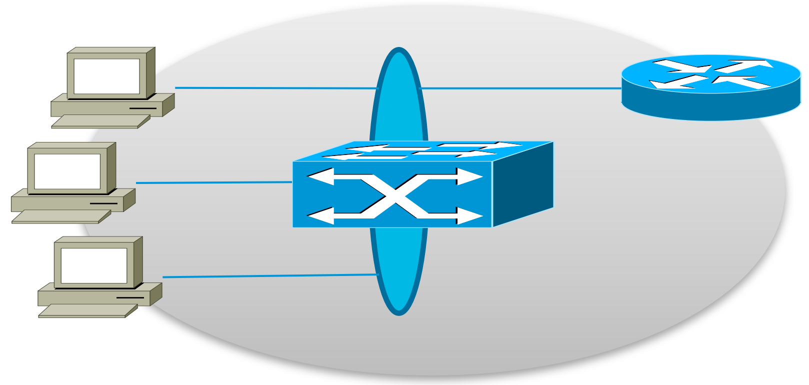

"A link"

Neighbour Discovery Protocol (NDP)

NDP (ARP replacement in IPv6)

- Discover other hosts & routers on local network
- Incorporates many features from older link-layer protocols
- Makes extensive use of IPv6 multicast addresses
- Operates using ICMPv6

NDP is also the protocol used to learn information:

- About other hosts
  - Address Resolution*
  - Duplicate Addresses
  - Neighbour Unreachable
  - Next Hop
- About routers
  - Discovery
  - Network Prefix
  - Network Parameters
  - Autoconfiguration

NDP and SLAAC

All can be used as attack vectors

- Primary ICMPv6 NDP Messages
  - Neighbour solicitation (NS)
  - Neighbour advertisements (NA)
  - Router solicitation (RS)
  - Router advertisements (RA)
  - Neighbour Unreachability Detection (NUD)
  - Duplicate Address Detection (DAD)
  - Redirects
- SLAAC
  - IPv6 Stateless Address Auto Configuration (SLAAC)

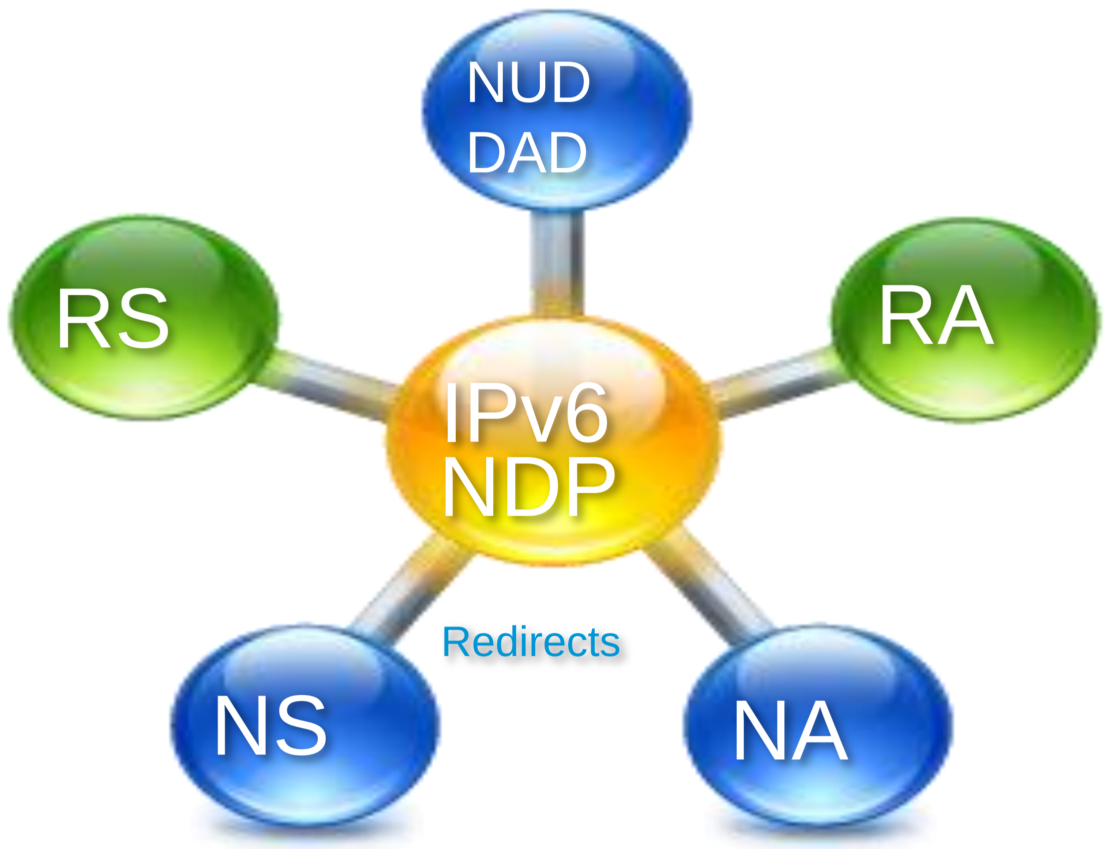

IPv6 First Hop Security

How

IPv6 FHS

IPv6 Binding Integrity Guard *

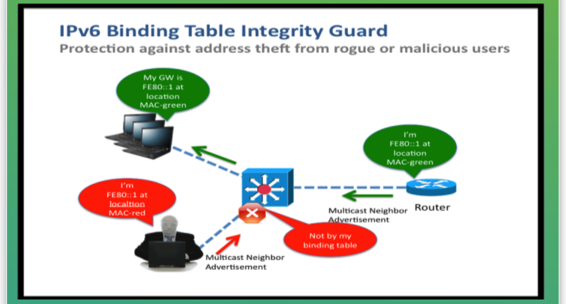

- Integrity protection for Neighbour Binding Cache & FHS binding table
- Protection against IPv6 address theft

IPv6 RA Guard

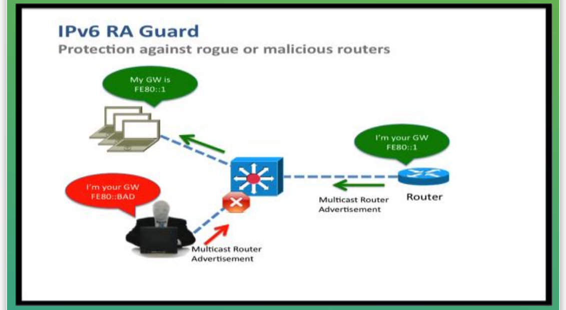

- Protection against MiM Attacks
- Protection against rouge or malicious Router Advertisement

IPv6 DHCP Guard

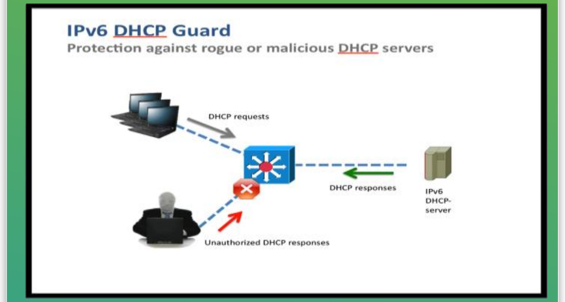

- Protection against MiM & DoS attacks
- Rejects invalid DHCP Offers

IPv6 Source Guard

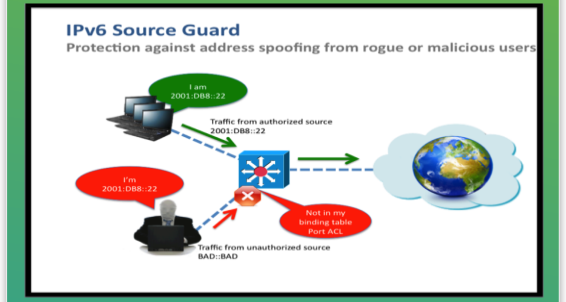

- Validate source address or prefix
- Protects against source address spoofing

IPv6 Destination Guard

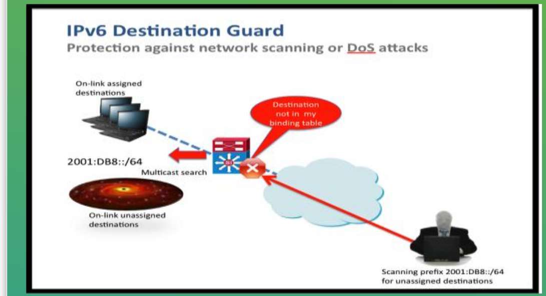

- Validates destination address of IPv6 traffic reaching the link
- Protects against scanning or DoS attacks

* Previously referred to as Address Glean/Watch

Rogue Router Advertisement

The Attack

RA w/o Any Authentication Gives Exactly Same Level of Security as DHCPv4 (None)

- Router Advertisements contains:
  - Prefix to be used by hosts
  - Data-link layer address of the router
  - Miscellaneous options: MTU, DHCPv6 use, …

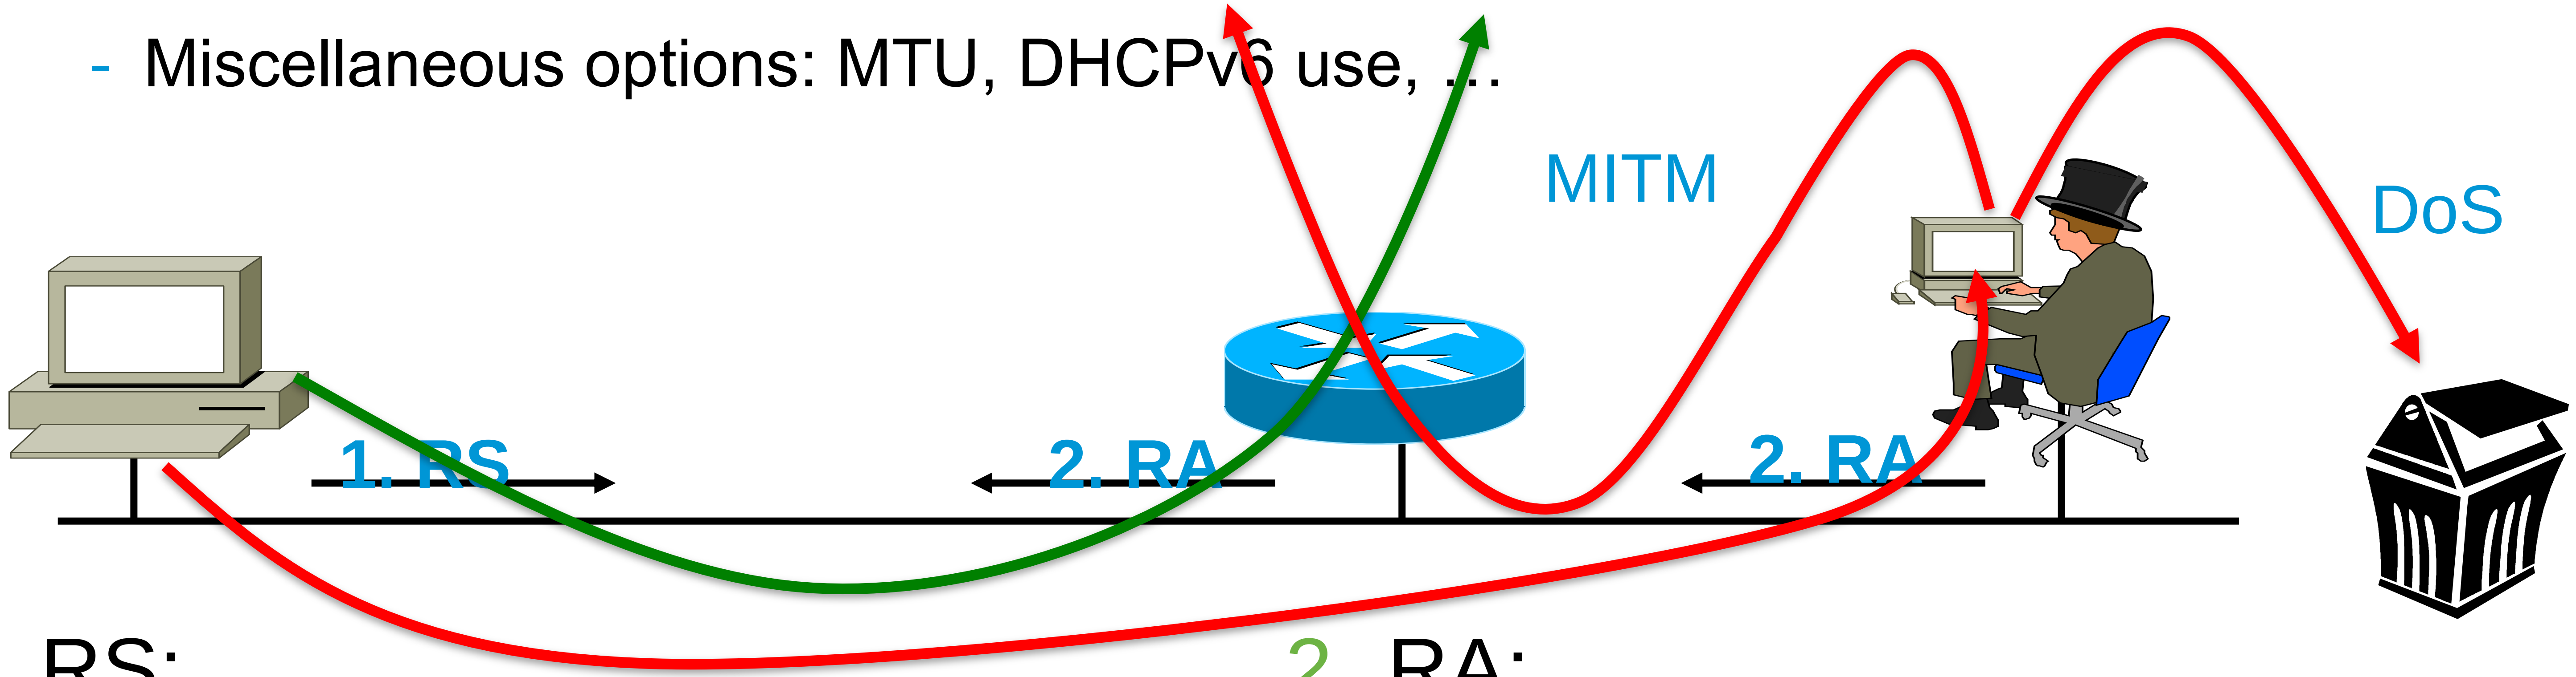

1. RS:

   Data = Query: please send RA

2. RA:

   Data= options, **prefix**, lifetime, A+M+O flags

Rogue Router Advertisement

Effect

- Devastating:
  - Denial of service: all traffic sent to a black hole
  - Man in the Middle attack: attacker can intercept, listen, modify unprotected data
- Also affects legacy IPv4-only network with IPv6-enabled hosts
- Most of the time from non-malicious users
- Requires layer-2 adjacency (some relief…)
- The major blocking factor for enterprise IPv6 deployment

Rogue Router Advertisement

Mitigation Techniques

Where	What
Routers	Increase “legal” router preference
Hosts	Disabling Stateless Address Autoconfiguration
Routers & Hosts	SeND “Router Authorisation”
Switch (First Hop)	Host isolation
Switch (First Hop)	Port Access List (PACL)
Switch (First Hop)	RA Guard

Secure Neighbour Discovery (SeND)

RFC 3971

- RFC 3972 Cryptographically Generated Addresses (CGA)
  - IPv6 addresses whose interface identifiers are cryptographically generated from node public key
- SeND adds a signature option to Neighbour Discovery Protocol
  - Using node private key
  - Node public key is sent in the clear (and linked to CGA)
- Very powerful
  - If MAC spoofing is prevented (port security)
  - No authentication, it does not replace 802.1X
  - But, not a lot of implementations: Cisco IOS 12.4(24)T, Linux, some H3C, 3rd party

Mitigating Rogue RA

RFC 6101

- **Port ACL** blocks all ICMPv6 RA from hosts

```
interface FastEthernet0/2
  ipv6 traffic-filter ACCESS_PORT in
  access-group mode prefer port
```

- **RA-guard lite** (12.2(33)SXI4 & 12.2(54)SG):

```
interface FastEthernet0/2
  ipv6 nd raguard
  access-group mode prefer port
```

- **RA-guard** (12.2(50)SY)

```
ipv6 nd raguard policy HOST device-role host
ipv6 nd raguard policy ROUTER device-role router
ipv6 nd raguard attach-policy HOST vlan 100
interface FastEthernet0/0
  ipv6 nd raguard attach-policy ROUTER
```

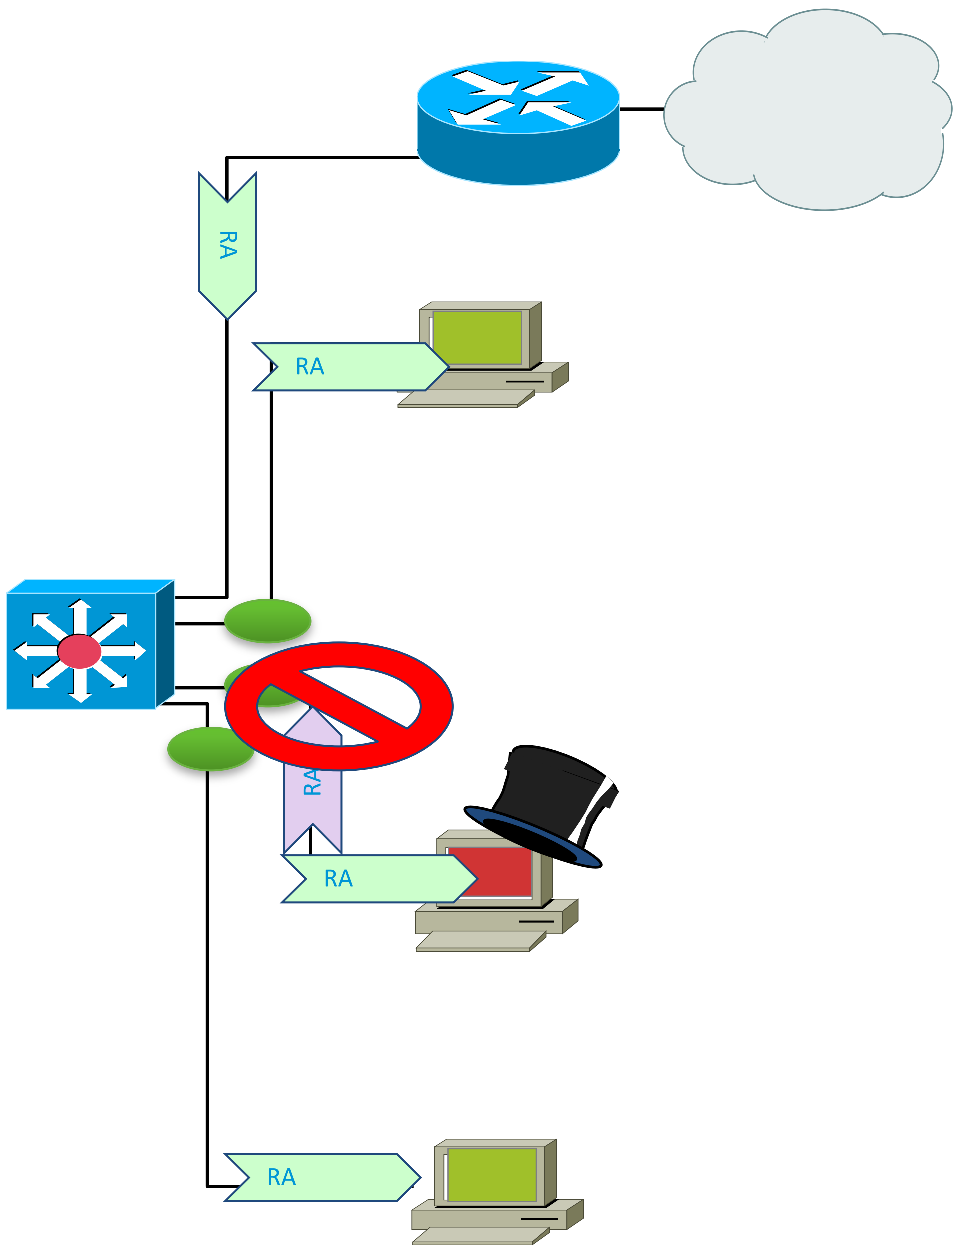

Mitigating Rogue RA

Dealing with Fragmentation

- Extension headers chain can be so large than it is fragmented!
- RFC 3128 is not applicable to IPv6
- Layer 4 information could be in 2nd fragment

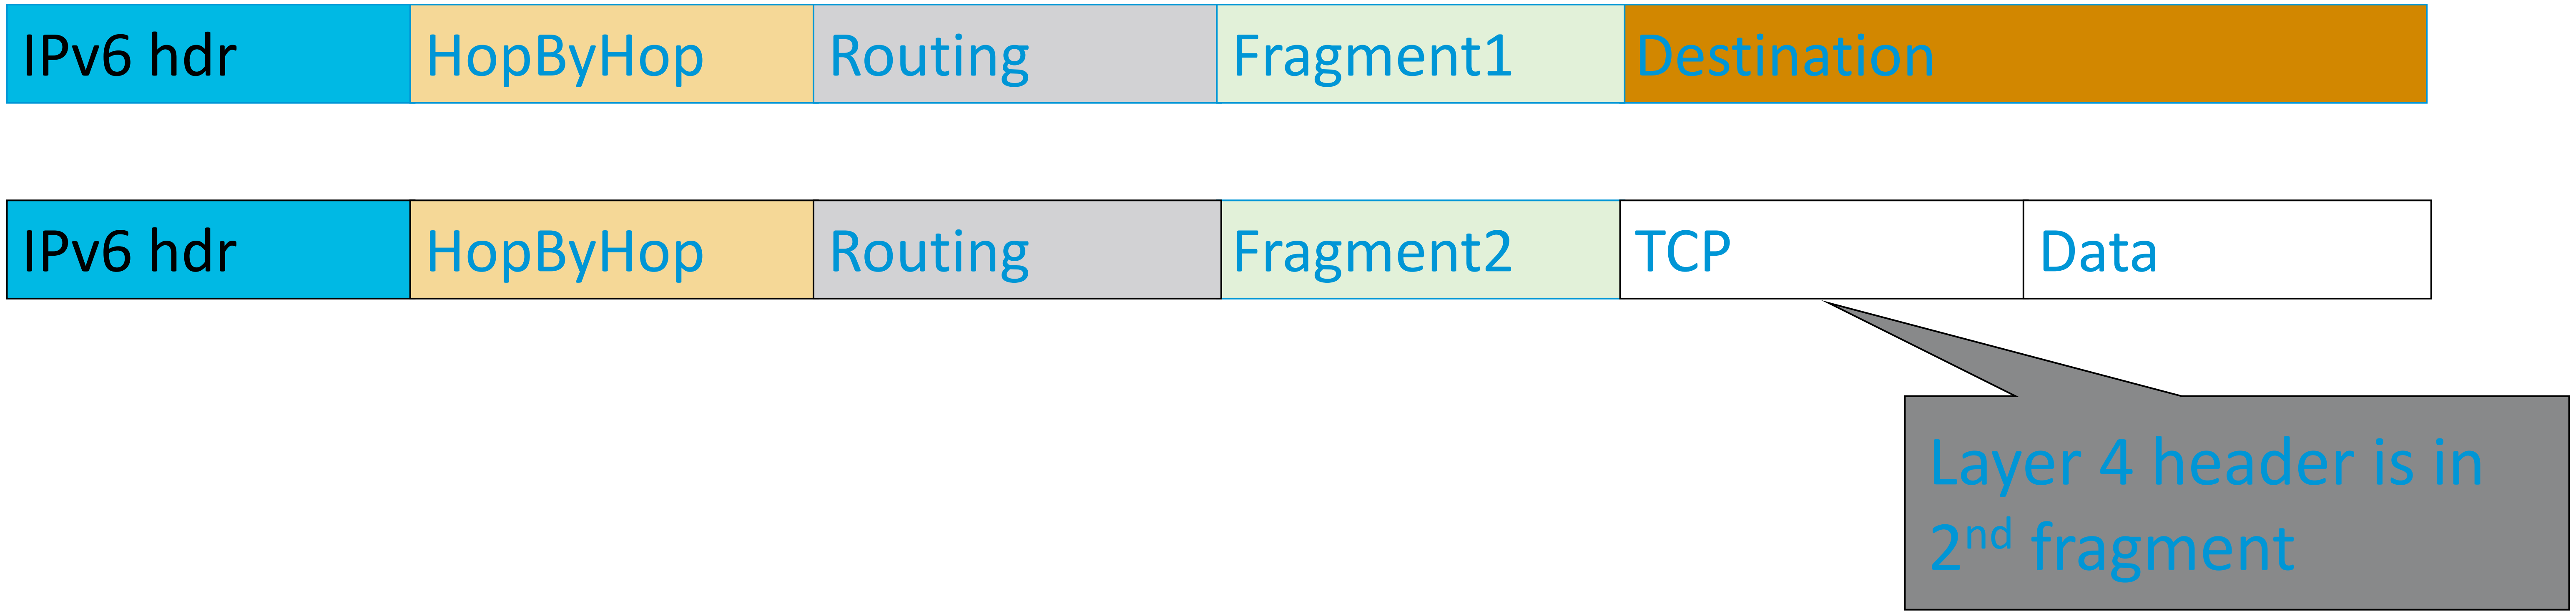

Parsing the Extension Header Chain

Fragments and Stateless Filters (RA Guard)

- RFC 3128 is not applicable to IPv6, extension header can be fragmented
- ICMP header could be in 2nd fragment after a fragmented extension header
- RA Guard works like a stateless ACL filtering ICMP type 134
- THC fake_router6 –FD implements this attack which bypasses RA Guard

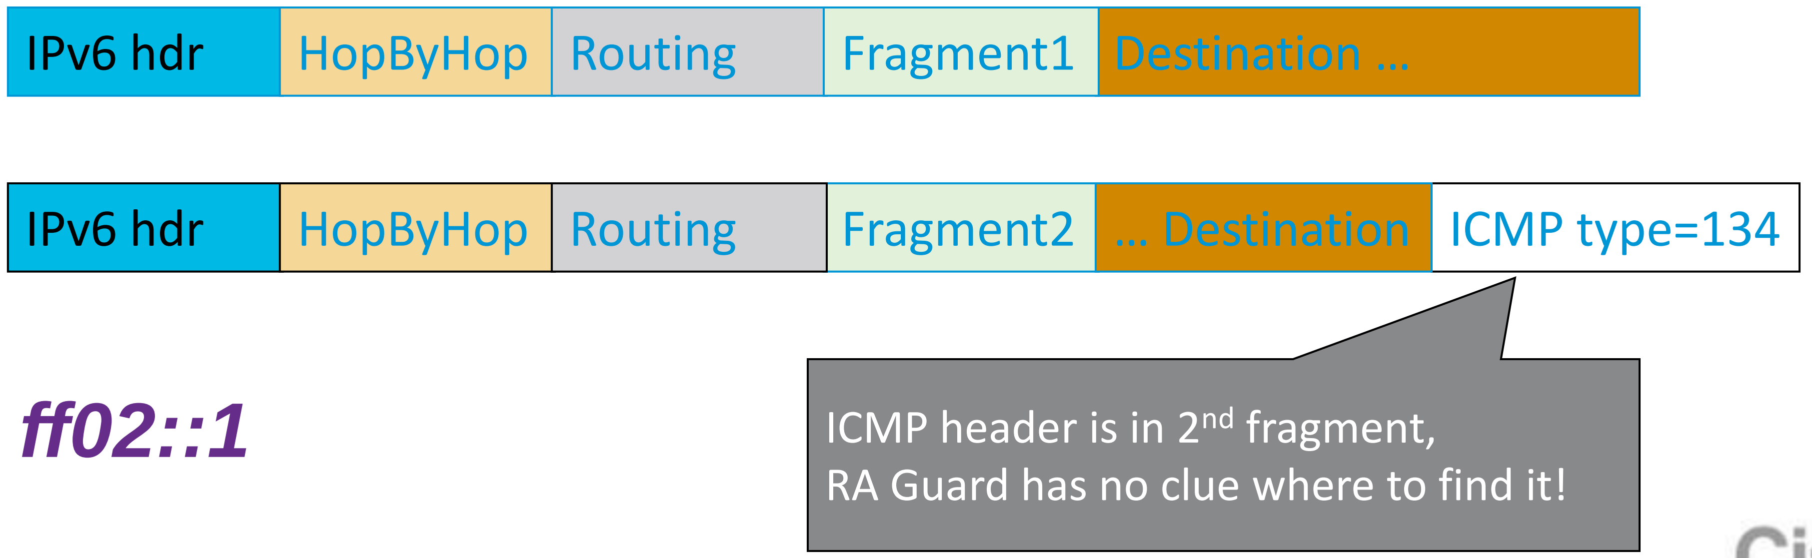

Partial work-around:
block all fragments sent to ff02::1

NDP Spoofing

The Attack

- Pretty much like RA: no authentication
  - Any node can ‘steal’ the IP address of any other node
  - Impersonation leading to denial of service or MITM
- Requires layer-2 adjacency
- IETF SAVI Source Address Validation Improvements (work in progress)

NDP Spoofing

Mitigation Techniques

Where	What
Routers & Hosts	configure static neighbour cache entries
Routers & Hosts	Use CryptoGraphic Addresses (SeND CGA)
Switch (First Hop)	Host isolation
Switch (First Hop)	Address watch • Glean addresses in NDP and DHCP • Establish and enforce rules for address ownership

SAVI

- If a switch wants to enforce the mappings < *IP address, MAC address*> how to learn them?
- Multiple source of information
  - SeND: verify signature in NDP messages, then add the mapping
  - DHCP: snoop all messages from DHCP server to learn mapping (same as in IPv4)
  - NDP: more challenging, but '*first come, first served*'
    - The first node claiming to have an address will have it

Mitigation : Binding Integrity Guard

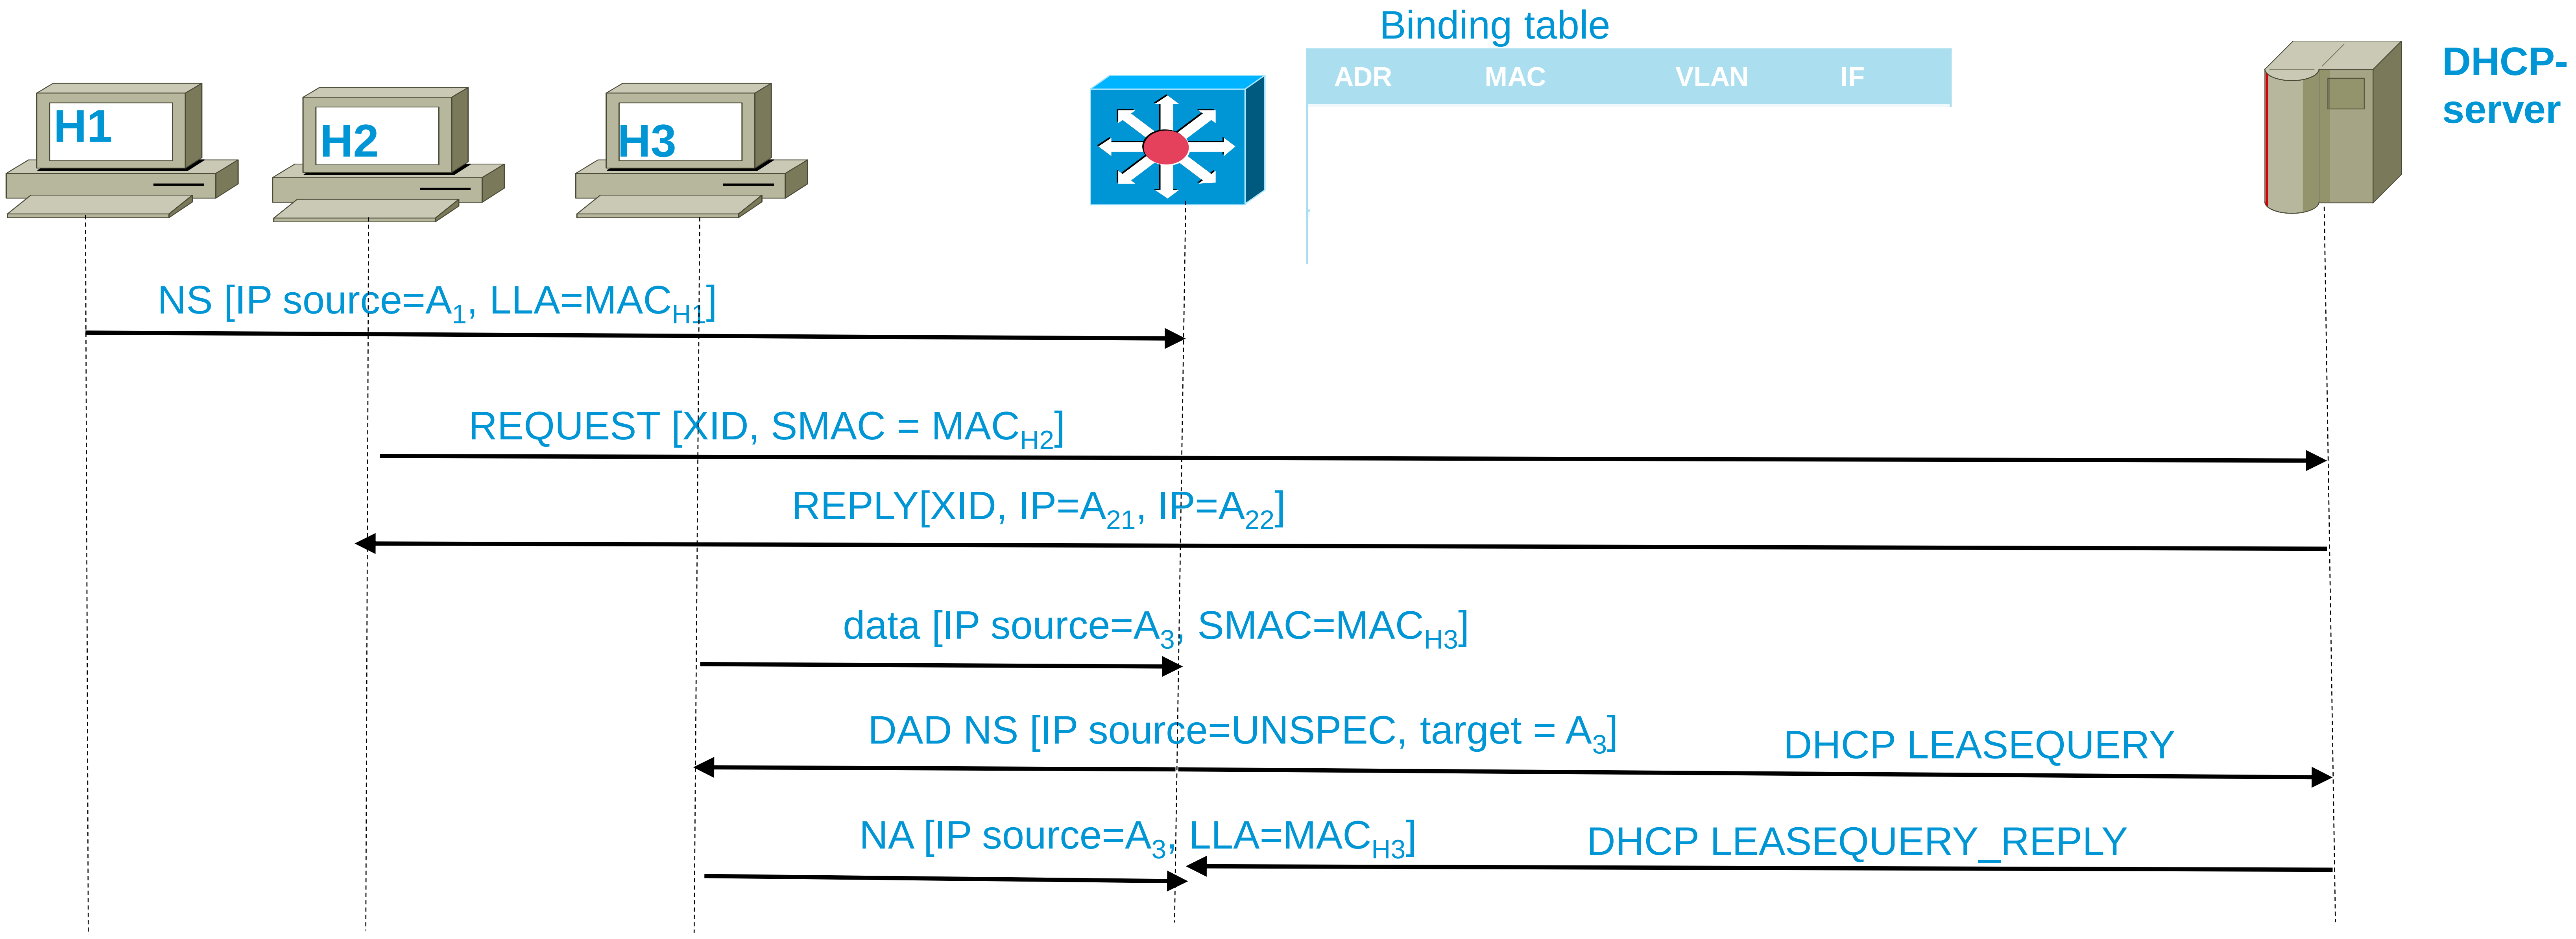

Then enforce the binding table in the TCAM

Exhausting the Neighbour Cache

The Attack

- Remote router CPU/memory DoS attack if aggressive scanning
  - Router will do Neighbour Discovery... And waste CPU and memory
- Local router DoS with NS/RS/…

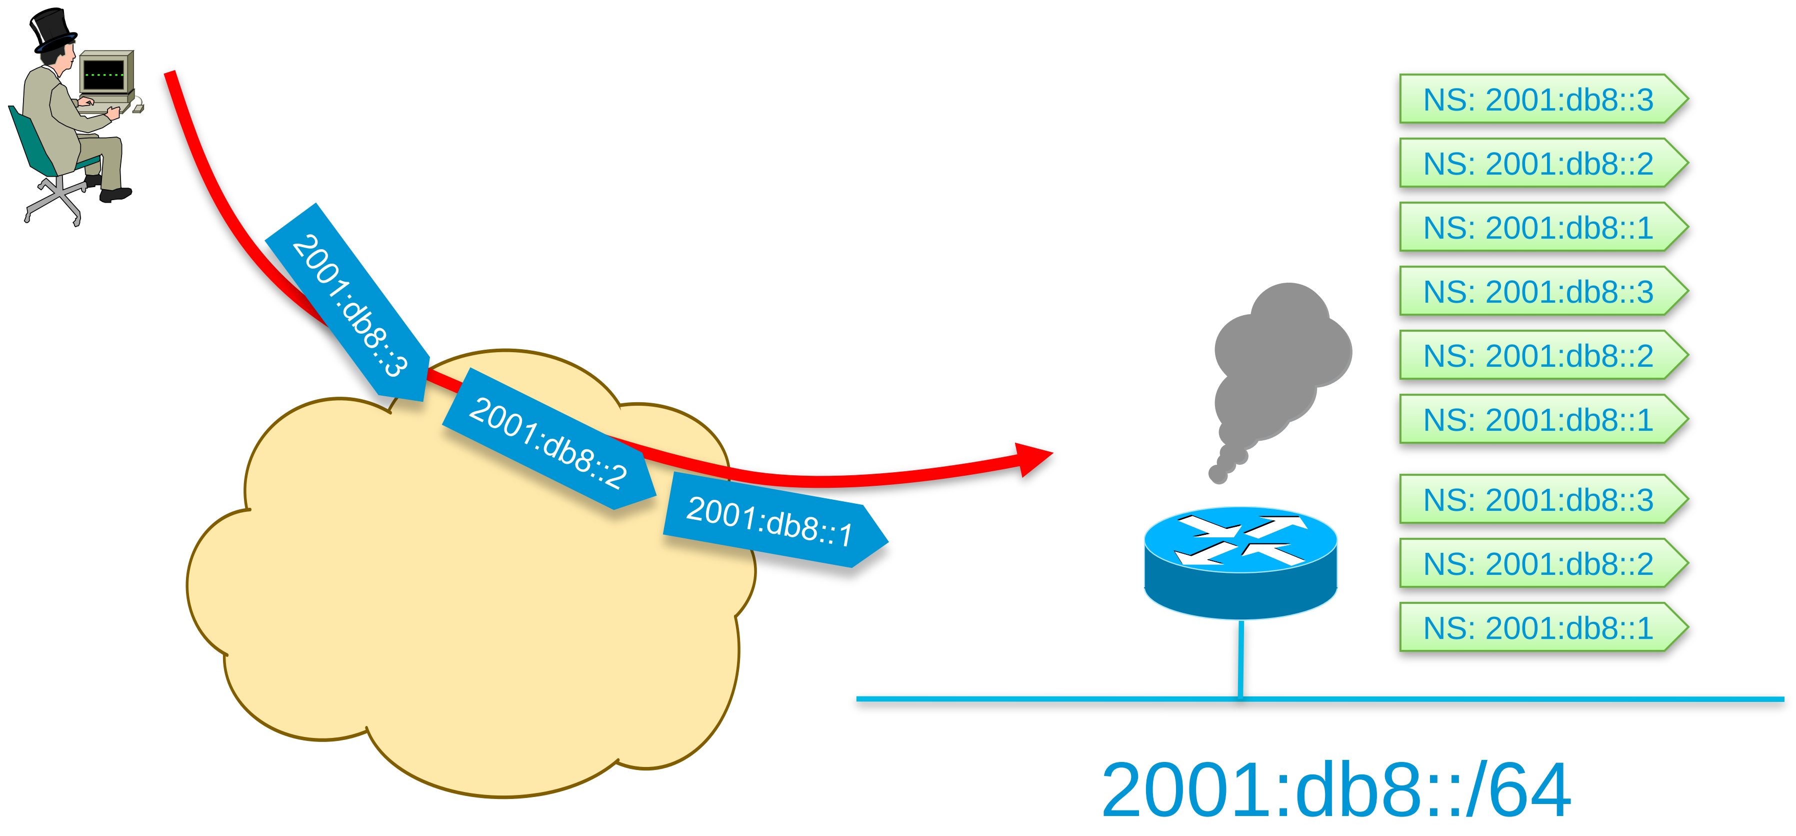

Exhausting the Neighbour Cache

Mitigation Techniques

- Mainly an implementation issue
    - Rate limiter on a global and per interface
    - Prioritise renewal (PROBE) rather than new resolution
    - Maximum Neighbour cache entries per interface and per MAC address
    - Since 15.1(3)T: **`ipv6 nd cache interface-limit`**
    - Or IOS-XE 2.6: **`ipv6 nd resolution data limit`**
- **Destination guard** (12.2(50)SY): drop all packets not matching an entry in the integrity binding table
- **Internet edge/presence**: a target of choice
    - Ingress ACL permitting traffic to specific statically configured (virtual) IPv6 addresses only
    - ⇒Allocate and configure a /64 but uses addresses fitting in a /120 in order to have a simple ingress ACL

Exhausting the Neighbour Cache

Simple Fix

- Ingress ACL allowing only valid destination and dropping the rest
- NDP cache & process are safe
- Requires DHCP or static configuration of hosts

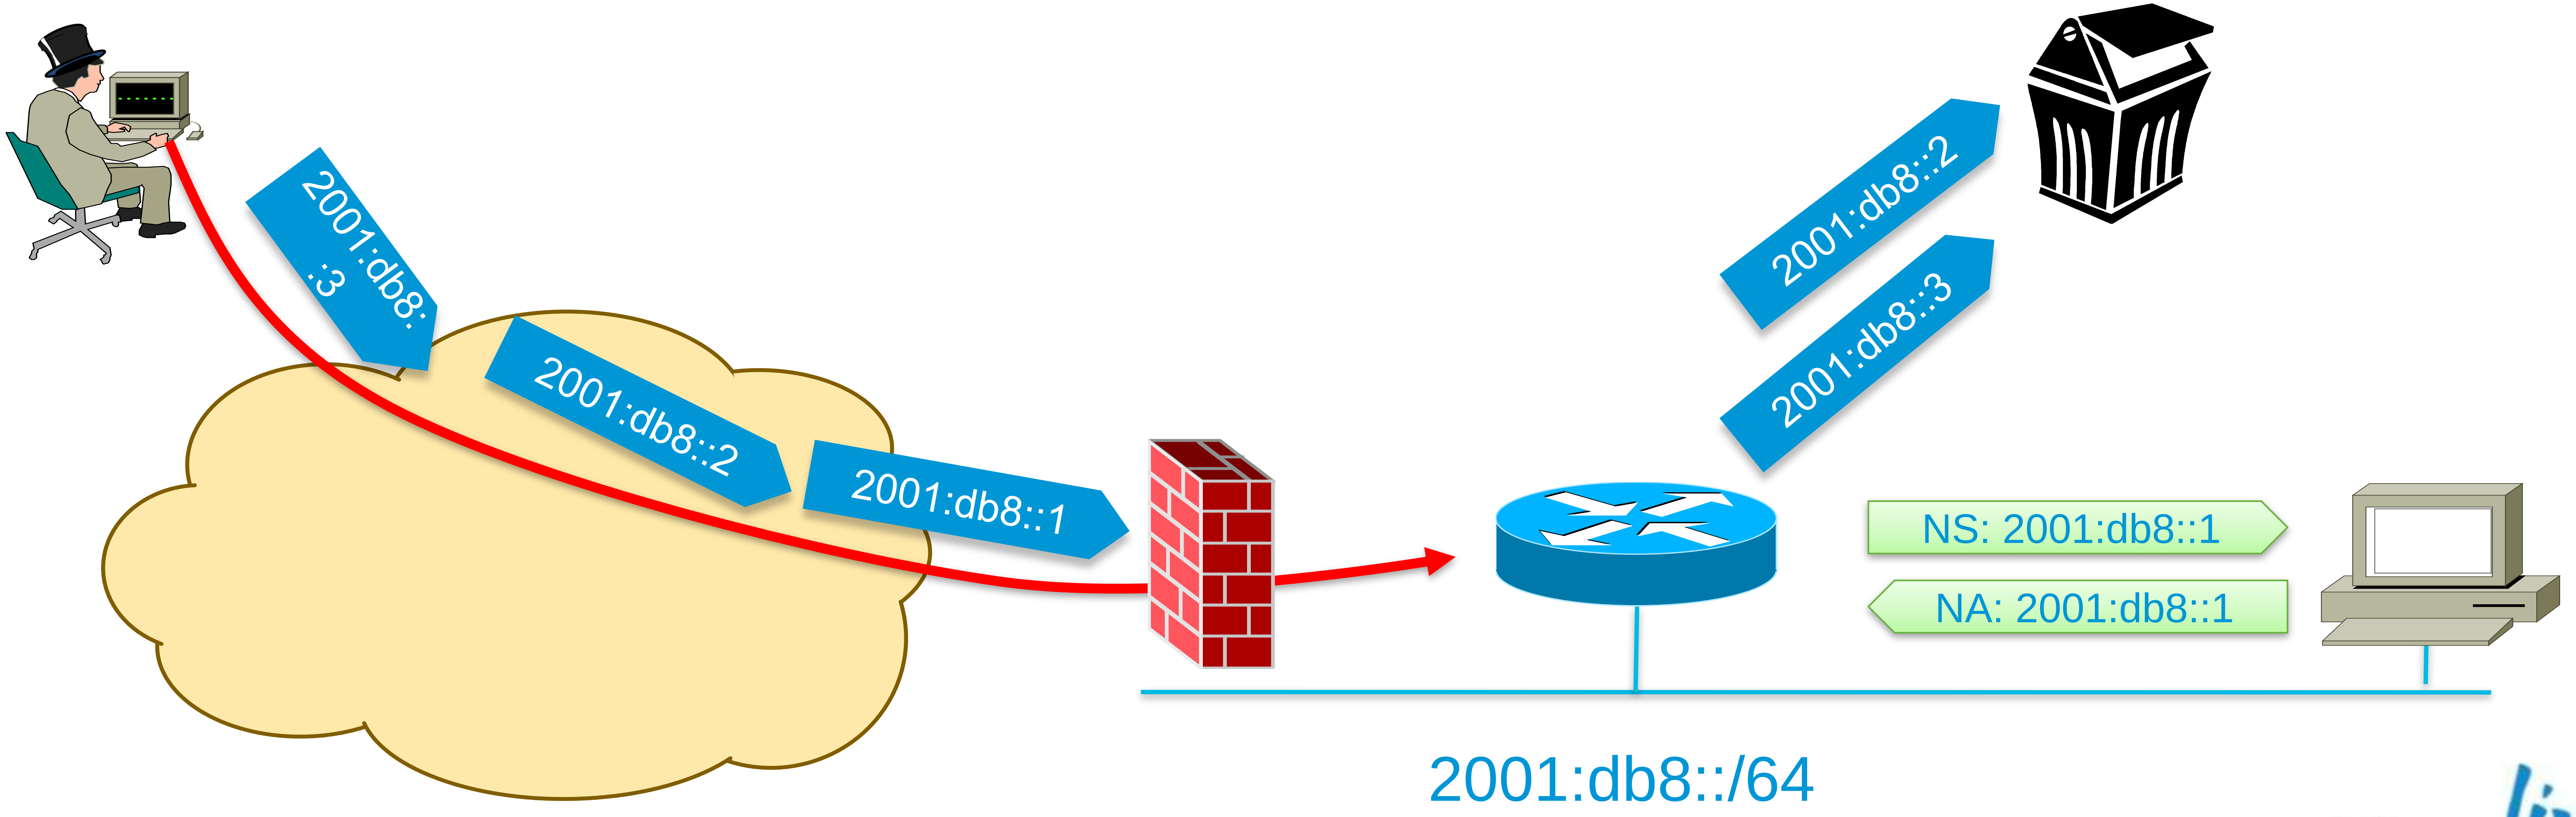

Key Takeaways

- Without a secure layer-2, there is no upper layer security
- Rogue Router Advisement is the most common threat
- Mitigation techniques
  - Host isolation
  - Secure Neighbour Discovery: but not a lot of implementations
  - SAVI-based techniques: discovery the ‘right’ information and dropping RA/NA with wrong information
  - Last remaining issue: (overlapped) fragments => drop all fragments…
- Neighbour cache exhaustion
  - Use good implementation
  - Expose only a small part of the addresses and block the rest via ACL

More Information

http://www.cisco.com/go/ipv6

http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-roadmap.html

http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-first-hop-security.html

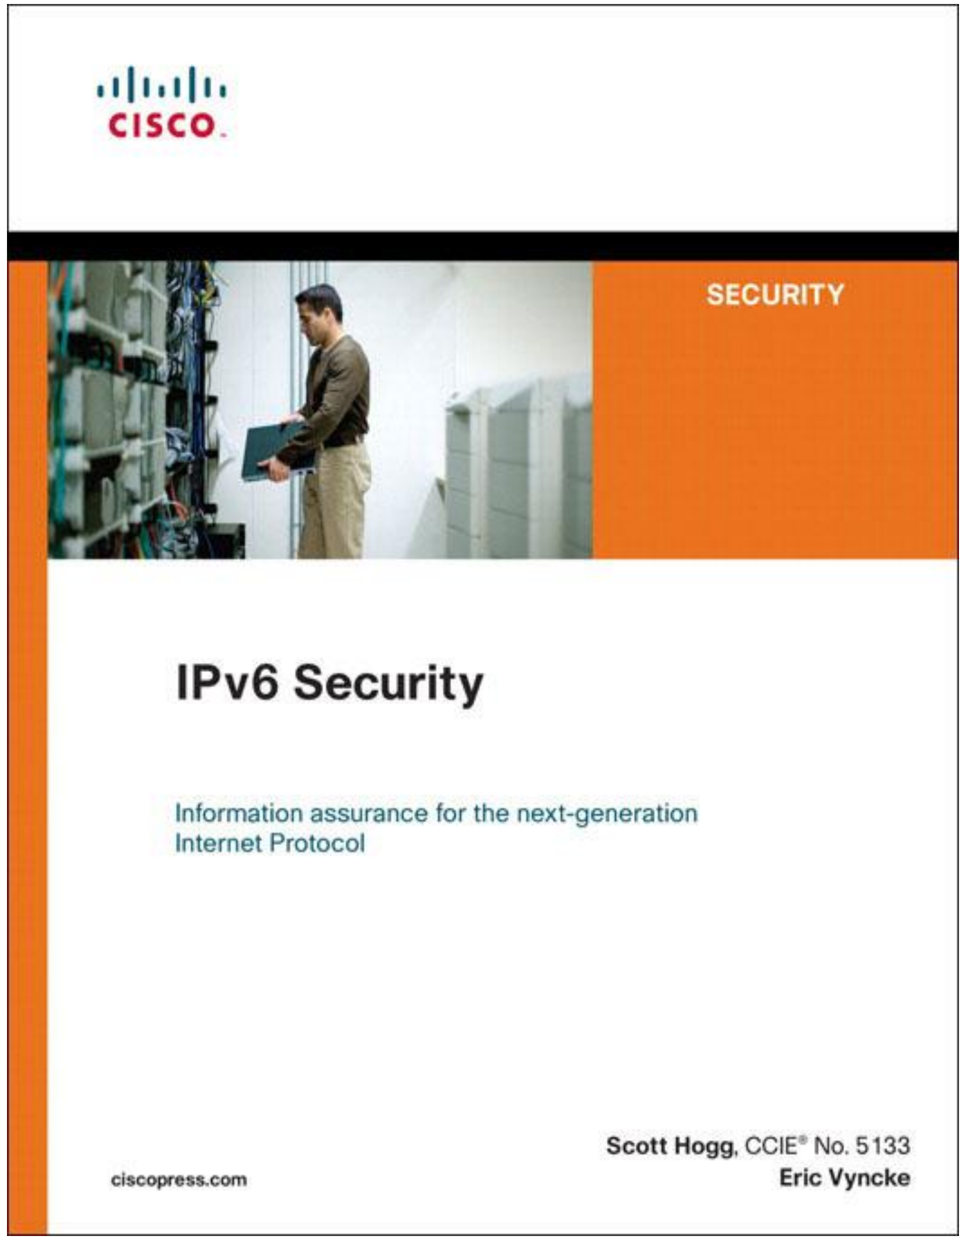

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 8 March 12:00pm-2:00pm

Don’t forget to activate your Cisco Live 365 account for access to all session material, communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.
www.ciscoliveaustralia.com/portal/login.ww