# Routing Security Strategy

**November 2009**

**Janne Östling**

# The Crystal Ball

Predicting Future Trends

# Networks in 2014

- Access / core bandwidth up by factor ~10
- Address exhaustion:
  - IPv6 widely deployed
  - Carrier-grade NAT widely deployed
- Any device, anywhere, any time
  - Mobility, diverse client devices
- Clients remain “intelligent”
  - Multiple OSs, running applications
  - Still “untrusted” (can run applets, etc)
- Consolidation of data centers in the Enterprise (cloud)
- P2P more widely deployed

# Mobile Networks Will Become Like Fixed Networks

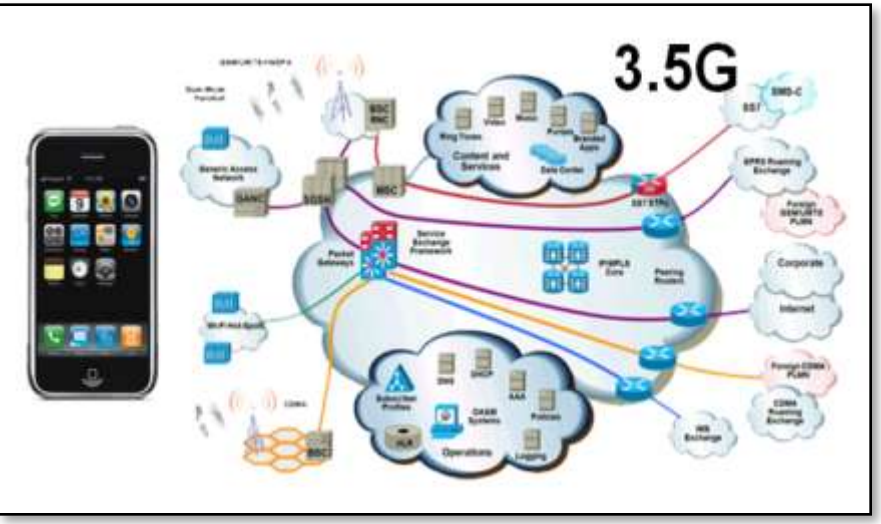


- Powerful clients
  - CPU, memory, bandwidth
- Same threats
  - Malware, bots, ...
- Firewalling towards mobile devices will remain in place
  - → This will change protocols (Skype)

# Complexity Is Increasing

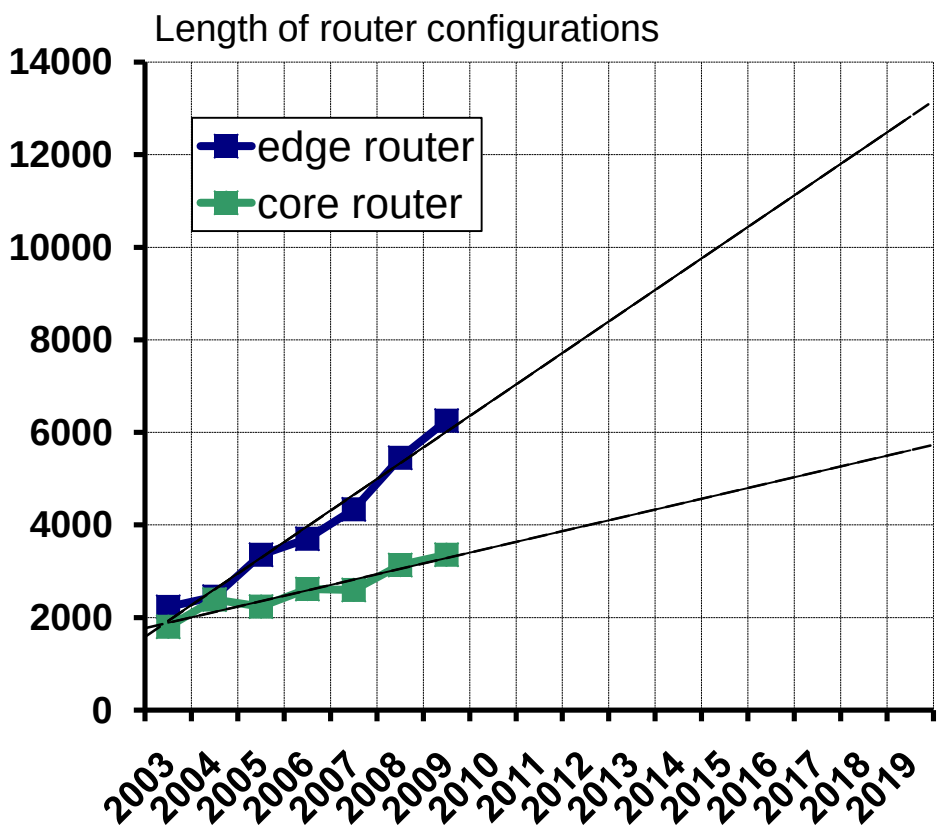


- Applications, OSs, networks will become more complex
  - More lines of code
  - More functions
- Lack of visibility
  - More “covert channels” (HTTP tunneling)
  - More encryption (SSL, etc)
- Double and triple NAT
- Home networks more complex
  - New endpoints: STB, Playstations, ...
- Router resource management and protection increasingly important

# Threats Evolution

- Economically based fraud will persist / increase
  - Potential other motivations: Political, military, ...
- Malware will use SSL, use covert channels
  - Will look like a banking application, or:
  - Piggyback on legitimate connections, as covert channel
- Security measures will increase
  - But so will hackers' capabilities
- Increasing internal threats: sabotage, misconfigurations

→ Basically, no change, except more obfuscation (SSL)

# The Future of the Data Plane

Packet Forwarding and Handling

# Data Plane Vision

"No data plane packet will have the potential to cause harm to the network"

# Data Plane Strategy

- Capability to view and control IPv4/v6 flows:
  - Monitor
  - Filter
  - Re-mark (QoS)
- Service Routing
  - Dynamically, policy based
  - Route flows to service points
- Based on:
  - L3/L4 parameters
  - Subscriber policy
- All at line rate

# Data Plane Execution, #1

- Identify and list known exceptions to the rule

  Where transit packets can cause harm

  E.g.: TTL expiry, IPv4 options, IPv6 HbH, prec 6/7, ...

- Audit all SP platforms (define all platforms):

| Platform | Issue | vulnerable by default | commands to fix/reduce issue | remaining issue |
|---|---|---|---|---|
| GSR (IOS) | TTL expiry | yes | CoPP (partial fix only) | LC can be DoSed |
| 7600 | TTL expiry | yes | mls-rate limiter | none |

"yes" means customer must do something; could we make this default?

if issue can't be fixed: Need to develop a solution

If no complete fix available, need to improve solution

# Data Plane Execution, #2

- Define minimum data plane feature set for core and edge platforms (and core and edge line cards), for IPv4/v6

  uRPF, ACLs, NetFlow, QoS, ... (define in detail)

- Audit all SP platforms

  Define scenarios (e.g., threat model, attack forms, etc)

  Develop features where required

# The Future of the Control Plane


Routing and Switching

# Control Plane Vision

"From the outside of a network, it is not possible to interfere with any control plane process"

"All control plane protocols are hardened on transport, content, and service level"

# Control Plane Strategy

- Priority 1:

  Complete isolation from outside

  BGP: GTSM: LPTS: iACL

  IGP: make unreachable

  Option: MPLS with all external services in a VPN

- Priority 2:

  Hardening

  Authentication

  Device level protection

  ( LPTS, CoPP, …)

# Control Plane Execution: BGP Security: Prefix hijacking

- Currently punctual events
  - will get worse, due to IPv4 address space exhaustion
- Monitor: Hijacking monitoring
- Currently, prefix hijacking cannot be prevented
  - Focus on fast detection and reaction
- Need SIDR, with a secure BGP variant
  - soBGP / S-BGP

# Validating Routing

Current SIDR Work

- Origin authentication only
- The RIRs maintain a database of all known address assignments
  - Route Origination Authorizations, or ROAs
  - X.509 certificates containing the assigned AS and a prefix block
- This database is distributed through rsync

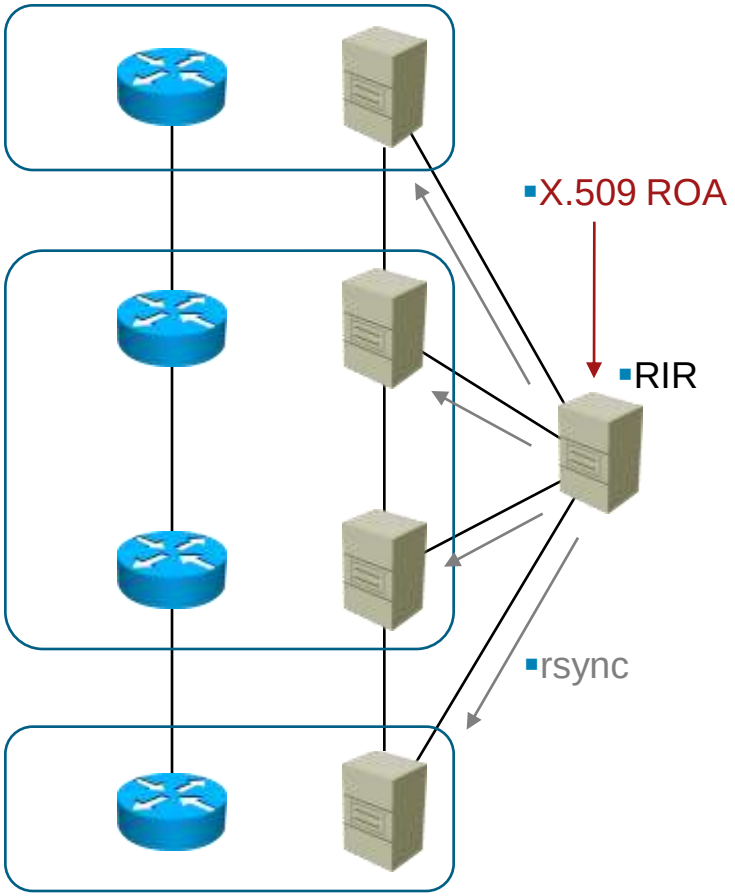

# Validating Routing

Current SIDR Work

- Each edge (eBGP) router in the network connects to a local server
- Some communication protocol allows the router to communicate with this server
- Through this, the router determines if each advertisement is valid or not

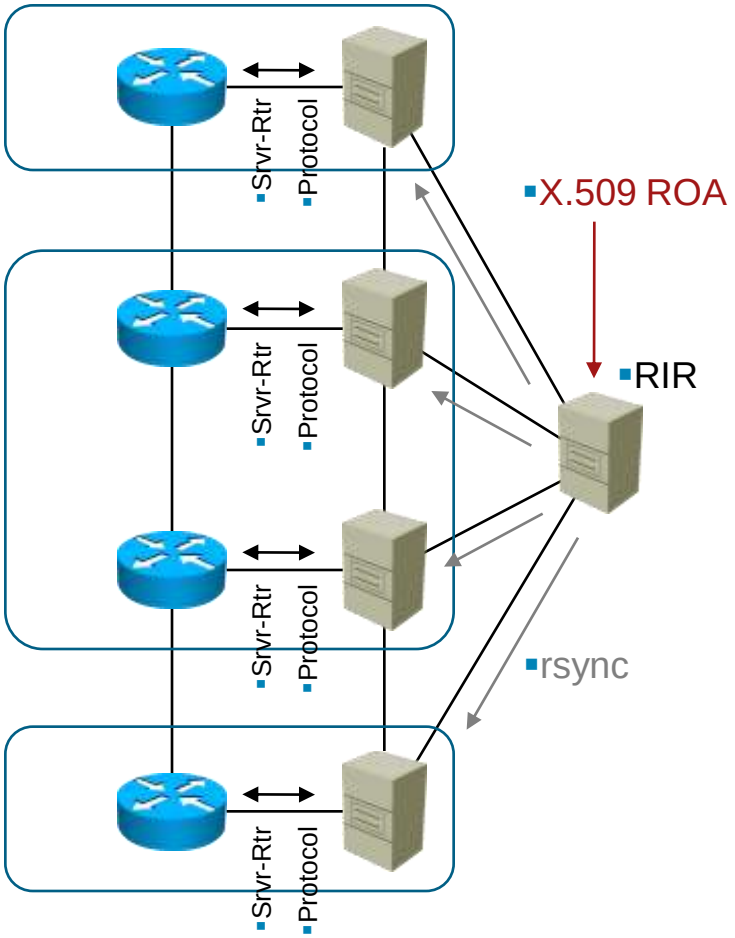

# BGP Bestpath Selection Modifications

- Path Validation States
  - –BGP_PFX_STATE_VALID (Lookup Successful)
  - –BGP_PFX_STATE_NOT_FOUND (Not in the table)
  - –BGP_PFX_STATE_INVALID (Lookup invalid -different origin AS or masklen not in the range)
- BGP Bestpath Modifications

  Input: Received Path, Current Bestpath

  If Received Path is an ibgplearnt path then skip the Prefix Origination check

  If Received Path's Prefix Origination Check state is BGP_PFX_STATE_INVALID then prefer the Current bestpath

  else If Received Path's Prefix Origination Check state > Current Bestpath Prefix Origination Check state, then prefer the Current bestpath

  else (they are equal) proceed to next bestpathcheck step

  Rest of the BGP BestpathSteps

# BGP Cli Modifications

- New CLI
  - –Disable Prefix Validation Globally
  - –Disable Prefix Validation per EBGP Peer
  - –Disable Prefix Validation per set of prefixes
- When disabled, the prefix origin validation state of EBGP Learnt routes will be set to BGP_PFX_STATE_NOT_FOUND

# BGP Cli Modifications (cont'd)
# Policy Knobs

- Allow routes with prefix origin validation state of BGP_PFX_STATE_INVALID for further BGP bestpathselection
- Disallow BGP_PFX_STATE_NOT_FOUND from further being evaluated in BGP bestpathselection
- Usage of communities for announcing validation states using outbound policies

# Dynamic Control Plane Policing

- (D)CoPP is made possible by Local Packet Transport Service
  - –LPTS enables distributed applications to reside on any or all RPs, DRPs, or LCs
  - –Filters local ‘for-us’ packets and sends them only to the nodes that need them
- LPTS has HW policers on line cards to limit traffic sent to (D)RPs – DCoPP
  - –LPTS entries in TCAM classifies packets to select a policer
  - –Polices on protocol (BGP, OSPF, SSH) and flow state (BGP established flows, BGP listen)
  - –Policing done on the LC ASIC before packets hit RP/LC CPU
  - –All filters are automatically and dynamically installed by the IOS XR infrastructure

# For-us packet path overview

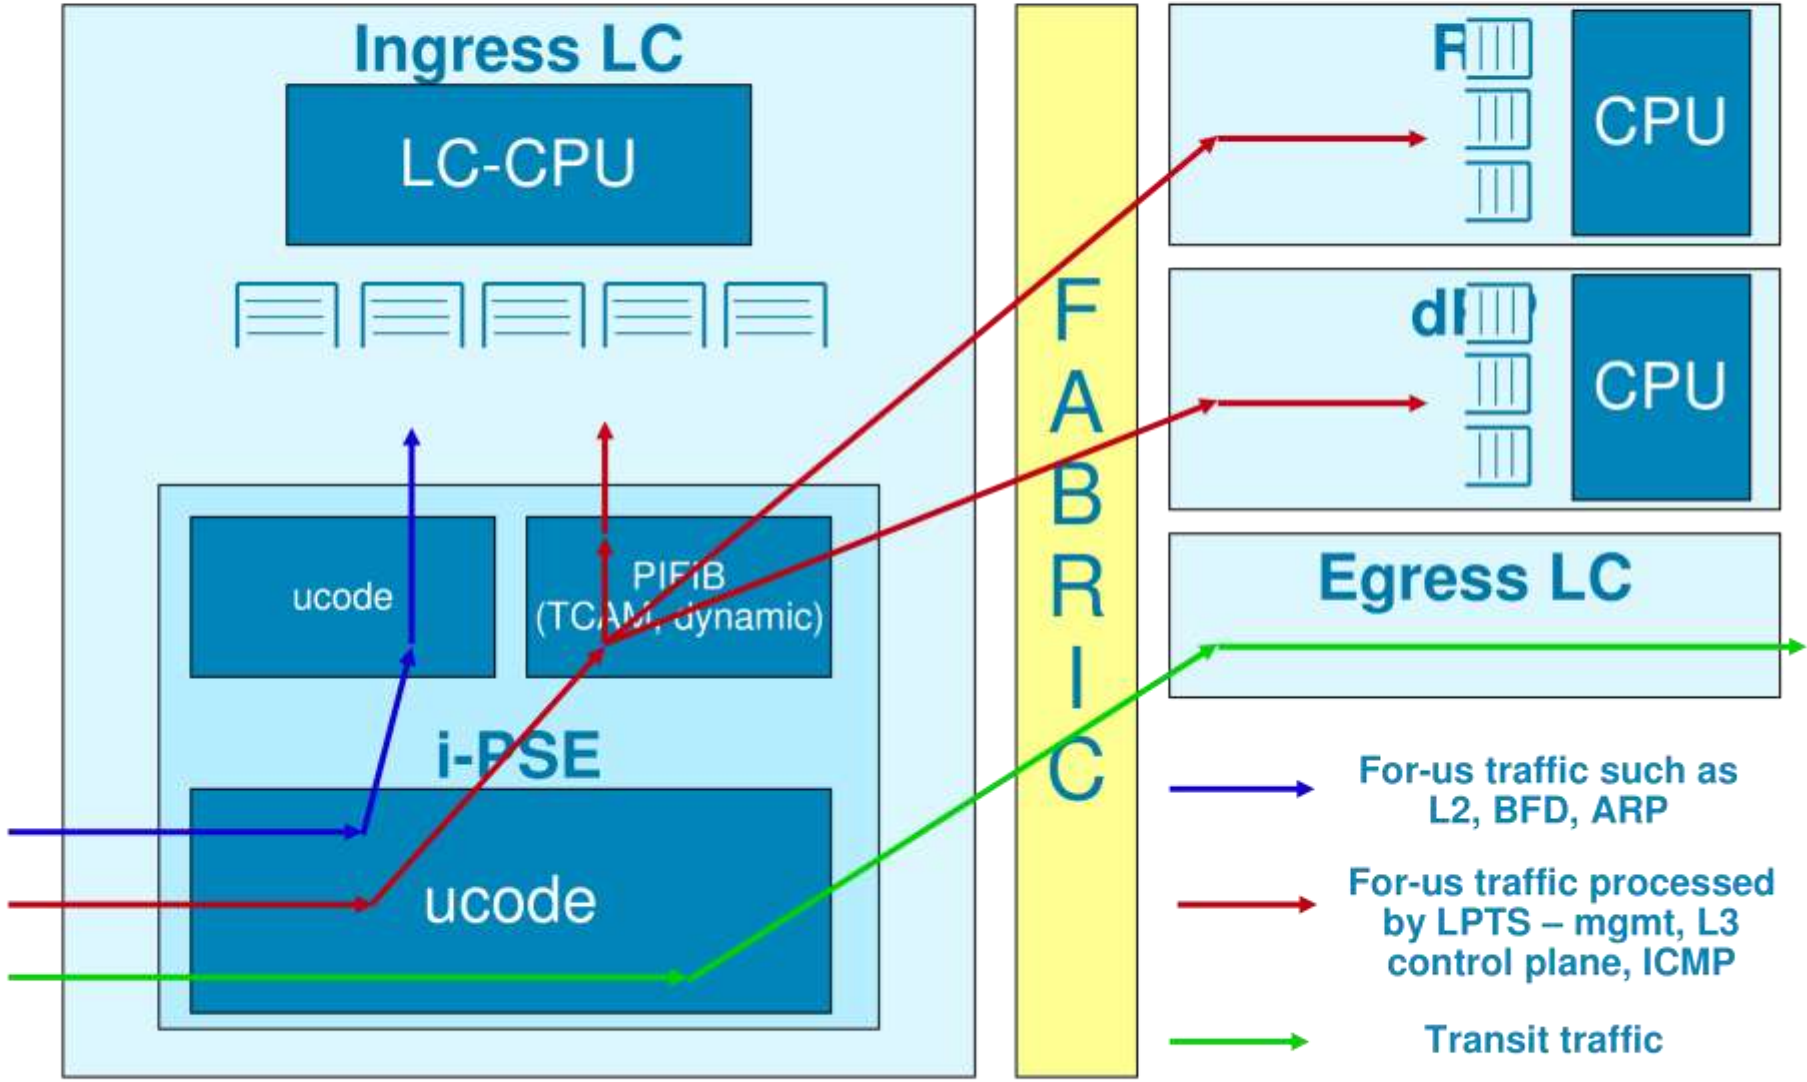

# PSE Policiers

| Punt-all enabled | Handle large | L2 low priority | L2 control | CDP | IPv6 PLU Punt |
|---|---|---|---|---|---|
| PLIM ASIC header error | Ethernet loopback | Bundle Control | Unknown OSI NLP | ARP | IPv6 frag needed |
| Diagnostic | IPv4 Options | IPv4 RSVP w/options | IPv4 IGMP w/options | RARP | IPv6 MC do all but forward |
| IPv4 PIM w/options | IPv4 TTL expiration | Iv4 PLU no match | IPv4 PLU punt | CGMP | IPv6 PLU no match |
| IPv4 PLU receive | IPv4 frag needed but DF set | IPv4 L3LI punt | IPv4 L2Li punt | SAP | IPv6 BFD Async |
| IPv4 L2LBE RP punt | IPv4 cannot frag, MTU too small | IPv4 BFD Async | IPv4 BFD Echo | IPv4 tunnel MTU exceeded | IPv6 PLU receive |
| IPv4 MC do all | IPv4 MC do all but forward | IPv4/IPv6 incomplete adjacency | IFIB lookup miss | ACL dency gen ICMP | IPv6 BFD Echo |
| ACL log | IPv6 link local | IPv6 routing extension hdr | IPv6 hop-by-hop options | IPv6 TTL expiration | IPv6 L3LI punt |
| MPLSL2LI punt | MPLS TTL expiration, IP payload | MPLS PLU | MPLS PLU receive | IPv6 L2LI punt | IPv6 MC directly connected |
| L2VPN VCCV | MPLS incomplete adjacency | MPLS IPv4 frag needed but DF set | MPLS IPv4 options frag needed | IPv6 MC do all | ILMI Packet |
| ATM LC Packet | Mac record | OAM Event | OAM Packet | MPLS L3LI punt | MPLS IPv6 frag needed |
| Service Card Punt | EOAM CFM non CCM Packets | EOAM EFM Packets | QNET | Biscuit DCAP TTL error | Biscuit MTU Violation and DF set |

# The Future of the Management Plane

Configuring and Monitoring Networks

# Management Plane Vision

"From the outside, management channels are unreachable."

"There are secure versions of all management protocols."

# Management Plane Strategy

- Complete isolation of management channels from the outside
  - Making devices unreachable from the outside
  - Management plane protection
  - Automatic, with minimum configuration
- Support of secure protocols for all management channels
  - Every management access is protected via strong authentication (e.g., AAA) and crypto;
- Support *unified* role based access control mechanisms

# The Overall Vision

The ideal SP architecture, from a security point of view

# Isolation between “inside” and “outside”

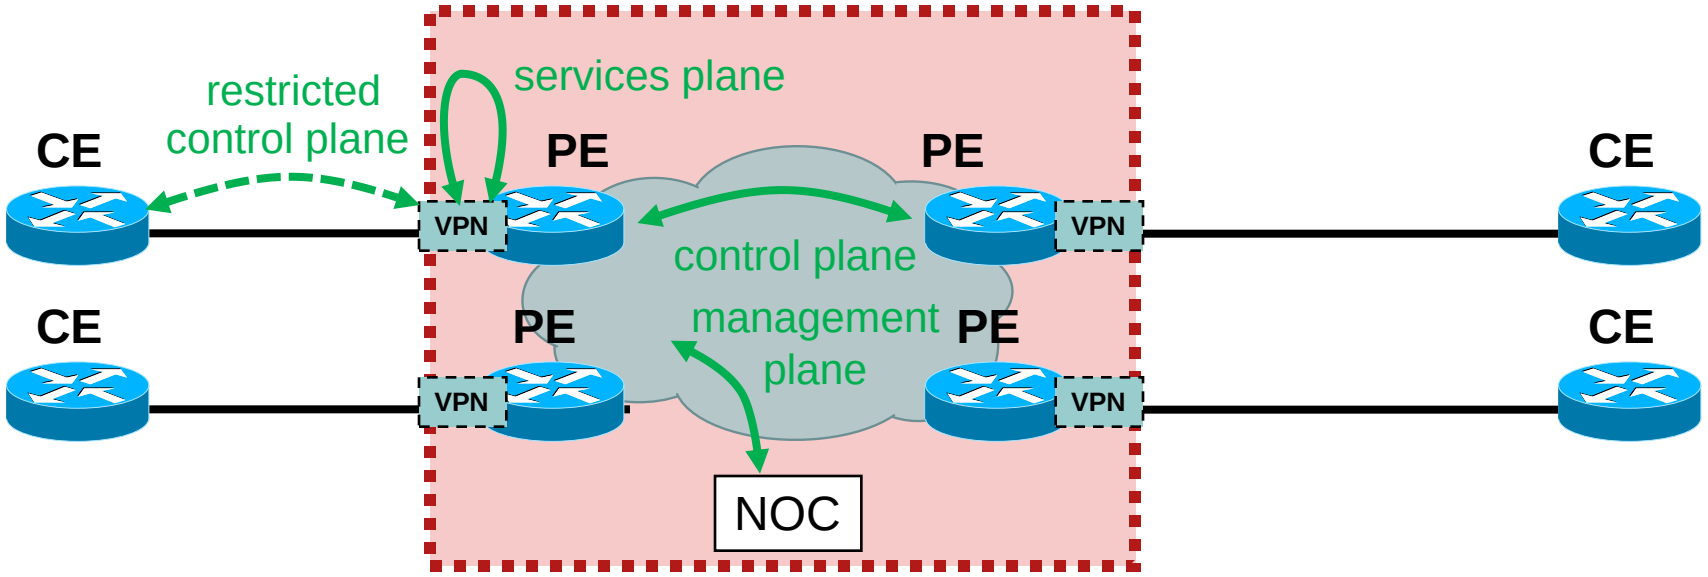


All external traffic will be a “VPN” on an SP core

→ Isolation towards outside world

→ Control, management plane remain “inside” only

→ Data plane traffic has no effect on network

# Target Architecture

- RFC 4364 networks ideally suited:
  - Strong separation outside / inside; but:
- Need to make Internet in VRF a reality!!
  - All PE platforms!
- Need to further secure the external facing interface
  - So far, manual security (iACL, etc) required
  - Need to automate
    - Good example: Management plane protection
  - Goal: By default, PE does not “receive” any packets coming in on an external i/f. Exceptions (for routing, ICMP, etc) need to be explicitly configured.

# Summary

- **“Be Careful or Be Roadkill” — Calvin**

- Important strategies:
  - Architectural isolation inside/outside
  - Internet in a VRF
  - Secure Inter Domain Routing
  - CoPP / LPTS