Dynamic Routing for IPsec VPN Manageability: Current IETF Standards Activities

Paul Knight, Nortel Networks

Gregory Lebovitz, Netscreen Technologies

Lars Eggert, USC/ISI

Agenda

- **Introductions**
- Why we need dynamic routing in IPsec
- Difficulty of doing dynamic routing in IPsec
- Quick Review: IPsec Transport and Tunnel Modes
- Current Implementations of dynamic routing in IPsec
- What's happening in IETF standards

Introductions

- Gregory M. Lebovitz
  - Architect, CTO Office, Netscreen Technologies
  - Design next generation feature sets and security solutions
  - gregory@netscreen.com
  - www.netscreen.com

Introductions

- Paul Knight
  - Standards Architect, Nortel Networks
  - Ensure that product plans incorporate standards, for interoperability
  - Promote innovative technologies as potential standards candidates
  - paul.knight@nortelnetworks.com
  - www.nortelnetworks.com

Introductions

- Lars Eggert <larse@isi.edu>
  - Ph.D. candidate, USC/ISI
- virtual networks since 1997
  - X-Bone, DynaBone, TetherNet
- other research
  - TCP, web caching, OS network issues
- http://www.isi.edu/larse/

Agenda

- Introductions
- **Why we need dynamic routing in IPsec**
- Difficulty of doing dynamic routing in IPsec
- Quick Review: IPsec Transport and Tunnel Modes
- Current Implementations of dynamic routing in IPsec
- What's happening in IETF standards

Virtual Network

- network equivalent of virtual memory
  - abstraction, protection, sharing
- network =
  hosts + routers + links
- virtual network =
  - virtual host → packet source/sink
  - virtual router → packet gateway
  - virtual link → tunnel X over Y
- **virtual Internet**: X = IP, Y = IP

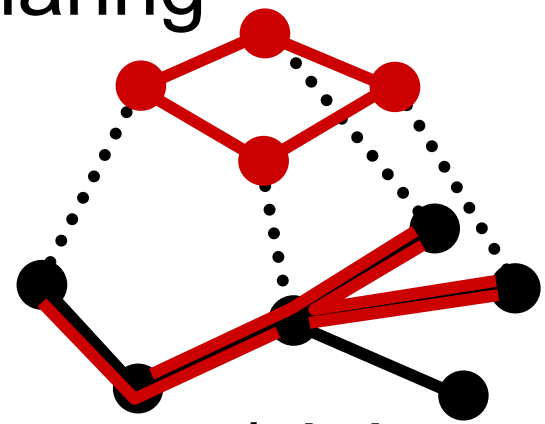

Virtual Private Network

- **private** = secure links
  - authenticate tunnel ends + encrypt
- virtual private Internet
  - secure IPIP tunnels hop-by-hop
- security is link property
  - decoupled from topology
- IPsec tunnel mode?

IPsec VPN – Frame Relay Replacement

- IPsec-based VPN as a frame relay replacement
- Business drivers
  - Lower monthly operating costs
  - ROI in 4 to 6 months
- Need equivalent functionality at lower cost

FR Features that Customers want: Can IPsec VPNs address them?

- Single physical connection with multiple virtual connections to remote sites
- **Privately** transport all internal networking information. Includes:
  - IP traffic
  - Private IP addressing schemes
  - non-IP traffic
  - IGP/EGP routing protocols
- CIR, assured level of performance (bandwidth)

IPsec VPN vs. Frame Relay

Feature	**IPsec VPN**	**Frame**
Single phy w/ multiple virtual connections to remote sites.	➨	➨
Private Transport	+	➨
Private Addressing Schemes	➨	➨
Non-IP Traffic	➨ **(in tunnels)**	+
IGP/EGP Routing Protocols	➨	➨
CIR	-	+
COST	+++	--

Use Case 1 – New Networks Added to a Remote Site

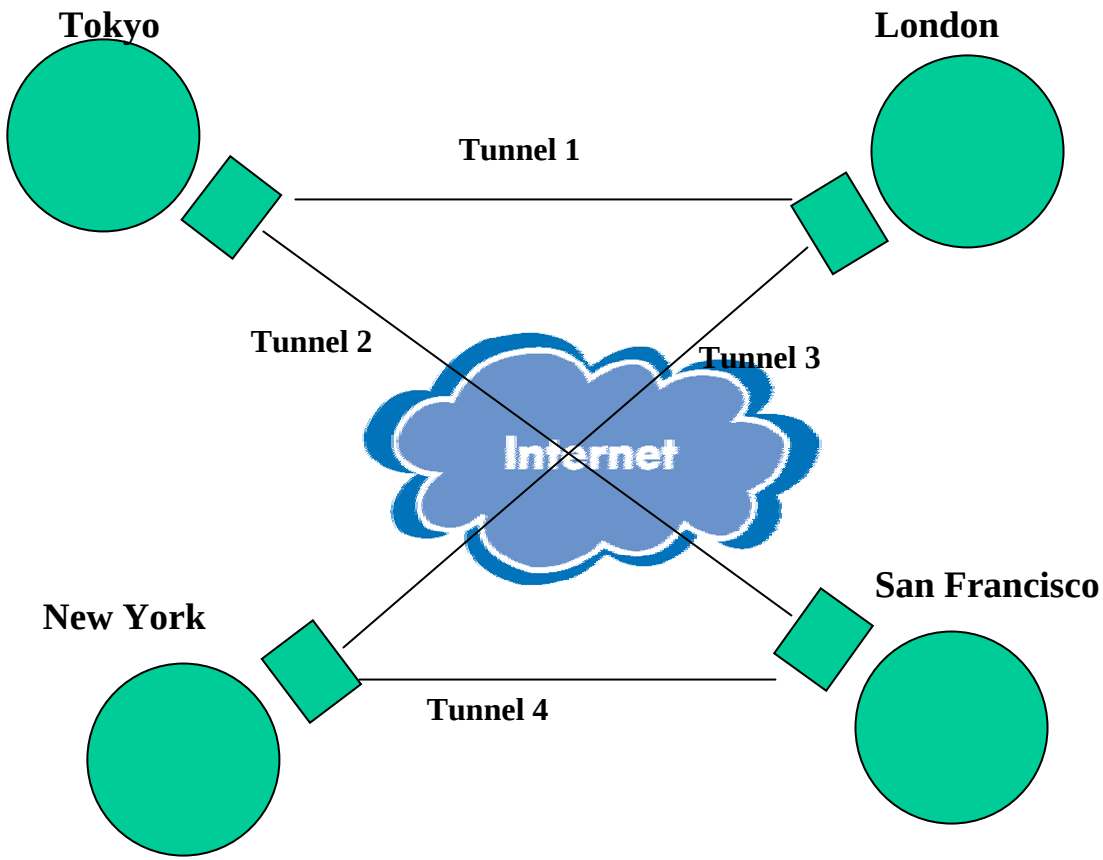

Use Case 1 – New Networks Added to a Remote Site

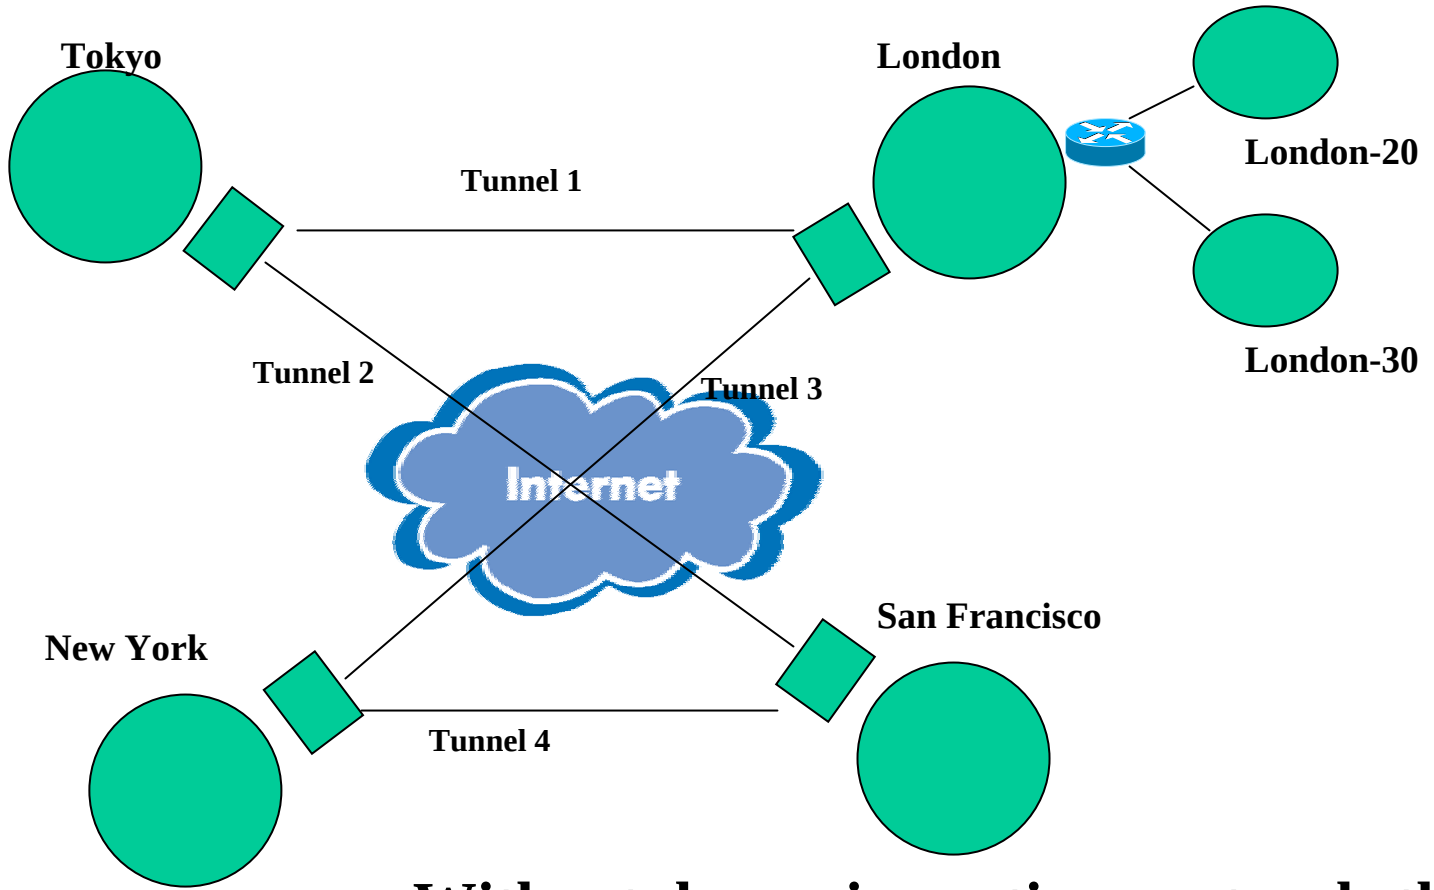

Without dynamic routing protocols through tunnels, none of the other Sites will be able to reach the London-20 and London-30 networks without configuration change.

Use Case 2 – Multiple Paths to Hanover

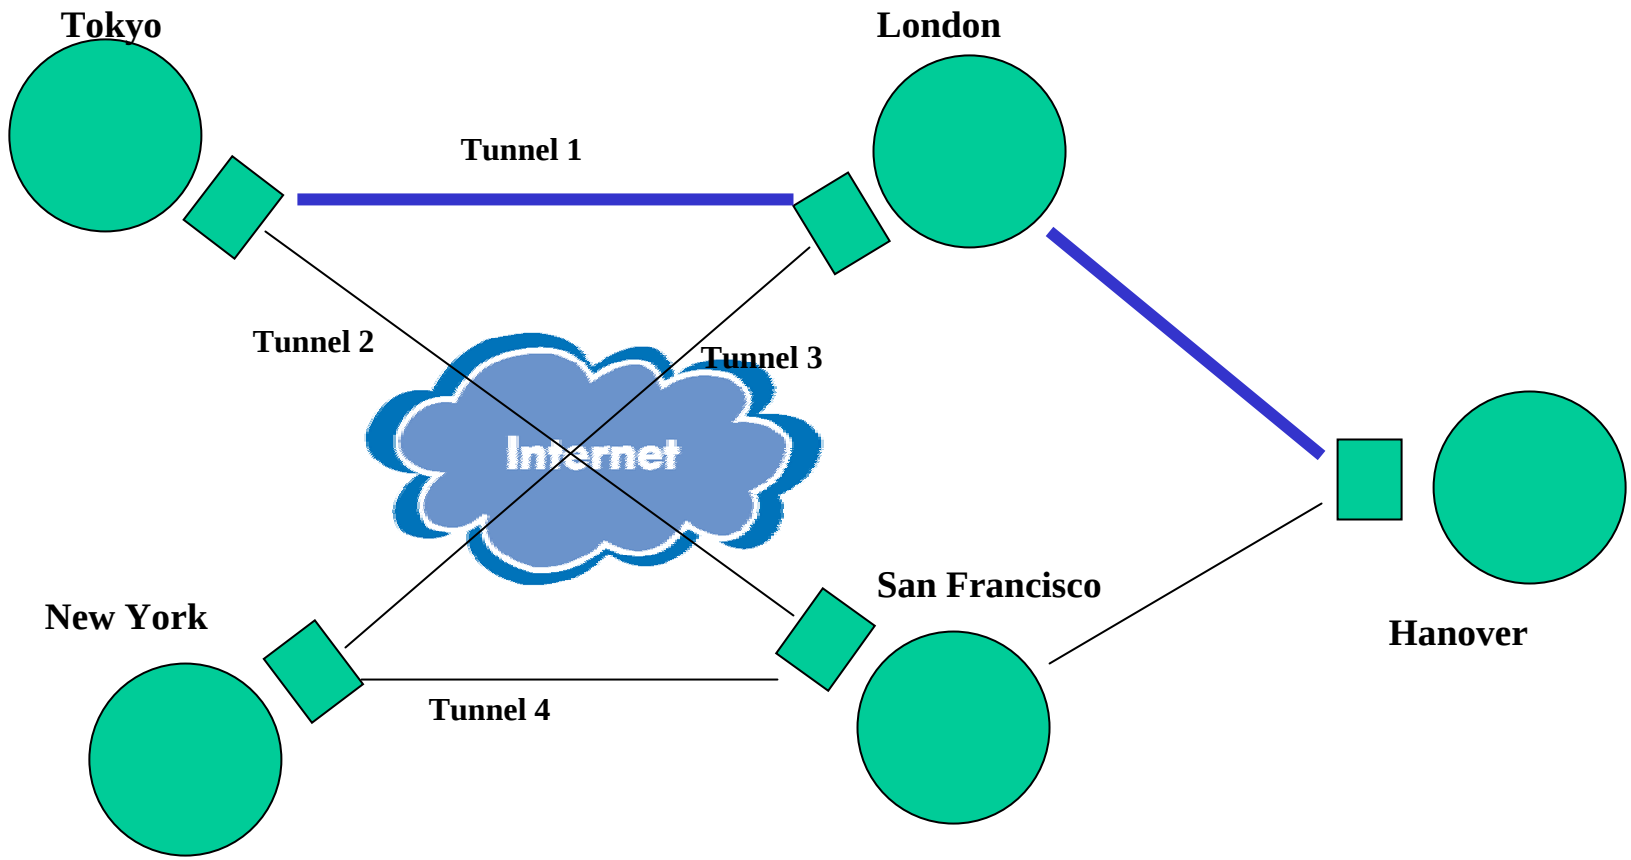

Use Case 2 – Without Dynamic Routing

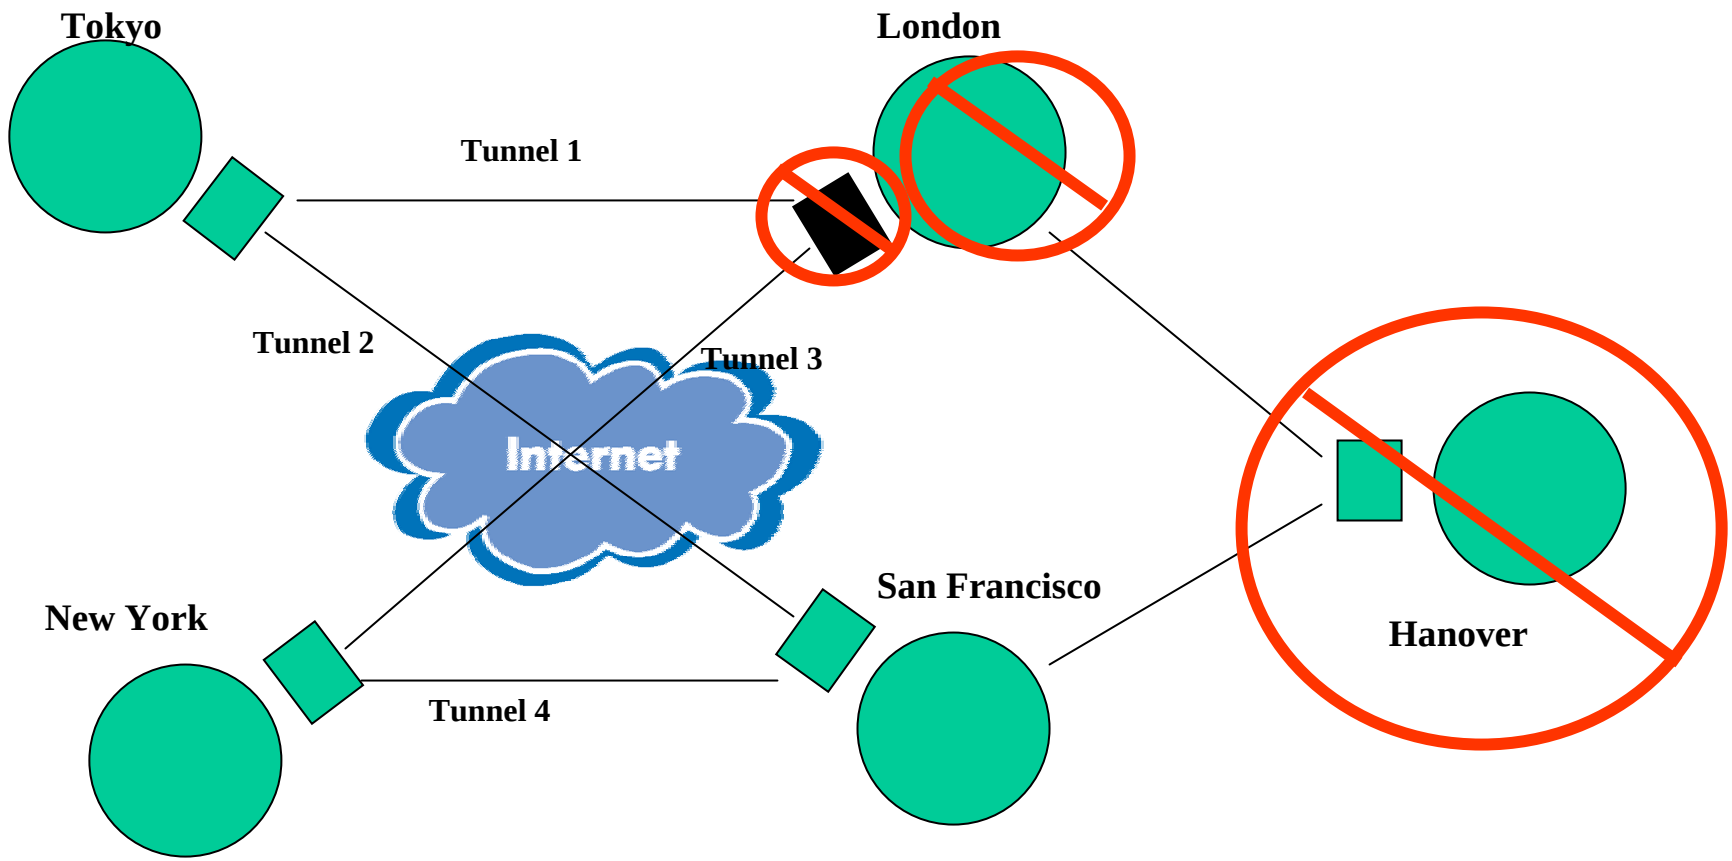

Connectivity to both London and Hanover are lost.

Use Case 2 – With dynamic routing protocols

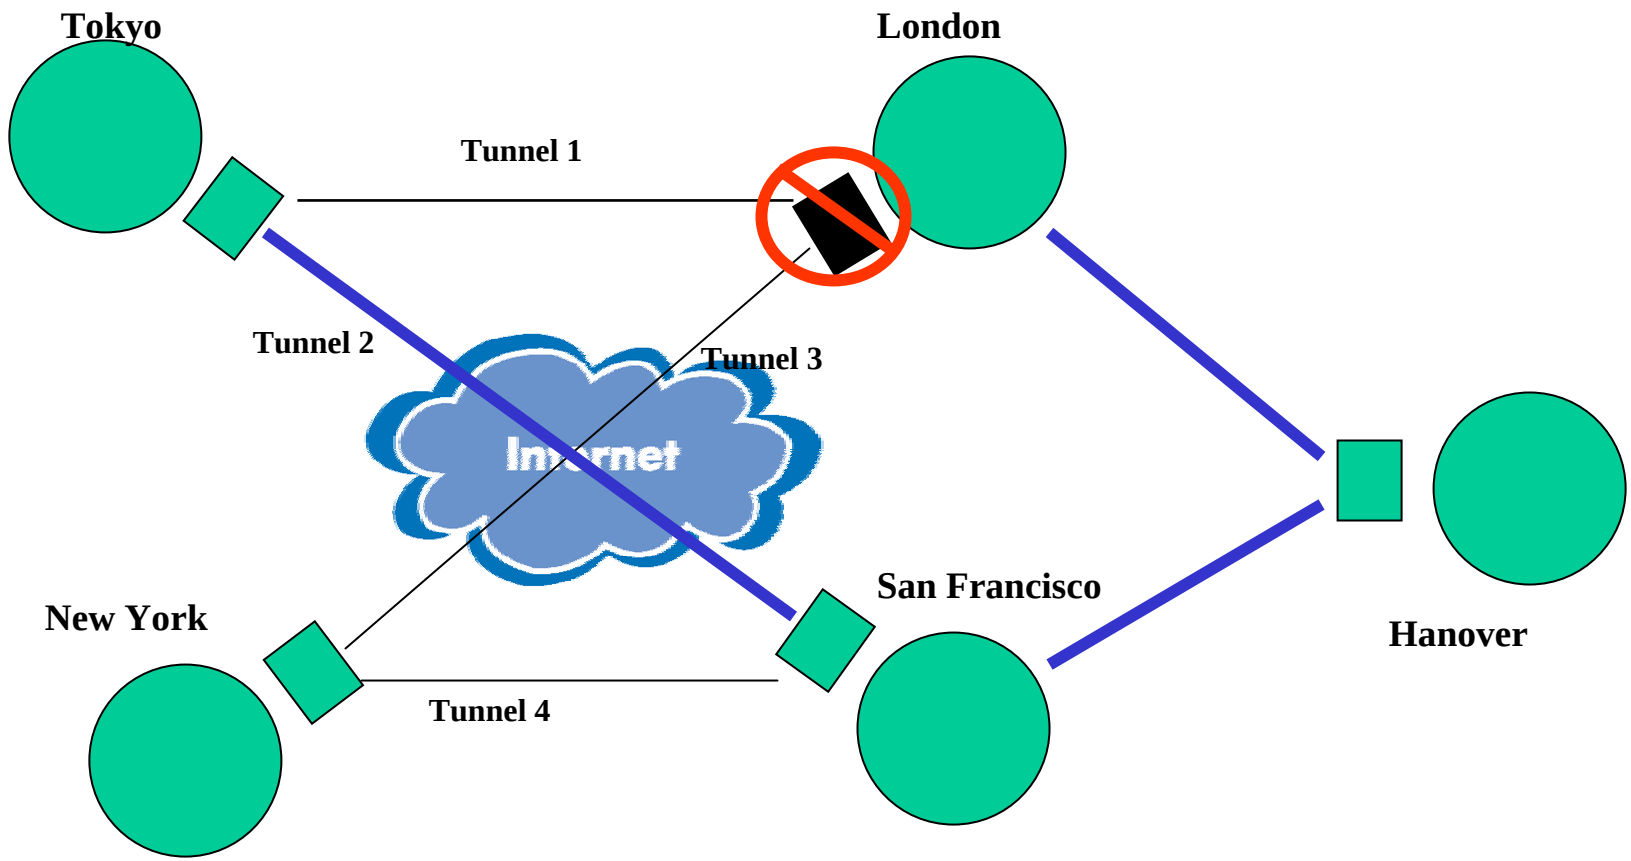

Connectivity to both London and Hanover are maintained.

Agenda

- Introductions
- Why we need dynamic routing in IPsec
- **Difficulty of doing dynamic routing in IPsec**
- Quick Review: IPsec Transport and Tunnel Modes
- Current Implementations of dynamic routing in IPsec
- What's happening in IETF standards

Typical Routing Environment

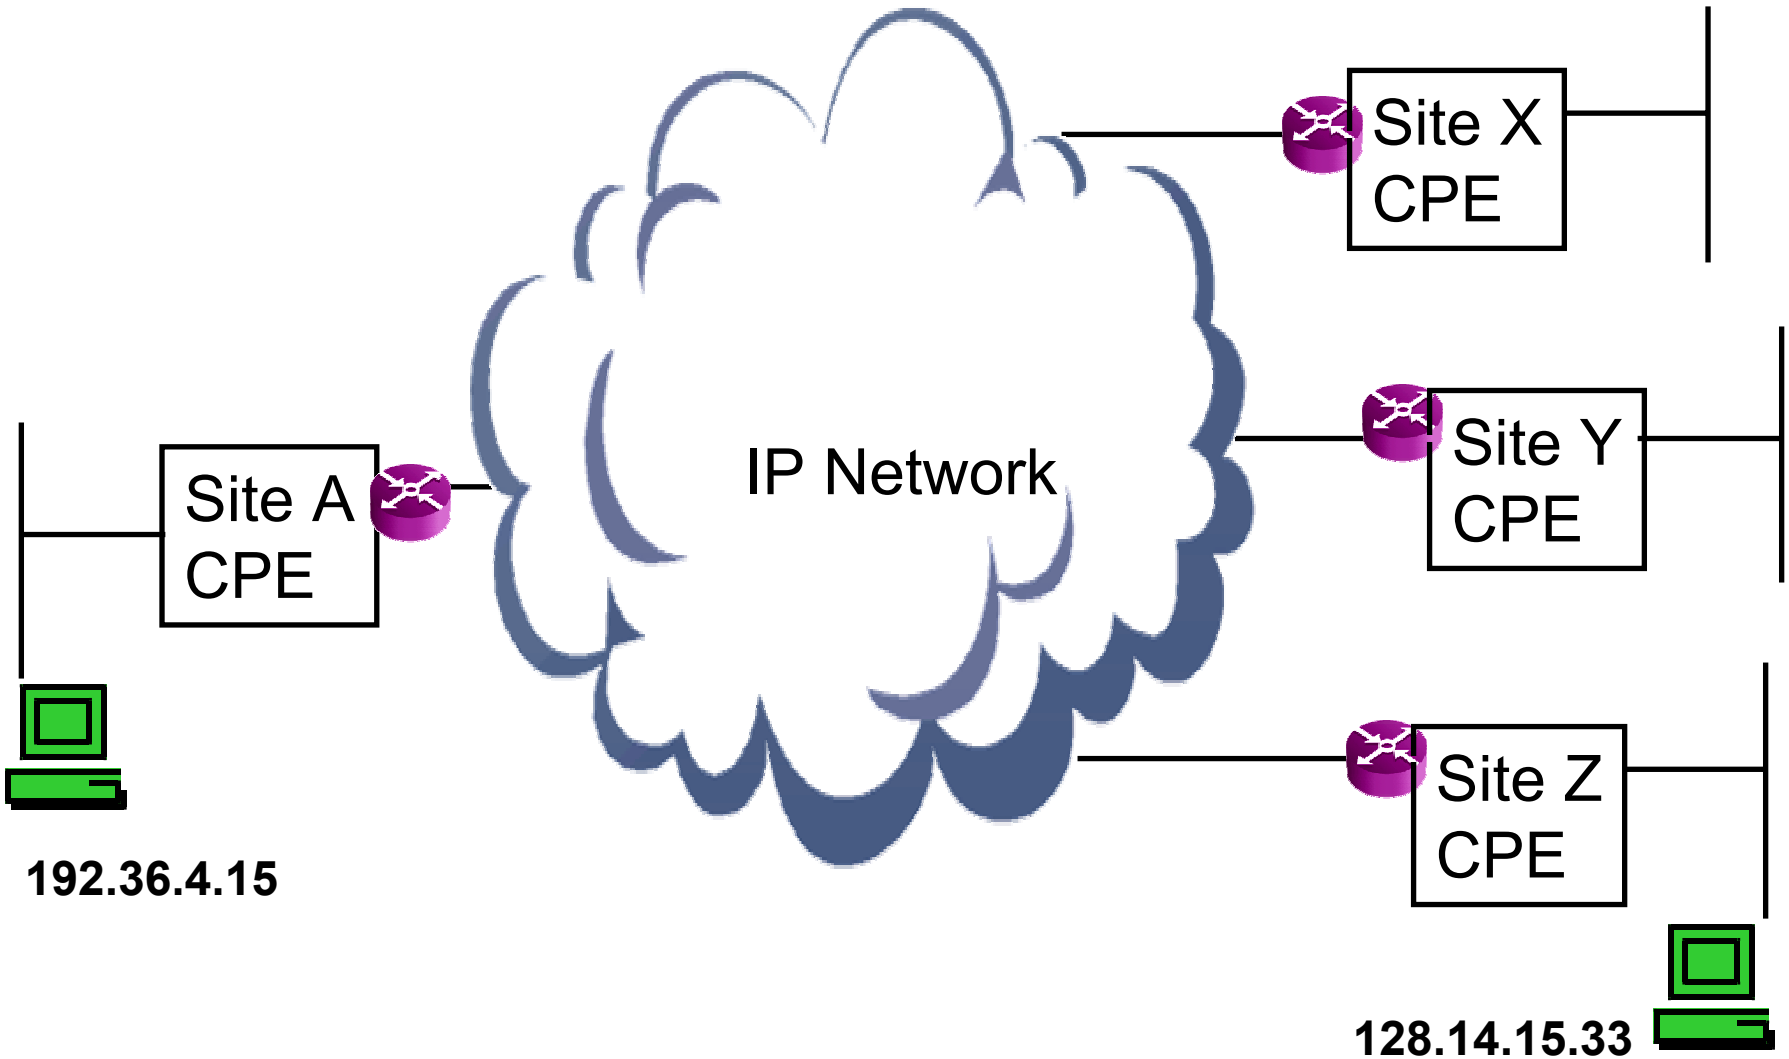

Typical Routing Environment

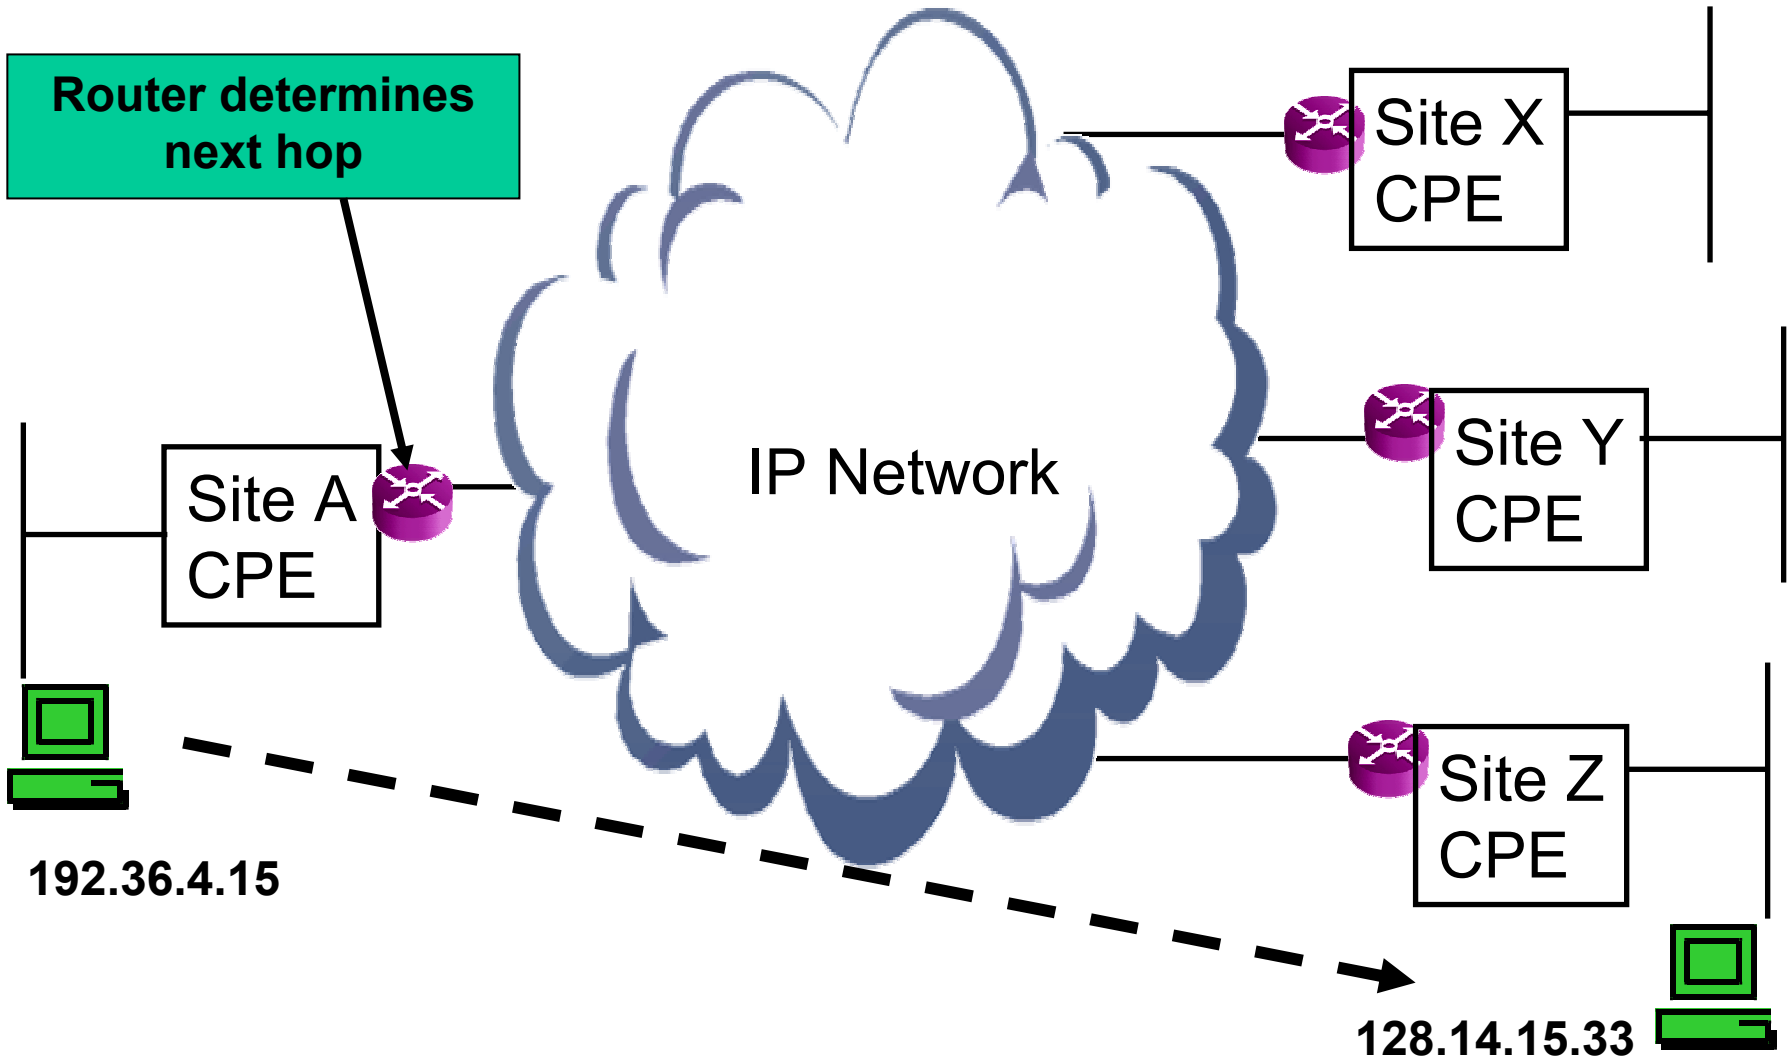

Typical IPSec VPN Environment

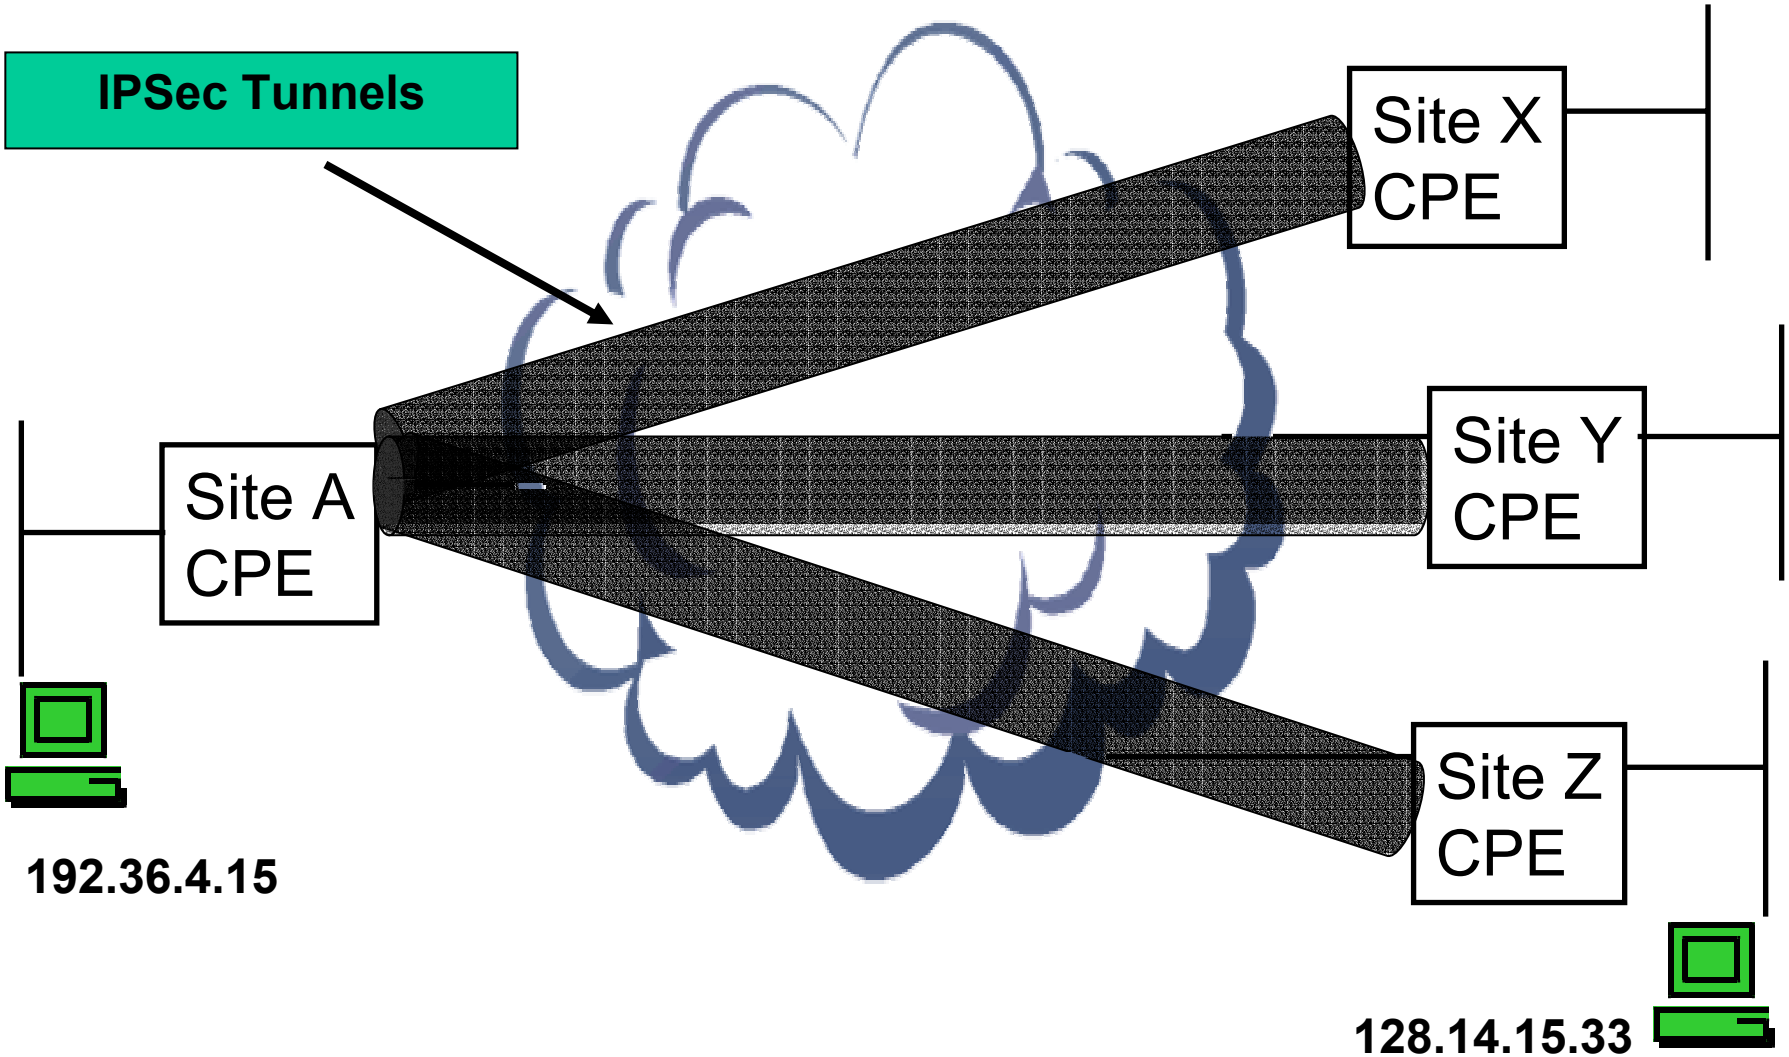

Typical IPSec VPN Environment

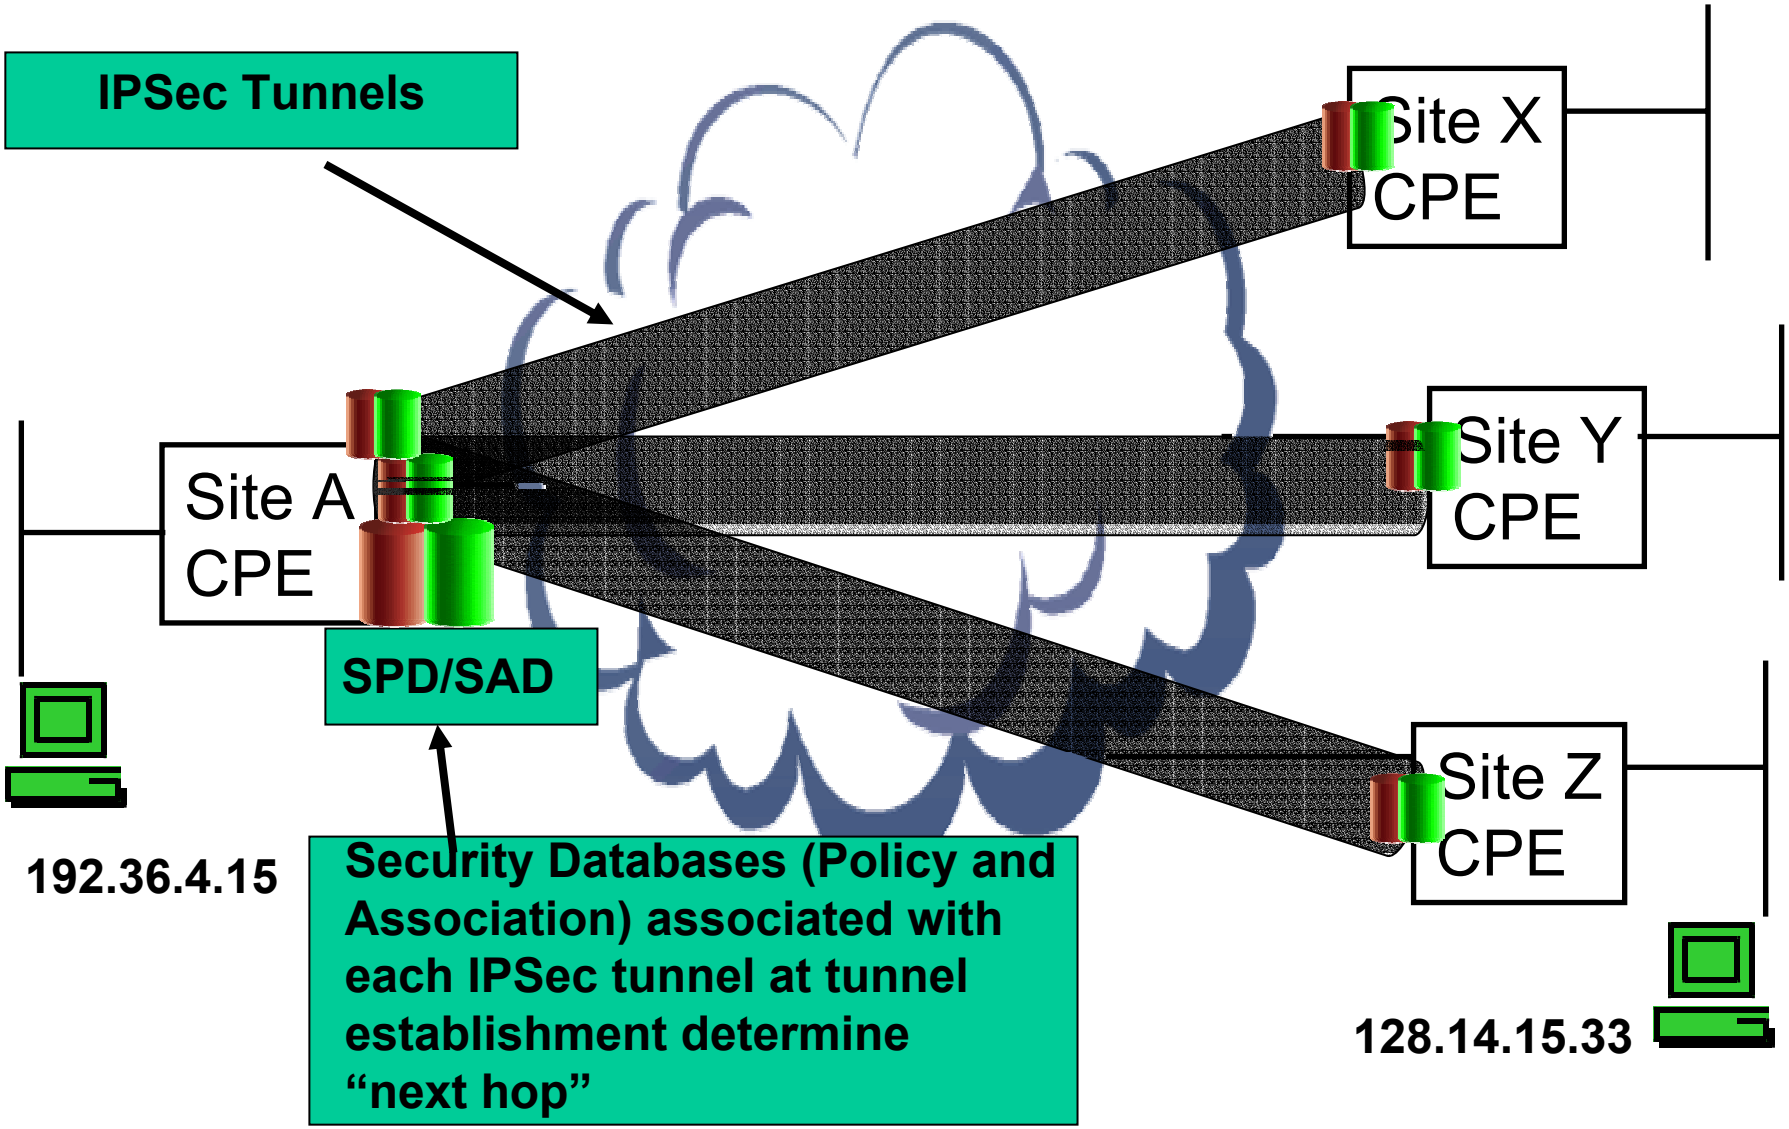

Typical IPSec VPN Environment

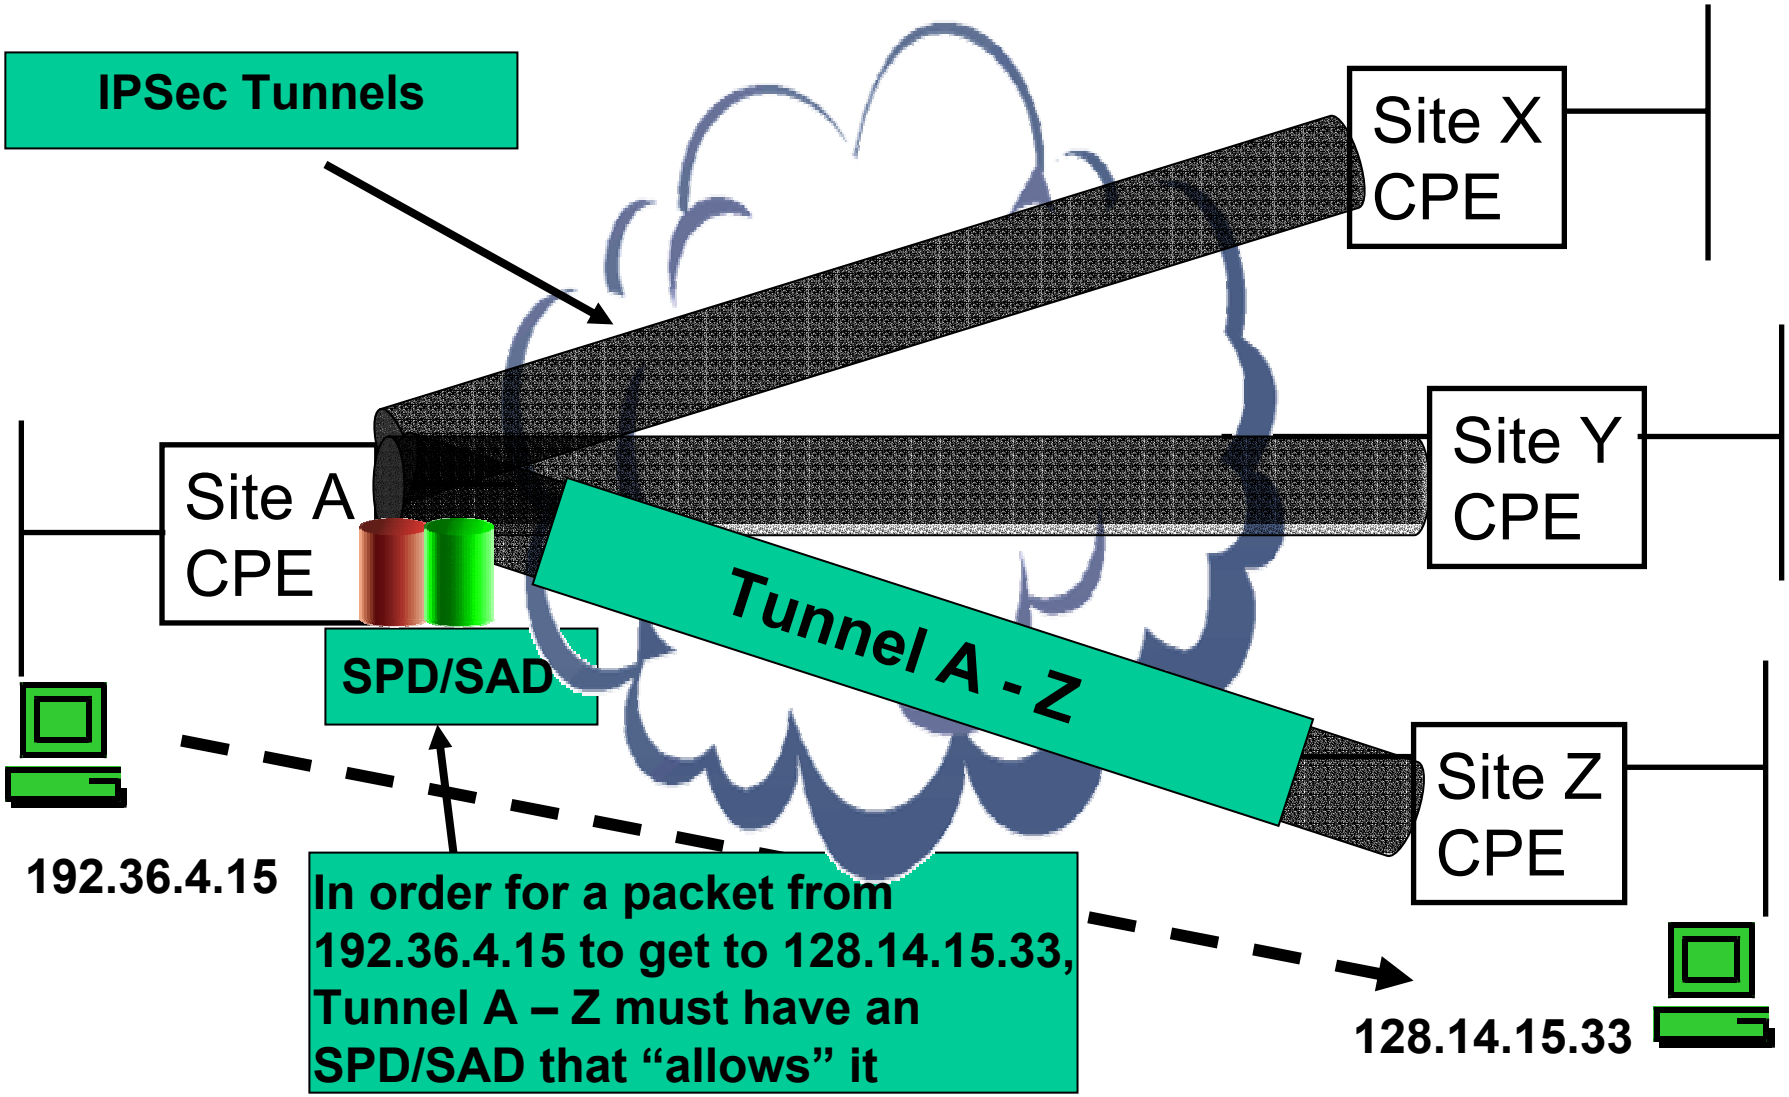

Typical IPSec VPN Environment

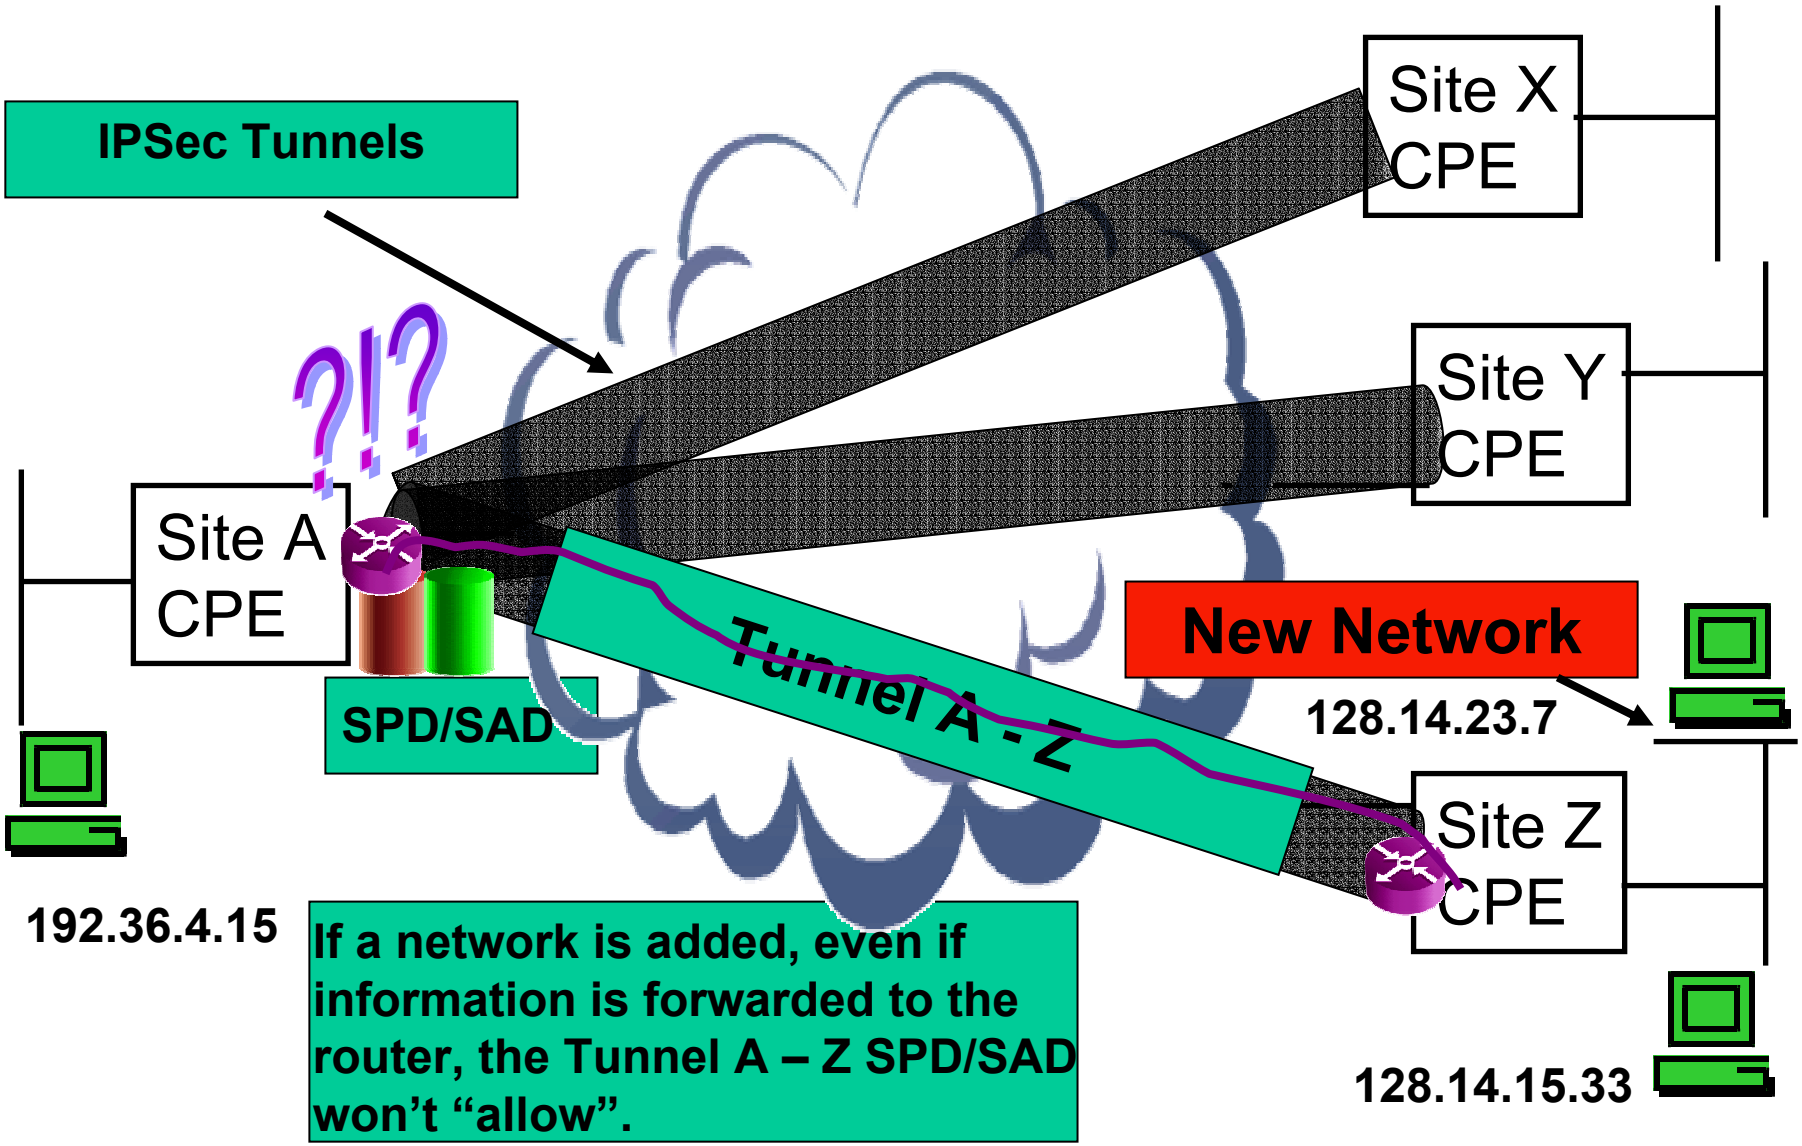

So: how do you do Dynamic Routing over IPSec tunnels?

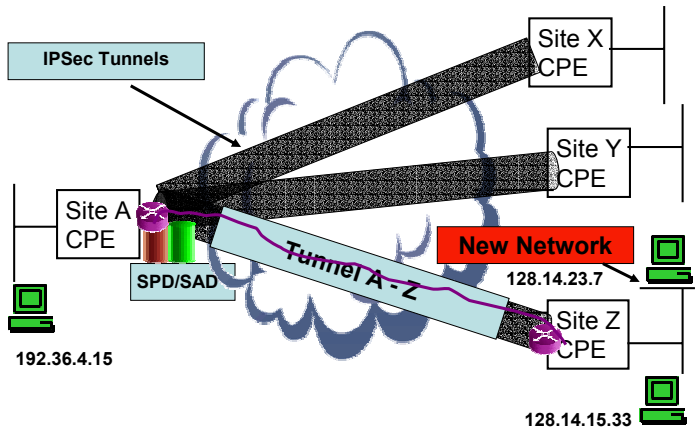

- Rebuild IPsec SA for each routing change?
- Make a “Wild Card” SPD/SAD for the IPsec Tunnels?
- Do the routing outside of IPsec?
- Current solutions incorporate the ideas of the last two points.

SA ≠ Interface

- tunnel SA = key, src, dst
- encapsulation: **interface operation**
- SAs not in IP forwarding table
- duplicate, separate forwarding mechanism

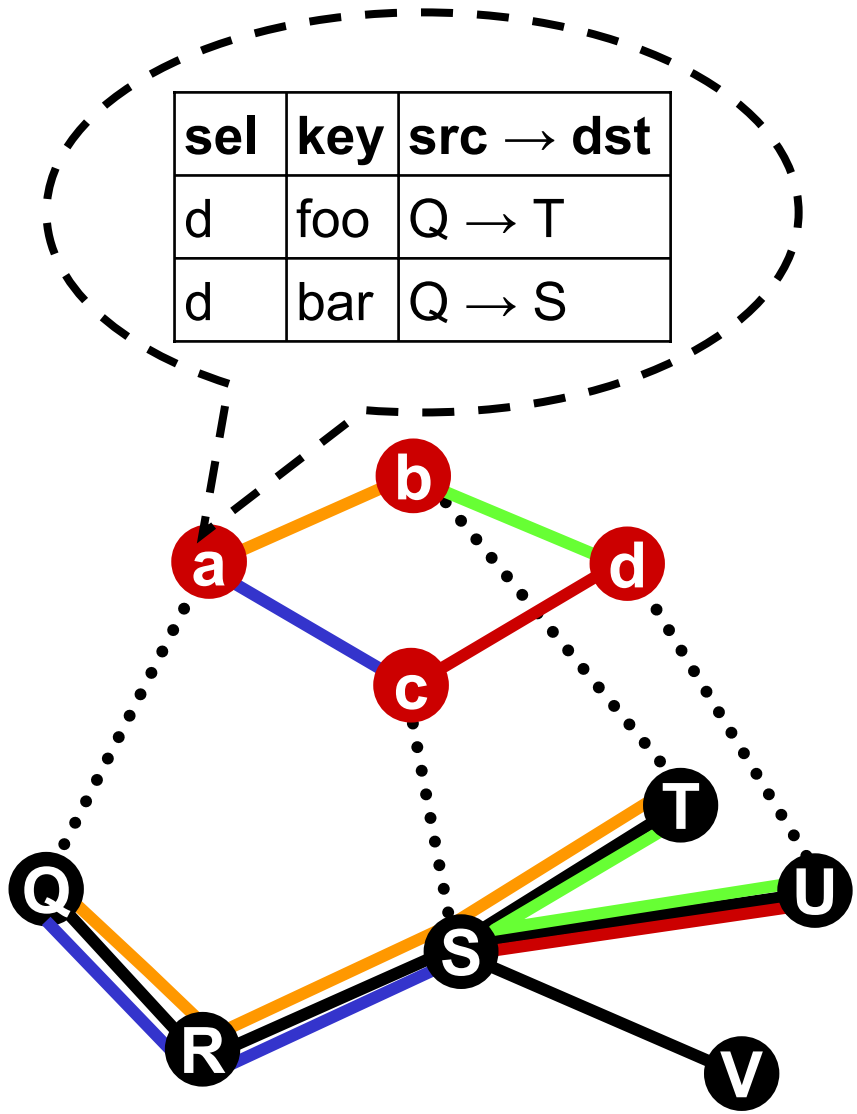

Source Address Selection

- 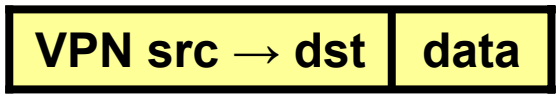

which source IP address?
- RFC 1122, section 3.3.4.3
  - uses notions of interface and route
- tunnel mode SA neither
- security implications!
  - replies in the clear
- result: special case for local traffic
  - must include in IPsec spec, bloat

Selectors and Routing

- selectors = tunnel firewall
- routing update → SA renegotiation
  - or valid traffic filtered
  - overhead, stabilization
  - couples routing + IPsec
- option: wildcard selectors
- selectors for tunnel mode less useful?

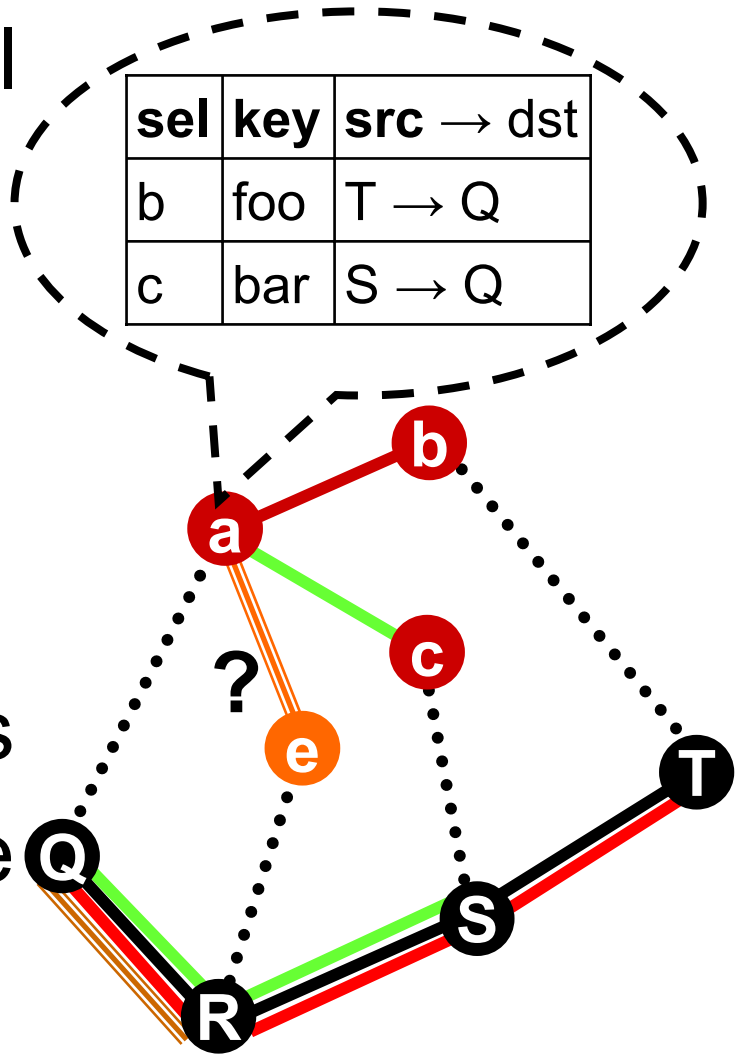

Routing Protocols via IPsec

- Tough when rtg protocol utilizes a L2 component
- OSPF – has multicast component used on broadcast networks, and NBMA
  - Solution: Use OSPF virtual links or pt-to-pt.
  - Must define neighbors. Good security anyway
- BGP – Works Great!
  - All peers pre-identified/pre-configured
  - All messages in IP. It's easy.
- RIP - L2 and IP level broadcast and can be carried w/o any trouble over the tunnel.
  - Gtwy on other side needs to act as a recipient of the RIP, and not just forward pkt into the internal network.
- ISIS – L2 component needed.

Agenda

- Introductions
- Why we need dynamic routing in IPsec
- Difficulty of doing dynamic routing in IPsec
- **Quick Review: IPsec Transport and Tunnel Modes**
- Current Implementations of dynamic routing in IPsec
- What's happening in IETF standards

Two IPSec Modes: Transport and Tunnel Mode

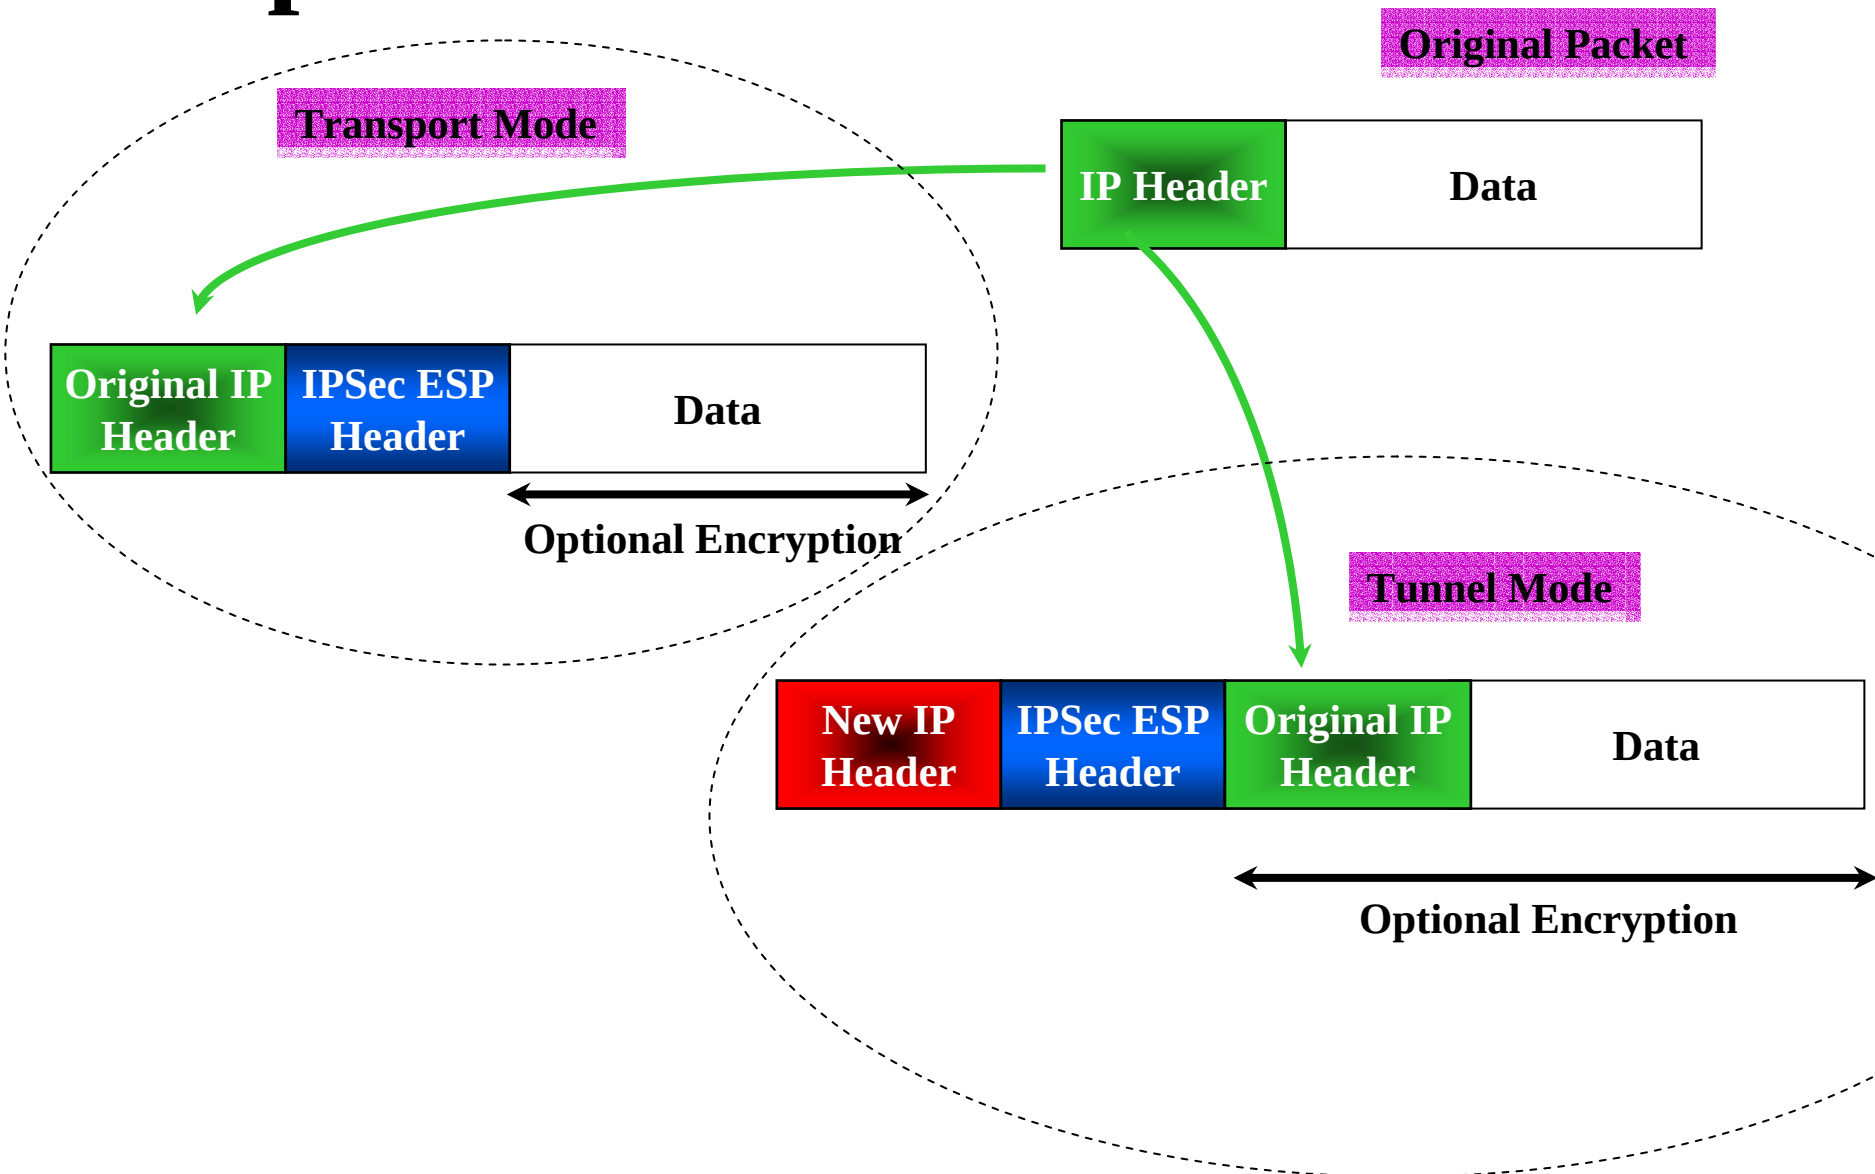

Application of the IPsec modes

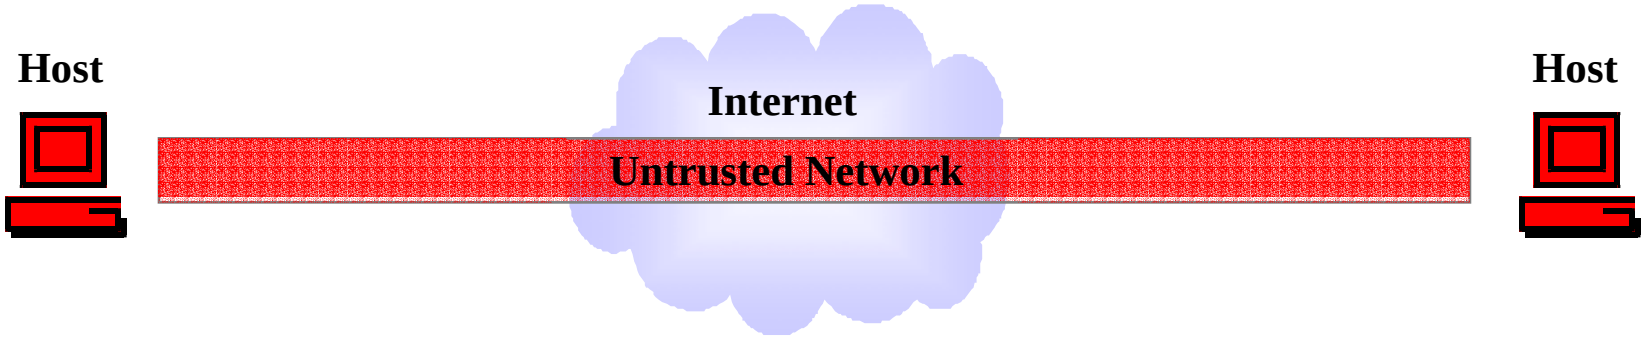

Can use **Transport (or Tunnel)** Mode between Hosts

Between Gateways: MUST hide IP addresses of trusted networks when traffic crosses the untrusted network.

- **Tunnel Mode… OR**
- **IP encapsulation within Transport Mode**

Agenda

- Introductions
- Why we need dynamic routing in IPsec
- Difficulty of doing dynamic routing in IPsec
- Quick Review: IPsec Transport and Tunnel Modes
- **Current Implementations of dynamic routing in IPsec**
- What's happening in IETF standards

Implementations: NetScreen

- How it works
  - Tunnel Mode using wild-card (0/0) Proxy-Ids
  - Remove Access control from tunnel decision
  - Explicitly separate Routing function from SPD function
  - Treat tunnel as routable interface (un- or numbered)
- Benefits
  - Less packet overhead
  - Works through NAT boundaries
  - Faster due to less encapsulation processes
- Drawbacks
  - Other side must support same method
  - IP traffic only

NetScreen – Functional Difference

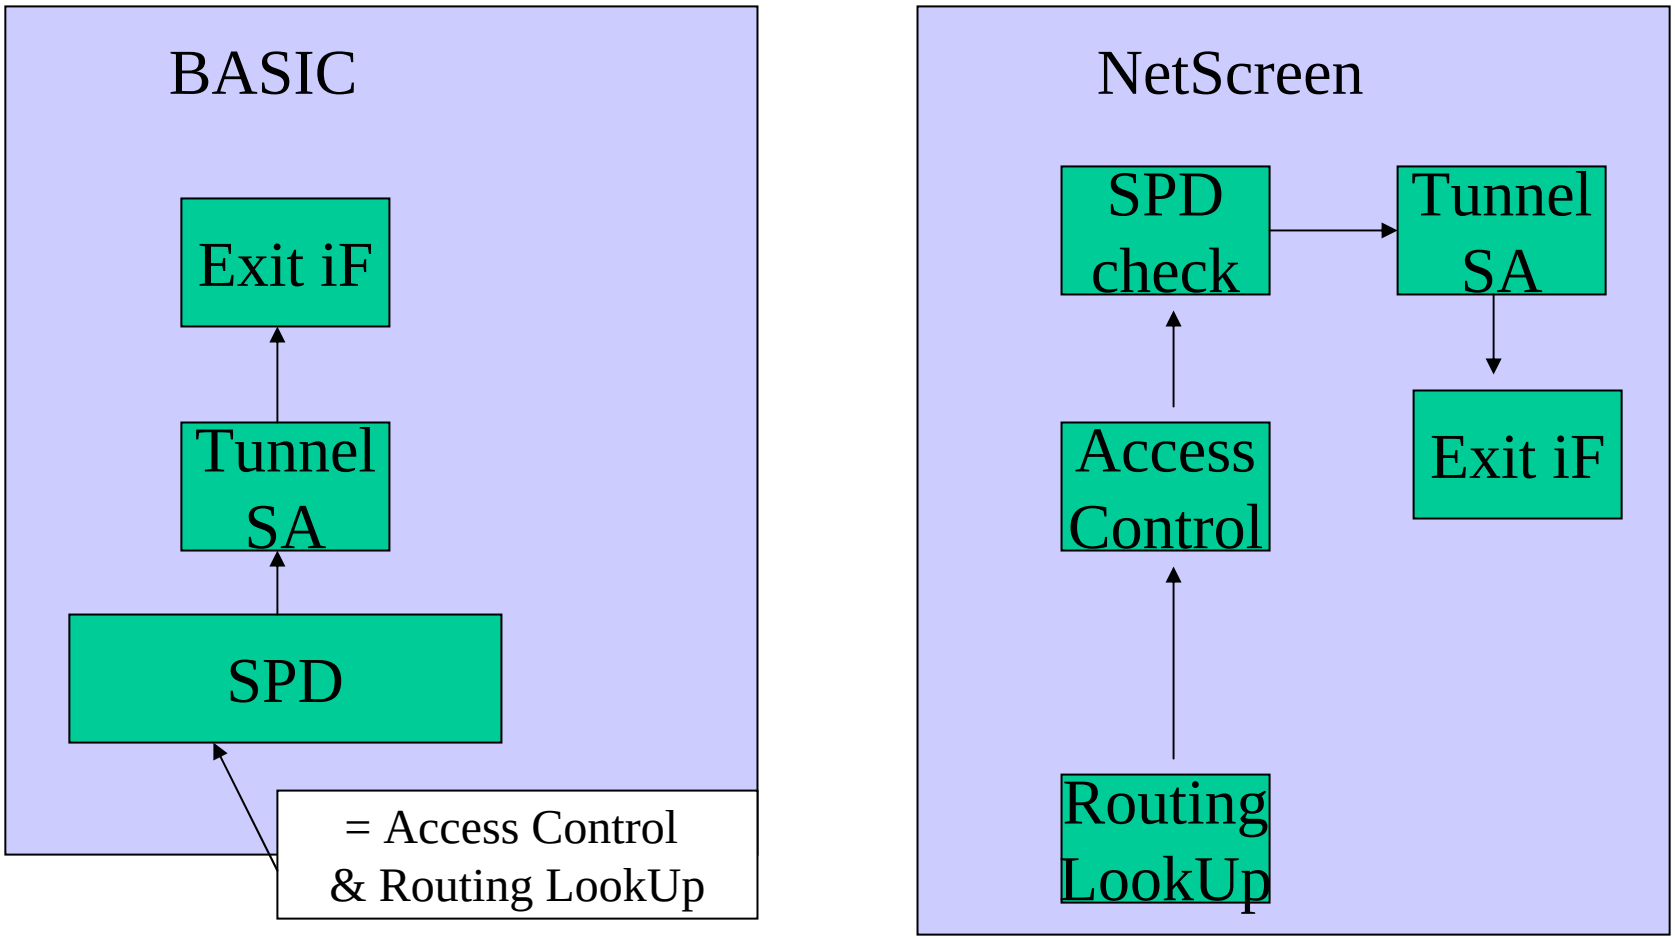

NetScreen's Solution

Tunnel Mode IPsec, Numbered or Unumbered "Tunnel" Interfaces,
Route decision and access control separate from IPsec processing

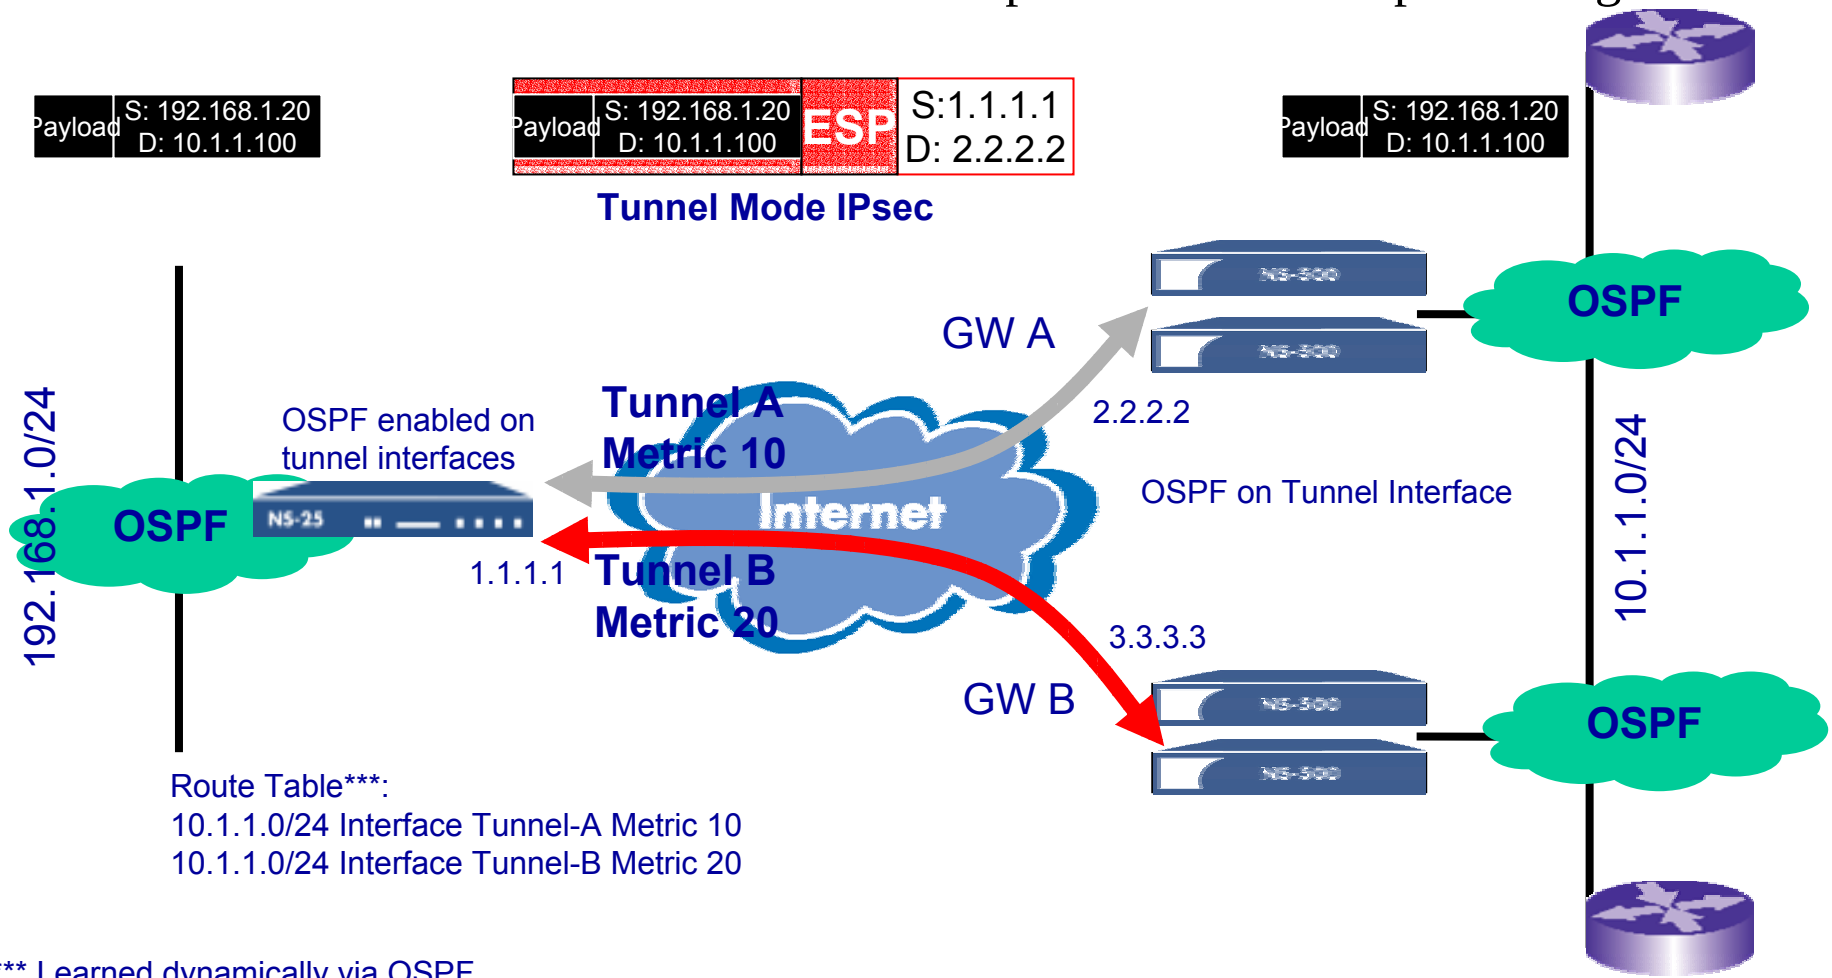

*** Learned dynamically via OSPF

Implementations: Nortel Networks Contivity

- How it works for dynamic routing*
  - Transport Mode IPsec Security Association is created, protecting IP-in-IP encapsulated traffic
  - IP-in-IP encapsulation assigns the tunnel endpoints based on routing table
  - Firewall, filtering, access control - applied outside IPsec
  - Contivity gateways see peers as next-hops for routing
- Benefits
  - Packets exactly same as Tunnel mode
  - Routing clearly separated from IPsec SPD processing; “Secure Routing Technology”
- Drawbacks
  - Other side must support same method

* IPsec Tunnel Mode is used with static routing

Transport mode + IP encapsulation

- Determine “next IPsec hop” of the packet, using any criteria the “routing engine” can handle:
  - route to destination (using dynamic information!)
  - protocol
  - port (socket)
  - even content analysis (URL, etc.)
- Construct new encapsulating IP header with source of own IPsec gateway address; destination of next IPsec hop
- Pass to IPsec process for TRANSPORT mode processing
- Resulting packet is equivalent to tunnel mode, but now it is routed using dynamic routing updates

Transport mode + IP encapsulation

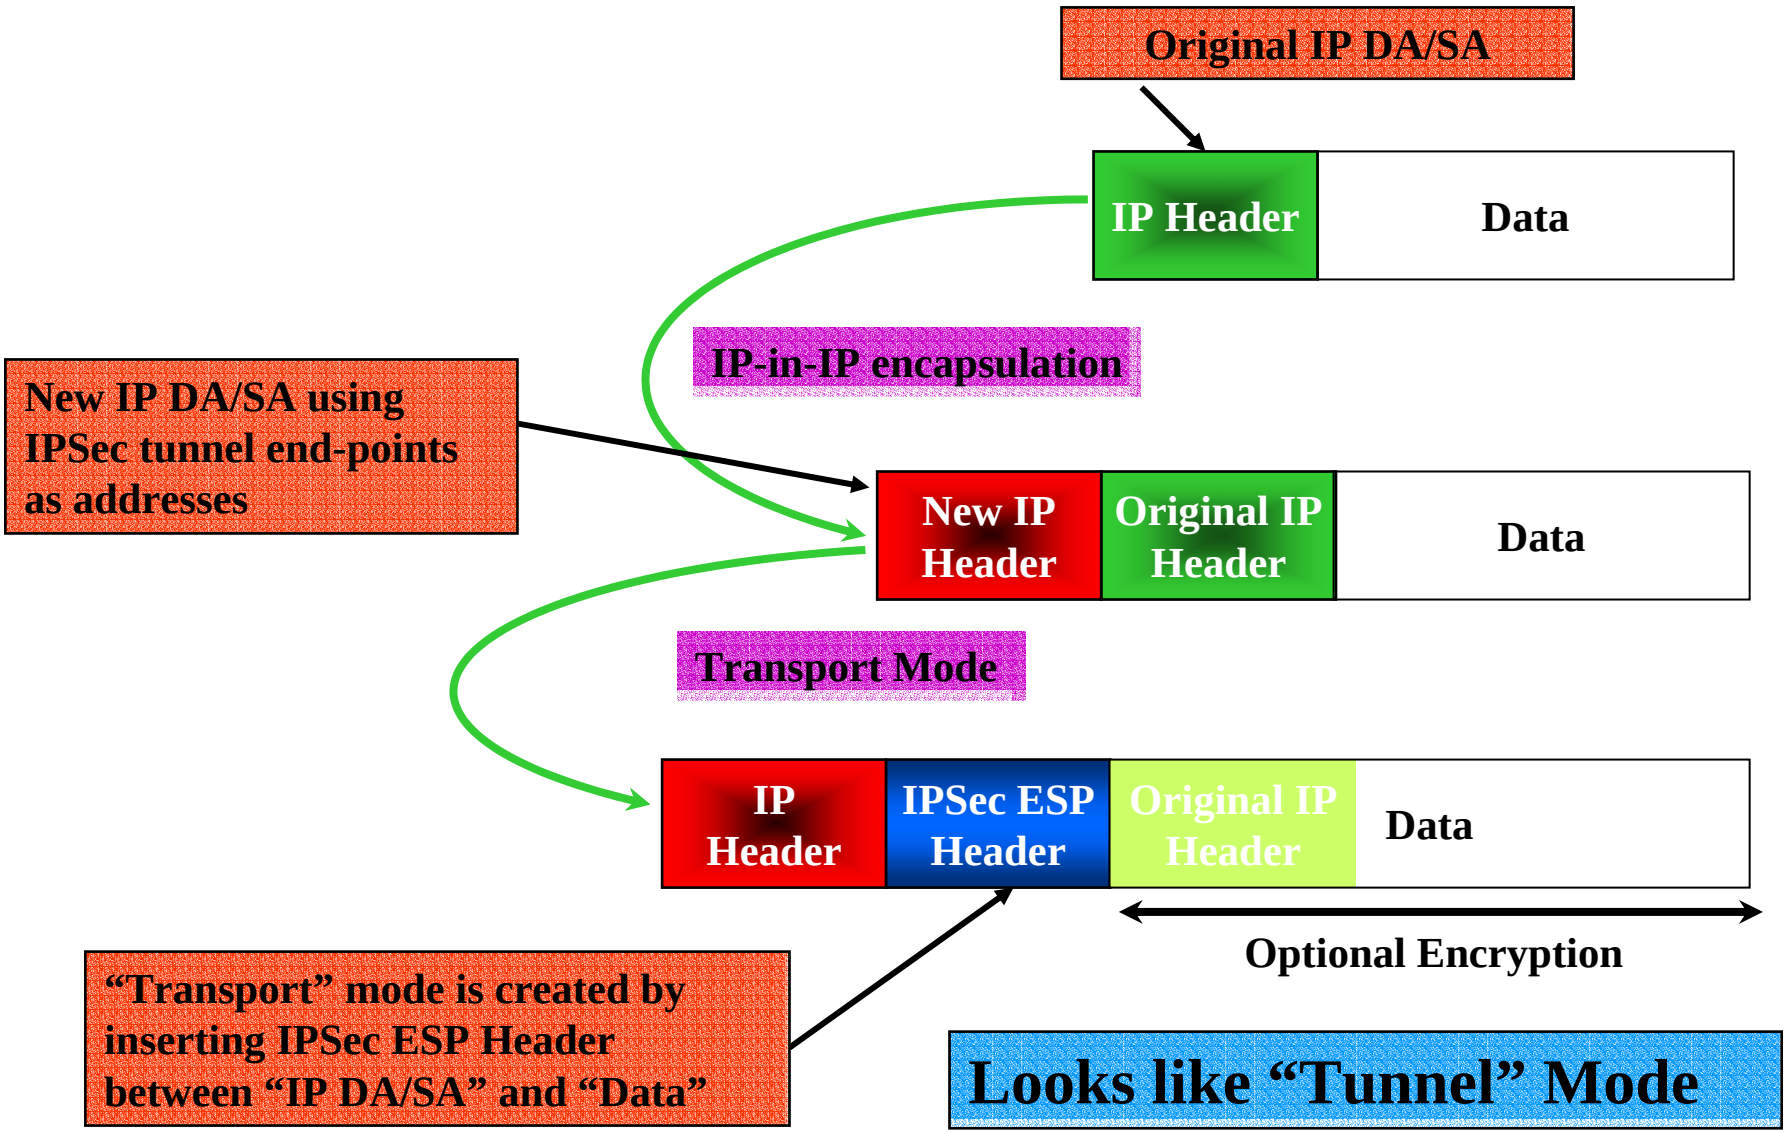

Implementations: ISI's X-Bone and TetherNet

Subjective IPsec History

- goal: secure end-to-end IP
  - everybody will do transport mode
- tunnel mode: **stopgap**
  - wrap packets from legacy boxes
  - one-hop topologies
- then virtual networks come along

Proposed Solution

- kill tunnel mode, instead combine:
  - RFC 2003 IPIP tunnel (step #1)
  - IPsec transport mode (step #2)
- route VPN IP → encaps → IPsec base IP

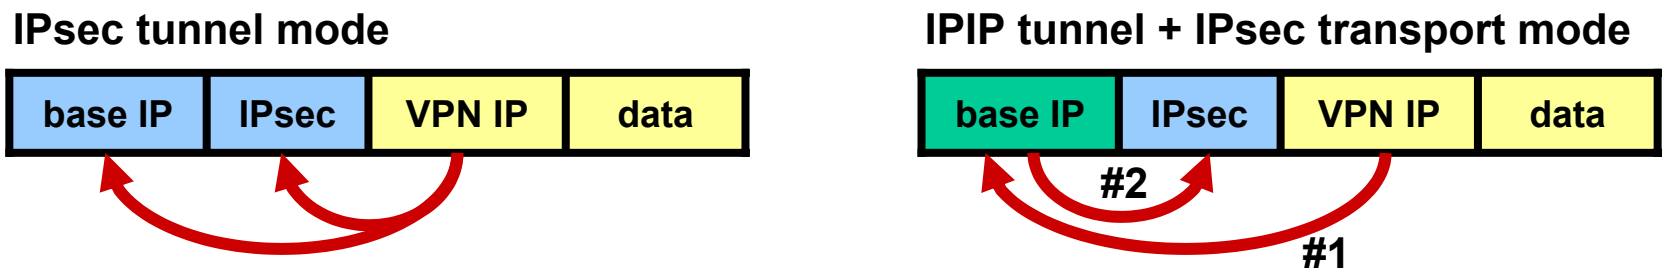

- draft-touch-ipsec-vpn-05.txt

Benefits

- IP tunnels: real interfaces with routes
  - explicit next hop
  - routing protocols and code just work
  - source address selection works
- simplifies spec
- decouples security from topology

Issues

- tunnel mode selectors more expressive
  - equivalent: policy routing + tunnel firewall
- IKE does 3 things:
  - key exchange → OK
  - tunnel management → factor out
  - policy negotiation → factor out
- NAT traversal
  - draft spec requires tunnel mode
  - equivalent: use UDP instead of IPIP

X-Bone

- parallel, secure, virtual Internets
  - IPv{4|6} with DNS, etc.
  - IPsec + dynamic routing
  - revisitation + recursion
  - web interface
- BSD, Linux (Cisco, Mac)
  - no OS changes: any IP app just works
- http://www.isi.edu/xbone/

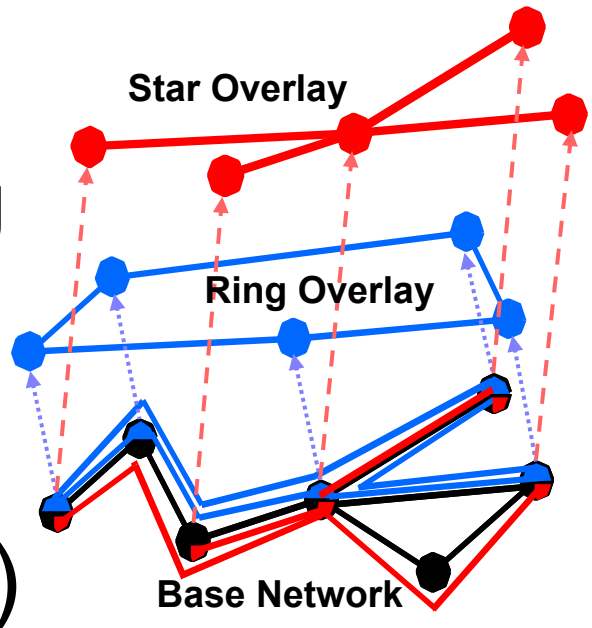

TetherNet

- true Internet behind NATs and firewalls
  - IPv{4|6}
  - multicast
  - fwd/rev DNS
  - traffic shaping
  - 802.11b AP
  - secure: IPsec for traffic, X.509 for user auth
  - web interface configuration

- http://www.isi.edu/tethernet/

Implementations: GRE-in-IPsec

- How it works
  - Creates virtual routing interface via Generic Routing Encapsulation (GRE), also called a tunnel interface
  - Makes SPD <Local GRE interface, Remote GRE Interface, GRE protocol type>
    - any traffic can pass in the IPsec tunnel w/o changing SPD
    - NEEDS ACCESS CONTROL ON GRE
  - Tunnel or Transport (more efficient) Mode

"S" = Source; "D" = Destination

GRE-in-IPsec

Transport Mode

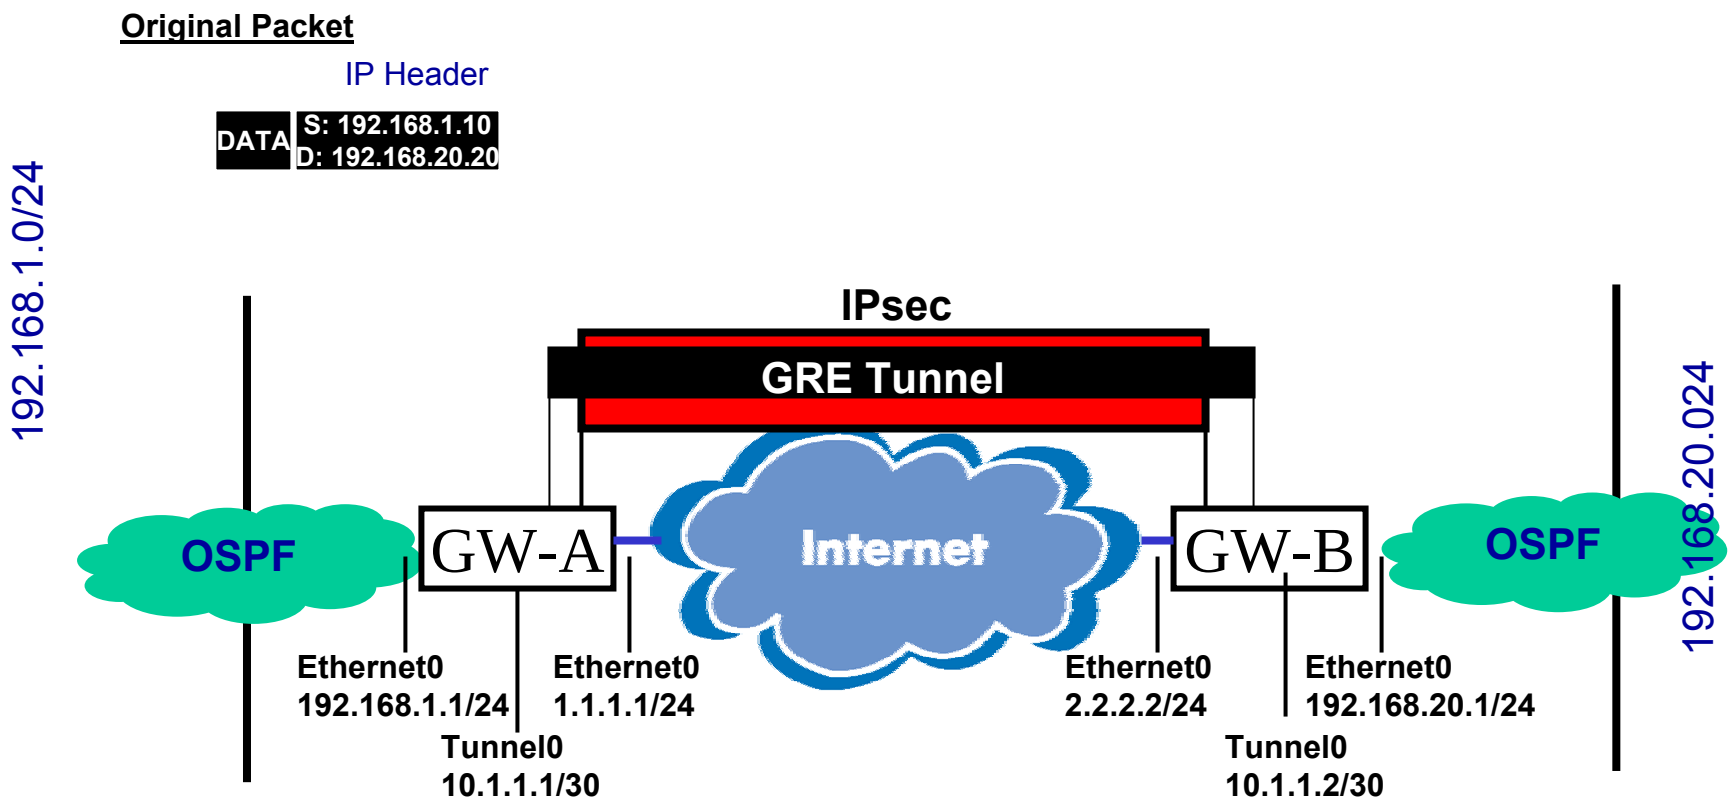

"S" = Source; "D" = Destination

GRE-in-IPsec

Transport Mode

Original Packet

IP Header

GRE Process

GRE header + New IP Header

GRE Encapsulates original IP Header and Data with a GRE header, and appends a New IP Header

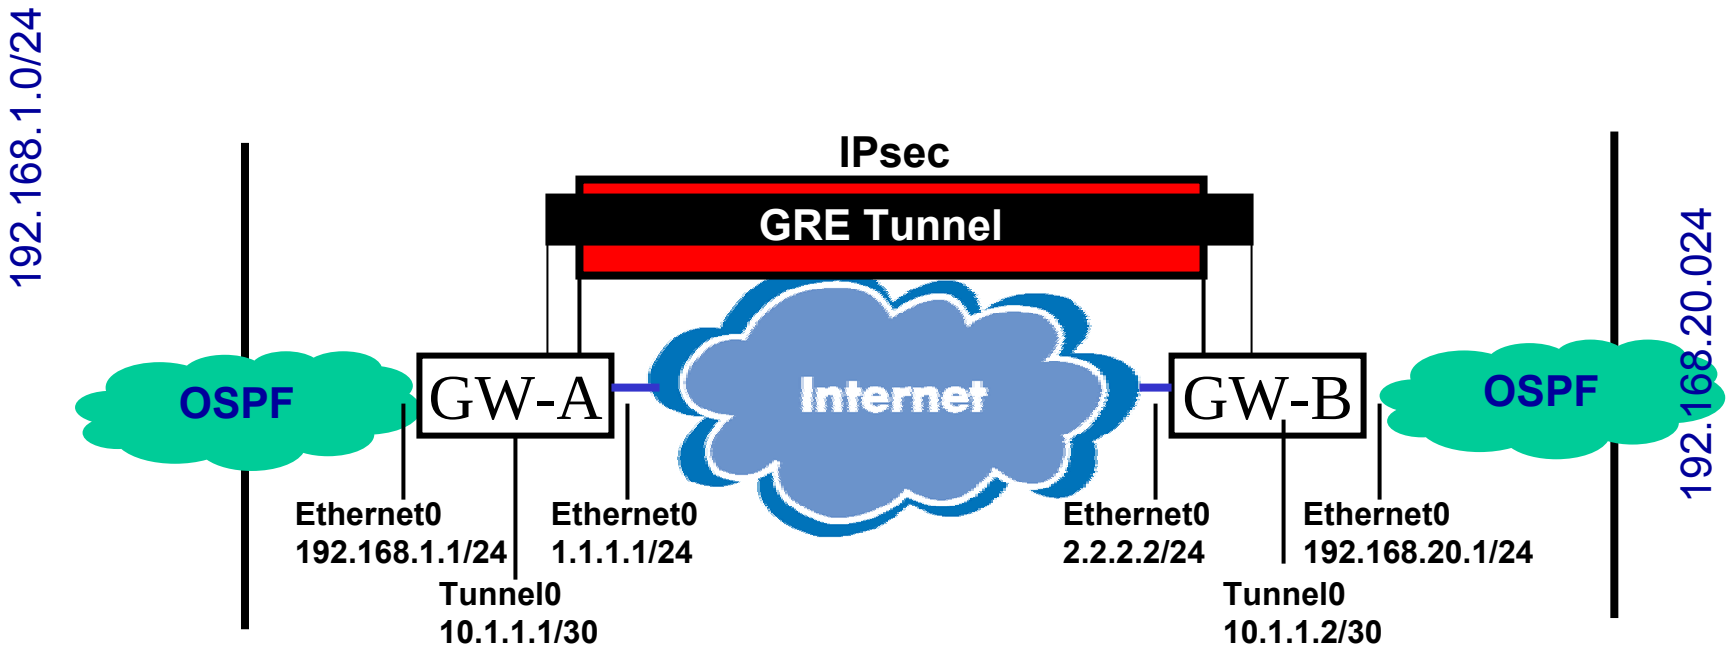

"S" = Source; "D" = Destination

GRE-in-IPsec

Transport Mode

Original Packet

IP Header

GRE Process

GRE header + New IP Header

GRE Encapsulates original IP Header and Data with a GRE header, and appends a New IP Header

IPsec Process

Encrypted

ESP inserted

IPsec Transport Mode Encrypts entire original packet + GRE header, and inserts the ESP header between encrypted payload and New IP Header

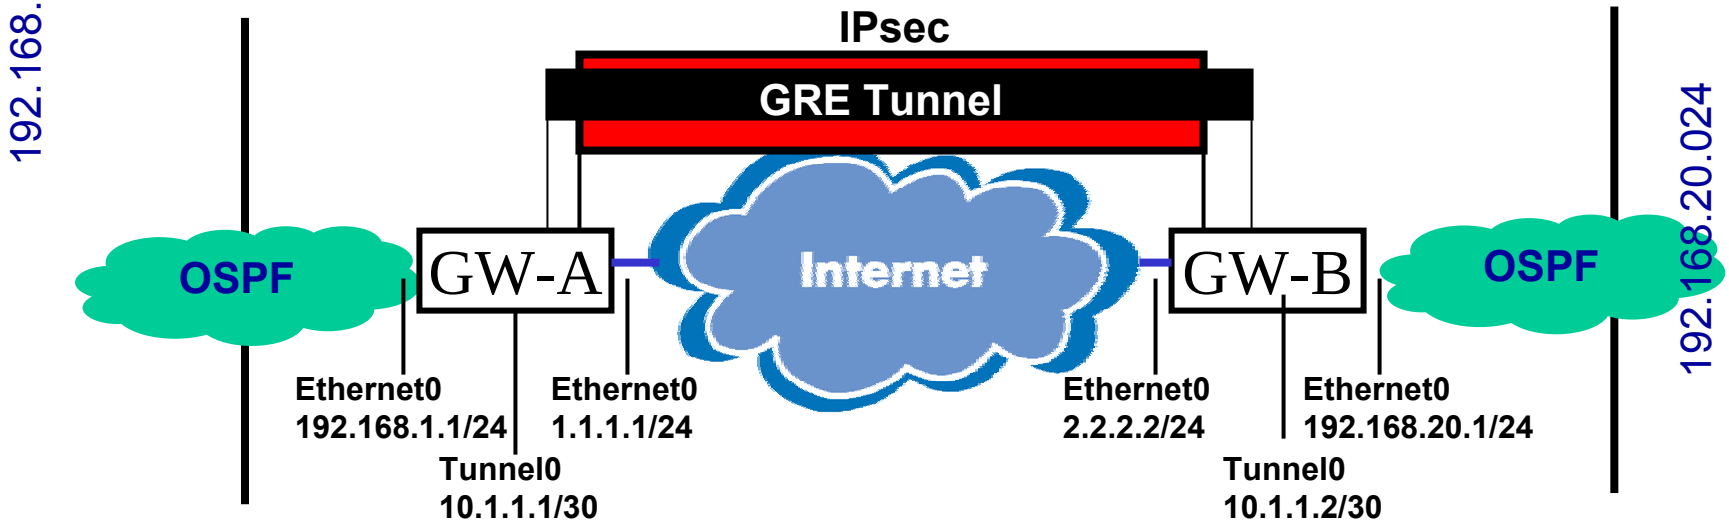

"S" = Source; "D" = Destination

GRE-in-IPsec

Transport Mode

Routing Statements:
Sent between Tunnel0's in GRE with original IP Header of S=10.1.1.1, D=10.1.1.2

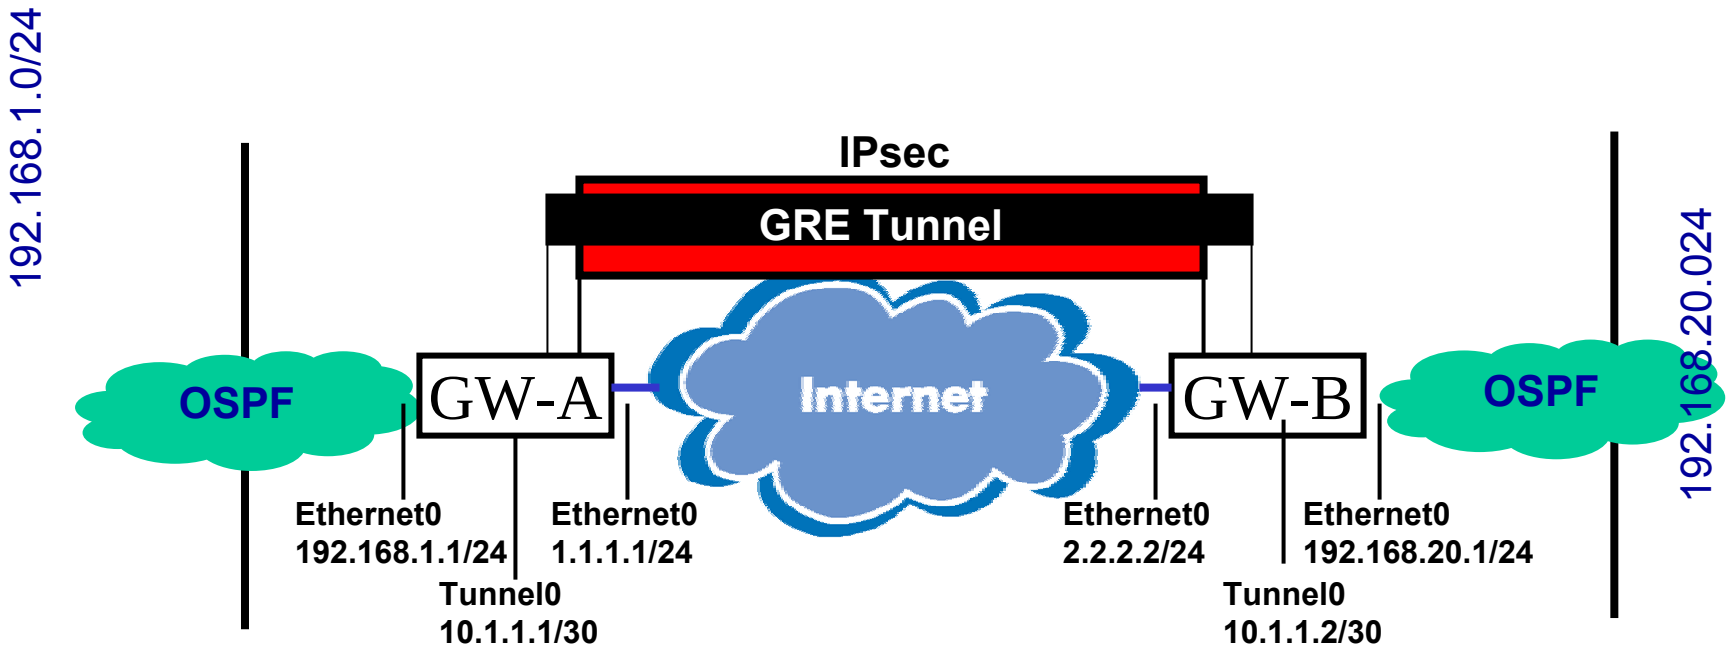

GRE Pro's & Con's

- Benefits
  - Carry non-IP traffic (only method for IS-IS)
  - Tunnel or Transport mode
- Drawbacks
  - Additional encapsulation, > overhead
    - 4 bytes for Transport
    - 20 bytes for Tunnel
  - Performance hit
    - another encapsulation to process
    - Fragmentation – offset by lowering MTU on GRE interface

Agenda

- Introductions
- Why we need dynamic routing in IPsec
- Difficulty of doing dynamic routing in IPsec
- Quick Review: IPsec Transport and Tunnel Modes
- Current Implementations of dynamic routing in IPsec
- **What's happening in IETF standards**

Changes in ESP (2401bis)

- ESP SPD lookup. Traffic Selectors in the SPD are only used to drop or permit traffic, but not used for a routing decision. Routing function exists outside of IPsec.

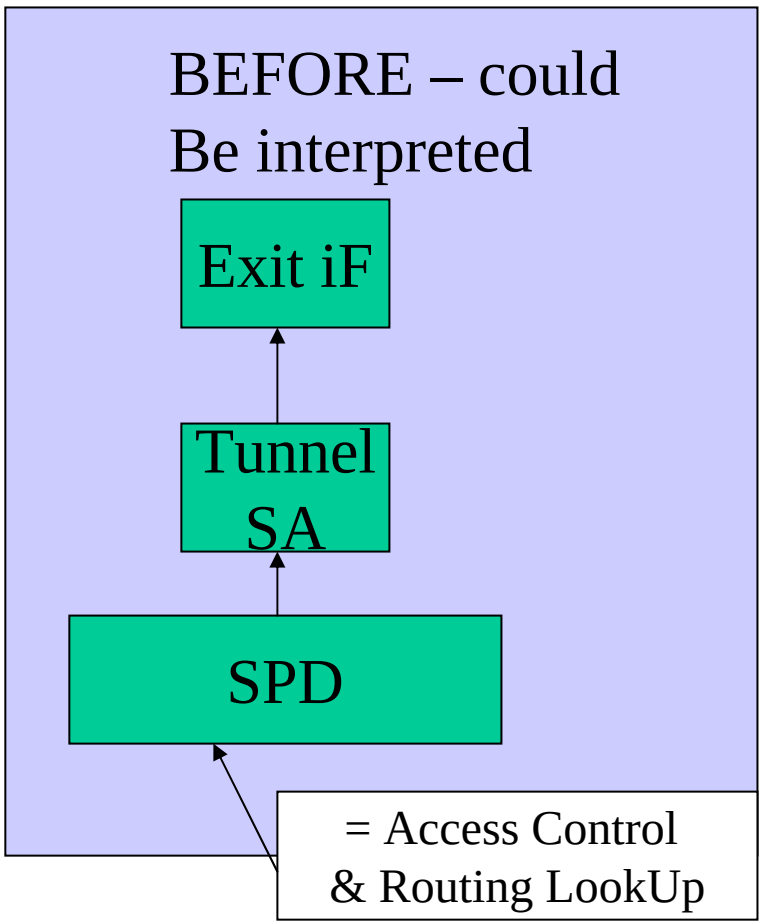

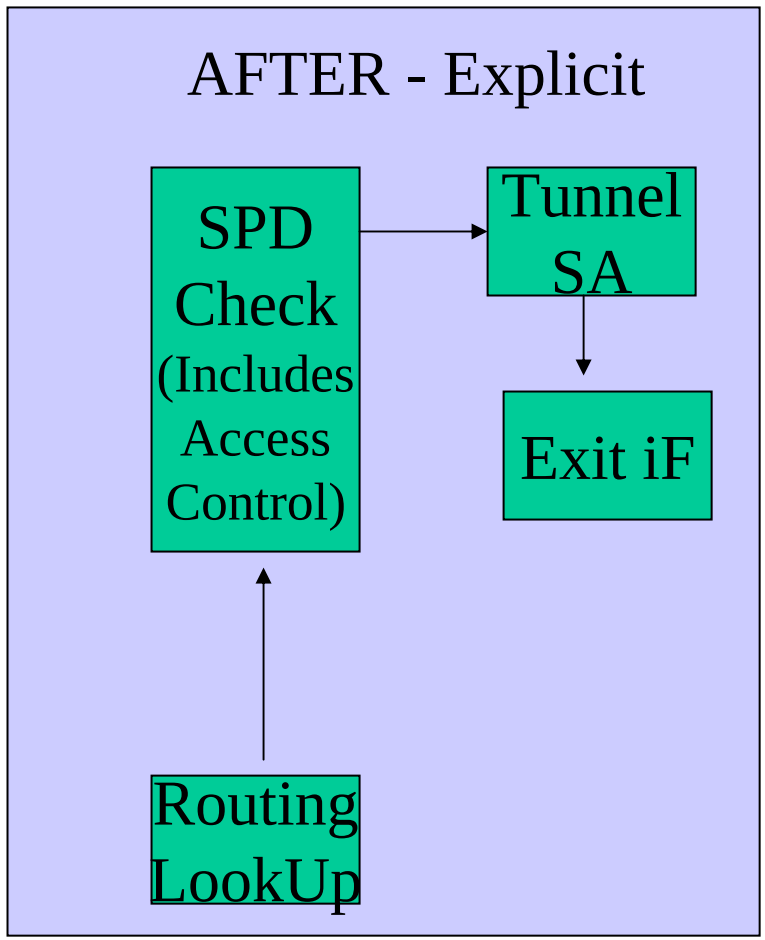

IETF - Dynamic Routing in IPsec

- Draft-knight-ppvpn-ipsec-dynroute-02.txt
- (http://www.ietf.org/internet-drafts/ *)
- Gives the gory details of using transport mode with IP-in-IP encapsulation for dynamic routing
- Describes transport of routing protocols within IPsec

Following IETF Activities

- Mailing lists and archives of Working Groups
  - IPSEC
    - General Discussion: ipsec@lists.tislabs.com
    - To Subscribe: ipsec-request@lists.tislabs.com
    - Archive: ftp://ftp.tis.com/pub/lists/ipsec
  - PPVPN (Provider Provisioned VPN)
    - General Discussion: ppvpn@nortelnetworks.com
    - To Subscribe: lyris@nortelnetworks.com
    - In Body: (UN)SUBSCRIBE ppvpn in message body
    - Archive: http://standards.nortelnetworks.com/ppvpn/index.htm
- Reading the drafts and RFCs
  - http://www.ietf.org/ID.html
  - http://www.ietf.org/rfc.html

Thank You!