NDSS'99

Network and Distributed Systems Security Symposium

Securing the Internet's Exterior Routing Infrastructure

Secure Border Gateway Protocol (S-BGP)

Dr. Charles Lynn

BBN Technologies

CLynn@BBN.Com

http://www.net-tech.bbn.com/sbgp/ndss99/index.html

Constraints and Goals

- BGP Implementation and Protocol Limitations
- Dynamic - must handle changes in topology and the addition of new networks, routers, and ASes
- Must handle current and projected usage
  - E.g., Aggregation, Communities, MPLS, Multi-Protocol, etc.
- Scalable - must handle foreseeable growth, i.e., IPv6
- Deployable - must use available technology that can be incrementally deployed
- Avoid Dependency Loops - cannot depend on inter-AS routing when initializing, e.g., non-local databases
- Leverage of off existing infrastructure
- No new trust relationships; least privilege design

Correct Operation of BGP

- Each UPDATE is intact, sent by the indicated sender, and is intended for indicated receiver
- Neighbor that sent the UPDATE was authorized to act on behalf of its AS to advertise the routing information in the UPDATE to the BGP speakers in the recipient AS
- Neighbor withdrawing a route is the advertiser for that route
- AS that generated the initial UPDATE is authorized by the Organization that owns the address space in the NLRI to represent the address space

Correct Operation of BGP (nice, but ...)

- Neighbor that sent the UPDATE, correctly applied BGP rules, local policies, etc.
- BGP speaker that received the UPDATE, correctly applied BGP rules, local policies, etc.
- Subscriber traffic forwarded by a BGP speaker is valid (not spoofed, duplicated, etc.)

We can’t enforce these aspects of correct operation because BGP affords speakers considerable latitude with regard to local policy, ASes do not tend to make public their local policies, and because validation and tracking of subscriber traffic is impractical

Design Overview

- IPsec --> authenticity, integrity, and anti-replay protection of peer-to-peer communication
- Public Key Infrastructures (PKIs) --> secure identification of BGP speakers and of owners of ASes and of address blocks
- Attestations --> authorization of the subject (by the issuer) to advertise the specified address blocks, or ASes in the AS Path
  - Can also protect other Path Attributes
- Validation of UPDATEs using certificates and attestations
- Distribution of countermeasures information --> certificates, CRLs, attestations

IP Address Allocation Example

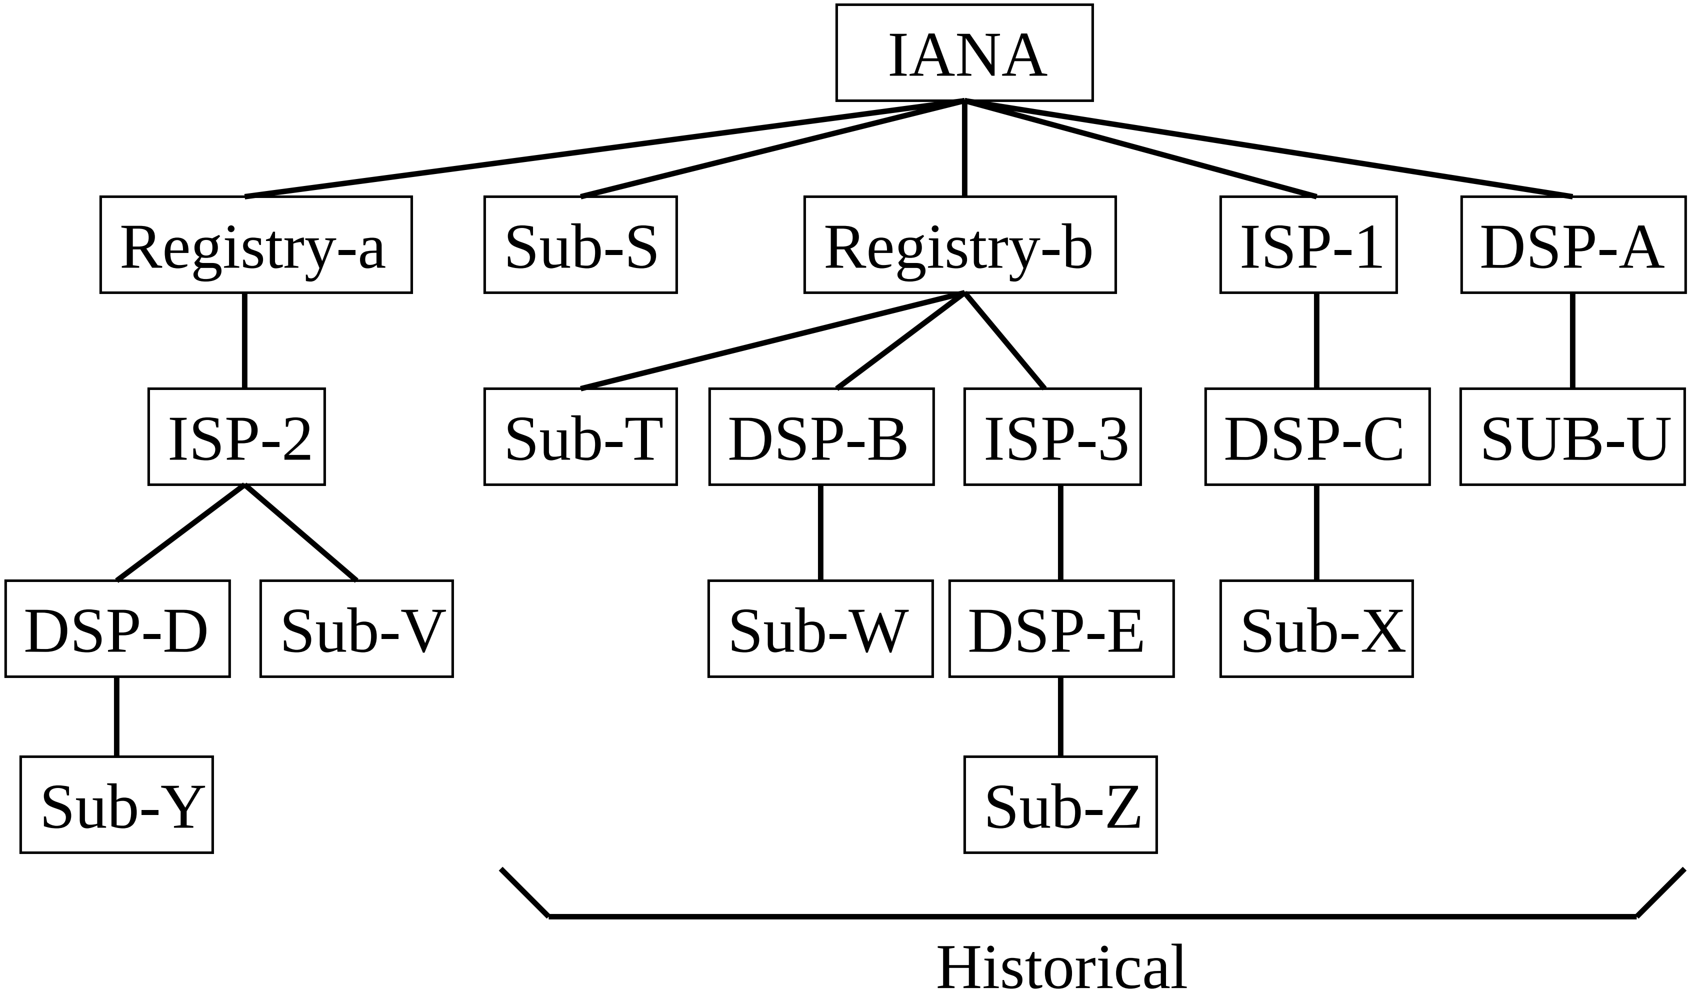

IP Address Allocation PKI Example

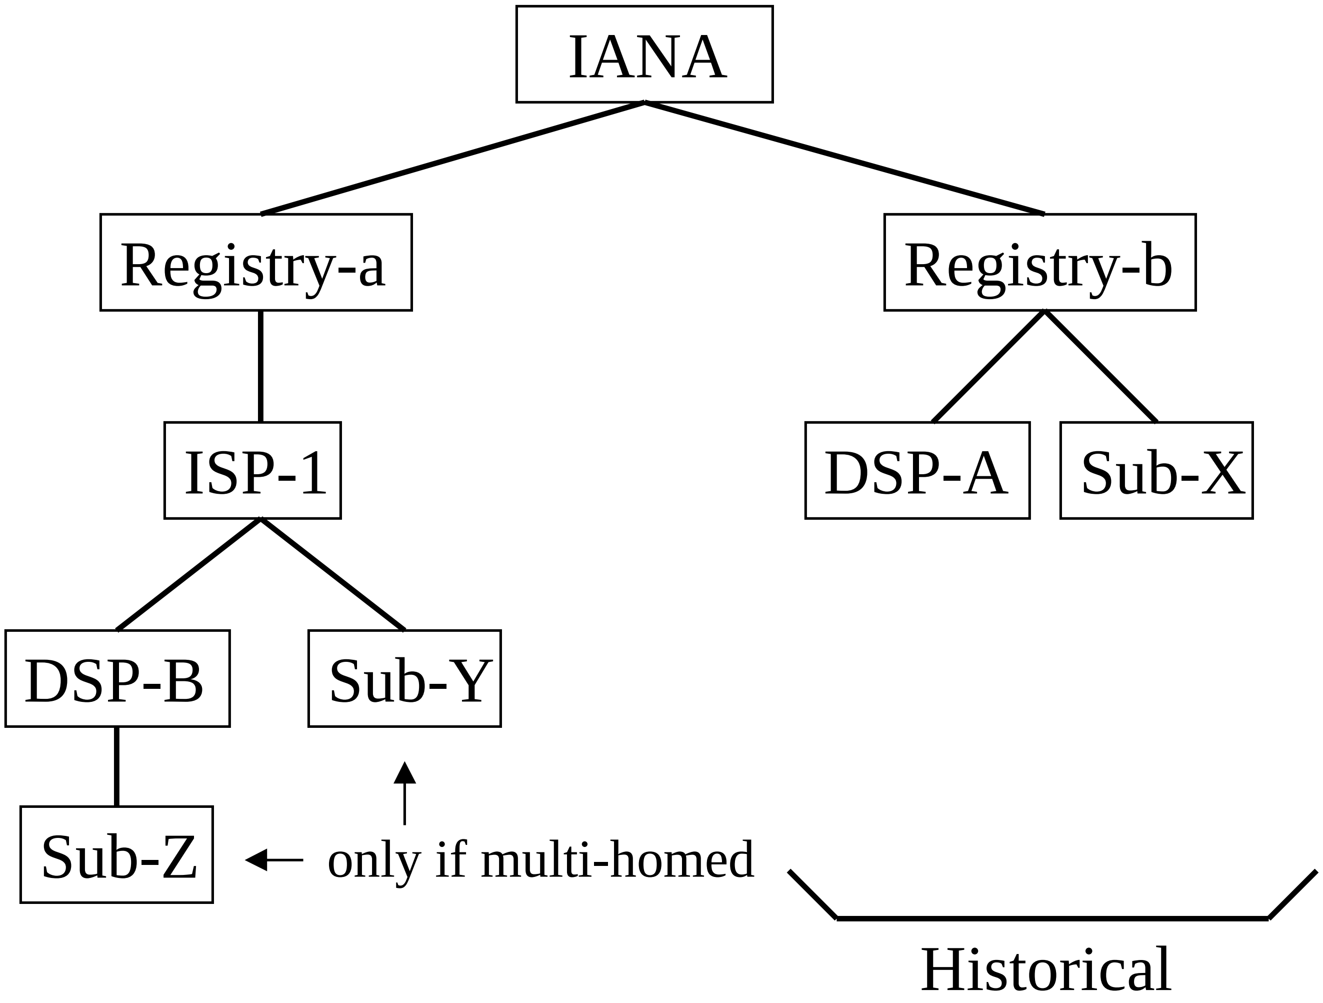

IANA

Registry-a

Registry-b

ISP-1

DSP-A

Sub-X

DSP-B

Sub-Y

Sub-Z

only if multi-homed

Historical

Address Certificates

	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all addr
Registry Certificate	IANA	Registry	addr blocks
ISP/DSP Certificate	**Registry (or IANA)**	ISP/DSP	addr blocks
Subscriber Certificate	**ISP/DSP (or Registry, IANA)**	Subscriber	addr blocks

AS # Allocation and Router Example

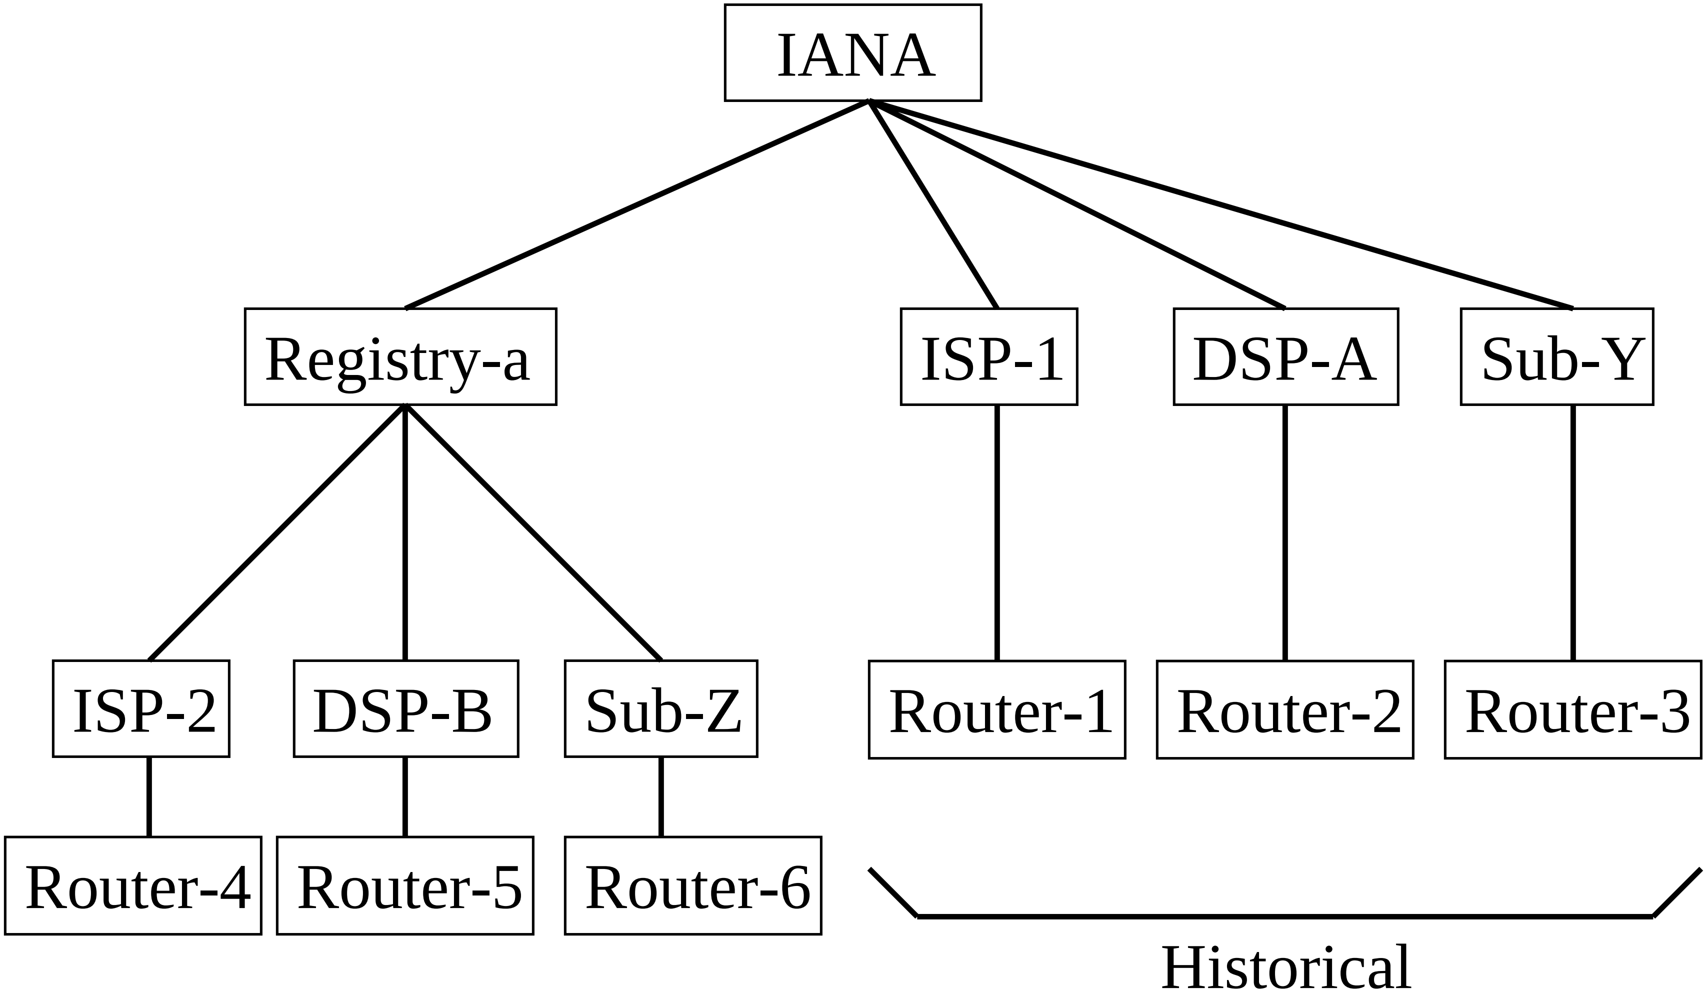

AS # Allocation and Router PKI Example

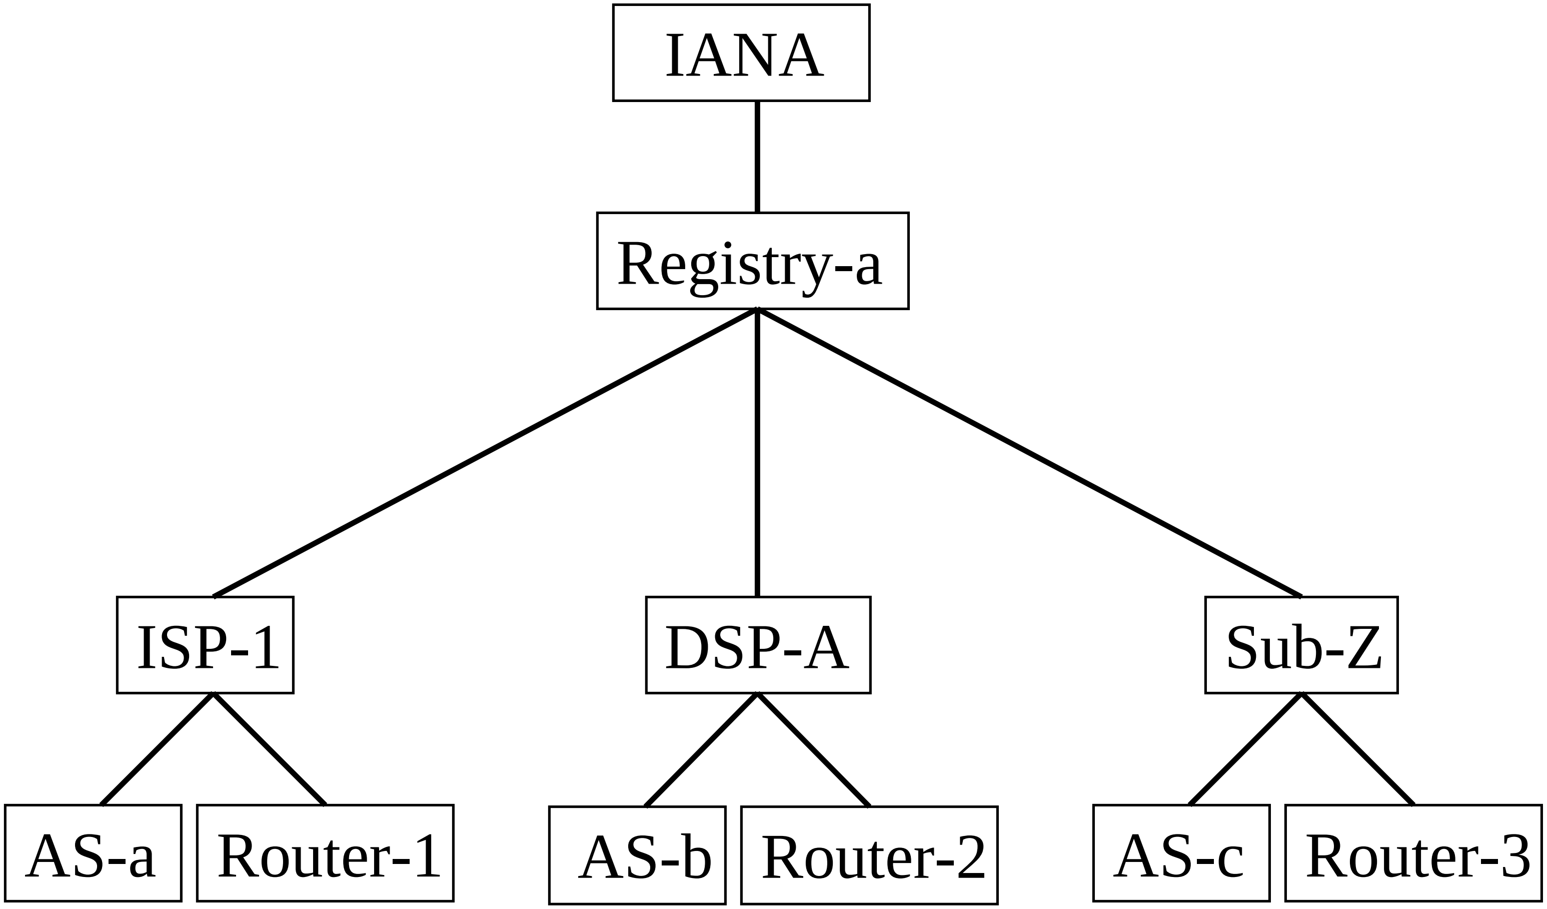

AS and Router Certificates

	Issuer	Subject	Extensions
Root Certificate	IANA	IANA	all ASes
Registry Certificate	IANA	Registry	ASes
AS Owner Certificate	Registry (or IANA)	ISP/DSP or Subscriber	ASes
AS Certificate	ISP/DSP or Subscriber	AS	
Router Certificate	ISP/DSP or Subscriber	Router*	AS, RtrId

* the subject name could be fully-qualified DNS name

Attestations -- Overview

- Address Attestations
  - Used to validate that a destination address block is being originated by an authorized AS
- Route Attestations
  - Used to validate that an AS is authorized to use an AS Path
- Each UPDATE includes one or more Address Attestations and a set of Route Attestations
- These are carried in a new, optional, transitive BGP Path Attribute

Address Attestation

- Includes identification of:
  - address blocks
  - their owner’s certificate
  - AS authorized to originate (advertise) the address blocks
  - expiration date/time
- Indicates that the AS listed in the attestation is authorized by the owner to originate/advertise the address blocks in an UPDATE
- Digitally signed by owner of the address blocks, traceable up to the IANA via a certification path
- Used to protect BGP from erroneous UPDATEs (authenticated but misbehaving or misconfigured BGP speakers)

Route Attestation

- Includes identification of:
  - AS's or BGP speaker's certificate issued by the AS owner
  - the address blocks and the AS Path (ASes) in the UPDATE
  - the AS number of the receiving (next) neighbor
  - expiration date/time
- Indicates that the BGP speaker or its AS authorizes the receiver's AS to use the AS Path & NLRI in the UPDATE
- Digitally signed by owner of the BGP speaker (or its AS) distributing the UPDATE, traceable to the IANA ...
- Used to protect BGP from erroneous UPDATEs (authenticated but misbehaving or misconfigured BGP speakers)

Encoding of Attestations

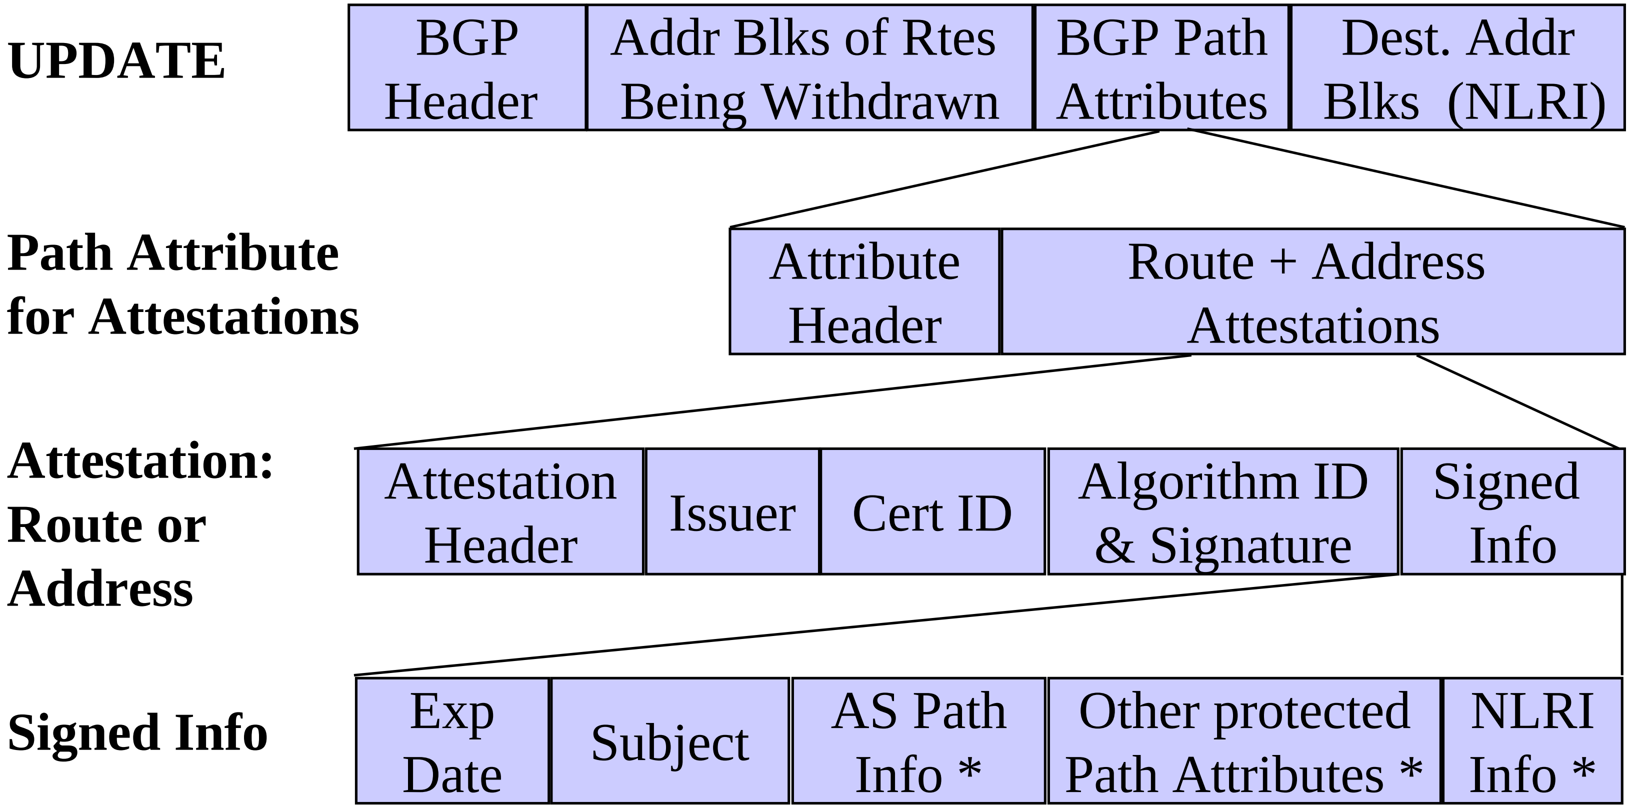

UPDATE: BGP Header | Addr Blks of Rtes Being Withdrawn | BGP Path Attributes | Dest. Addr Blks (NLRI)

Path Attribute for Attestations: Attribute Header | Route + Address Attestations

Attestation: Route or Address: Attestation Header | Issuer | Cert ID | Algorithm ID & Signature | Signed Info

Signed Info: Exp Date | Subject | AS Path Info * | Other protected Path Attributes * | NLRI Info *

*explicit in the aggregation case, or if Path Attribute changes unpredictably

Detail of Attestation Path Attribute

Path Attribute	Flags	Attestation	Len : (Ext Len)					
RA=8 or AA=9	Type	Attestation Len	L	# Attestations following				
Issuer	0 0 0 1	Issuer Len	AF/Type					
	AS Number	IPv4 Address						
	IPv6 Address	DNS Name (nil terminated)						
Signature	0 0 1 0	Signature Len	Signature Algorithm Identifier					
	Actual Data Length Signed	Key Hash	Coverage Len					
Coverage	N 1 1 0 0 0 1 1 1 0 0 1 0 0 1 0 … Path Attribute bit mask							
	Signature bytes (DSA 40 bytes)							
Not Valid Before	0 0 1 1	Len	R	YrBef	Month	Day	Hour	Minutes
Not Valid After	Year	Month	Day	Hour	Minutes			
Subject	0 1 0 0	Subject Len	AF/Type					
	AS Number	IPv4 Address						
	IPv6 Address	DNS Name (nil terminated)						
Explicit Path Attributes	0 1 0 1	Len present	Actual covered Path Attributes					
	from UPDATE, including Flags and Attribute Number							
	(NLRI encoded as: x50, x00, 16-bit length, Prefixes)							
	Additional Attestations							

Signed Data

Signature
covers
all the
Attributes
Identified
By Coverage,
not just
explicit
data, if any

Propagation of an S-BGP UPDATE

# seq:5432221,nlri:a,b

% seq:43222,nlri:a,b

signed data (implicit)

$ seq:32221,nlri:a,b

= seq:21,nlri:a,b

Hdr
seq:432221
RA:r7 5 #
RA:r5 4 %
RA:r3 3 $
RA:as1 2 =
AA:orga 1 a
AA:orgb 1 b
NLRI:a,b

Hdr
seq:32221
RA:r5 4 %
RA:r3 3 $
RA:as1 2 =
AA:orga 1 a
AA:orgb 1 b
NLRI: a,b

Hdr
seq:2221
RA:r3 3 $
RA:as1 2 =
AA:orga 1 a
AA:orgb 1 b
NLRI: a,b

Hdr
seq:1
RA:as1 2 =
AA:orga 1 a
AA:orgb 1 b
NLRI: a,b

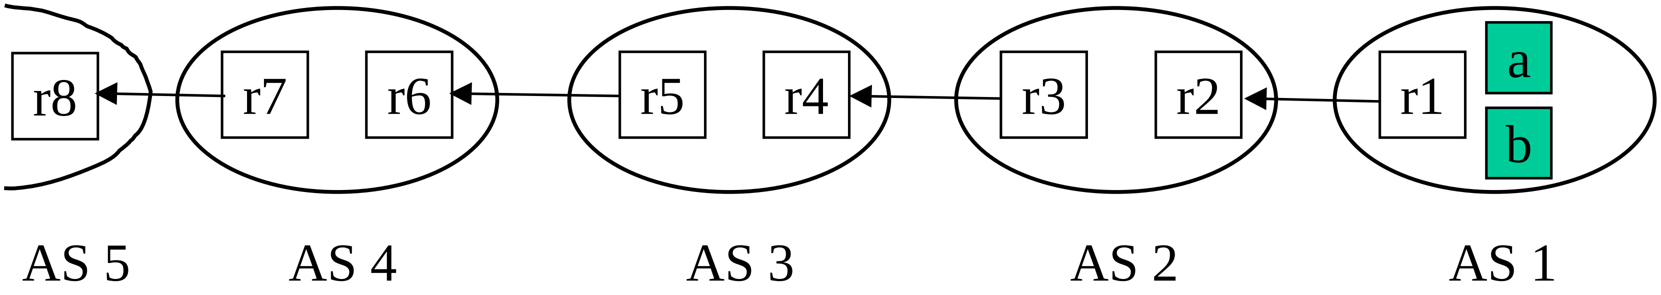

Validating a Route

To validate a route from ASn, ASn+1 needs:

- 1 address attestation from each organization owning an address block(s) in the NLRI
- 1 address allocation certificate from each organization owning address blocks in the NLRI
- 1 route attestation from every AS along the path (AS1 to ASn), where the route attestation for ASk specifies the NLRI and the AS Path up to that point (AS1 through ASk+1)
- 1 certificate for each AS or router along the path (AS1 to ASn) to use to check signatures on the route attestations
- and, of course, all the relevant CRLs must have been verified

S- BGP Path Aggregation Example

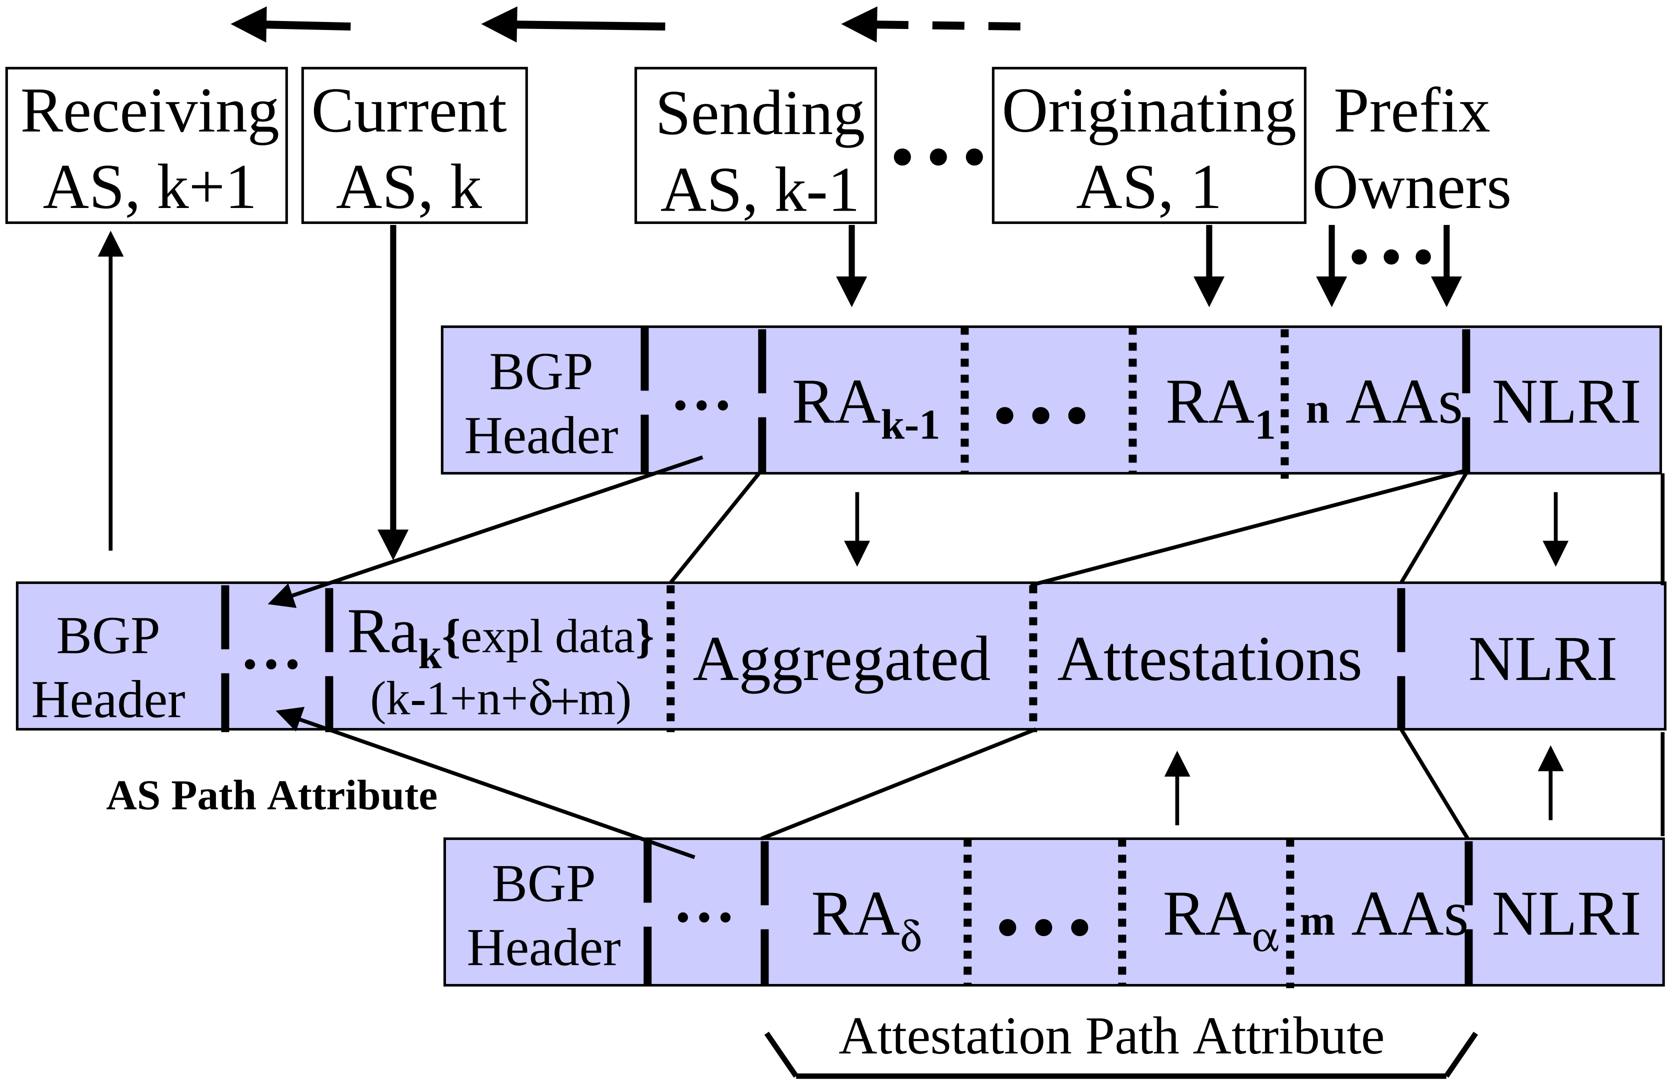

Performance Issues -- Resources

- Certificates (generation and signing done offline)
  - disk space for storing certificates
  - CPU resources for validating certificates
- CRLs (generation and signing done offline)
  - disk space for storing CRLs
  - CPU resources for validating CRLs
- Attestations
  - RIB memory space for storing attestations
  - disk space for faster recovery from router reboot (optional)
  - CPU resources for signing and validating attestations
  - resources for transmitting attestations (to make this a dynamic system)

Performance -- Certificates

- Processing -- certificates and CRLs are signed infrequently; this should be done off-line (and not by routers)
- Storage:
  - ~30 Mbytes for ~65K Certificates
  - ~2 Mbytes for ~3K CRLs
  - DNS or Certificate server -- 1 entry/address block, 1 entry/AS, 1 entry/BGP-speaker in an AS
- Transmission bandwidth -- An UPDATE will not hold the certificates needed to validate an average route. Therefore, certificates will have to be cached. Certificates will be transmitted at a low frequency except at startup (or preloaded from the NOC).

*** Estimates are based on observed MRT data from Jan 1998**

Performance -- Attest.’s (worst case)

- Transmission bandwidth
  - countermeasures information adds ~400 bytes to a typical (2.6 ASes in path) UPDATE of 63 bytes, but UPDATEs represent a very small portion of all traffic
- Processing (using DSA/SHA-1, ~ 0.15 second each/75 MHz)
  - Per boot: ~7.5 hours to validate LOC-RIB with 50000 NLRI
  - ~1.8 hours to generate and sign route attestations
  - Per day: 5-6 minutes to validate UPDATEs; 5-6 minutes to generate/sign attestations (assuming 10 UPDATEs/second)
- Storage
  - address attestations -- ~7 Mbytes
  - route attestations -- ~20 Mbytes for ADJ-RIB per BGP peer; 80 Mbytes (4 peers) to 500 Mbytes (25 peers)

Optimizations

- Cache previously validated routes (attestations) to avoid re-validation, e.g., if router crashes or link to neighbor is lost
- Required BGP caching covered ~89% of UPDATEs
  - caching last two distinct paths from each peer hit another 6% of UPDATEs
- Mark routes “withdrawn” for use later, e.g., if link flapped, to speed up validation for reinstated routes
- Defer verification until route is put into LOC-RIB
- Background verification of alternate routes
- Exploit previously validated common tail attestations

Optimizations (continued)

- Attestation design eliminates redundant information in common case, i.e., just prepending local AS to path
- Keep only needed certificate fields in S-BGP databases
- Offload generation/signing of route attestations, e.g., from routers to AS
- Heuristics to guide which prefixes are aggregated in an UPDATE -- this is aimed at reducing the amount of information that has to be made explicit when an aggregate encounters a preferred more-specific route

Other Performance Savings

- Most organizations will obtain their address blocks from their provider --> they do not need to provide address attestations and do not need certificates
- Most organizations/users are singly homed (see above)
- Limit where UPDATE validation occurs, e.g., not needed for
  - Organizations that use default routing, e,g,, singly homed "leaf" organizations or singly homed DSPs
  - multi-homed DSP -- check only if receive >1 route to same address blocks
- Cryptographic hardware for signing and verification

Proof of Concept

- DARPA and NSA are funding prototype development
  - Prototype code will be available, e.g., in GateD
- Replay some of the Merit historical data
- Deploy in the wide-area DARPA CAIRN testbed
  - PC-based routers running FreeBSD/mrtd/gated
  - Partition testbed into several ASes or a confederation
  - Peer with (import from) diverse BGP speakers in the Internet
  - Insert “missing” attestations on the fly, e.g., from local cache
  - Do nasty things to routers, links, BGP sessions (w/IPSec off :-)
  - Find performance problems and devise optimizations for them
- Have some ISPs evaluate it
  - Monitoring mode vs. enforcement mode

Benefits of S-BGP

- Secure identification of entities participating in global internet routing, e.g., prefix owners, AS number owners, AS administrations, routers speaking for an AS
- Secure authorization of:
  - AS to advertise an address prefix
  - AS to use a route
- Integrity of peer communications and advertised routes

Questions?

An Example BGP Topology

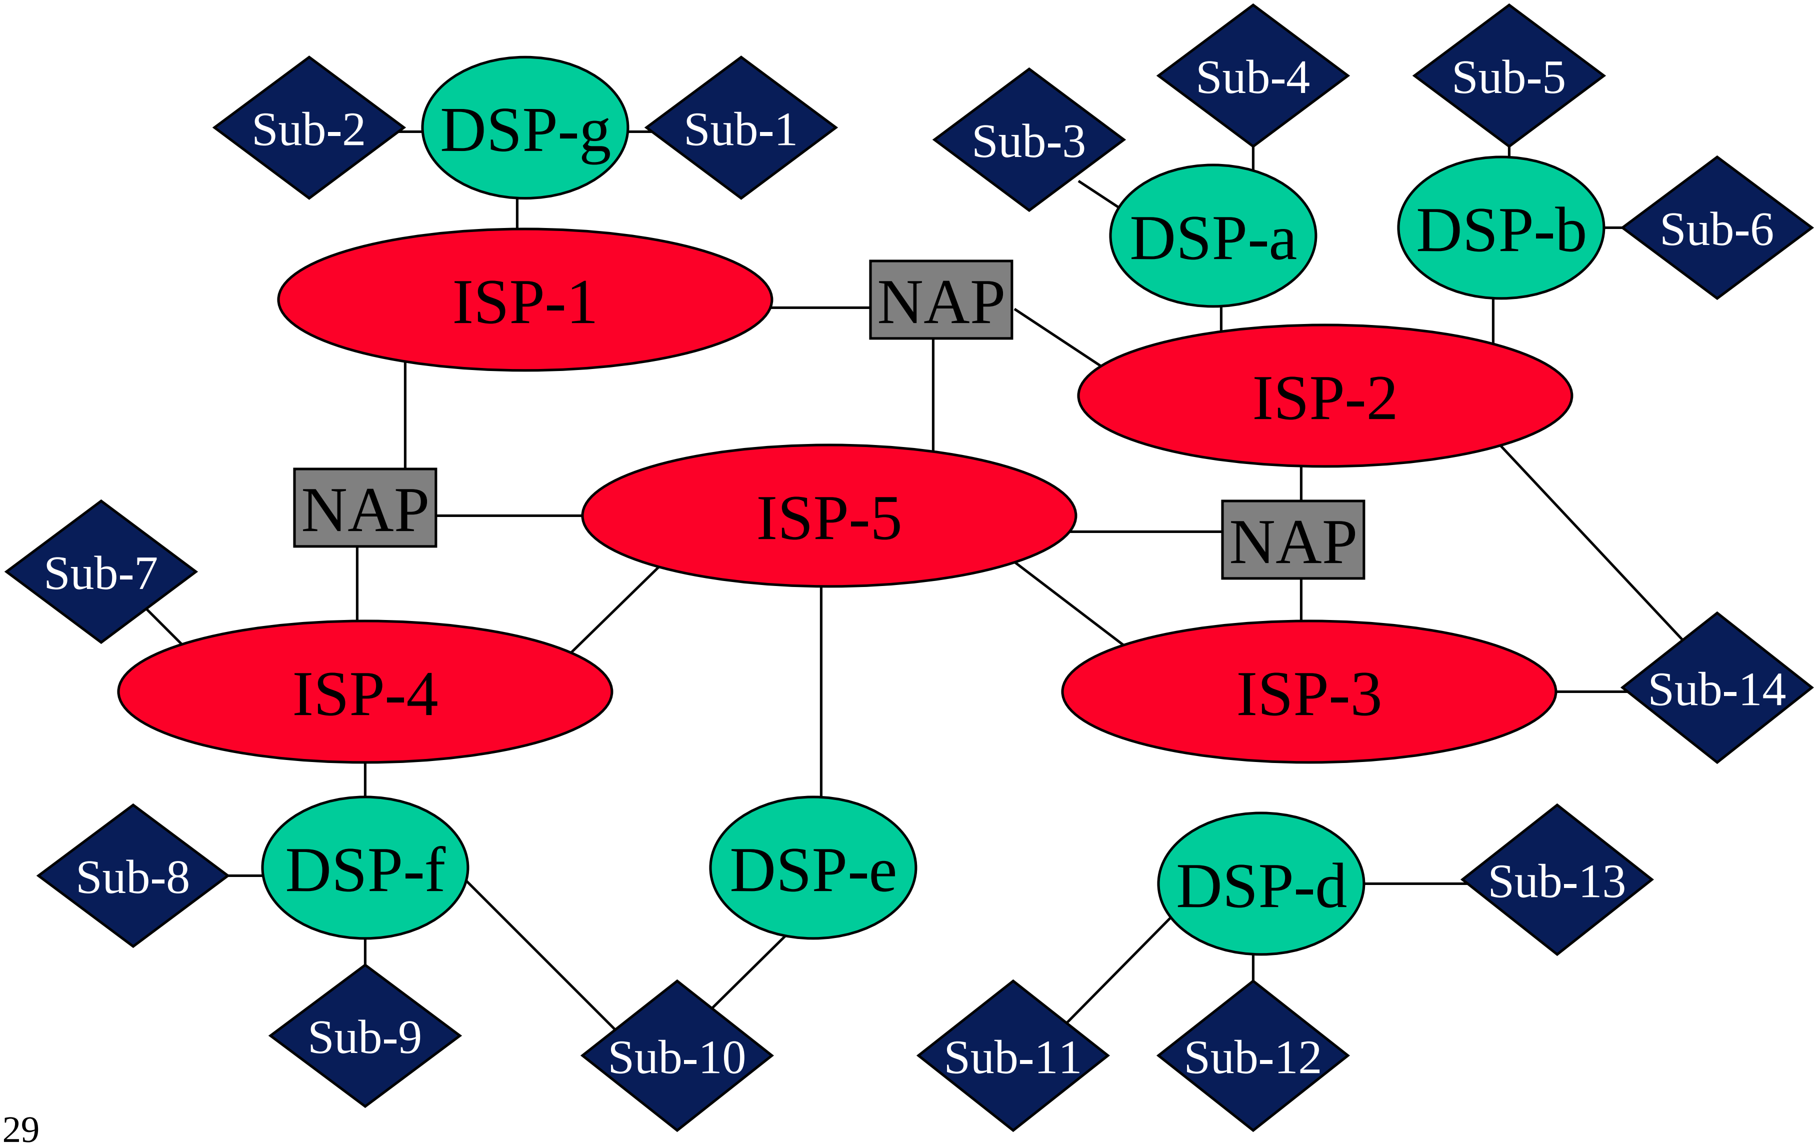

Basic BGP Model

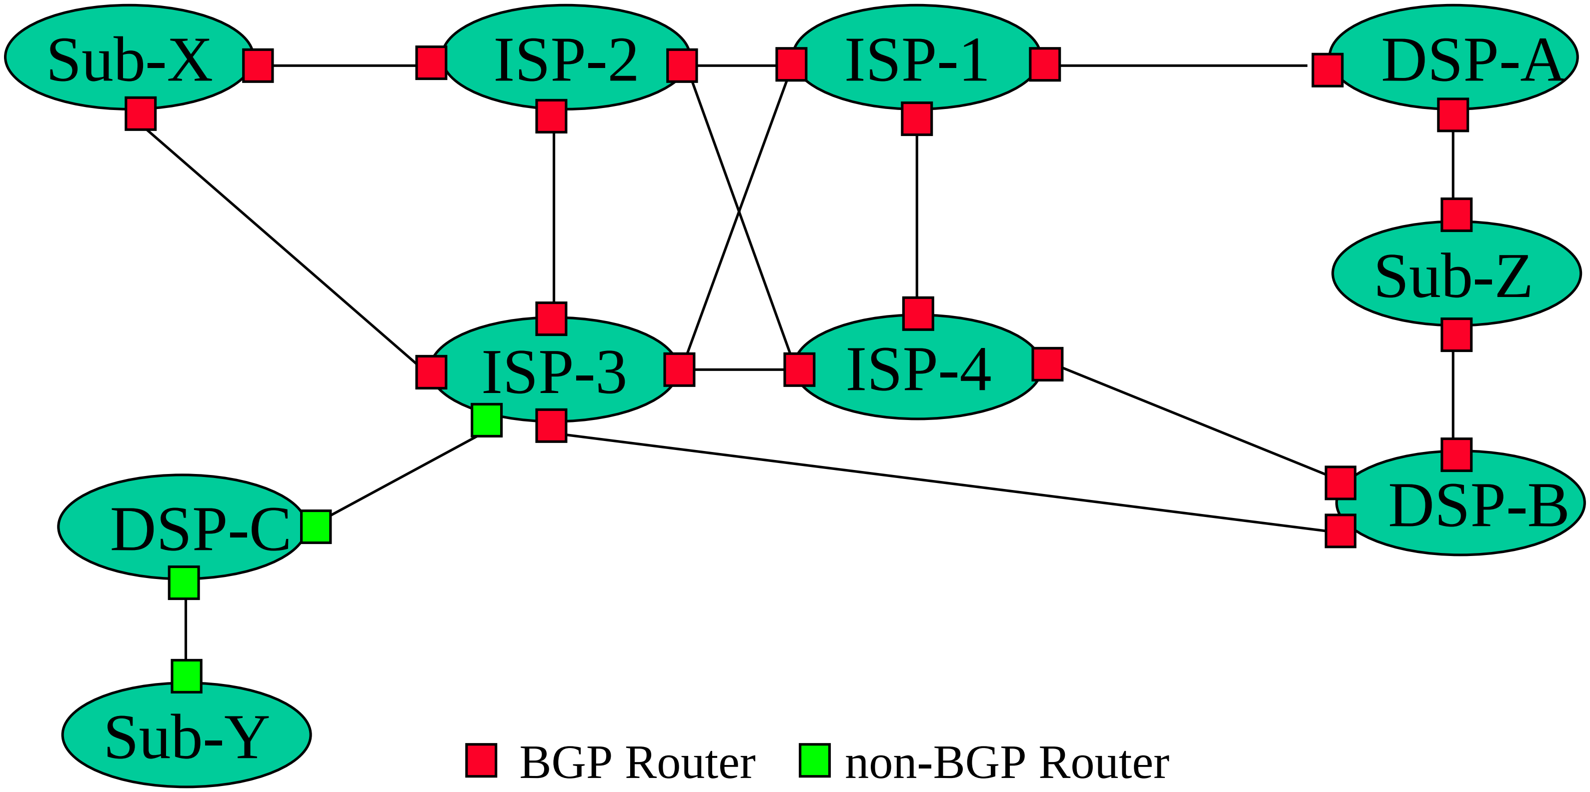

Format of Attestation Path Attribute

```
        +---------------+---------------+---------------+---------------+
PathAtt |     Flags     |     Type      |      Len      :   (ExtLen)    |
        +=======+=======+===============+=+=============+===============+
Type    |RA8/AA9|    Attes Len          |L| # Attestations following    |
        +-------+-------+---------------+-+-------------+---------------+
Signer  |0 0 0 1|      Len              |          AF/Type              |
        +-------+-------+---------------+---------------+---------------+
        |            AS Num             :              IPv4             :
        +---------------+---------------+---------------+---------------+
        :             IPv6              DNS Name (nil terminated)       |
        +-------+-------+---------------+---------------+---------------+
Sig     |0 0 1 0|      Len              |        Sig Alg Id             |
        +-------+-------+---------------+---------------+---------------+
        | Actual Data Length Signed     |   Key Hash    | Coverage Len  |
        +---------------+---------------+---------------+---------------+
(cvrg)  |N 1 1 0 0 0 1 1 1 0 0 1 0 0 1 0 … Path Attribute bitmask      |
        +---------------+---------------+---------------+---------------+
        :            Signature Bytes          (DSA 40 bytes)            :
        +-------+-------+-+-----+-------+---------+-----+---+-----------+  ----
NVB     |0 0 1 1| Len   |R|YrP:3|Month:4| Day:5   |Hour:5 | Minute:6  |  ^
        +-------+-------+-+-----+-------+---------+-----+---+-----------+  |
NVA     |      Year:12          |Month:4| Day:5   |Hour:5 | Minute:6  |  |
        +-------+-------+-------+-------+---------+-----+---+-----------+  |
Subj    |0 1 0 0|      Len              |         AF/Type               |Signed
        +-------+-------+---------------+---------------+---------------+  |
        |            AS Num             :              IPv4             :  |
        +---------------+---------------+---------------+---------------+  |
        :             IPv6              DNS Name (nil terminated)       |  |
        +-------+-------+---------------+---------------+---------------+  |
ExpPA   |0 1 0 1|      Len (here)       |Actual Covered Path Attributes:   :
        +-------+-------+---------------+---------------+---------------+  :
        :from UPDATE, including Flags and Attribute Number (may omit);  :  :
        +---------------+---------------+---------------+---------------+  :
        :    "x50 x00, 16-bit len, NLRI (may omit);                     |  v
        +===============+===============+===============+===============+ ----
        :                Additional Attestations                        :
        +===============+===============+===============+===============+
```

PKI: Certificates

	Issuer	Subject	Public Key	Extensions	Signer
1.	IANA	IANA	IANA		IANA
2.	IANA	Registry	Registry		IANA
3.	Reg/Org	Org/Sub	Org/Sub	addr blks	Reg/Org
4.	Registry	Organiz.	Organiz.	AS	Registry
5.	Organiz.	AS	AS	AS**	Organiz.
6.	Organiz.	Router	Router	AS	Organiz.

UPDATE Processing

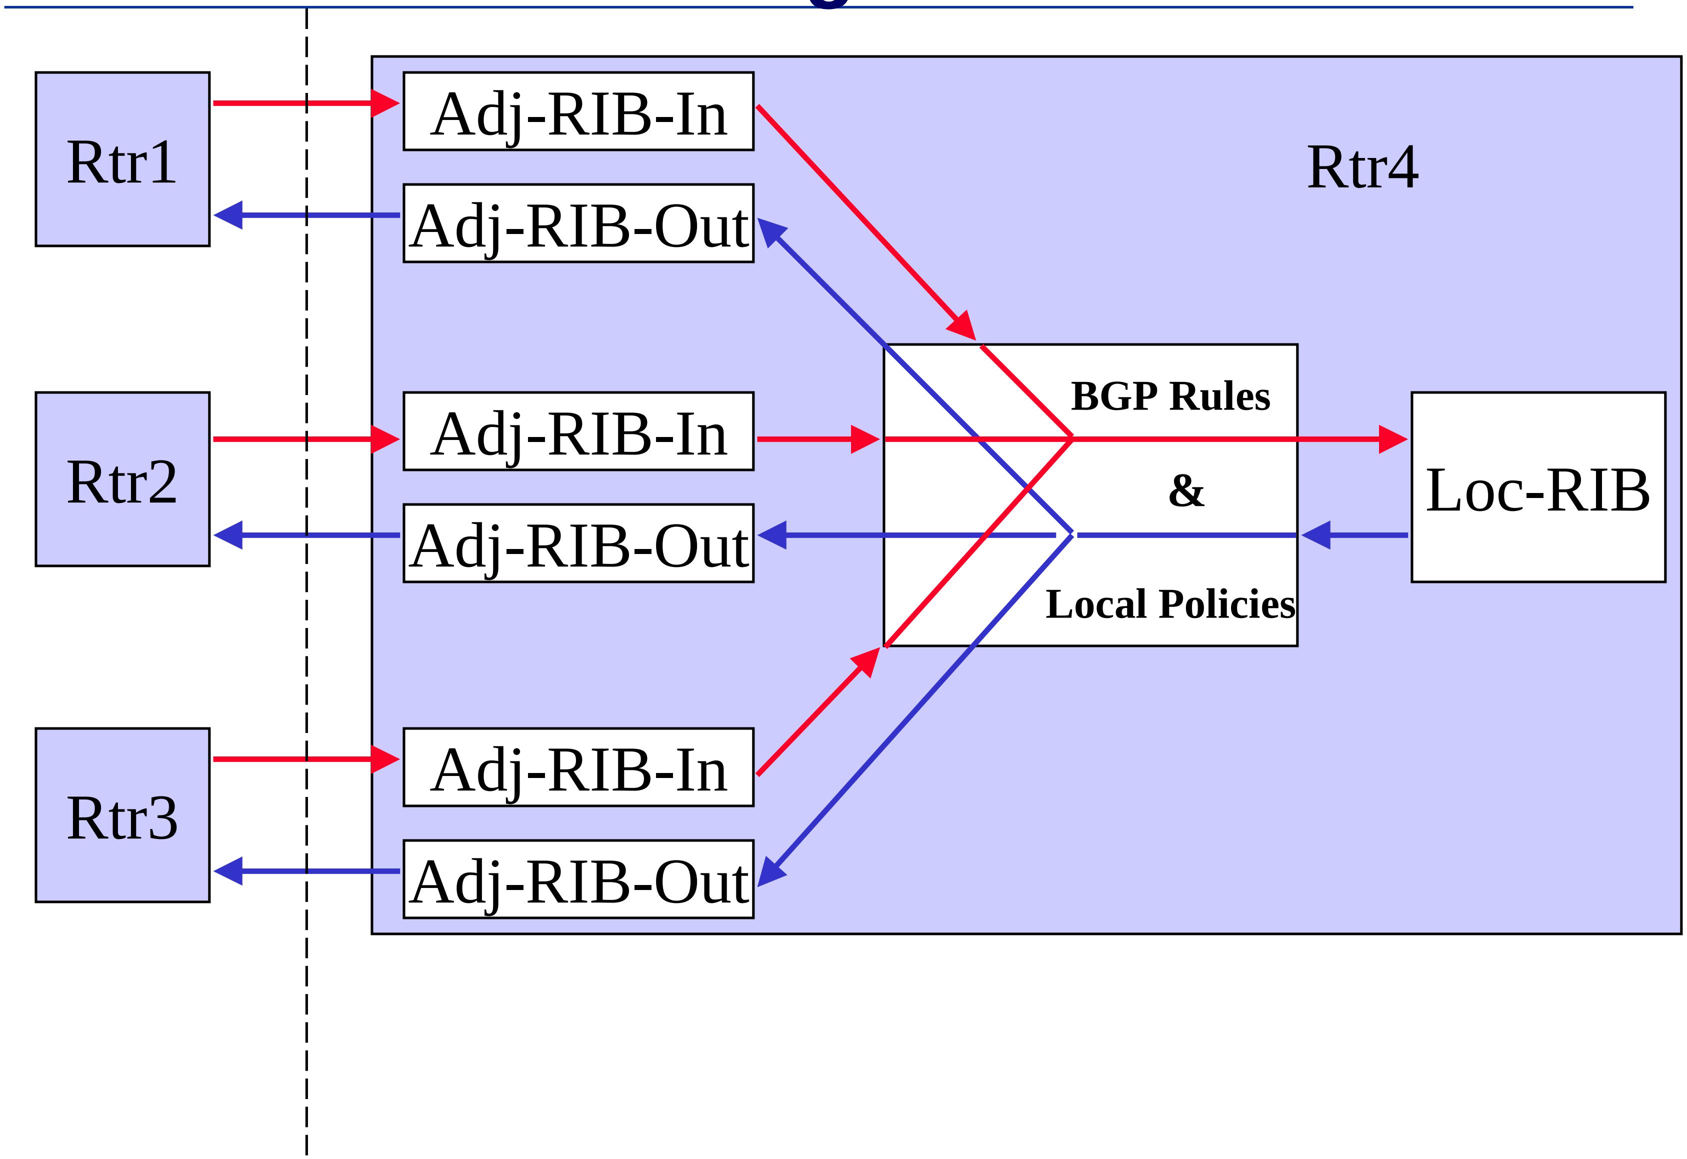

Sample Auxiliary Box Topology

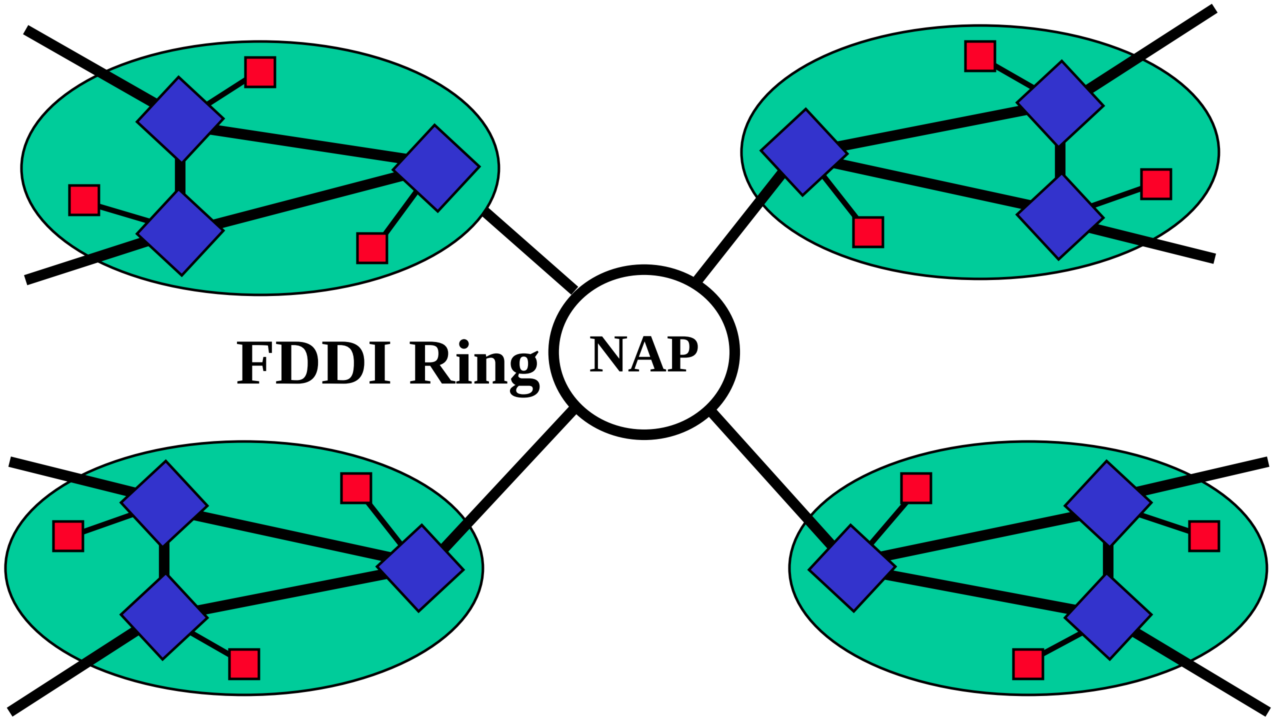

◆ Border Router

🟥 S-BGP box

⬭ AS

Sample Peering of one Auxiliary Box

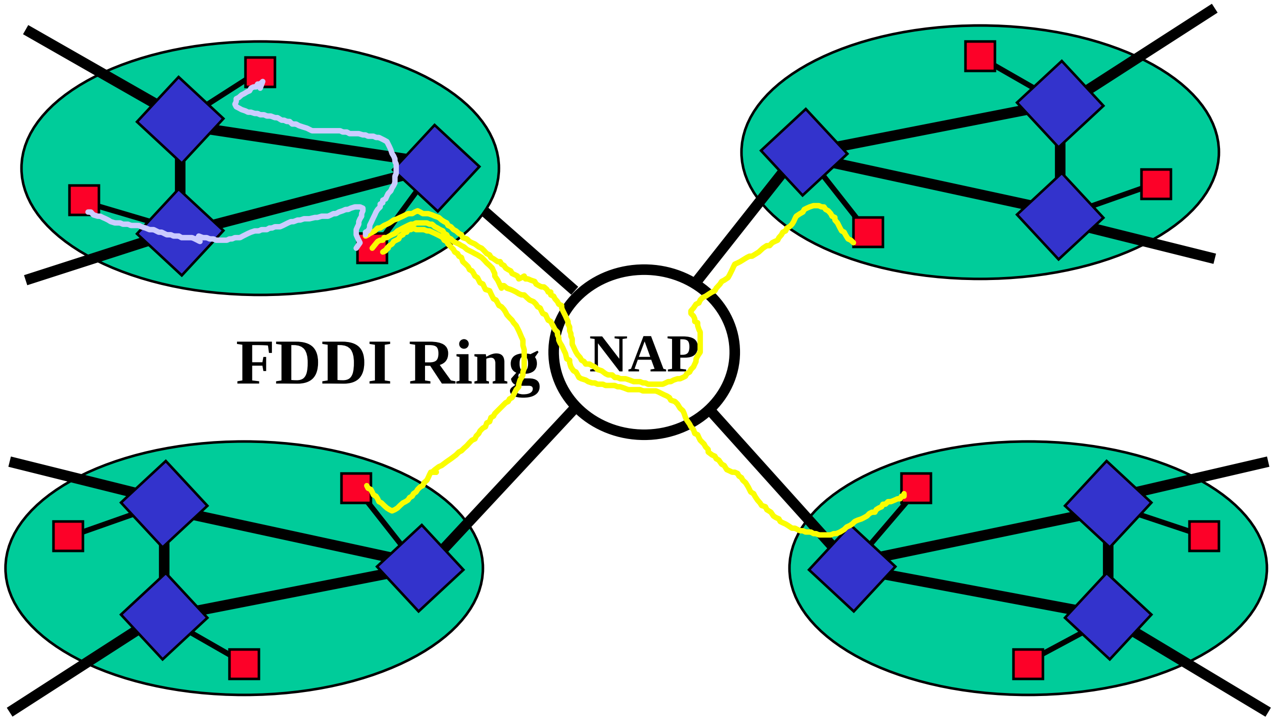

FDDI Ring

NAP

◆ Border Router

🟥 S-BGP box

⬭ AS

— External Peer

— Internal Peer