Computer Networks and Data Systems

Exterior Gateway Protocols (EGPs)

* primarily BGP *

One of two critical systems

Routing (most importantly BGP) and naming (solely DNS) are, by far, the two most critical subsystems of the Internet infrastructure. Participation in and access to the routers themselves are generally, or rather should be, restricted to network administrators.

Agenda

- Routing Refresher
- BGP Overview
- Advanced topics and exercises

BGP Overview

- The routing protocol for connecting *domains*
- Besides the *network prefix* the path is the key component of a BGP route
- Autonomous system numbers (ASNs) define path
  - generally an ASN == domain
    - NOTE: this is not a reference to DNS!
- Even if you don't use it for actual Internet routing, it might be handy for other things (e.g Team Cymru bogon route server, IP addr to ASN mapping)

Exchanges and Peering

- Networks need to connect to each other
- Question: Who pays who?
- Question: Where do they physically connect at?
- Peering is an entire “ecosystem” unto itself
- Paid versus settlement-free peering
- Transit versus peering and exchanges
- Peering requirements and network types

Path Vector Routing

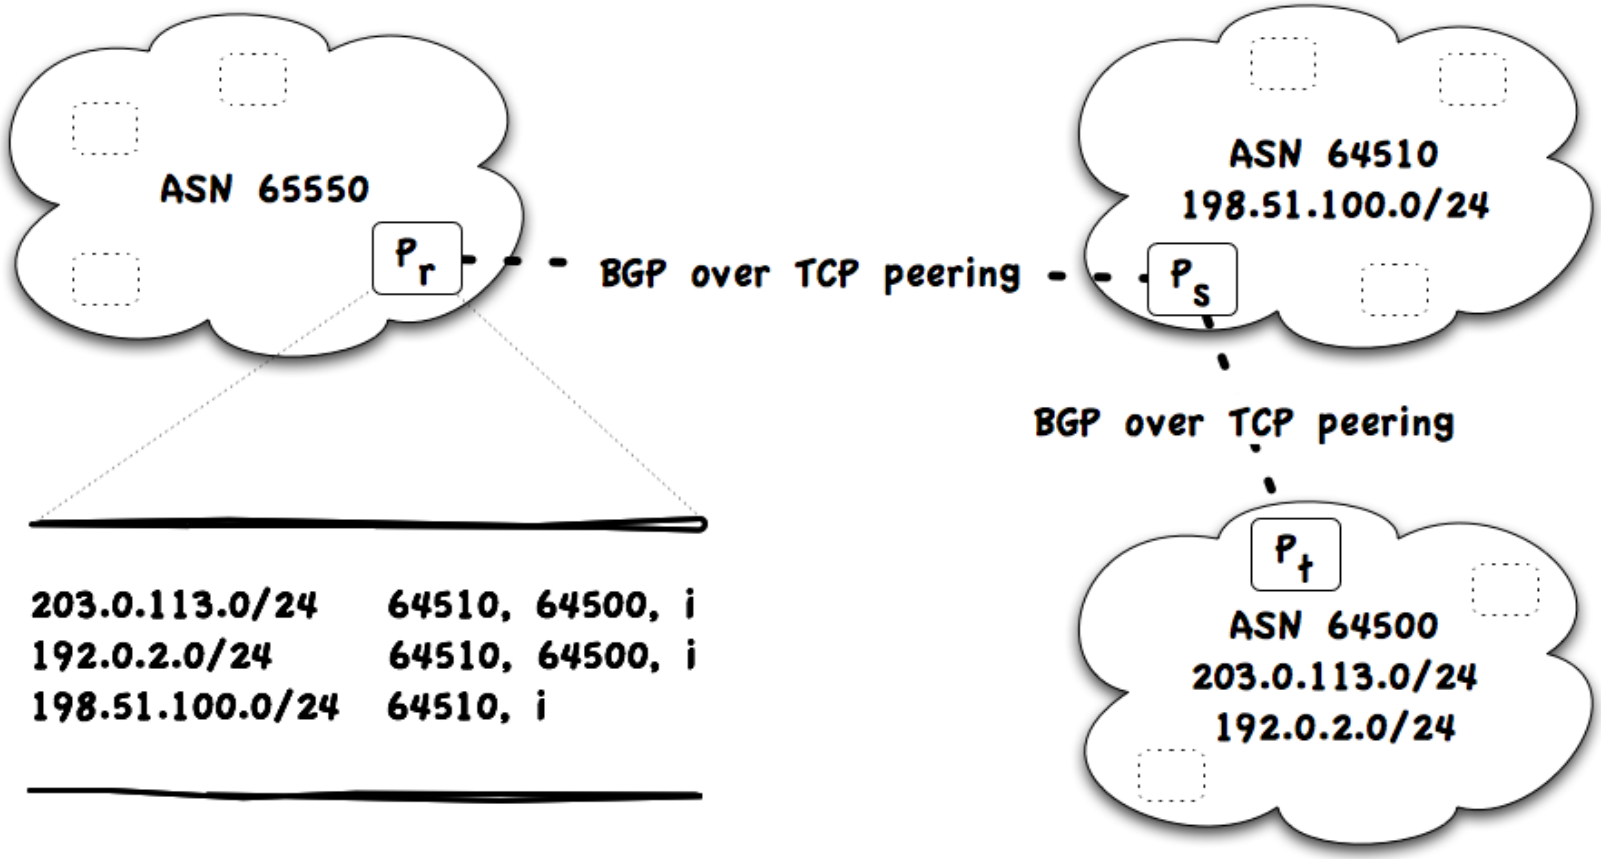

BGP over TCP port 179

- One-to-one peering relationship
- Inherit TCP behaviors, advantages and threats

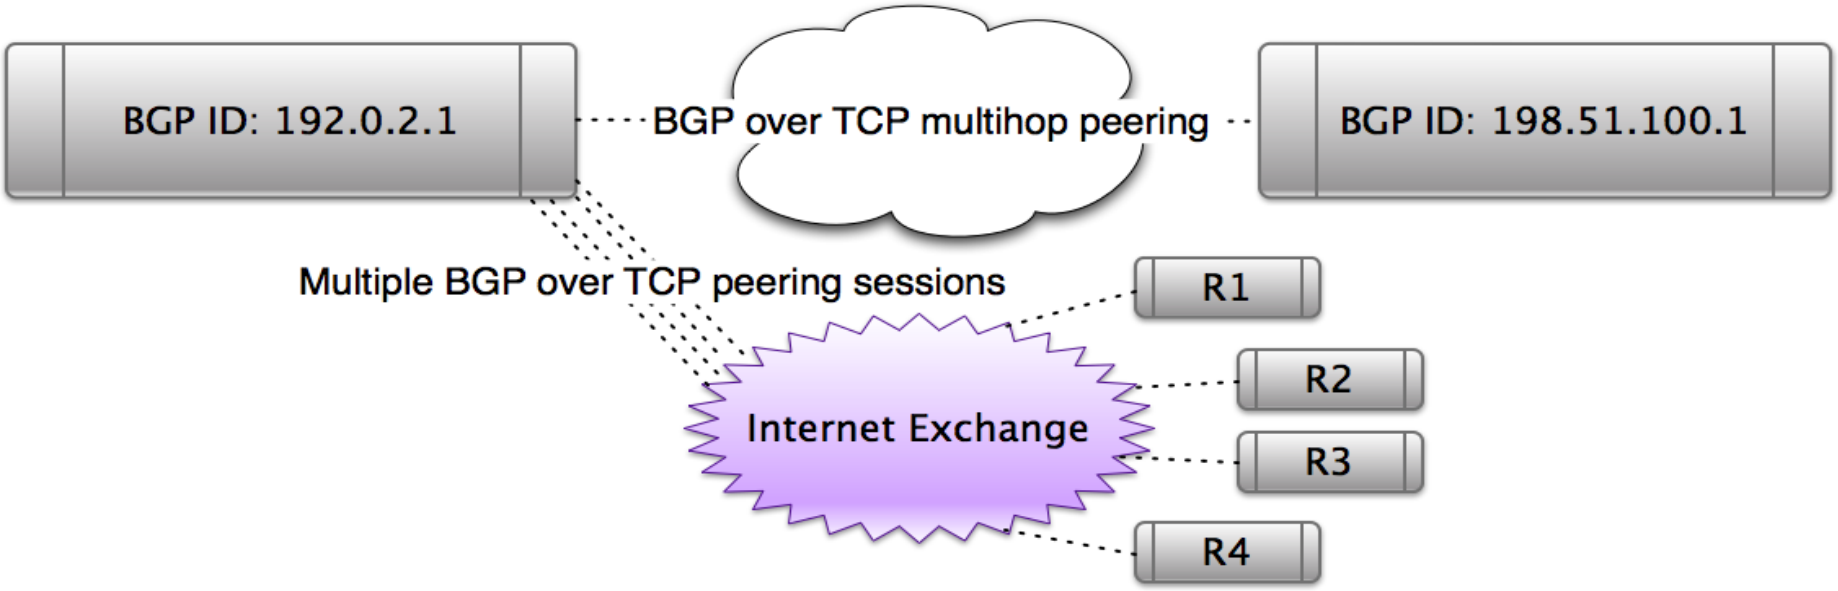

Common BGP Header

BGP message types

1 – OPEN

2 – UPDATE

3 – NOTIFICATION

4 – KEEPALIVE

5 – ROUTE-REFRESH

BGP OPEN

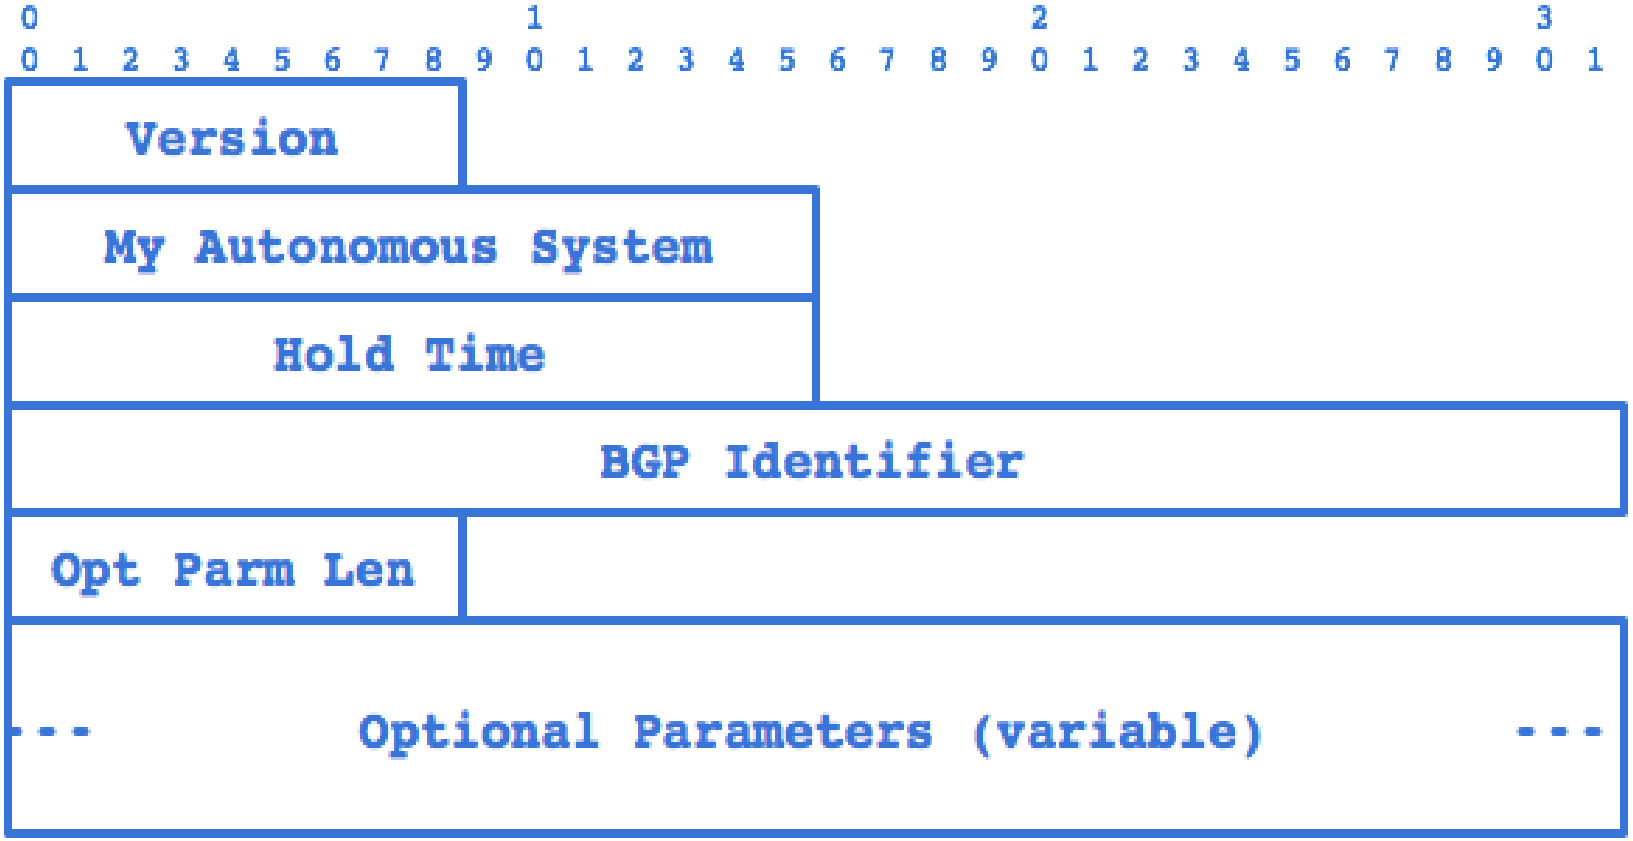
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Version
My Autonomous System
Hold Time
BGP Identifier
Opt Parm Len
Optional Parameters (variable)

BGP UPDATE

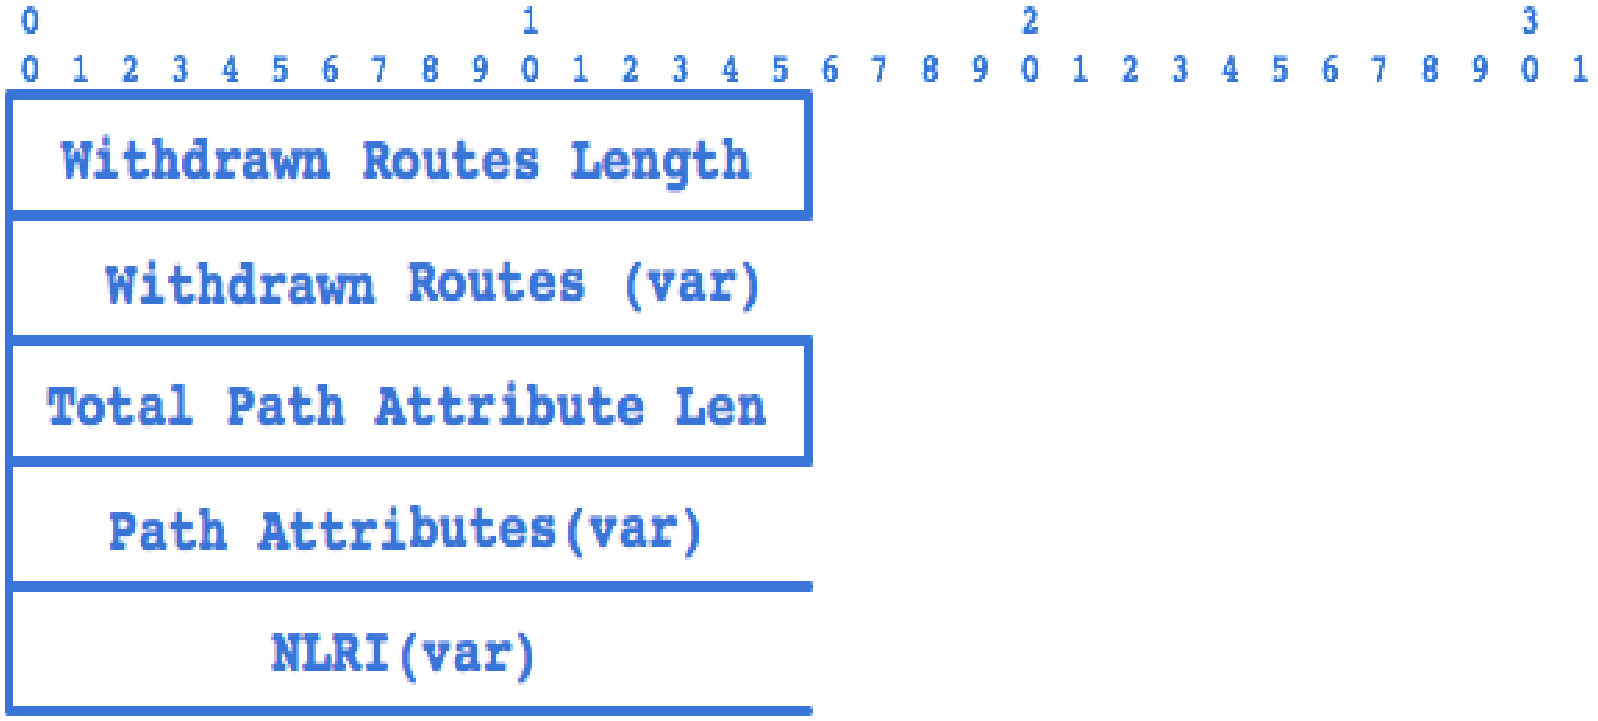
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Withdrawn Routes Length
Withdrawn Routes (var)
Total Path Attribute Len
Path Attributes(var)
NLRI(var)

BGP NOTIFICATION

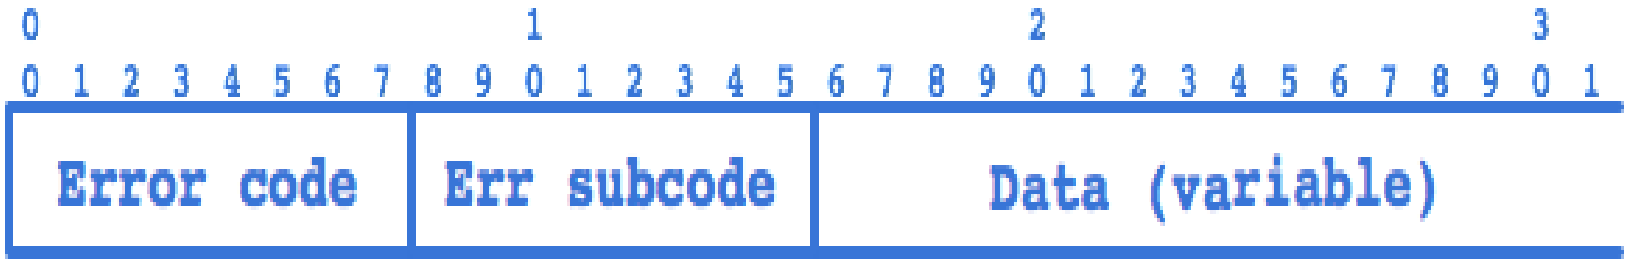

Common BGP Path Attributes

Attribute	Well-known mandatory	Well-known discretionary	Optional transitive	Optional non-transitive
ORIGIN	X			
AS_PATH	X			
NEXT_HOP	X			
MULTI_EXIT_DISC				X
ATOMIC_AGGREGATE		X		
AGGREGATOR			X	
COMMUNITY			X	
MP_REACH_NLRI				X

Affecting BGP Route Decisions

- Prefix length
- LOCAL_PREF
- ORIGIN
- AS_PATH length
- MULTI_EXIT_DISC
- Router import and export policies
- … and more ...

BGP Operational Challenges

- Each AS operates autonomously
- Implicit trust (“routing by rumor”)
- Configuration and policy intensive
- In-band control traffic

BGP Security Matters

- Availability
- Confidentiality
- Integrity

Threats to Availability

- TCP and lower layer attacks
- Packet floods and control path congestion
- Route instability and churn
- Route flap dampening
- Disaggregation and route table exhaustion
- Implementation bugs and configuration errors
- Route hijacking and black holes
- Policy disputes

Threats to Confidentiality

- Clear text communications
- Routing leaks
- Policy configuration leaks
- Route hijacking

Threats to Integrity

- Implementation bugs
- Protocol design weaknesses
- Compromised systems
- Route hijacking
- Path editing
- Overt or covert transit theft
- Divergence

A Quantitative Analysis of the Insecurity of Embedded Network Devices: Results of a Wide-Area Scan

- "...we have identified over 540,000 publicly accessible embedded devices configured with factory default root passwords."
- "...range from enterprise equipment such as firewalls and routers to consumer appliances such as VoIP adapters, cable and IPTV boxes to office equipment..."
- "Vulnerable devices were detected in 144 countries, across 17,427 unique private enterprise, ISP, government, educational, satellite provider as well as residential network environments."

Mitigation

- Protecting the transport
- Router BCPs
- Route monitoring
- Policies and Defensive filtering
- RPKI and BGPSEC

Protecting the transport

- TCP MD5 signature option and TCP-AO
  http://tools.ietf.org/html/rfc2385
  http://tools.ietf.org/html/rfc5925
- IPSec
  http://tools.ietf.org/html/rfc4301
- RFC 5082 Generalized TTL Security Mechanism
  http://tools.ietf.org/html/rfc5082

Router BCPs

- Configuration templates
  http://www.team-cymru.org/ReadingRoom/Templates/
  https://www.iad.gov/iad/library/ia-guidance/security-configuration
- Control plane protection
- Limited and protected remote access
- Current software
- Configuration management

Route monitoring

- http://bgpmon.net
- http://bgplay.routeviews.org/bgplay/
- http://puck.nether.net/bgp/leakinfo.cgi
- http://www.ripe.net/data-tools/stats/ris/
- http://bgp.he.net

Policies and Defensive Filtering

- Document policy with peers
- Internet Routing Registries (IRRs)
- Max prefix and path length limits
- Limiting disaggregation
- Remote triggered black hole filtering (RTBH)
  http://tools.ietf.org/search/rfc5635
- Dissemination of Flow Specification Rules
  http://tools.ietf.org/search/rfc5575
  http://www.cymru.com/jtk/misc/community-fs.html

BGP Remote Triggered Blackhole

- Goal: Have remote router /dev/null certain traffic
- Trick: use next-hop address that points to /dev/null
- Trick: Using policy, set next-hop for matching traffic
- Team Cymru bogon route server does this
- Many ISPs offer this as a DDoS relief service
- IPaddr getting packeted? Have upstream null it
- Also see IETF RFC 5635

unicast Reverse Path Check (uRPF)

- Goal: mitigate source address spoofing
- Trick: Validate source address to ingress interface
  - is there a route back via that interface?
- Loose versus strict mode
- Easier than ACLs (filters)? Maybe
  - doesn't work for everyone
- What do you do if you have a default route?

NetFlow

- Key router technology for analysis and monitoring
- NetFlow is not like pcap
- A unidirectional summary of traffic for a “flow”
- Flow is a unique tuple of addrs, proto, ports
- Router “exports” flows at timer, RST, FIN, etc
- Data may be limited due to sampling
- Scales very well and is very popular

Flow specification (flow-spec)

- Using BGP, exchange “flow-spec” to act on
- Largely used as a distributed firewall filter
- Can be more precise than a BGP RTBH
- Besides filtering, you can rate limit, log, pass
- Not widely implemented

RPKI and BGPSEC

- Observation:

  There is no official, and consequently, no strong association between address assignment and routing announcements

- Problem:

  How do you guard against routing threats, such as hijacks, without a means to verify the routing announcements?

Resource Public Key Infrastructure (RPKI)

- RIRs maintain RPKI infrastructure
- Prefixes linked to authorized AS
- In-band via soBGP or S-BGP
- Out-of-band for near-time validation and monitoring
- IETF SIDR WG spearheading RPKI work
  - Only for origin validation
  - BGPSEC working on path validation
- Similar in scope to DNSSEC architectural changes

A Future of IRR with RPKI

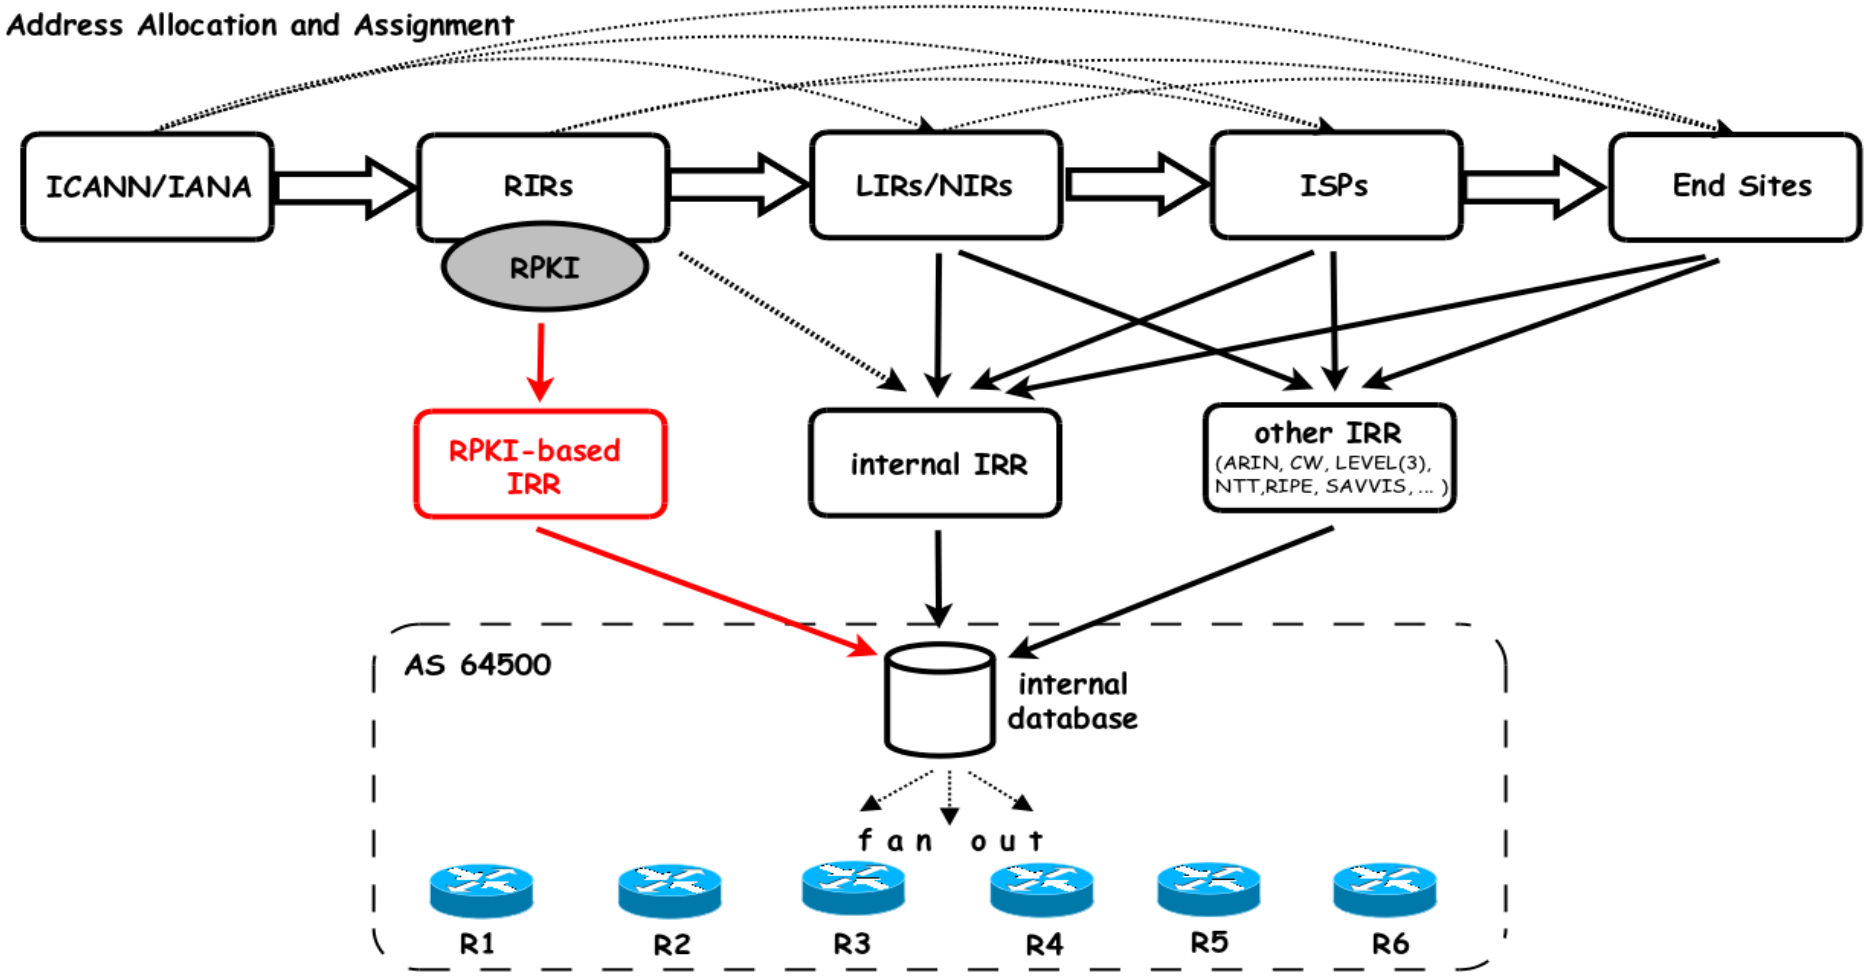

In Closing

- RFC 4271 A Border Gateway Protocol 4 (BGP-4)
- RFC 4593 Generic Threats to Routing Protocols
- IETF Secure Inter-Domain Routing (sidr) WG
- Academic papers:
  - Securing BGP – A Literature Survey
  - A Survey of BGP Security Issues and Solutions
  - Securing BGP with BGPsec