Management: Fault Detection and Troubleshooting

Nick Feamster
CS 7260
February 5, 2007

Today's Lecture

- Routing Stability
  - Gao and Rexford, *Stable Internet Routing without Global Coordination*
  - Major results
  - Business model assumptions (validity of)
- Network Management
  - "State-of-the-art": SNMP
  - Research challenges for network management
  - Routing configuration correctness
    - *Detecting BGP Configuration Faults with Static Analysis*

Is management really *that* important?

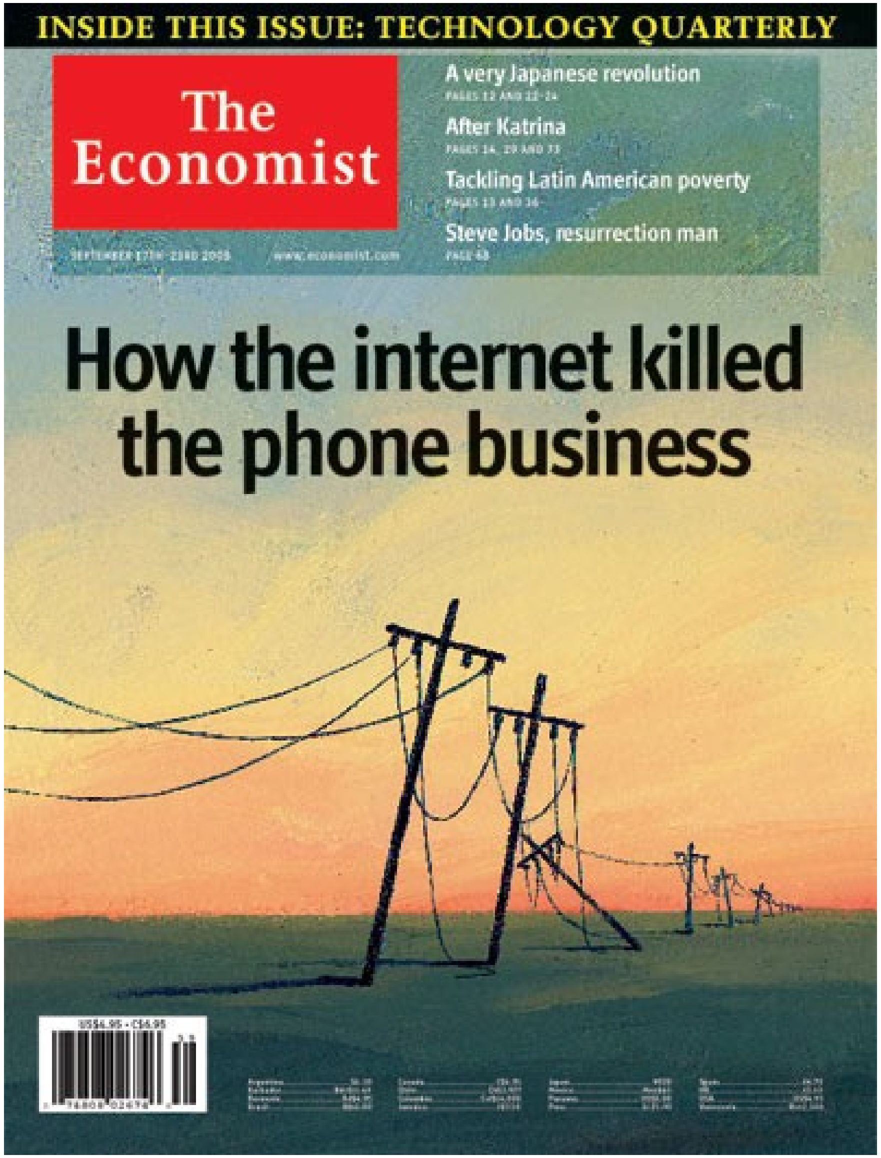

Is management really *that* important?

- The Internet is increasingly becoming part of the mission-critical Infrastructure (a public utility!).

FCC Requires VoIP Providers to Offer E911 Service

Emergency service call ability to be mandatory within six months.

Grant Gross, IDG News Service

Thursday, May 19, 2005

WASHINGTON--Voice over Internet Protocol carriers that connect to the U.S. public telephone network will be required this year to provide their customers with enhanced 911 emergency calling service, the U.S. Federal Communications Commission ruled Thursday.

Big problem: Very poor understanding of how to manage it.

Simple Network Management Protocol

- Version 1: 1988 (RFC 1065-1067)
- Management Information Base (MIB)
  - Information store
  - Unique variables named by OIDs
  - Accessed with SNMP
- Three components
  - *Manager:* queries the MIB ("client")
  - *Master agent:* the network element being managed
  - *Subagent:* gathers information from managed objects to store in MIB, generate alerts, etc.

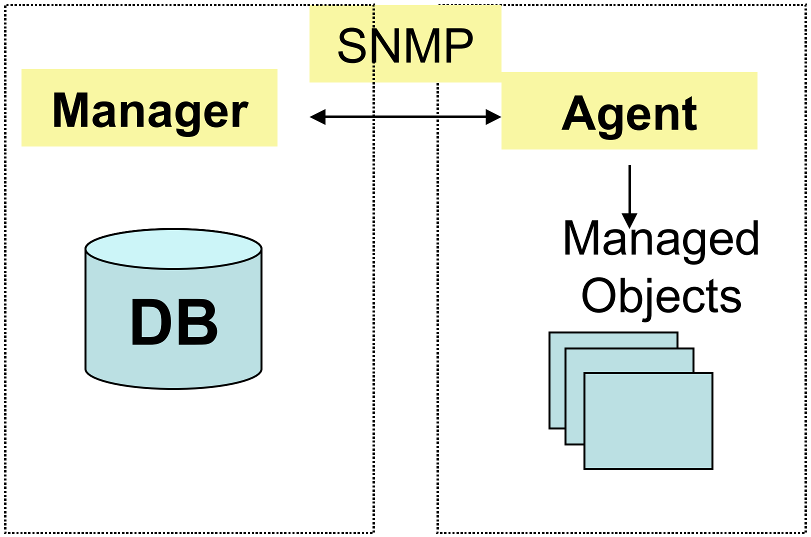

Naming MIB Objects

- Each object has a distinct object identifier (OID)
  - Hierarchical Namespace
- **Example**
  - *BGP:* **1.3.6.1**.2.1.15 (RFC 1657)
    - bgpVersion: "1.3.6.1.2.1.15.1"
    - bgpLocalAs: "1.3.6.1.2.1.15.2"
    - bgpPeerTable: "1.3.6.1.2.1.15.3"
    - bgpIdentifier: "1.3.6.1.2.1.15.4"
    - bgpRcvdPathAttrTable: "1.3.6.1.2.1.15.5“
    - bgp4PathAttrTable: "1.3.6.1.2.1.15.6"

Tables are sequences of other types

MIB Structure

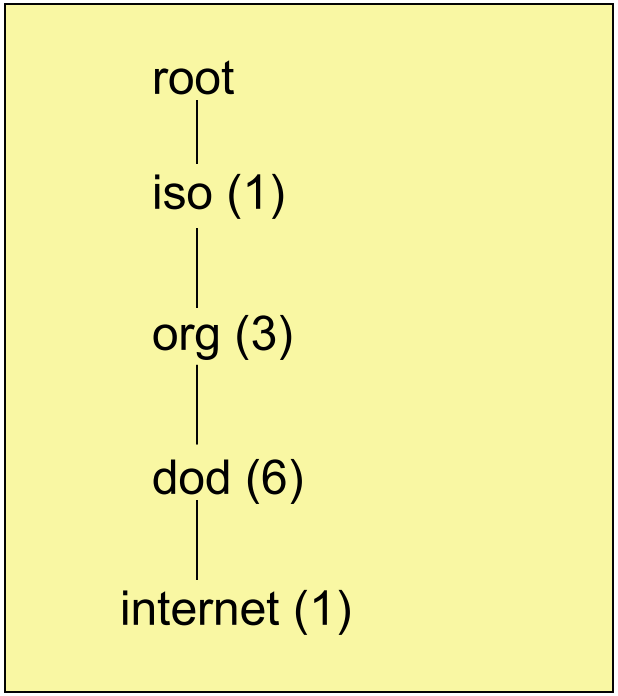

MIB Definitions

Example from RFC 1657

```
bgpVersion OBJECT-TYPE
    SYNTAX     OCTET STRING (SIZE (1..255))
    MAX-ACCESS read-only
    STATUS     current
    DESCRIPTION
            "Vector of supported BGP protocol version
            numbers.  Each peer negotiates the version
            from this vector.  Versions are identified
            via the string of bits contained within this
            object.  The first octet contains bits 0 to
            7, the second octet contains bits 8 to 15,
            and so on, with the most significant bit
            referring to the lowest bit number in the
            octet (e.g., the MSB of the first octet
            refers to bit 0).  If a bit, i, is present
            and set, then the version (i+1) of the BGP
            is supported."
    ::= { bgp 1 }
```

"1.3.6.1.2.15.1"

MIB Definitions: Lots of Them!

ADSL	RFC 2662
ATM	Multiple
AppleTalk	RFC 1742
BGPv4	RFC 1657
Bridge	RFC 1493
Character Stream	RFC 1658
CLNS	RFC 1238
DECnet Phase IV	RFC 1559
DOCSIS Cable Modem	Multiple
…	

Interacting with the MIB

- Four basic message types
  - **Get:** retrieving information about some object
  - **Get-Next:** iterative retrieval
  - **Set:** setting variable values
  - **Trap:** used to report
- Queries on UDP port 161, Traps on port 162
- Enabling SNMP on a Cisco Router for BGP

```
# snmp-server enable traps bgp
# snmp-server host myhost.cisco.com informs version 2c public
```

- Notifications about state changes, etc.

SNMPv2c (1993)

- Expanded data types: 64-bit counters
- Improved efficiency and performance: **get-bulk**
- Confirmed event notifications: inform operator
- Richer error handling: errors and exceptions
- Improved sets: especially row creation/deletion
- Transport independence: IP, Appletalk, IPX
- *Not widely-adopted:* security considerations
  - Compromise: SNMPv2u (commercial deployment)

Common Use of SNMP: Traffic

- Routers have various counters that keep byte counts for traffic passing over a given link
  - Periodic polling of MIBs for traffic monitoring
- **Problem:** these measurements are device-level, not flow-level
  - Detect a DoS attack by polling SNMP?!
  - *Trend:* end-to-end statistics

More Problems with SNMP

- Can’t handle large data volumes
  - SNMP “walks” take very long on large tables, especially when network delay is high
- Imposes significant CPU load
- **Device-level, not network-level**
- Sometimes, implementation issues
  - Counter bugs
  - Loops on SNMP walks

http://www.statseeker.com/pdf/snmp.pdf

Management Research Problems

- Organizing **diverse data** to consider problems across different time scales and across different sites
  - Correlations in real time and event-based
  - How is data normalized?
- Changing the focus: **from data to information**
  - Which information can be used to answer a specific management question?
  - Identifying root causes of abnormal behavior (via data mining)
  - How can simple counter-based data be synthesized to provide information eg. “something is now abnormal”?
  - View must be expanded across layers and data providers

Research Problems (continued)

- **Automation** of various management functions
  - Expert annotation of key events will continue to be necessary
- **Identifying traffic types** with minimal information
- Design and deployment of measurement infrastructure (both passive and active)
  - Privacy, trust, cost limit broad deployment
  - Can end-to-end measurements ever be practically supported?
- Accurate **identification of attacks** and intrusions
  - Security makes different measurements important

Overcoming Problems

- Convince customers that measurement is worth additional cost by targeting their problems
- Companies are motivated to make network management more efficient (*i.e.*, reduce headcount)
- Portal service (high level information on the network's traffic) is already available to customers
  - This has been done primarily for security services
  - Aggregate summaries of passive, netflow-based measures

Long-Term Goals

- Programmable measurement
  - On network devices and over distributed sites
  - Requires authorization and safe execution
- Synthesis of information at the point of measurement and central aggregation of minimal information
- Refocus from measurement of individual devices to measurement of network-wide protocols and applications
  - Coupled with drill down analysis to identify root causes
  - This must include all middle-boxes and services

Why does routing go wrong?

- Complex policies
  – Competing / cooperating networks
  – Each with only limited visibility

- Large scale
  – Tens of thousands networks
  – …each with hundreds of routers
  – …each routing to hundreds of thousands of IP prefixes

What can go wrong?

Some things are out of the hands of networking research

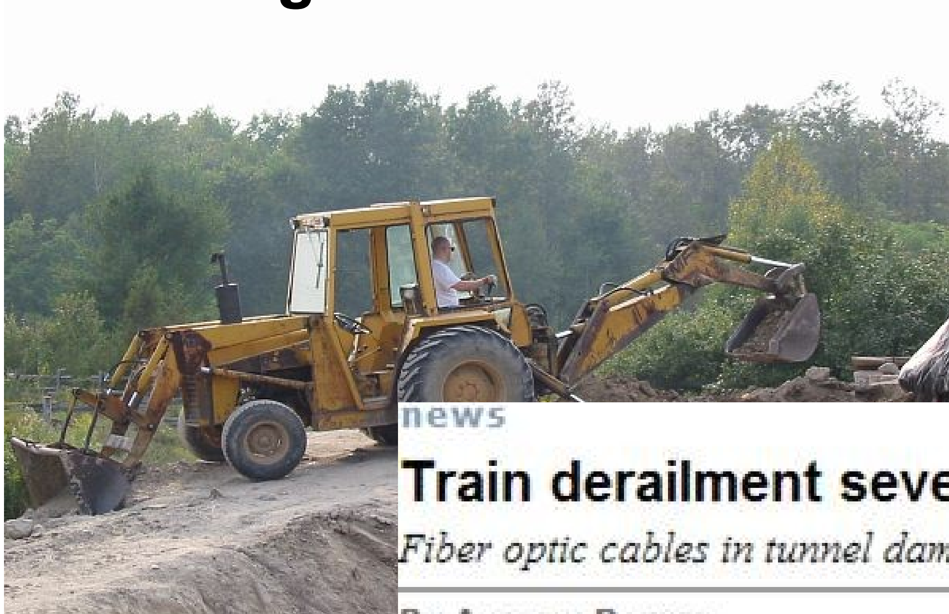

news

Train derailment severs communications

Fiber optic cables in tunnel damaged; flood knocks out phone service

By Andrew Ratner
Sun Staff
Originally published July 20, 2001

When a train falls in Baltimore, it knocks out e-mail halfway around the world.

But…

Two-thirds of the problems are caused by *configuration* of the routing protocol

Complex configuration!

Flexibility for realizing goals in complex business landscape

- **Which neighboring networks can send traffic**
- **Where traffic enters and leaves the network**
- **How routers *within* the network learn routes to external destinations**

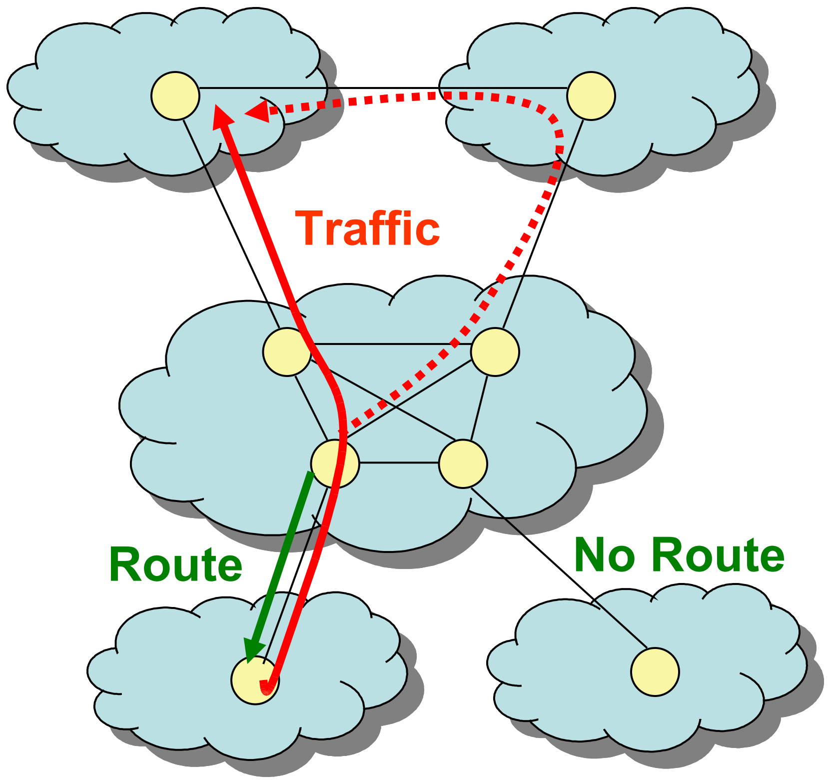

Flexibility ⟶ Complexity

Configuration Semantics

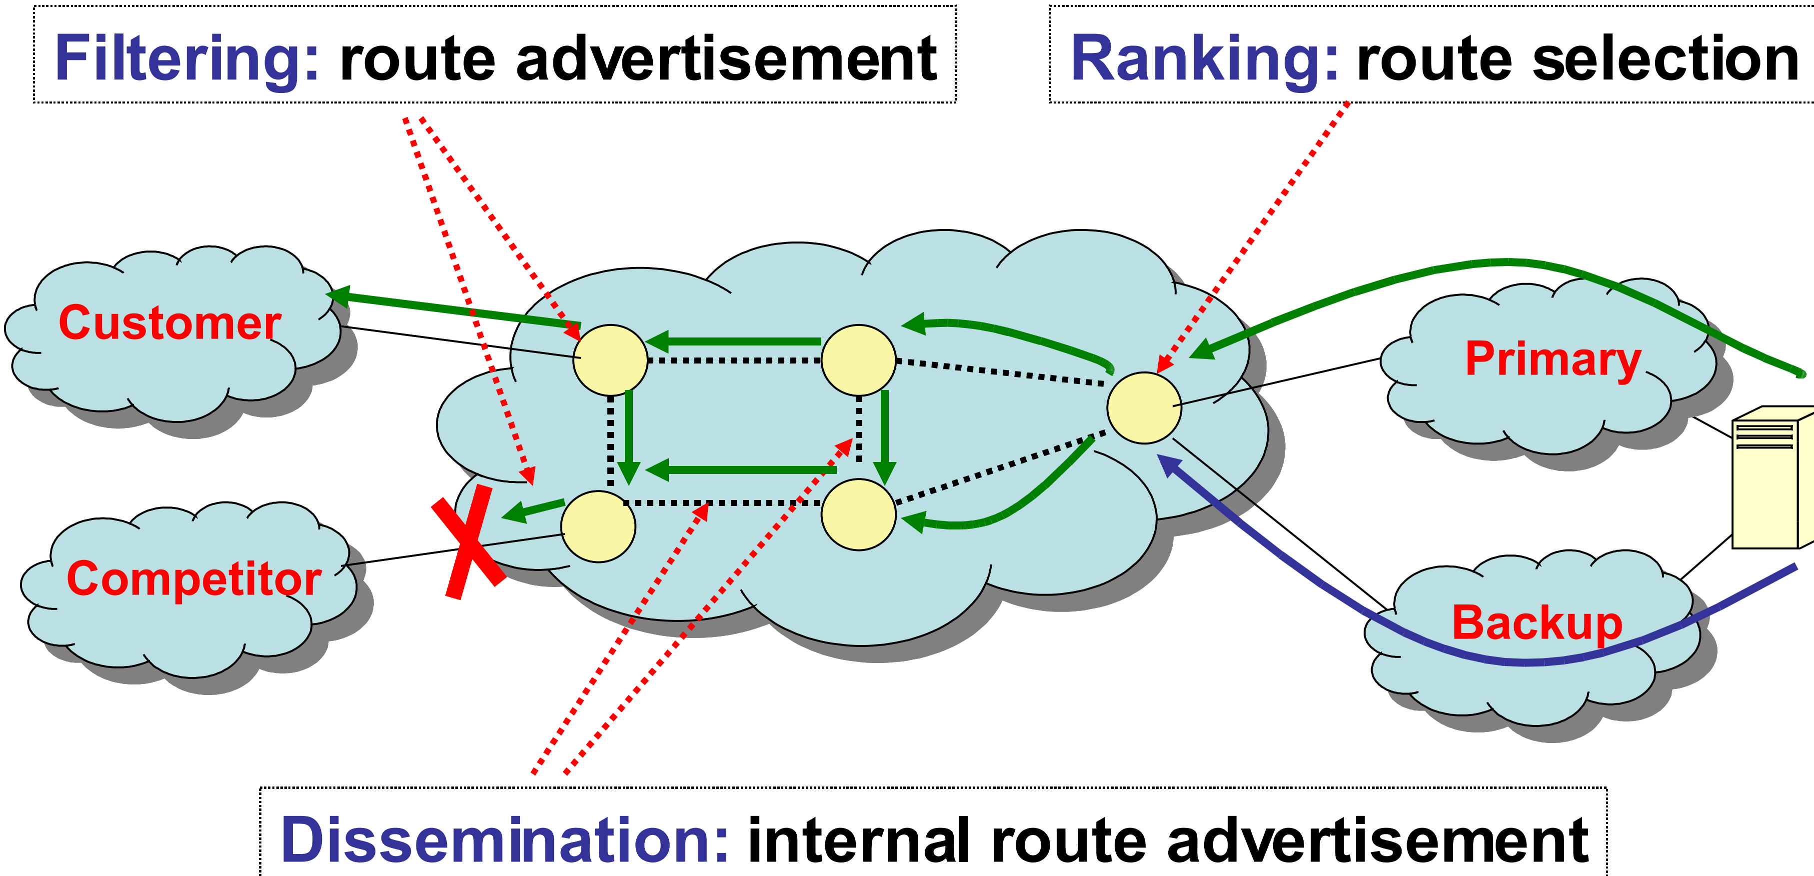

What types of problems does configuration cause?

- Persistent oscillation *(last time)*
- Forwarding loops
- Partitions
- “Blackholes”
- Route instability
- …

Real Problems: "AS 7007"

"…a glitch at a small ISP… triggered a major outage in Internet access across the country. The problem started when MAI Network Services...passed bad router information from one of its customers onto Sprint." -- *news.com*, April 25, 1997

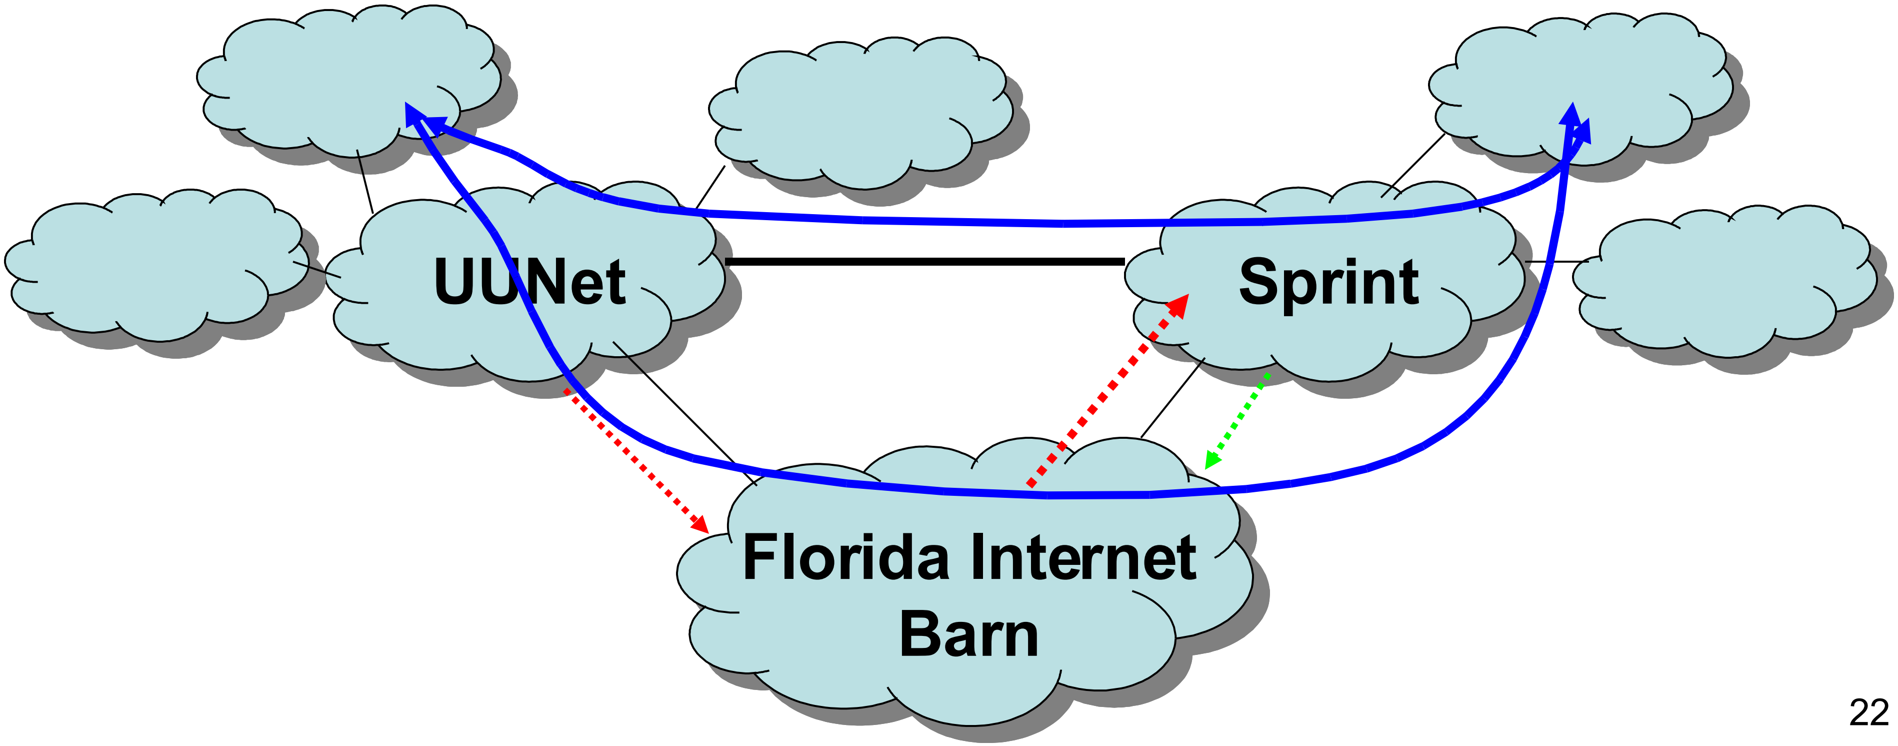

Real, Recurrent Problems

"…a glitch at a small ISP… triggered a major outage in Internet access across the country. The problem started when MAI Network Services...passed bad router information from one of its customers onto Sprint."

-- *news.com*, April 25, 1997

"Microsoft's websites were offline for up to 23 hours...because of a [router] misconfiguration…it took nearly a day to determine what was wrong and undo the changes." -- *wired.com*, January 25, 2001

"WorldCom Inc…suffered a widespread outage on its Internet backbone that affected roughly 20 percent of its U.S. customer base. The network problems…affected millions of computer users worldwide. A spokeswoman attributed the outage to "a route table issue."

-- *cnn.com*, October 3, 2002

"A number of Covad customers went out from 5pm today due to, supposedly, a DDOS (distributed denial of service attack) on a key Level3 data center, which later was described as a route leak (misconfiguration)."

-- *dslreports.com*, February 23, 2004

January 2006: Route Leak, Take 2

Con Ed 'stealing' Panix routes (alexis) Sun Jan 22 12:38:16 2006

All Panix services are currently unreachable from large portions of the Internet (though not all of it). This is because Con Ed Communications, a competence-challenged ISP in New York, is announcing our routes to the Internet. In English, that means that they are claiming that all our traffic should be passing through them, when of course it should not. Those portions of the net that are "closer" (in network topology terms) to Con Ed will send them our traffic, which makes us unreachable.

“Of course, there are measures one can take against this sort of thing; but it's hard to deploy some of them effectively when the party stealing your routes was in fact once authorized to offer them, and its own peers may be explicitly allowing them in filter lists (which, I think, is the case here). “

Several "Big" Problems a Week

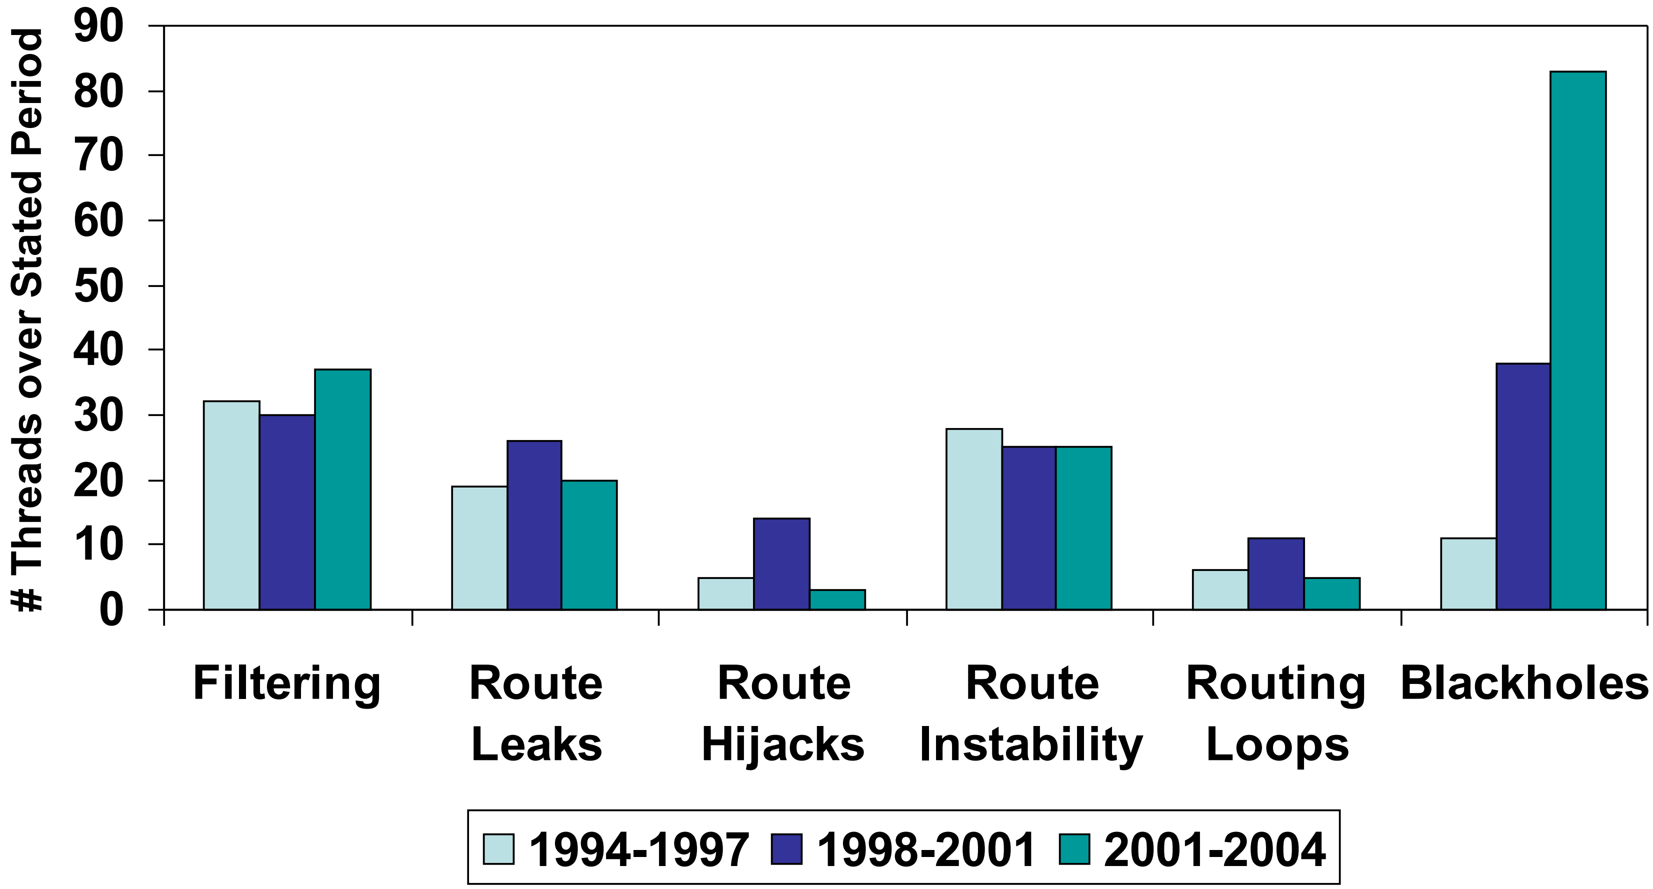

Why is routing hard to get right?

- **Defining correctness is hard**
- **Interactions cause unintended consequences**
  - Each network independently configured
  - Unintended policy interactions
- **Operators make mistakes**
  - Configuration is difficult
  - Complex policies, distributed configuration

Correctness Specification

Safety

The protocol does not oscillate

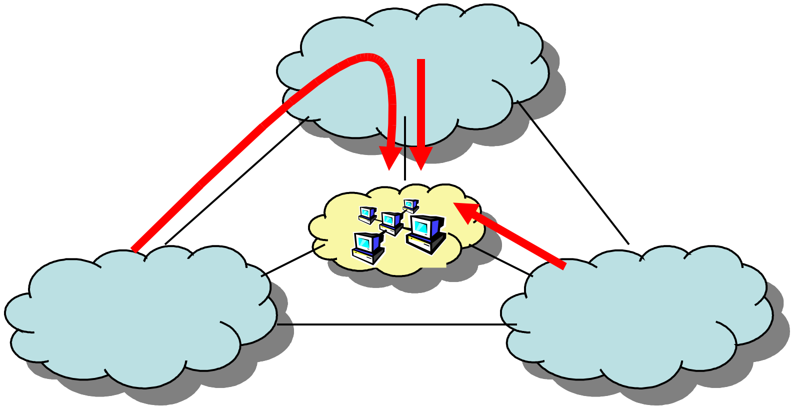

What about properties of resulting paths, after the protocol has converged?

We need additional correctness properties.

Correctness Specification

Safety

The protocol does not oscillate

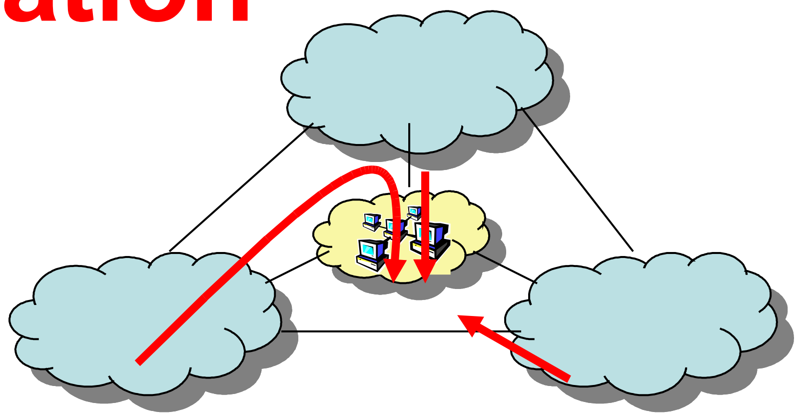

Path Visibility

If there exists a path,
then there exists a route

Example violation: Network partition

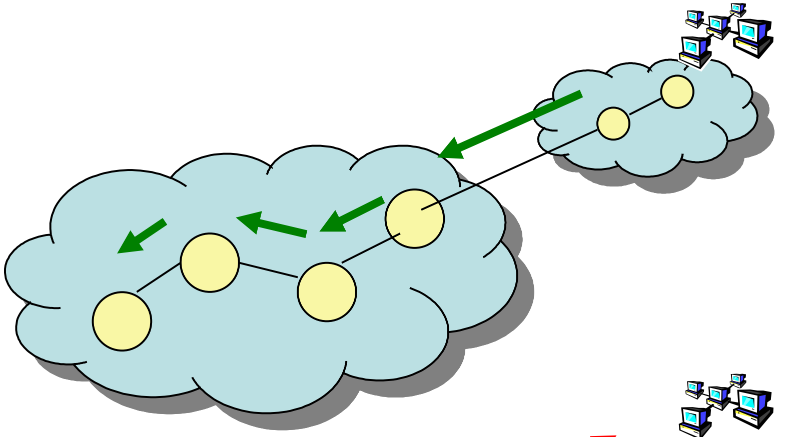

Route Validity

If there exists a route,
then there exists a path

Example violation: Routing loop

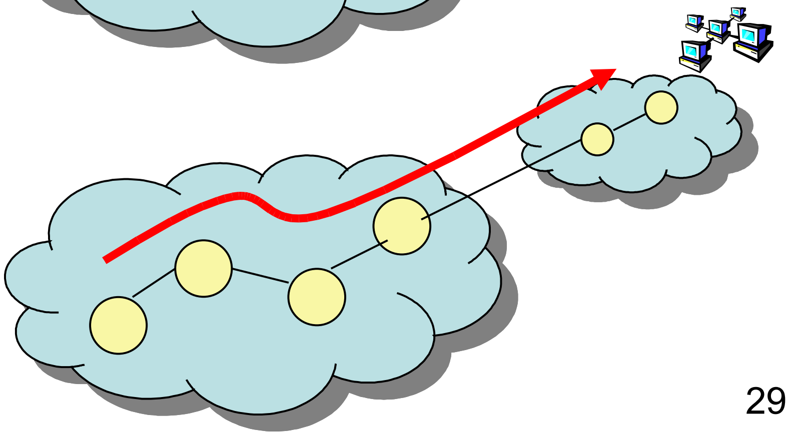

Path Visibility: Internal BGP (iBGP)

Default: “Full mesh” iBGP.
Doesn’t scale.

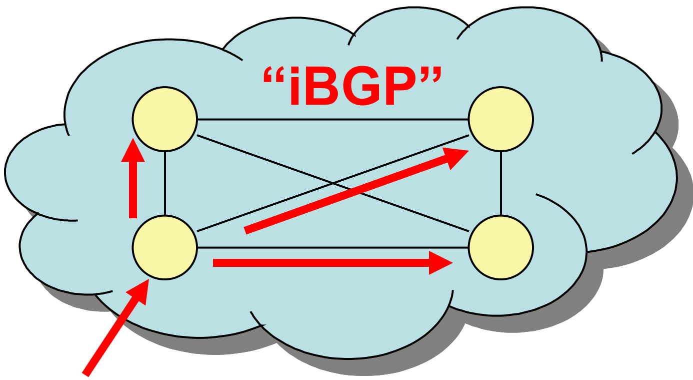

Large ASes use **“Route reflection”**
Route reflector:
non-client routes over client sessions;
client routes over all sessions
Client: don’t re-advertise iBGP routes.

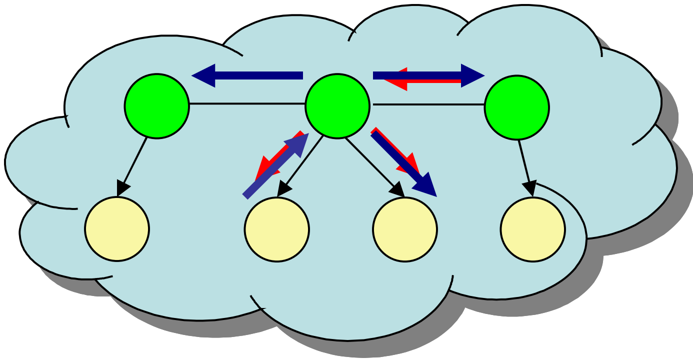

iBGP Signaling: Static Check

Theorem.
Suppose the iBGP reflector-client relationship graph contains no cycles. Then, path visibility is satisfied if, and only if, *the set of routers that are not route reflector clients forms a clique*.

Condition is easy to check with static analysis.

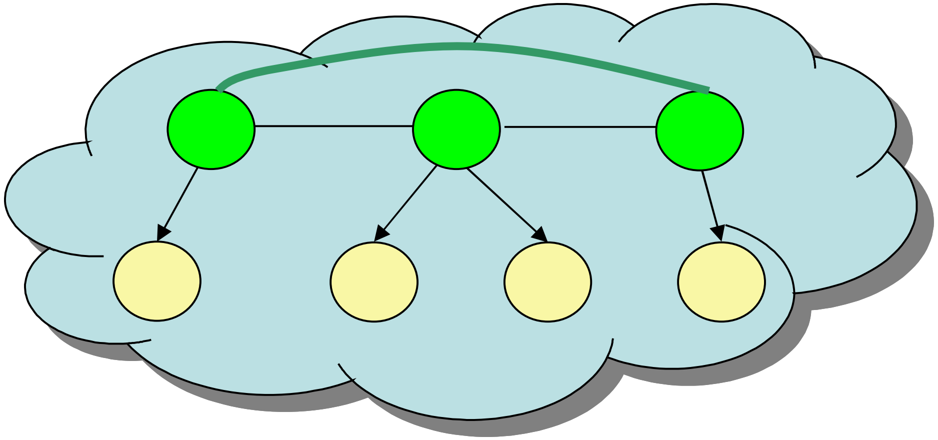

How do we guarantee these additional properties in practice?

Today: Reactive Operation

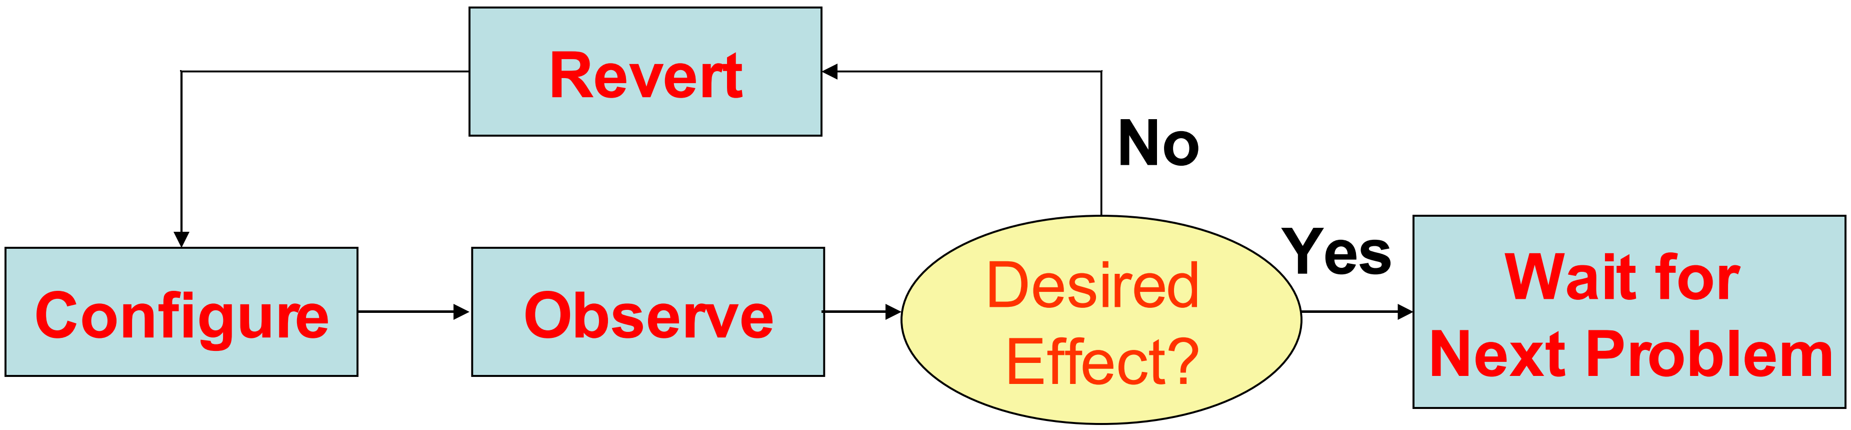

- Problems cause downtime
- Problems often not immediately apparent

Goal: Proactive Operation

- **Idea:** Analyze configuration *before* deployment

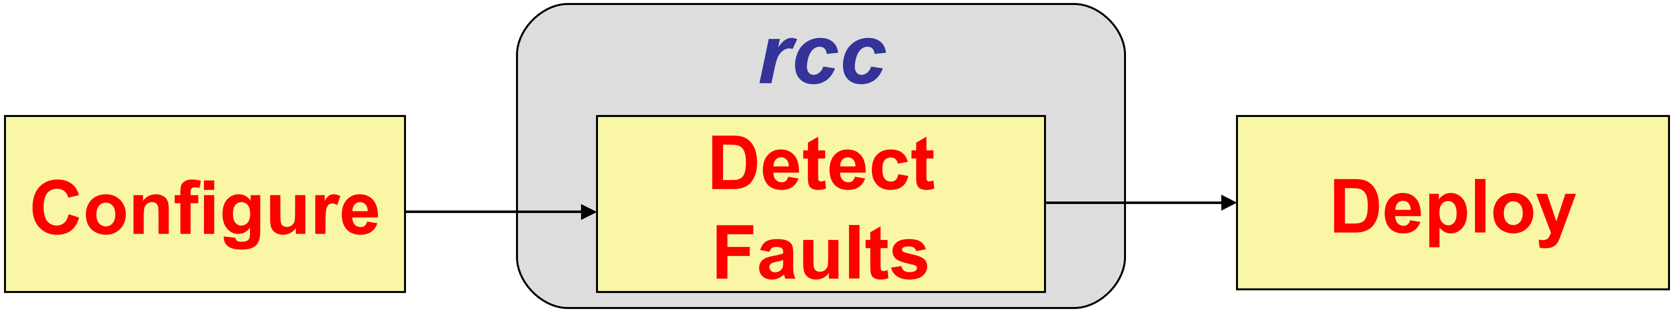

Many faults can be detected with static analysis.

rcc Overview

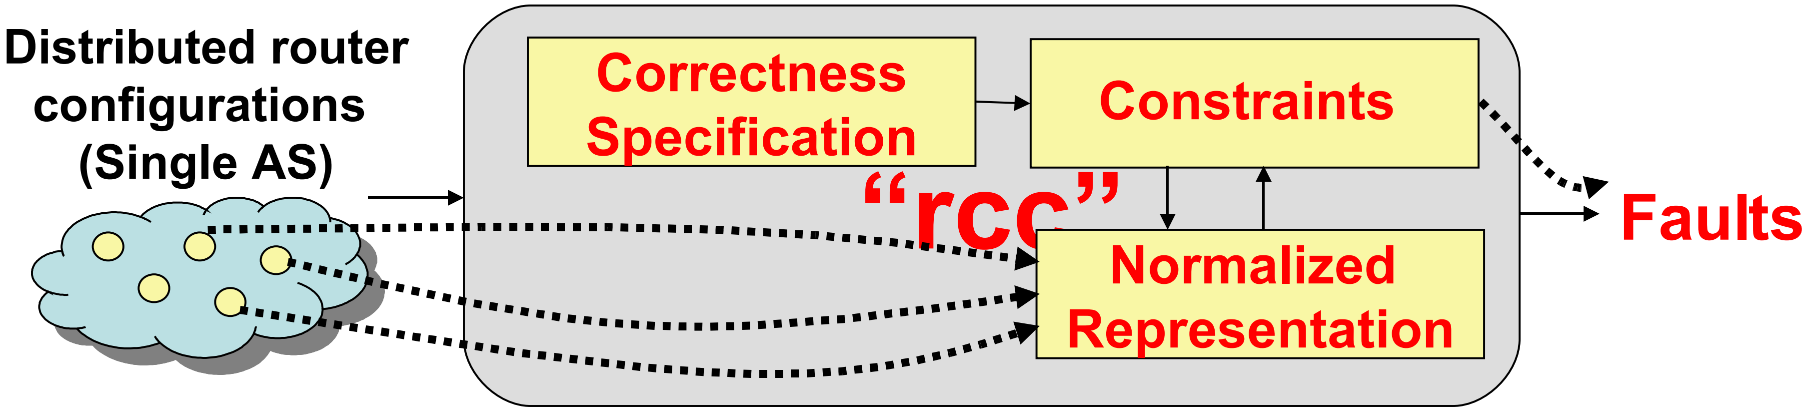

Challenges

- Analyzing complex, distributed configuration
- Defining a correctness specification
- Mapping specification to constraints

rcc Implementation

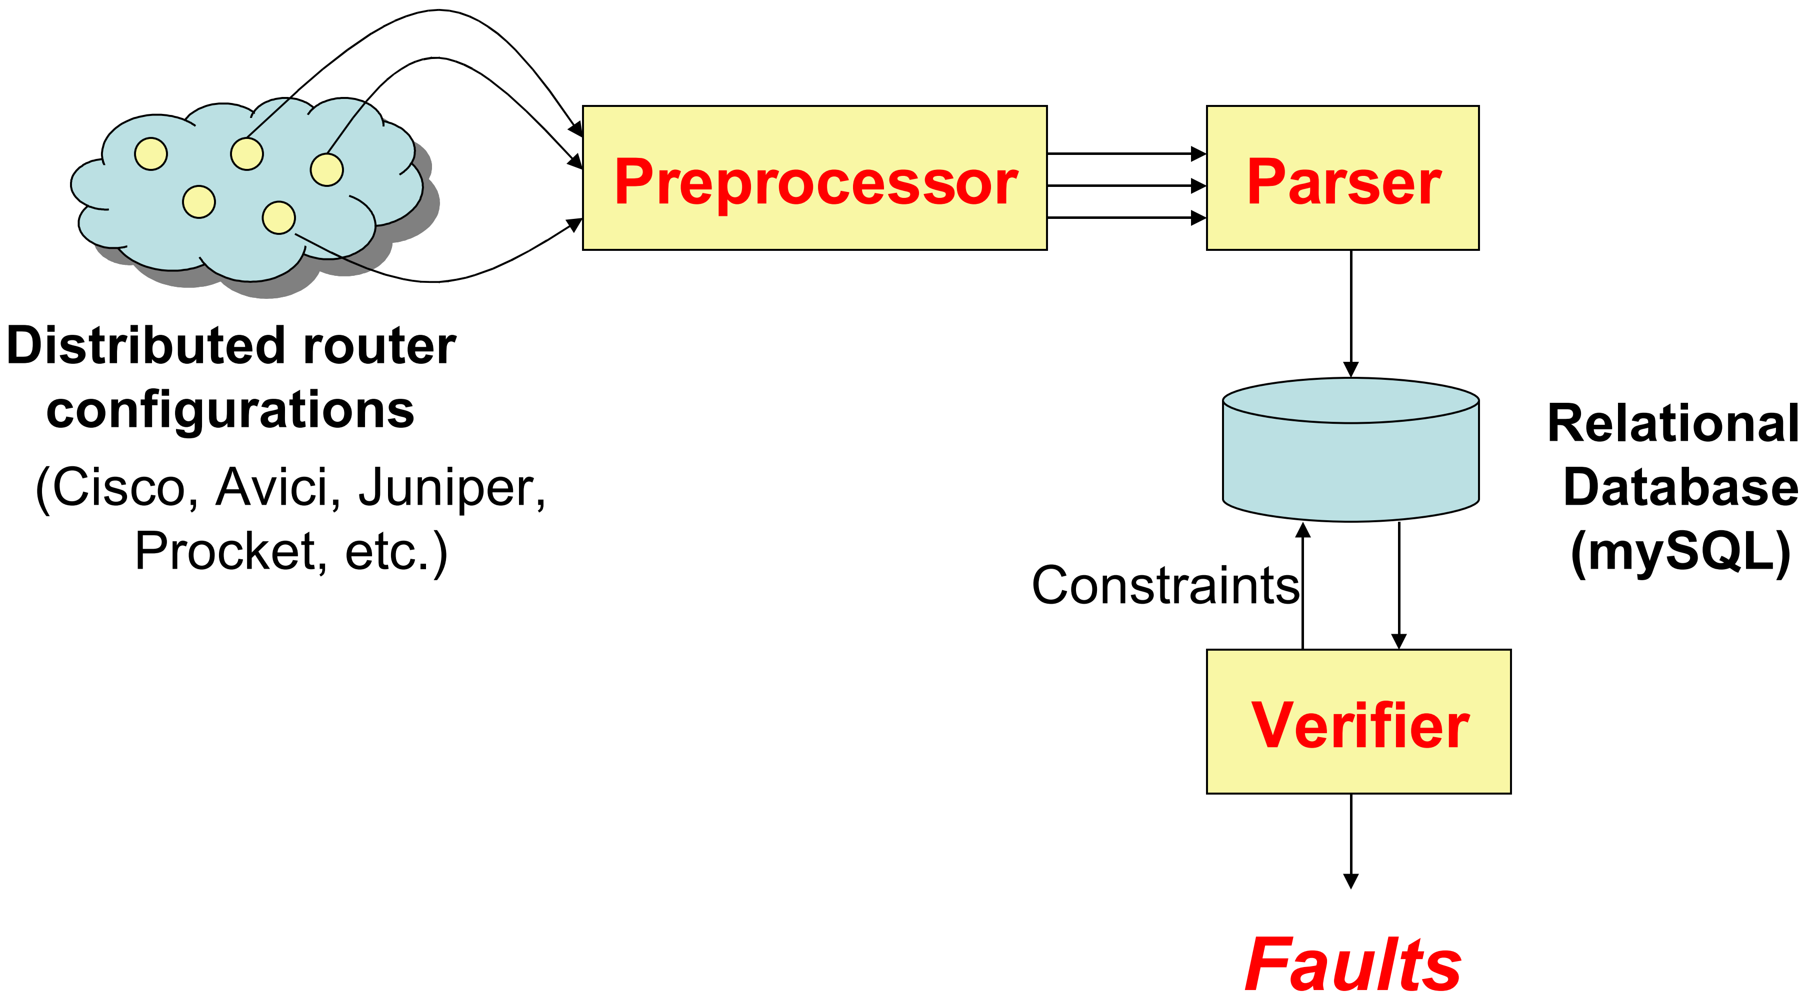

Summary: Faults across 17 ASes

Every AS had faults, regardless of network size
Most faults can be attributed to distributed configuration

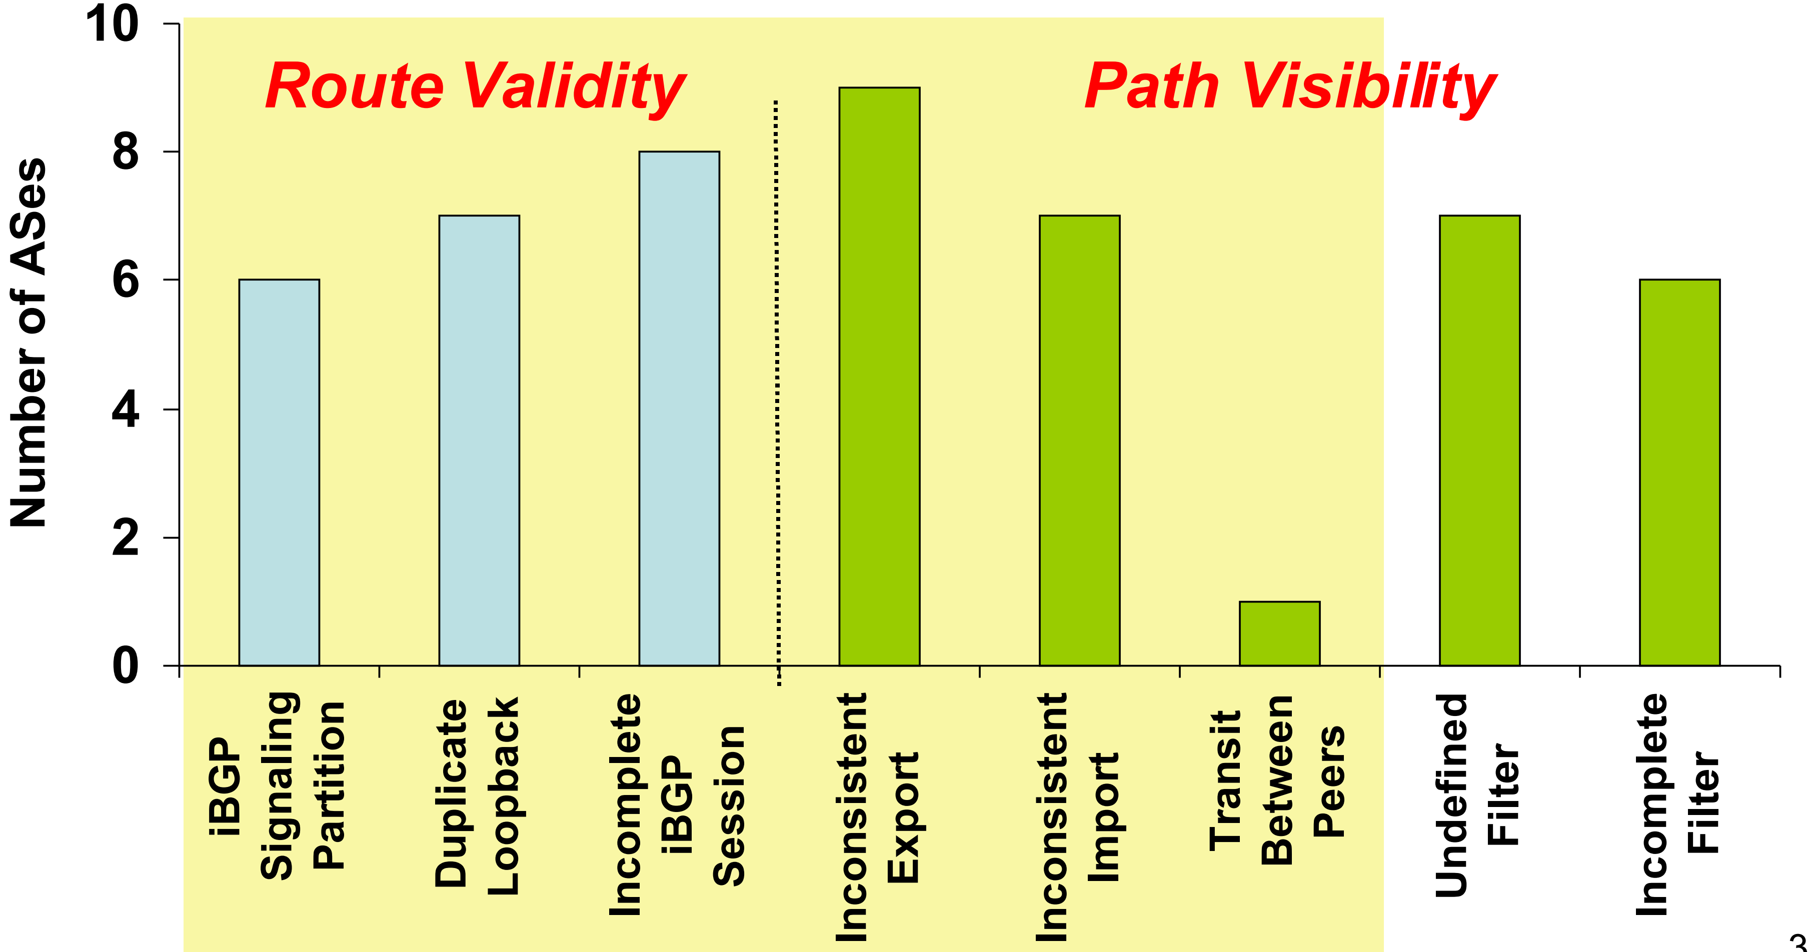

rcc: Take-home lessons

- Static configuration analysis uncovers many errors
- Major causes of error:
  - Distributed configuration
  - Intra-AS dissemination is too complex
  - Mechanistic expression of policy

Two Philosophies

- **The "rcc approach":** Accept the Internet as is. Devise "band-aids".

- **Another direction:** Redesign Internet routing to guarantee safety, route validity, and path visibility

Problem 1: Other Protocols

- Static analysis for MPLS VPNs
  - Logically separate networks running over single physical network: ***separation is key***
  - Security policies maybe more well-defined (or perhaps easier to write down) than more traditional ISP policies

Problem 2: Limits of Static Analysis

- **Problem:** Many problems can't be detected from static configuration analysis of a single AS

- Dependencies/Interactions among multiple ASes
  - Contract violations
  - Route hijacks
  - **BGP "wedgies" (RFC 4264)**
  - Filtering

- Dependencies on route arrivals
  - Simple network configurations can oscillate, but operators can't tell until the routes actually arrive.

BGP Wedgie Example

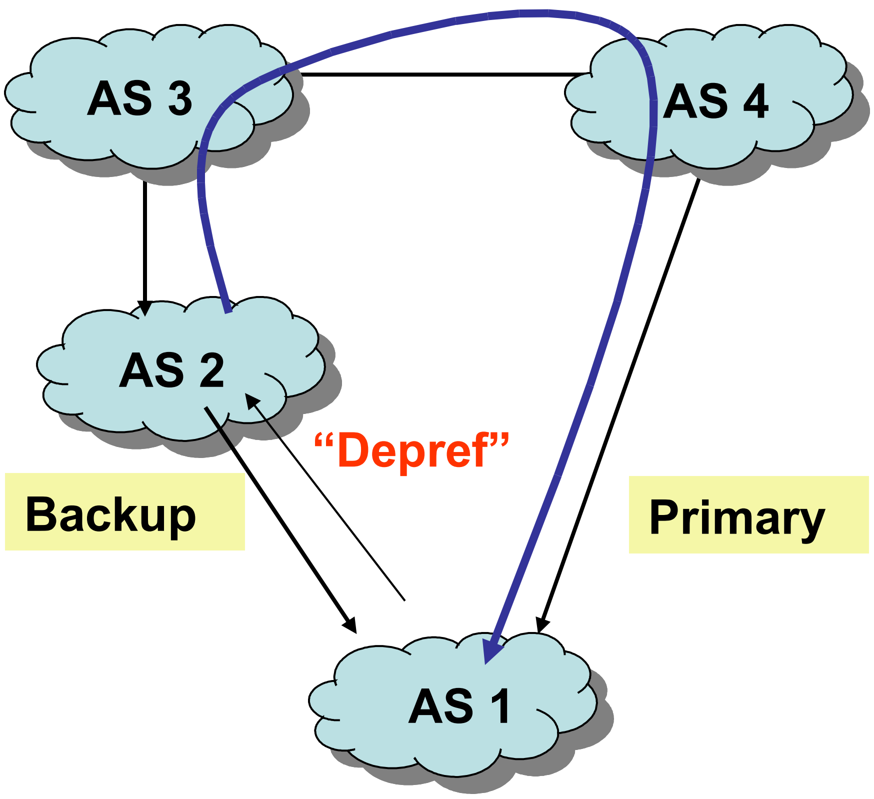

- AS 1 implements backup link by sending AS 2 a "depref me" community.

- AS 2 sets localpref to smaller than that of routes from its upstream provider (AS 3 routes)

Failure and "Recovery"

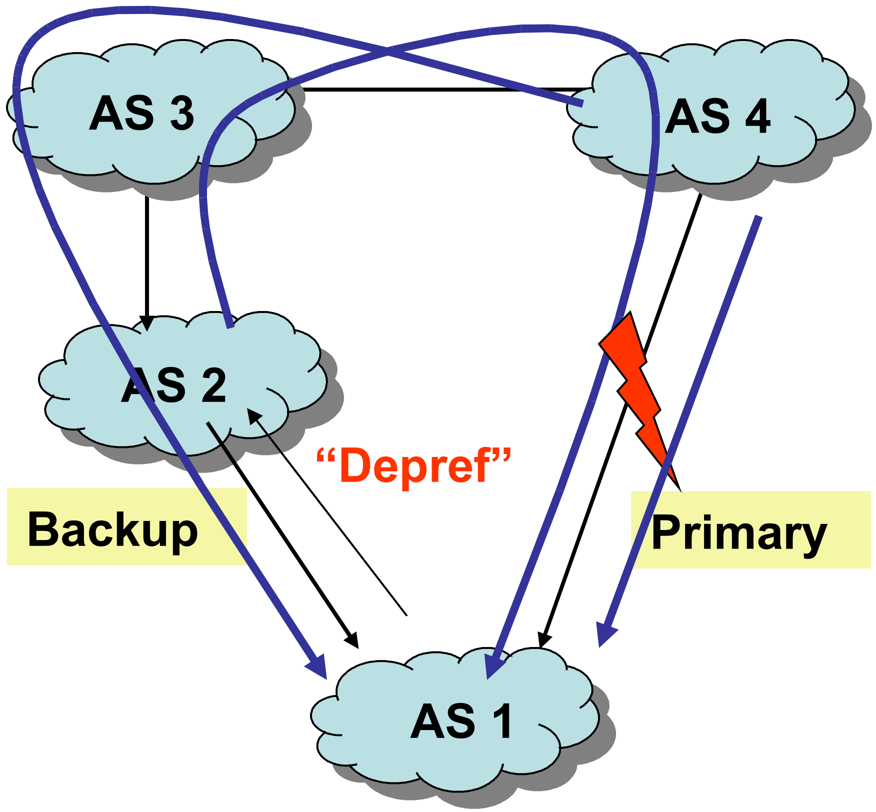

- Requires manual intervention

Detection Using Routing Dynamics

- Large volume of data
- Lack of semantics in a *single* stream of routing updates

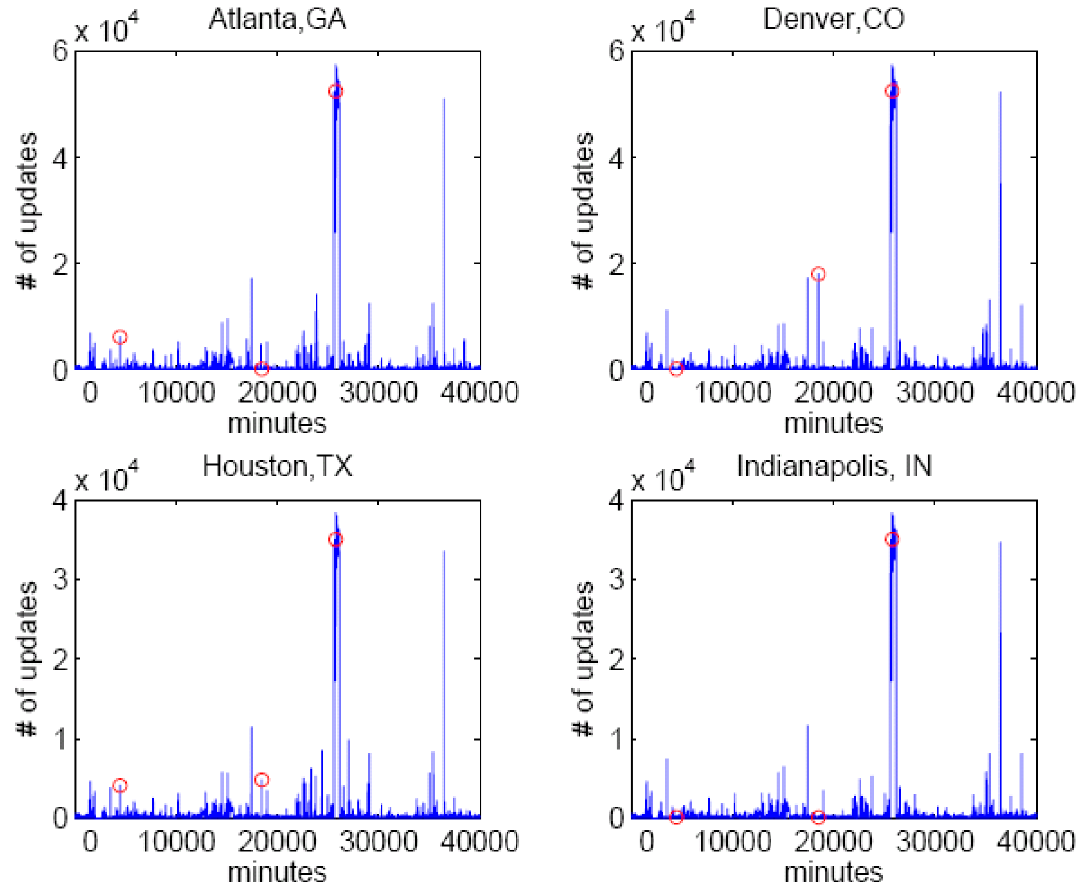

Idea: Can we improve detection by mining network-wide dependencies *across* routing streams?

Problem 3: Preventing Errors

Before: conventional iBGP

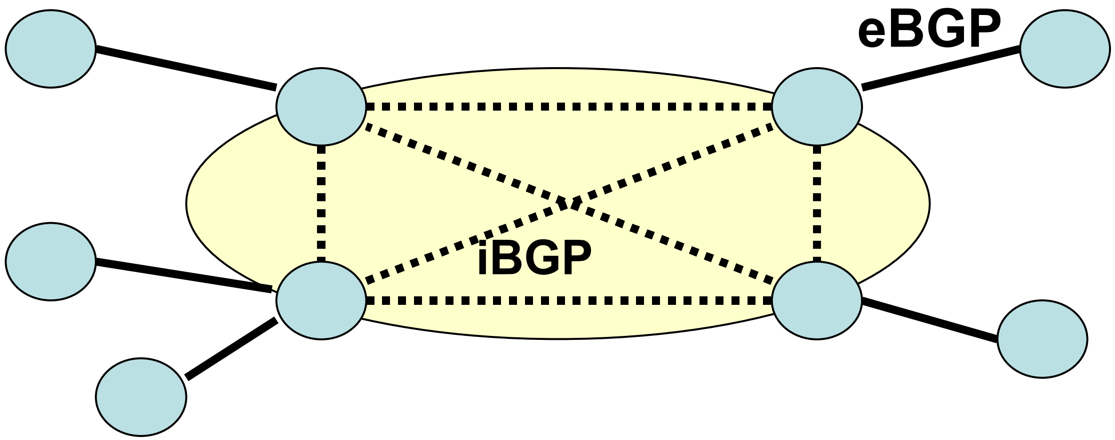

After: RCP gets "best" iBGP routes (and IGP topology)

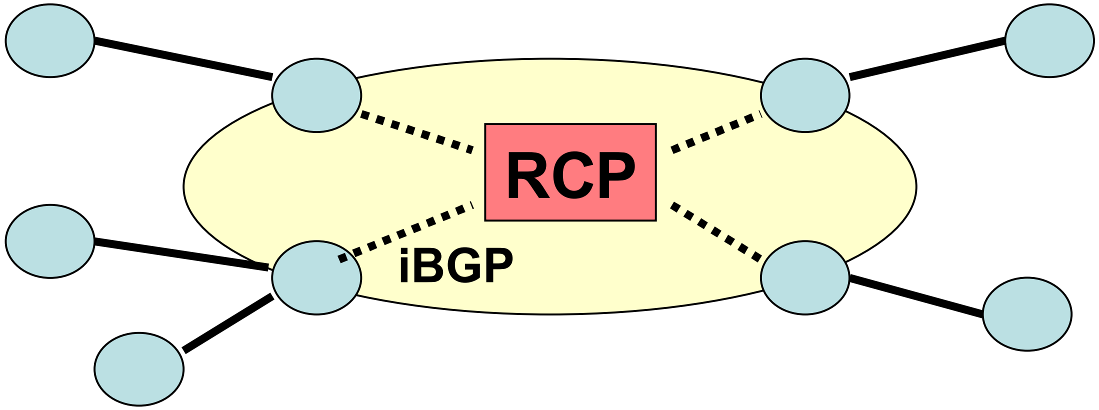

Caesar *et al.*, "Design and Implementation of a Routing Control Platform", *NSDI*, 2005