

## Part 2: Network Performance Evaluation and Management

### Goals:

- Learn techniques, strategies, pitfalls to evaluate and manage network performance

### Overview:

- simulation; ns-2
- measurement: techniques, strategies, and pitfalls
- routing behavior
- traffic estimation and engineering
- fault detection and diagnosis

**Credit:** most slides are based on Nick Feamster's rcc talk.

Part 2.5 1

## Is management really *that* important?



Part 2.5 2

## 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.

3

## 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

Part 2.5 4

## 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

**But...**

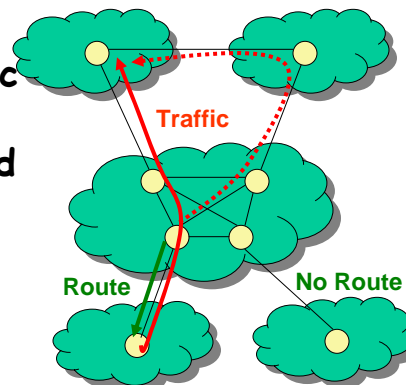
When a train falls in Baltimore, it knocks out e-mail halfway around the world.

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

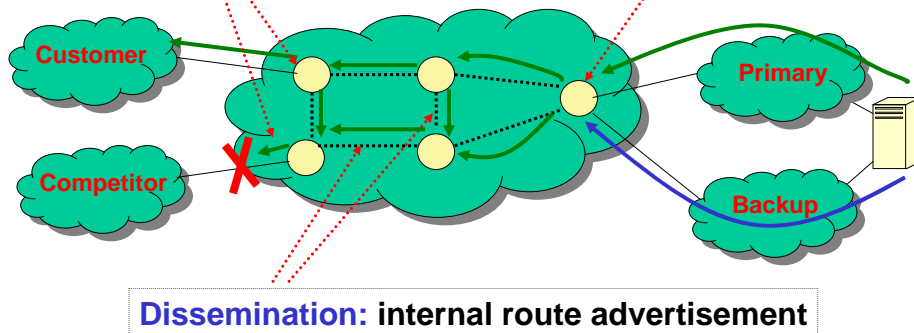


**Flexibility** → **Complexity**

## Configuration Semantics

**Filtering:** route advertisement

**Ranking:** route selection



**Dissemination:** internal route advertisement

Part 2.5 7

## What types of problems does configuration cause?

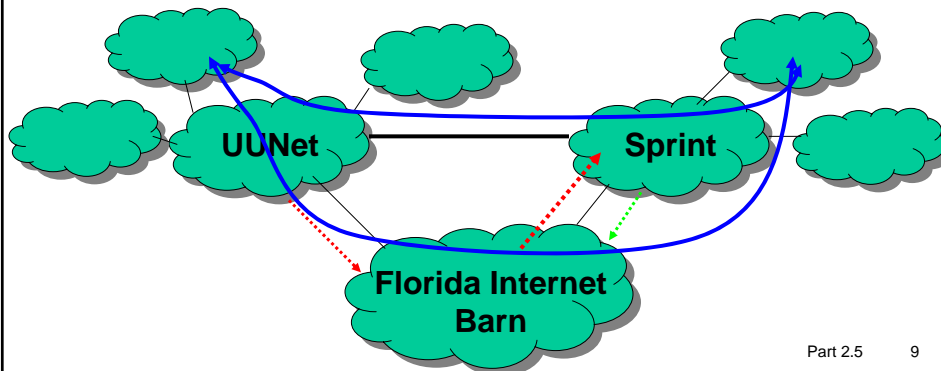
- Persistent oscillation (*last time*)
- Forwarding loops
- Partitions
- "Blackholes"
- Route instability
- ...

Part 2.5 8

## 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



Part 2.5 9

## 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

Part 2.5 10

## January 2006: Route Leak, Take 2

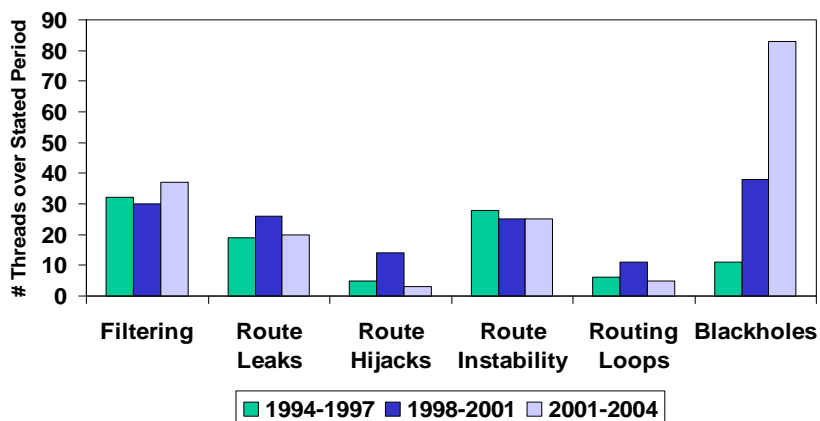
### Con Ed 'stealing' Panix routes (alexis) Sun Jan 22 12:38:16 2006

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). "

Part 2.5 11

## Faults on NANOG Mailing List



Note: Only includes problems openly discussed on this list.

Part 2.5 12

## Example: Blackhole

Date: Thu, 18 Jul 2002 06:05:10 -0400 (EDT)  
From: Chad O'leary <col@pobox.com>  
Subject: Re: problems with 701  
To: <nanog@merit.edu>

We're starting to see the same issues with UUNet, again. Anyone else seeing this? Trying to reach Qwest...

```
traceroute to 63.146.190.1 (63.146.190.1), 30 hops max, 38 byte packets
 1  esc-lp2-gw.e-solutionscorp.com (63.118.220.1)  1.167 ms  1.163 ms  1.142 ms
 2  500.Serial2-10.GW1.TPA2.ALTER.NET (157.130.149.9)  1.097 ms  1.059 ms  1.044 ms
 3  161.at-1-0-0.XL4.ATL1.ALTER.NET (152.63.81.190)  13.839 ms  14.108 ms  16.638 ms
 4  0.so-3-1-0.XL2.ATL5.ALTER.NET (152.63.0.238)  14.370 ms  14.587 ms  14.553 ms
 5  POS7-0.BR2.ATL5.ALTER.NET (152.63.82.193)  13.928 ms  14.099 ms  14.053 ms
 6  * * *
 7  * * *
 ...
```

Part 2.5 13

## 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

Part 2.5 14

## Correctness Specification

### Safety

The protocol does not oscillate



Part 2.5 15

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


***We need additional correctness properties.***

Part 2.5 16

## 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



Part 2.5 17

## 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.




Part 2.5 18

## 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.***



Part 2.5 19


How do we guarantee these additional properties in practice?

Part 2.5 20

## Today: Reactive Operation



What happens if I tweak this policy...?




- ❑ Problems cause downtime
- ❑ Problems often not immediately apparent

Part 2.5 21

## Goal: Proactive Operation


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



Many faults can be detected with static analysis.

Part 2.5 22

## rcc Overview

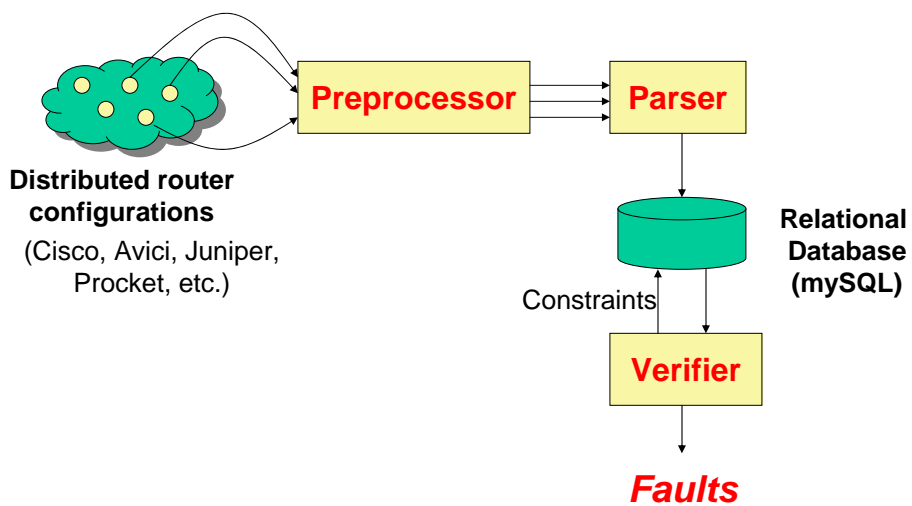


### Challenges

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

Part 2.5 23

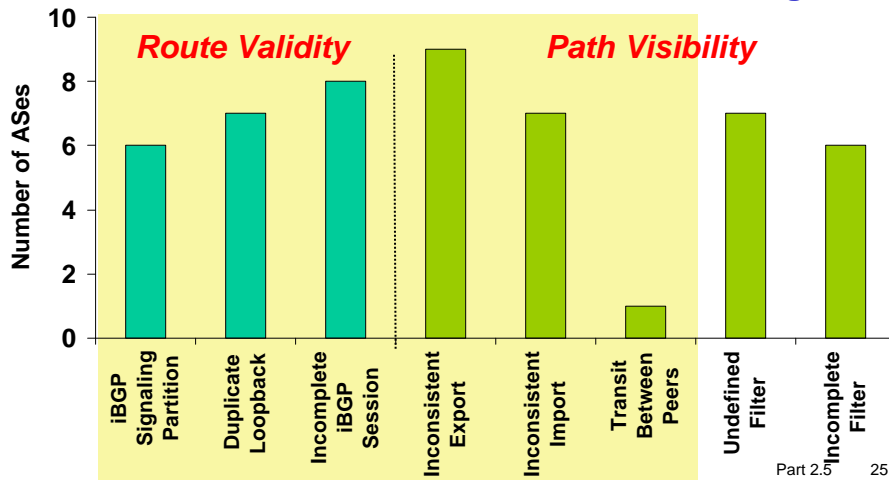
## rcc Implementation



Part 2.5 24

## 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


Part 2.5 27

## Example: 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

Part 2.5 28