ΟΙΚΟΝΟΜΙΚΟ ΠΑΝΕΠΙΣΤΗΜΙΟ ΑΘΗΝΩΝ

ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

Information-Centric Networks

Section # 4.2: Routing Issues

Instructor: George Xylomenos

Department: Informatics

European Union
European Social Fund

OPERATIONAL PROGRAMME
EDUCATION AND LIFELONG LEARNING
investing in knowledge society
MINISTRY OF EDUCATION, LIFELONG LEARNING AND RELIGIOUS AFFAIRS
MANAGING AUTHORITY

NSRF 2007-2013
programme for development
EUROPEAN SOCIAL FUND

Co- financed by Greece and the European Union

Funding

- These educational materials have been developed as part of the instructors educational tasks.
- The **“Athens University of Economics and Business Open Courses”** project only funded the reformatting of these educational materials.
- The project is being implemented as part of the Operational Program “Instruction and Lifelong Learning” and is co-financed by the European Union (European Social Fund) and national funds.

Licencing

- These educational materials are subject to a Creative Commons License.

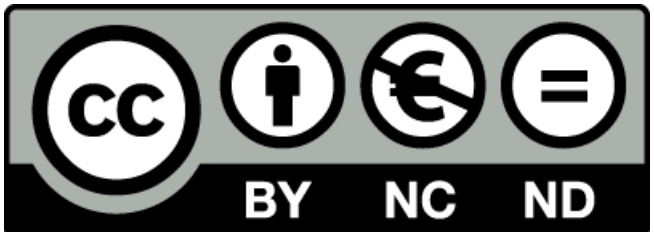

Week 4 / Paper 2

- Understanding BGP Misconfiguration
  - Rahil Mahajan, David Wetherall, Tom Anderson
  - ACM SIGCOMM 2002
- Main point
  - BGP misconfiguration can disrupt Internet connectivity
  - How often does it occur? Why does it occur?
  - Observation from multiple vantage points
  - 200-1200 prefixes misconfigured each day
  - Users are affected by very few of them

Introduction

- Focus on two types of misconfiguration
  - Accidental injection of routes into BGP tables
  - Accidental export of routes in violation of policy
- Goals of the study
  - How often are misconfigurations?
  - What is their impact on connectivity and load?
  - Why do they occur?
  - How can they be reduced?
- Observation study
  - 23 vantage points during 21 days
  - Use of simple heuristics to identify errors
  - Polling of operators to verify causes

Misconfiguration

- Focus on two types of BGP misconfiguration
  - Origin misconfiguration: erroneous injection in BGP tables
    - Failure to summarize prefixes
    - Announcing someone else's address space
    - Propagation of private prefixes
  - Export misconfiguration: advertisement of policy violating routes
  - There are many other types of misconfiguration
    - These are externally visible and clearly against policy
- Adverse impacts of misconfiguration
  - Increase of routing load due to unnecessary updates
  - Partial or global connectivity disruption
  - Routing policy violations

Methodology

- Analysis of data from the RouteViews BGP listener
  - 45% of new routes last for less than a day
  - 30% of new routes last for more than 7 days
  - Inference: misconfigurations last for less than a day
    - Requires verification by operator polling
  - Result: a lower bound on actual misconfigurations
- Origin misconfiguration analysis
  - Examination of new routes (not reappearing ones)
    - Self deaggregation: possible aggregation error
    - Related origin: possible backup route
    - Foreign origin: possible address hijacking
  - Look for routes that disappear quickly
    - More likely to be an error that was noticed

Methodology

- Export misconfiguration analysis
  - Paths are normally valley free
    - Up to the core, through the core and down to the destination
  - We can only infer the AS relationships via BGP tables
  - Result: a lower bound on actual misconfigurations
  - Types of misconfiguration
    - Provider->AS->Provider
    - Provider->AS->Peer
    - Peer->AS->Provider
    - Peer->AS->Peer
- Verification: email to operator and connectivity testing
  - Emails often bounced due to erroneous data in registries
  - Test reachability of suspect AS's from multiple vantage points

Results

- Origin misconfiguration
  - Short lived routes were clustered into incidents
    - Sets of prefixes from the same AS that appear/disappear together
  - Up to 72% of new routes seen in a day are misconfigurations
    - Extrapolation from the e-mail answers for incidents
    - Connectivity tests matched well with e-mail responses
  - 13% of the incidents impact connectivity
    - Some of the connectivity problems were not noticed by operators!
    - Extrapolation: 25 incidents per day disrupt connectivity
  - 50% of misconfigurations last less than 10 minutes
    - 80% less than an hour, 90% less than 10 hours
    - Connectivity disruptions are fixed sooner

Results

- Export misconfiguration
  - Segments with policy violations were clustered into incidents
  - Most incidents do not affect connectivity, only load
  - Provider->AS->Provider is the most common violation
    - Followed by Provider->AS->Peer
  - Impact on load is normally low
    - But it can even double load in some incidents

Causes

- Classification of human errors
  - Slips: errors in executing a correct plan
  - Mistakes: correct execution of an erroneous plan
- Origin misconfigurations
  - Mistakes
    - Initialization bugs: bug in a specific vendor's product
    - Reliance on upstream filtering: response to attacks of load balancing
    - Old configuration: unsaved changes or backup routers
  - Slips
    - Redistribution: of internal routes
    - Community: incorrect scoping of routes
    - Hijack: of addresses prefixes (attack or typing error)
    - Forgotten filter: error in filtering
    - Incorrect summary: larger or smaller address blocks

Causes

- Export misconfigurations
  - Mistakes
    - Prefix based configuration: a backup path leads to transit violations
    - Old configuration: as in origin misconfigurations
    - Initialization bug: as in origin misconfigurations
  - Slips
    - Bad ACL or route map: obvious
    - Typo: obvious
    - Forgotten filter: as in origin misconfigurations
    - Community: as in origin misconfigurations

Discussion

- What can we do to reduce misconfigurations?
- User interface design
  - Many CLIs are problematic and should be improved
  - Often operators do not really understand the CLI
- High-level languages and checking
  - Router configuration is a very low level task
  - At least high level configuration checking would be good
- Database consistency and replication
  - Registries are very outdated, leading to errors
- Protocol extensions
  - Secure BGP guards against hijacks
  - Better error reporting would reveal many other errors

End of Section # 4.2

Course: Information-Centric Networks, **Section # 4.2: Routing Issues**

Instructor: George Xylomenos, **Department:** Informatics