

HELSINKI UNIVERSITY OF TECHNOLOGY 2/28

TE in action

S-38.3192 Verkkopalvelujen tuotanto
S-38.3192 Network Service Provisioning

Networking laboratory

HELSINKI UNIVERSITY OF TECHNOLOGY 3/28
Visa Holopainen, M.Sc. (Tech.)

Concept of Traffic Engineering (TE)

- Traffic Engineering (TE) (Traffic Management) is a field of communications engineering that tries to make network operations more **effective** and **reliable** while at the same time **optimizing resource utilization**
- “Application of technology and scientific principles to the measurement, characterization, modeling, and control of Internet traffic”
- Traffic engineering ~ Traffic measurements + Traffic classification + Bandwidth management + Traffic Shaping + Protocol tuning
- RFC 3272: *Overview and Principles of Internet Traffic Engineering*

HELSINKI UNIVERSITY OF TECHNOLOGY 4/28
Visa Holopainen, M.Sc. (Tech.)

Some problems that TE tries to solve

- Effective bandwidth utilization on the path that packets are *currently using*
- Effective bandwidth utilization *within an Autonomous System*
- Optimal policy usage *between Autonomous Systems* (BGP TE)
- Fast connectivity restoration after a component breakdown (IGP Fast Convergence, MPLS Fast Re-Route, etc.)
- Result: Happier users = more money

5/28
HELSINKI UNIVERSITY OF TECHNOLOGY
Visu Holopainen, M.Sc. (Tech.)

TE in practice

- High-level TE concepts (ISP revenue models, conceptual frameworks, policies, flowcharts, heuristic functions, etc.) must be mapped (coded) into the network's components

6/28
HELSINKI UNIVERSITY OF TECHNOLOGY
Visu Holopainen, M.Sc. (Tech.)

Load balancing - problem

- Networks may have such properties that traffic is sometimes distributed very unevenly within the network -> bottlenecks
- It may be expensive (or impossible) to add links into the network
- In some situations it may be possible to distribute the offered load more evenly within the network, so that bottlenecks are avoided

7/28
HELSINKI UNIVERSITY OF TECHNOLOGY
Visu Holopainen, M.Sc. (Tech.)

Load balancing: ECMP

- IGP routing protocol maintains multiple equal-cost routes to all destinations
- Traffic is distributed (somewhat) evenly for all equal-cost routes
- An implementation may choose to keep only a fixed number of routes to any given destination

8/28
HELSINKI UNIVERSITY OF TECHNOLOGY
Visu Holopainen, M.Sc. (Tech.)

ECMP in practise

- Easy setup, "configure and play" (at least in principle), however, some things need to be considered...
- Per packet load balancing with ECMP is not usually a good thing
 - Variable latency and packet re-ordering
 - Solution: Cache nexthop for each flow
 - Problem: Need to have much more cache entries than with traditional routing
- Some routing protocol implementations will not cause the Forwarding Information Base (FIB) to change if one of the equal-cost interfaces goes down
 - Need to have some kind of L2 detection that takes care of removing the links from FIB
- ECMP is sometimes not enough, since it doesn't use "intelligence" in routing

HELSINKI UNIVERSITY OF TECHNOLOGY 9/28
Visa Holopainen, M.Sc. (Tech.)

Traffic Measurements

- If we want more intelligent load balancing than ECMP, measurements are needed
 - Passive measurements
 - Listen to network traffic
 - Active measurements
 - Send probes to the network and analyze responses
 - Benefit: Possibly more accurate measurements
 - Drawback: Extra traffic into the network

HELSINKI UNIVERSITY OF TECHNOLOGY 10/28
Visa Holopainen, M.Sc. (Tech.)

Traffic Measurements

- Many decisions need to be made
 - What to measure?
 - Bits, Bytes, packets, flows, topology?
 - Where to measure?
 - Networks edges, all links, certain links, certain nodes?
 - How to measure?
 - What types of measurements does the platform enable?
 - How to export the measurements?
 - What methods does the platform enable?
 - File formats?
 - How often to measure?
 - Too often may cause overhead
 - How often to export the measurement data?
 - Too often may cause overhead & oscillation

HELSINKI UNIVERSITY OF TECHNOLOGY 11/28
Visa Holopainen, M.Sc. (Tech.)

Traffic Measurements

- Internet Protocol Flow Information eXport (IPFIX) IETF working group
 - Created from the need for a common, universal standard of exporting IP flow information from routers
 - RFC 3917: Requirements for IP Flow Information Export
- SNMP
 - Query the devices (Request-Response) about their load, or configure the devices to send the data on triggers (SNMP traps)
 - Nice principle, however sometimes just doesn't work (buggy implementations)

HELSINKI UNIVERSITY OF TECHNOLOGY 12/28
Visa Holopainen, M.Sc. (Tech.)

Traffic Measurements


- Cisco NetFlow
 - Cisco routers that have NetFlow-feature enabled, generate netflow records
 - These are exported from the router using UDP or SCTP packets and collected using a netflow collector
 - Understands flows in addition to packets and bytes
- Juniper Cflow, Huawei NetStream
 - Basically same things that NetFlow
- Do It Yourself: Packet/Byte counters
 - Possible on open platforms (Linux, BSD, ...)
 - Self-made code that monitors interface packet counters
 - Export the statistics whenever you want



Traffic Measurements

- Wire tap

- Place a passive Ethernet tap inline between a host machine and the Ethernet switch using the two outside positions labeled "HOST"
- Tap A will show half-duplex traffic and Tap B will show the remaining traffic. You will need to use two Ethernet interfaces to examine both halves of the full-duplex signal



Improving load balancing

- Traditional routing protocols (OSPF, IS-IS) consider a link to be unusable only when it is physically broken
- A link may also be effectively unusable when there is too much load on it
- How do we know when there is too much load on a link
 - Obviously measurements are needed
 - Also a standardized way to tell others about the results of measurements has been developed – *TE extensions*
 - Note: any other way can also be used...



TE extensions

- Similar in OSPF (RFC 3630) and IS-IS (only drafts available – work in progress)
- The information made available by TE extensions can be used to build an extended link state database just as router LSAs are used to build a "regular" link state database
- The difference is that the extended link state database (Traffic Engineering Database, TED) has additional link attributes (e.g. free BW)
- Uses of the TED include:
 - Monitoring the extended link attributes
 - Local constraint-based source routing
 - Global traffic engineering



Traffic Engineering LSA

- TE attributes are carried by TE LSAs (Opaque LSAs)
- The LSA payload consists of one or more nested Type/Length/Value (TLV) triplets for extensibility



TE extensions in practise

- Ambitious goal: make *routing protocols* traffic-aware
- This is not an easy task:
 - Routing protocol code has to be altered fundamentally in order to use TE extensions –hard work
 - Router LSAs are easy to handle – no LSAs in time period t from router Y to X via link Z means that router X should decide that link Z is down
 - The handling of TE LSAs is not so well-defined
 - What does it mean if a TE LSA has value X or Y in some of its fields
- Not "configure and play"



Local constraint-based source routing

- Router A computes path to B
 - This path may be subject to various constraints on the attributes of the links and nodes that the path traverses, e.g., use only links that have unreserved bandwidth of at least 10Mbps
- One means of instantiating these paths is using MPLS tunnels
- Constraint-based routing can be NP-hard, or even unsolvable, depending on the nature of the attributes and constraints, and thus many implementations will use heuristics



Global TE

- A device (TE server) can build its own traffic engineering database (TE extensions not needed), input a traffic matrix and an optimization function, crunch on the information, and thus compute optimal or near-optimal routing for the entire network
- The device can subsequently monitor the traffic engineering topology and react to changes by recomputing the optimal routes



Global TE

- The TE server (Policy Decision Point) can either calculate optimal LSPs (MPLS TE) or optimal IGP costs (IGP metrics optimization)
 - Configurations are then sent to routers (Policy Enforcement Point)





MPLS TE

- RFC 2702 - Requirements for Traffic Engineering Over MPLS
- Benefits:
 - Enables explicit (optimal) routing – optimal load balancing
 - Protection can be incorporated in the computation
- Drawbacks:
 - Additional layer of complexity



IGP Metrics Optimization

- Benefits:
 - Simple solution
 - Fallback to traditional routing easy
- Drawbacks:
 - Does not enable optimal routing
 - However, in many real networks gets very close to optimal




TE tools

- Can be used at a centralized TE server (PDP)
- Input = measured data from the network
- Output = (near) optimal configurations



TE tools - MATE

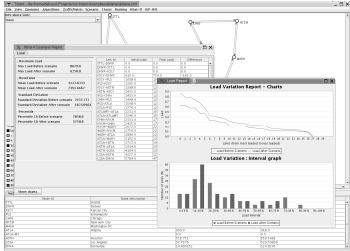
- Cariden (www.cariden.com) MATE-tool
 - Commercial
 - IGP Metric Optimization Package, MPLS Simulation Package, Explicit Routing Optimization Package, Demand Deduction Module, BGP simulation, ...



HELSINKI UNIVERSITY OF TECHNOLOGY 25/28
Vesa Holopainen, M.Sc. (Tech.)

TE tools - TOTEM

- TOTEM (TOolbox for Traffic Engineering Methods) (www.totem.run.montefiore.ulg.ac.be)
 - Open-source



HELSINKI UNIVERSITY OF TECHNOLOGY 26/28
Vesa Holopainen, M.Sc. (Tech.)

TOTEM

- Traffic matrix, topology and scenario XML files are input of TOTEM
- Output: Optimal IGP costs, MPLS LSPs, ...

```

<!-- Example XML snippet for TOTEM input -->
<!-- Topology -->
<!-- Node definition -->
<!-- Interface definition -->
<!-- Link definition -->

```

HELSINKI UNIVERSITY OF TECHNOLOGY 27/28
Vesa Holopainen, M.Sc. (Tech.)

TE tools - Opnet SP Guru

- Commercial (www.opnet.com)
- “SP Guru MPLS Traffic Engineering feature automatically defines explicit routes that minimize maximum link utilizations under normal conditions, and secondary routes that will survive link and node failures.”

HELSINKI UNIVERSITY OF TECHNOLOGY 28/28
Vesa Holopainen, M.Sc. (Tech.)

TE tools - Wandl IP/MPLSView

- Commercial (www.wandl.com)

