Traffic management

Lecture for QoS in the Internet –course S-38.180

16.10.2003 Mika Ilvesmäki

Contents

- Traffic management
  - Terminology and system structure
- Traffic management components
  - Classification
  - Traffic handling mechanisms
  - Bandwidth management
  - Policy systems
  - Monitoring
  - Billing
- Traffic management applied in DiffServ

Traffic management

- TM systems consist of a set of high-level rules that are propagated out to enforcement points using a policy system
  - Policy must be enforced to ensure that the users are behaving properly
- Network should classify, handle, police and monitor the traffic
  - operator should also be able to bill the customer

Terminology (RFC 3198)

- Policy is either:
  - A definite goal, course or method of action to guide and determine present and future decisions. "Policies" are implemented or executed within a particular context (such as policies defined within a business unit).
  - a set of rules to administer, manage, and control access to network resources [RFC3060].
- Policies are built with policy rules
  - Policy rule is a basic building block of a policy-based system. It is the binding of a set of actions to a set of conditions - where the conditions are evaluated to determine whether the actions are performed [RFC3060].
- Policy condition is usually a filter
  - A set of terms and/or criteria used for the purpose of separating or categorizing. This is accomplished via single- or multi-field matching of traffic header and/or payload data. "Filters" are often manipulated and used in network operation and policy. For example, packet filters specify the criteria for matching a pattern (for example, IP or 802 criteria) to distinguish separable classes of traffic.

Policy system structure

- Policy systems as such are pretty straightforward
  - Policy clients at routers ask the policy parameters from the policy server
  - Policy servers get the policy data from the information store
- Key question rarely given thought: How do you *create* the policy rules and the corresponding actions?

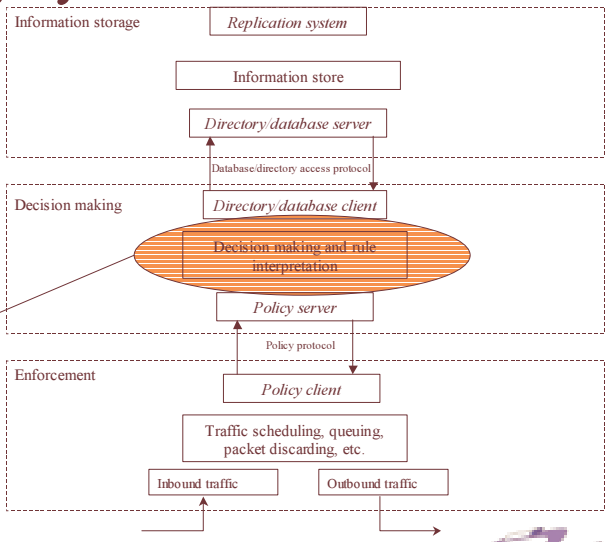

Traffic classification

- The main idea is to determine the packet class
- Based on experience and scalability studies the easiest way to bring service differentiation into the Internet is to use a limited amount of traffic classes (DiffServ).
  - But how many? 2, 3, 8 or more?
- Different traffic classes represent different priority levels

User decisions

- Users may inform the network on the service level (class) of the packet.
  - resource restrictions -> admission control
  - malicious users may want to misuse the network capacity
  - users want to measure the service level they get -> added complexity/software/traffic
  - and... do all the users _really_ have the expertise to make the decisions?!
- Users should be required to provide only minimum of information on the traffic characteristics!

Network decisions

- Network determines the service level (class) of the packet
  - feedback from the resource usage
    - SLAs do not promise anything absolute in terms of network service
  - AAA (Authentication, Accounting and Administration) guarantees the service levels to appropriate users
- If network decides individual packet treatment it should know what kind of packet it is classifying
  - This requires knowing the application characteristics
    - by examining the packet headers and/or content
    - by information obtained from other network devices that know the packet's type

Where's the info on the packet contents?

- Packet header information
  - layers 1 and 2 do not contain any information on packet content
  - layer 3 (IP) identifies the sending source and receiving destination the upper layer 4 protocol (TCP/UDP)
    - oversimplification: who sends packets where
  - layer 4 (UDP/TCP) identifies the port numbers used at source and destination
    - oversimplification: what application is used
    - source identifies the application that originates the packet and the destination tells us where the packets are headed
- Layers 3 and 4 are the first ones that contain any information on the application that the user is using to create packets in the network.
  - Aim is to limit the processing on the packet

HELSINKI UNIVERSITY OF TECHNOLOGY Mika Ilvesmäki, Lic.Sc. (Tech.)

Design guidelines for classification

- Plan for scalability
  - do not associate port numbers to QoS classes (-> potentially 65535 classes), instead bind the port numbers to DiffServ Codepoints (DSCP), for instance.
    - Port number have nothing to do with QoS identification whereas DSCP is designed just for that
- Do not imply policy within design
  - Use as value-neutral design as possible and leave room for freedom of choice
- Preserve end to end principle: "If possible do everything at the edges."
  - Profiling and marking should be done and used at the edges of the network
    - although measurements may, of course, be done anywhere in the network

Some classification problems

- NAT
  - User-based classification impossible
    - Pre-translation packet marking
- Stateful traffic
  - Upper-layer negotiates traffic (FTP)
    - Traffic monitoring
- VPN
  - Hides (as does NAT) the “true nature” of the traffic
    - Pre-VPN-entry packet marking

Traffic handling

- In a device
  - Shaping and queuing traffic
    - Leaky and token buckets, FIFO, PQ, CBQ, WFQ…
    - RED, WRED etc. for queue management
      - What are the correct parameter values?
- By path selection (QoS routing)
  - IntServ and DiffServ do not choose or resolve routes
    - the “best” routes chosen by current protocols are used
      - OSPF, BGP, etc.
    - problems: route oscillation, path capacity

Bandwidth Broker

- Outside intelligence which controls the network provisioning
  - Makes possible to offer end-to-end semantics
    - **Domain wide**
    - **Inter-domain**
      - We need to
        - » translate domain specific service attributes at the border of two domains (pretty fixed)
        - » Dynamically adjust resource requests to the other domain...

Enabling IntServ / DiffServ co-existence

- Bandwidth Broker interprets RSVP messages to modify the domain specific weights and filters
- We need to be able to pass reservation attributes to and from IntServ cloud.
  - IntServ cloud may be
    - Corporation
      - Outbound / inbound traffic is delivered as guaranteed traffic
        - » Mapping to DiffServ classes based on policy
    - Other ISP having IntServ as backbone
      - Mapping between IntServ and DiffServ classes

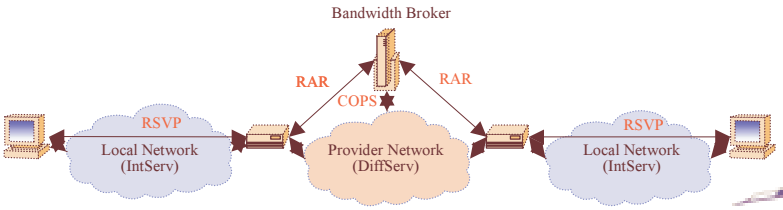

Bandwidth Brokers vs. IntServ routers

- Are we rotating things back to IntServ ?
  - BB:s require knowledge from the network (offered load, provisioning)
    - By measuring the network
    - By signaling from the users
  - BB:s modify conditioning and forwarding actions of network routers
- What is the difference to the IntServ ?
  - If we provide end-to-end service we need fixed routes and resources that at the minimum match the requirements
    - We need state information somewhere
      - Centralized - DiffServ BB:s
      - Distributed - IntServ routers

HELSINKI UNIVERSITY OF TECHNOLOGY Mika Ilvesmäki, Lic.Sc. (Tech.)

Policy systems

- RADIUS
  - Remote Access Dial Up User Services
  - Stateless protocol for authenticating dial-up users
- DIAMETER (extended RADIUS ☺)
  - Extensibility and statefulness
- COPS
  - A client/server model where Policy Enforcement Point (PEP) sends requests, updates and deletes to Policy Decision Point (PDP) and where PDP sends its decisions back to PEP.
  - TCP based
  - Stateful
  - Provides a way to distribute policy configuration to devices
    - No monitoring

Inter-domain issues

- Inter-domain traffic forwarding is based on bilateral of multilateral peering agreements
  - These tend to be rather static (due to the fact that there are probably lawyers etc. dealing with the issues)
    - Rule of thumb: more money -> more lawyers -> more static
  - However, demand is varying rapidly and therefore we more flexible peering agreements
    - We need to brake that rule by defining peering more dynamically
      - Could inter- operator billing be based on the aggregate traffic in the classes and rate of change requests

Evaluation of the policy systems

- Evaluate the network (element)
  - Use of transmission capacity, architecture dependent router resources (connection setup / class, packet forwarding / class etc.)
- Evaluate the effect on user
  - What applications are classified to priority
    - Relevance, application type, application count
    - How good the user feels?
    - Is she getting her money's worth?

Monitoring

- Measurements need to be 2-way
- Passive measurements
- Active measurements
  - May affect the network status
- The measured properties may be sorted, or otherwise analyzed against
  - absolute boundaries (particular packet sizes, certain variance limits)
  - each other (all packets smaller/larger than the average packet size are classified/not classified)

What to measure?

- The basic data block in the Internet is the IP packet
  - packets are made of bytes that are made of bits.
    - no info on the overlaying application
- IP packet identifies the overlaying protocol
  - TCP or UDP as far as user apps are concerned
    - TCP/UDP port numbers identify, to some degree, the application used

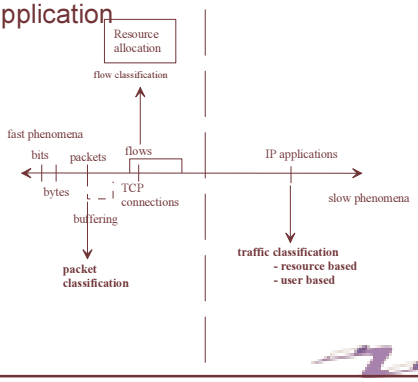

Pricing/Billing alternatives

- Flat rate (even sum/month)
- Usage based
  - received data
  - sent data
  - use of resources (Bandwidth etc.)
- Billing based on user profile
  - Being a member of user group
  - Using certain applications (VoIP-phone vs. Web-browser)
- In practise Internet routers and Internet in general has not been designed to collect and update the network usage of an individual user (scalability)
- Combination of any and all of the above
- How complicated can an Internet-bill be so that the user may verify it and accept it?!

Case: End-to-end service in DiffServ

- Obstacles
  - Structure of DiffServ is based on local control (policies)
    - Classification based on the policies at the edge of the network
    - Forwarding based on the policies in the core of the network
  - We **can** stretch through single domain (ISP) with EF
  - We **may** stretch through single domain (ISP) with AF
- End-to-end
  - **Is not** within single ISP
  - It **is** between source and destination

What a customer wants...

- Lets face the music
  - Customer **is** only interested in the **perceived quality**
    - How things are rolling compared
      - Minute ago
      - Year ago
  - Customer **is not** interested in the novel **technology** which is behind the service
  - This means end-to-end service quality

HELSINKI UNIVERSITY OF TECHNOLOGY
Mika Ilvesmäki, Lic.Sc. (Tech.)

Expedited Forwarding

- 'End-to-end' service
  - Actually single domain end-to-end
  - Quality is defined by two constraints:
    - Provisioning
      - Class should be provisioned with enough resources to handle worst case aggregate
    - Sharing
      - No resource reservation for individual flows.
      - Under and overflows possible
      - Timing and delays can not be held or guaranteed

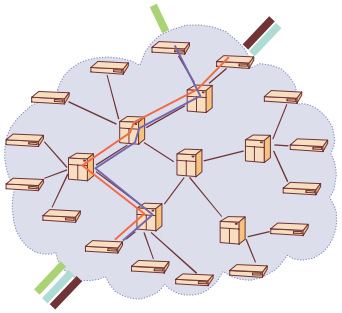

Assured Forwarding

- No end-to-end semantics
  - Service can be deployed
    - Point-to-point
    - Any-to-any
  - Uncontrollable resource usage inside the network
    - Problem of commons

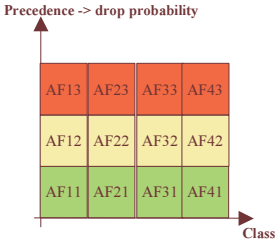

Making AF end-to-end

- To make AF end-to-end we need to control resources and offered load hand in hand
  - Adjust scheduling (to control resources)
  - Reroute some of the classes (to control offered load)
    - » Class and constraint based routing

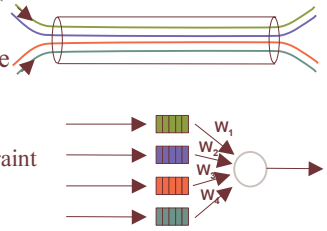

- Adjust scheduling by modifying CBQ heuristics:
  - If class green is unsatisfied and class turquoise is unsatisfied but at the scale of the network only class green is unsatisfied we allow only green to borrow.
    - Not possible with the logic we have today in DiffServ, because a single router does not know network scale situation (stateless)

A possible solution

- Have intelligence (bandwidth broker) outside the network which would control the scheduling of classes adjust scheduling parameters.

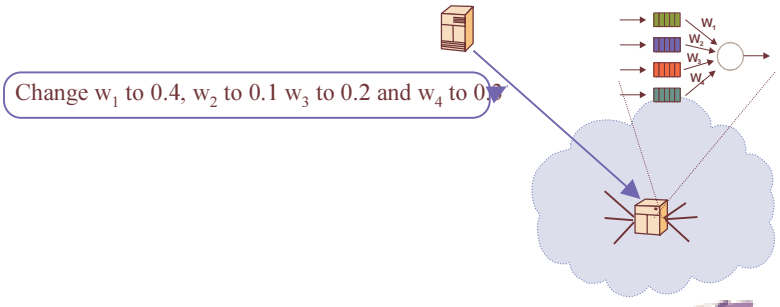

What DiffServ offers…

- Differentiated Services is service architecture which allows to build N logically separated Best Effort networks into a single physical network
- Differentiated Services provides tools to offer QoS which is only assured
- Differentiated Services does not provide end-to-end semantics to the services which are built upon it
- End-to-end semantics are only achieved with outside intelligence - like bandwidth brokers

Summary: General model for QoS traffic management

- Differential handling of traffic
  - Explicit reservations or implicit or administrative differentiation
- Making decisions to handle incoming packets
  - Local, preconfigured or on-line admission control
- Packet forwarding
  - Queuing, shaping, discarding etc.
- Removal of obsolete policy information