NSWI090: **Computer Networks**

http://www.ksi.mff.cuni.cz/~svoboda/courses/202-NSWI090/

Lecture 8

Internetworking II

Martin Svoboda
svoboda@ksi.mff.cuni.cz

26. 4. 2021

Charles University, Faculty of Mathematics and Physics

Lecture Outline

Internetworking

- **Broadcast** domains
- Interconnection devices at L3
- **Virtual LANs**
  - Motivation and deployment
- **Firewalls**

Internetworking Principles

80/20 rule

- Traditionally…
  - Usually ≈ **80% of traffic was local** within a given network
  - And only ≈ 20% was leaving such a network

20/80 rule

- Things significantly changed with the Internet…
  - Usually only ≈ 20% is still local
  - Even ≈ **80% of traffic crosses the border of a local network**
- Routers may no longer be able to handle increasing data flows
- **Solutions**
  - **Virtual Local Area Networks** (**VLAN**)
    - Harness fast interconnection at L2, but limit broadcast domains
  - L3 **Switches**
    - Increase overall efficiency and throughput of traditional routers

Broadcast Transmissions

L2 broadcast

- Intended **recipients**
  - **All nodes within a given local network** = **broadcast domain**
    - I.e., all nodes residing in the same network as the sender node
- Frame **destination address**
  - `FF:FF:FF:FF:FF:FF`
    - Special address with binary ones only
- **Delivery** process
  - **Bridges and switches**: forwarding based on **flooding**
  - **Routers** (in our network): further propagation is stopped
- Natural motivation
  - **Limiting the size** of broadcast domains

Broadcast Transmissions

Local L3 broadcast

- Intended **recipients**
  - Once again, **all nodes within a given local network**
    - Only this time in the context of IP datagrams at L3
- Datagram **destination address**
  - `255.255.255.255`
    - Once again special address with binary ones only
- **Delivery** process
  - **Sender**: IP datagram is requested to be sent using L2 **broadcast**
  - **Routers** (in our network): further propagation is stopped

Broadcast Transmissions

Targeted L3 broadcast (**Directed L3 broadcast**)

- Intended **recipients**
  - **All nodes within a given particular network**
    - Usually **foreign network** (but also works for the local one)
- Datagram **destination address**
  - E.g.: `192.168.1.255`
    - **Network prefix at the beginning**, binary ones at the end
- **Delivery** process
  - IP datagram is first **routed and forwarded** using **standard unicast delivery**
  - Once the **router** serving as the **entry point to the target network** is reached, local L2 **broadcast** is then utilized
- Security considerations
  - Incoming targeted broadcasts are usually ignored nowadays

Network Layer Devices

Possible alternatives for L3 interconnection devices

- **Router**
  - **Traditional** complex device allowing for **routing and forwarding**
  - Suitable for **transition** between **heterogeneous** environments
- **L3 Switch**
  - Newer **integrated device** combining L2 **and** L3 **functionality**
    - Standard L2 switch for local network
    - Simplified but more efficient L3 router
  - Suitable for **interconnection** of **homogeneous** environments
- **Multilayer switch**
  - Basically L3 switch allowing to take into account information from higher layers L4 and / or even L7 for **routing decisions**
    - In particular, **L4 Switch** and **L7 Switch**

Network Layer Devices

Router

- Optimized for **logical functions** (and not only the core ones)
  - **Routing and forwarding**
  - Network Address Translation (**NAT**)
    - Allows to use **private IP addresses** in private networks
  - Assignment of IP addresses (**DHCP**)
  - Security: **firewall**, access rights, …
  - **Monitoring**, **management**, …
  - …
- Speed and throughput are not critical
  - As router was originally designed for 80:20 environments
  - **Implemented at the software level**
    - On top of a dedicated operating system (Cisco IOS)

Network Layer Devices

Router (cont'd)

- Suitable for transition between heterogeneous environments
  - **Bigger routing tables**
  - Usually bigger buffers
  - Can have **physical interfaces** with **different technologies**
    - Ethernet, EuroDOCSIS, xDSL, SDH, …
  - Can support **multiple routing protocols**
- Used for **connection to other networks**
  - Usually smaller networks (LAN, MAN) to larger ones (WAN)
  - Emphasis is put on…
    - Adaptation, logical separation, correct decision-making, …

Network Layer Devices

L3 Switch

- Optimized for **speed and throughput**
  - As L3 switch was originally designed for 20:80 environments
  - **Implemented at the hardware level**
    - So that it can match the wire speed
  - Focuses only on the **core functionality**
    - I.e., routing and forwarding
- Suitable for interconnection of homogeneous environments
  - Usually smaller routing tables and smaller buffers
  - Usually **Ethernet** physical interfaces only
- Used for **interconnection of related networks** (LAN, MAN)
  - Also allows to limit broadcast domains
    - Analogously to routers, but more efficiently

L4 and L7 Switches

L4 Switch

- L3 switch which can take L4 **information** into account
  - I.e., **routing decisions** can also be based on…
    - Transport **protocols** (TCP, UDP, …) and / or **port numbers**
- Different kinds of traffic can thus be treated differently
  - E.g., port 80 (HTTP requests), port 53 (DNS queries), …

L7 Switch (Content Switch)

- L3 switch which can take L4+L7 **information** into account
  - I.e., **routing decisions** can also be based on L4 and…
    - Application **protocols** (HTTP, SMTP, …) and their data
- Analogous utilization as above
  - E.g., port 80 HTTP requests to specific URLs in GET headers, …

L4 and L7 Switches

Use cases: **diversified routing**

- **Distribution** of requests
  - Requests to different services (e.g., HTTP, FTP, …) are in fact forwarded to different servers each providing just one of them
- Simulation of **anycast** transmissions
  - Requests to the same service are in fact split between multiple standalone serves (stickiness may be required)
- **Load balancing**
  - Exploitation of more different routing paths
- **Transparent caching**
  - HTTP requests are redirected to a dedicated cache server
- **Redirection** of DNS queries
- …

L4 and L7 Switches

Use cases: **traffic management**

- Traffic **prioritization**
  - Multimedia data may be handled preferentially
- Traffic **blocking**
  - Certain kinds of traffic may be strictly prohibited
    - E.g., VoIP communication, …
- Traffic **limitation**
  - Introduction of **volume quotas** for various kinds of traffic
    - E.g., **Fair Use Policy** (**FUP**)

Virtual Local Area Networks

Motivation

- L3 **network** = **set of end nodes** residing in one or more L2 segments interconnected using bridges / switches
  - All involved nodes are **mutually visible** and **directly reachable**
    - And so all L2 **traffic is also visible** to the entire network
  - This is not always desirable
    - Especially in buildings with systematic cabling deployed
    - Since individual users (end nodes) may not be related at all
- And so what if **membership of end nodes to networks** would be determined differently?
  - I.e., **independently on physical locations**
  - Separate switches and physical rewiring could then help
    - But this approach is not flexible enough
  - And so the concept of **VLAN** was introduced

Virtual Local Area Networks

VLAN (**Virtual LAN**)

- Principle: coexistence of **multiple different virtual networks on top of one physical `L1+L2` infrastructure**
  - Allows to decouple…
    - **Physical users locations** from **logical network memberships**
  - And so individual VLANs can reflect different…
    - Organizational needs, groups or categories of users, access or other privileges, usage of services and servers, …
- Whole concept is generic
  - Both older proprietary and newer standardized solutions exist
  - **Implemented in several technologies**
    - **Ethernet**, ATM, …

VLAN Principles

Requirements

- **Additional logic** needs to be added into the infrastructure
  - Primarily **VLAN-aware switches** at L2
  - But also **routers** at L3
- Practical expectations
  - **End nodes should remain ignorant** to the whole concept
    - I.e., they should not need to know what VLAN they are part of, nor whether VLANs are being deployed and utilized at all
    - Thus their interfaces / software do not need to be upgraded
  - $\Rightarrow$ only network administrators should concern themselves
- Fundamental requirement
  - **Traffic belonging to a given VLAN** must stay within that VLAN
    - I.e., it must be guaranteed that it will not leak to a different one
    - And so **VLAN hopping** must be avoided

VLAN Principles

Consequences and features

- Limiting **broadcast domains**
  - Broadcasts and unknown unicasts are flooded everywhere
- Improving **security and privacy**, minimizing external threats
- Enabling **Quality of Service**
  - Kind of VLAN side-effect, based on traffic prioritizing
- Simplifying **network administration** and **fault management**

VLAN concepts

- Two basic types of virtual networks can be distinguished
  - **Local VLANs** and **End-to-End VLANs**
- They both differ in the primary motivation and objectives
  - However, their mutual **boundaries are not defined strictly**

VLAN Concepts

Local VLANs

- Aim at **<u>separating</u> geographically <u>close</u> nodes**
  - In the reach of just one switch (or a small group of switches)
  - This allows for easier implementation of the whole concept
- Primary goal: **limiting broadcast domains**

End-to-End VLANs

- More generic concept
- Aim at **<u>interconnecting</u> geographically <u>remote</u> nodes**
  - Individual nodes are dispersed throughout the whole network
  - And so **VLANs span multiple switches** across the network
    - Special links between the switches are therefore needed
    - So that they can carry traffic of several different VLANs at a time
- Primary goal: **grouping users with similar interests**

Logical Model

Set of VLANs, each associated with…

- Distinct **integer** VLAN Identifier (**VID**)
- Optional **name** allowing for user-friendly management

Types of **segments** involved in the infrastructure

- **VLAN-unaware segments**
  - Contain nodes from **exactly one VLAN**
    - Actually just a single node in case of microsegmentation
    - Transmitted frames do not need to be mutually distinguished
  - Correspond to **switch-to-host** links
- **VLAN-aware segments**
  - Carry traffic from **several different VLANs**
    - And so such frames must be tagged to be mutually recognizable
  - Correspond to **switch-to-switch** or **switch-to-router** links

Logical Model

Operation principles

- VLAN can actually be seen simply as kind of a projected network consisting of only segments where it is activated
  - From this point of view, everything works as expected
  - I.e., **filtering and forwarding**
    - Including Spanning Tree Protocol (STP), etc.

VLAN configuration

- Expressed via **association of switch ports to VLANs**
  - I.e., not directly in terms of the intended usage of segments
- In particular, **each port** is labeled with a **set of permitted VIDs**
  - Obviously, network administrator must ensure **consistency**
    - I.e., corresponding ports on switches containing a given segment must be configured identically

Types of Ports

Access port (**untagged port**)

- Connects a VLAN-unaware segment
  - Labeled with **exactly one VID**
    - If not specified, **default VLAN** is assumed (usually VID 1)
  - This very VID **determines the VLAN membership** of nodes
- All frames (are expected to) belong to this single VLAN
  - **Incoming frame** is altered by **tagging** it with a given **port VID**
    - So that it becomes prepared to enter VLAN-aware segments
    - Already tagged frame is only accepted if it matches the port VID
  - **Outgoing frame** is altered by removing its tag
- **Tagging mechanism** is required
  - Open standard **IEEE 802.1q** (**Dot1q**)
  - Proprietary approaches: Cisco ISL (Inter-Switch Link), …

Types of Ports

Trunk port (tagged port)

- Connects a VLAN-aware segment
  - Labeled with **one or more VIDs**
    - By default, **all VLANs**
    - Or enumeration of only **selected VLANs**
- Frames of all involved VLANs are carried alongside each other
  - And so they must be **tagged so that they can be distinguished**
  - **Incoming frame** is only accepted if it matches the allowed VIDs
- Native VLAN may optionally be specified
  - Its frames may **remain untagged**
    - This allows to have VLAN-unaware devices in trunks as well
  - Configured on a per-port and per-device basis
    - Must hence be consistent within the entire trunk segment
    - Typically the same value everywhere (for sanity)

VLAN Configuration

Static (port-based) approaches

- Each port is **configured manually** by network administrator
- Relatively small overhead, higher security, not flexible enough

Dynamic approaches

- **VLAN membership** is resolved dynamically
  - Based on **MAC addresses** (deprecated, not a good idea anyway)
  - Or **IEEE 802.1X authentication** (based on user credentials)
- Information needs to be shared between switches
  - **Multiple VLAN Registration Protocol** (**MVRP**) (IEEE 802.1ak)
    - L2 protocol allowing to de/registers VIDs on ports, …
  - Proprietary approaches: Cisco VTP (VLAN Trunking Protocol)
- Greatly **simplifies network design and deployment**

Ethernet Frames

IEEE 802.1q (**Dot1q tagging**)

- **VLAN tag** is added into the original **Ethernet frames**
  - Between Source MAC and Type / Length header fields
  - TPID = Tag Protocol Identifier = `0x8100`
    - So that tagged frames can be distinguished from untagged ones
  - TCI = Tag Control Information
    - Contains 12-bit long **VLAN Identifier** (**VID**) $\approx$ **4094 VLANs**
    - Certain values are reserved (at least `0x000` and `0xFFF`)
- Adding and removing tags also involves recalculating the CRC

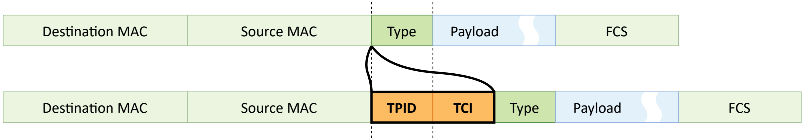

Routing Between VLANs

Observation

- **IP traffic between VLANs must normally go through routers**

Routing options

- **VLAN-unaware router** with **separate physical interfaces**
  - One separate port is needed for each VLAN on the router
  - They are all connected to different **access ports** on a switch
  - Obviously working, but not efficient enough and scales poorly
- **VLAN-aware router** with **sub-interfaces**
  - Physical interface is split up into multiple virtual sub-interfaces
    - Each corresponds to one particular VLAN
    - Frames outgoing from the router are tagged appropriately
  - Connected to a **trunk port** on a switch
- **VLAN-aware L3 switch**

Firewalls

Firewall

- General security system permitting to **monitor and control** both **incoming and outgoing** network traffic
  - Allows to **block unauthorized** / **allow authorized** access
    - So that users (their traffic) can only get where they are allowed
- Forms a **barrier** between a **trusted** and an **untrusted network**
  - I.e., between the inner (LAN) and outer (Internet) networks

Firewalls

Possible **deployments**

- **Network-based firewall**
  - Protects the whole inner corporate / school / home network
    - And so all its nodes / users
- **Host-based firewall** (individual, personal)
  - Protects just a single node / user

Possible **implementations**

- **Dedicated device** (combination of hardware and software)
- Purely **software solution**
- Set of **organizational measures**

Firewalls

Possible **strategies**

- **Prohibited unless permitted**
  - Everything is by default prohibited
  - Only something is explicitly permitted via **positive exceptions**
    - Having the nature of permissions
  - Approaches
    - Demilitarized Zones, Packet Filters
- **Permitted unless prohibited**
  - Everything is by default permitted
  - Only something is explicitly prohibited via **negative exceptions**
    - Having the nature of prohibitions
  - Approach
    - Packet Filters

Demilitarized Zones

Demilitarized Zone (**DMZ**) (Perimeter Network)

- Physical or logical **network acting as a barrier separating the inner and outer networks / zones**
  - Serves as kind of a gateway to the public Internet
    - Neither as secure as the inner zone,
      nor as insecure as the outer zone
  - Provides additional security especially from external attacks
- Permitted traffic
  - **Outer zone** $\leftrightarrow$ **inner zone**
    - This kind of communication is **entirely prohibited**
    - I.e., **no traffic can directly pass through DMZ**
  - **Outer zone** $\leftrightarrow$ **DMZ** and **DMZ** $\leftrightarrow$ **inner zone**
    - Possible in principle
    - But can also be partially restricted if need be

Demilitarized Zones

Demilitarized Zone (cont'd)

- Means of **implementation**
  - Simply via appropriate **configuration of routing tables** in both the routers separating the zones (i.e., at L3)
    - **Only traffic commencing / terminating in DMZ is allowed**
    - Which is detectable using **source / destination IP addresses**
- **DMZ contains…**
  - **Public servers** providing services to external users
    - E.g.: HTTP, SMTP, POP3, DNS, …
    - These are the hosts that are most **vulnerable to attacks**
    - And so when any of them gets compromised, inner zone is still likely to remain protected
  - Application Gateways
    - **Mediate** otherwise impossible **outer** $\leftrightarrow$ **inner communication**

Demilitarized Zones

Application Gateway (L7 Gateway, Application Proxy)

- **Server mediating communication with the outer zone**
  - E.g.: HTTP Proxy Gateway for requesting web pages, …
- Principle
  - (1) Inner node sends an intermediate request to the gateway
    - I.e., not directly to the intended target node
    - And so the sender must be aware of the gateway existence!
    - $\Rightarrow$ application **gateways are not transparent**
  - (2) Gateway then generates and sends its own request
  - (3) Response from the target node is received by the gateway
  - (4) It is then forwarded to the original node in the inner zone
- Observation
  - Gateways are always **application-dependent**
    - I.e., specifically designed for a given particular L7 protocol

DMZ Architectures

Dual Firewalls (**Back-to-Back DMZ**)

- **Two routers** (firewalls) are needed
  - Front-end (**perimeter**) between the outer zone and DMZ
  - Back-end (**internal**) between DMZ and the inner zone
- **Higher security**
  - Because two devices would need to be compromised at a time
    - Especially when devices from **different vendors** are used
    - Since it is not likely they would have the same vulnerabilities
- **Relatively costly** solution
  - And so suitable only for larger corporate networks

DMZ Architectures

Single Firewall (**Three-Legged DMZ**)

- **Only one router** (firewall) with (at least) **3 network interfaces**
- Represents a single point of failure
  - Since it must be able to handle all of the traffic

Integrated DMZ

- DMZ on a **software basis** without even a single router device
  - I.e., within a node directly separating the outer / inner zones

DMZ Host – not a true DMZ!

- Solution frequently appearing in **small home routers**
  - One server in the inner network can be specified
    - It then receives all unrecognized incoming traffic
  - This server is not isolated from the inner network at all
    - And so this solution has nothing to do with the DMZ concept

Packet Filters

Packet Filter

- Inspects and **filters** both **incoming and outgoing traffic** based on a set of configured **rules**
  - Works at L3
    - In terms of both blocking and permitting
    - In contrast, DMZ blocks at L3 and permits at L7
- Both **positive and negative** strategies are possible
  - Individual **rules** are described via **Access Control Lists**
- Available information
  - **Source / destination IP addresses** by default
  - But also information from higher layers
    - Such as **transport protocols** or **port numbers** at L4, …

Packet Filters

Modes of operation

- **Stateless Packet Inspection** (**Static Packet Filtering**)
  - Each packet is treated **independently** on each other
  - Easier to implement
  - Less computationally demanding
- **Stateful Packet Inspection** (**Dynamic Packet Filtering**)
  - Each packet is treated with regard to the recent **history**
    - I.e., also with respected to the previously handled packets
  - And so **more undesirable situations can be detected**
    - Especially various **concurrencies**
    - Can help to prevent DOS / DDOS attacks

Packet Filters

Access Control List (**ACL**)

- **List of rules** to be applied
  - Based on **positive permissions** or **negative exceptions**
- Standard **ACL**
  - Only **source IP address** is considered
  - Recommended deployment
    - Usually as **close to the target nodes** as possible
- Extended **ACL**
  - Other information is considered as well
    - **Destination IP address**, port number, …
  - Recommended deployment
    - Usually as **close to the source nodes** as possible

Lecture Conclusion

Broadcasts

- L2, local L3, targeted L3
- Broadcast domains

L3 **interconnection devices**

- Routers, L3 / L4 / L7 switches

VLANs

- VLAN-aware / VLAN-unaware segments
- Access (untagged) / trunk (tagged) ports
- Static / dynamic configuration

Firewalls

- **Demilitarized zones**, application gateways, **packet filters**