COSC 301
Network Management and Security

Lecture 21: Network Accounting & Visualization

Today's Focus

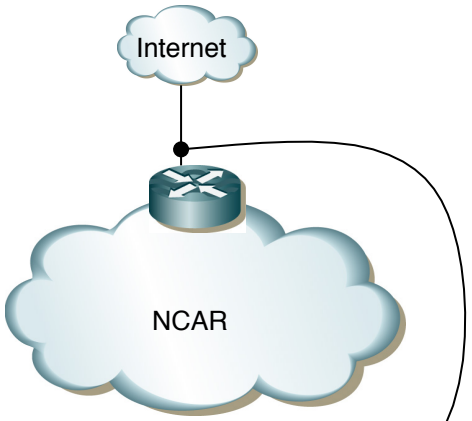

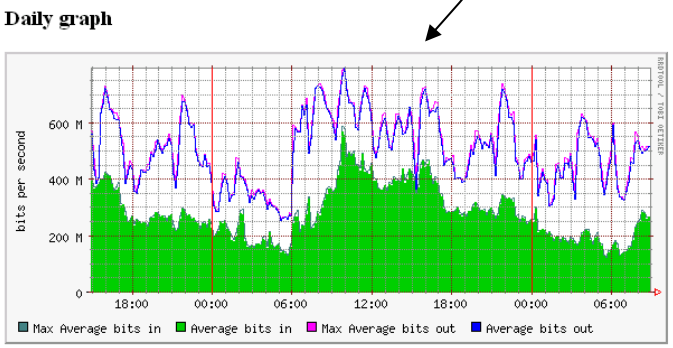

Network Accounting & Visualization

-- Why network accounting?
-- How to perform accounting?
-- Accounting tools

Why Account?

- Usage-based Billing
  - Charge groups/people for used bandwidth.
- Peering agreements
- Security analysis
  - Provide audit trail for connections, including src/dest addr, protocol, port, time, duration
- Network monitoring & anomaly detection
- Network/capacity planning
- Application monitoring and profiling
- User monitoring and profiling.

Where are we accounting?

- Datalink Layer
  - Bad
  - All Ethernet frames, including broadcast and management
- Network Layer
  - Common
  - includes traffic that may be unwanted.
  - Charges for IP headers too.
  - Makes protocols such as SSH very expensive.
- Application Layer
  - Common at proxies
  - fairest from users' point of view
  - does not charge for LAN/IP overhead traffic.

Caching and Charging

- If a user's request goes through a proxy, do they still get charged for cache hits?
  - Is it fair that the first requester gets charged if subsequent users do not?
  - Similar problems with multicast.
  - Are you charging for a data product (bytes), or a service (connectivity)?
- Charge provider or consumer?
- Consumers want predictable charging.

International/Domestic

- Commercial links may be charged at different rates for different types of traffic.
- How can we tell whether traffic is international or domestic?
  - Use a table of known national-IP ranges.
  - Hard to come by, no standard mechanism.
  - Processor / memory intensive.
- Best results comes from routing tables for national routers.

Getting the Data (1)

- Method 1: Use firewall counters
  - Put rules at the start of your firewall that match only (no ACCEPT or DROP).
  - Each rule has byte and packet counters.
  - What about traffic that would be dropped?
    Most useful for client-requested data.
  - Adds to latency.
  - Cannot acquire a post-capture breakdown of traffic.

Getting the Data (2)

- Method 2: Capture packet headers
  - Either listen on a router, or a switch's mirror port
  - Flexibility in processing of the packet headers
    - As in Method 1, there can be problems with respect to NAT. Do you get the packets pre/post NAT?
    - Again, don't know if packets get dropped.

Capturing Packets

- Modern (usually managed) switches have a mirror port, in which a copy of every frame that goes through the switch also gets forwarded out the mirror port.
- For optical networks, fibre splitters can be used.
- A traffic probe would be attached to the copied data.
- Unlike router methods, that can be useful for measuring link-local activity, although this is less useful for most accounting.

NetFlow

- Developed by Cisco originally.
- Primary accounting technology used in industry today.
  - IPFIX is IETF's standardisation of NetFlow
  - Different versions export different sorts of values.
  - Version 5 most common for IPv4.
  - Version 9 for IPv6.

- Use UDP as transport

Flow Concept in NetFlow

- A flow is a unidirectional sequence of packets between a given source and destination, defined by a 7-tuple key consisting of the following fields:
  - Source IP address
  - Destination IP address
  - Source Port
  - Destination Port
  - IP Protocol
  - Ingress interface
  - IP Type of Service

NetFlow Architecture (1)

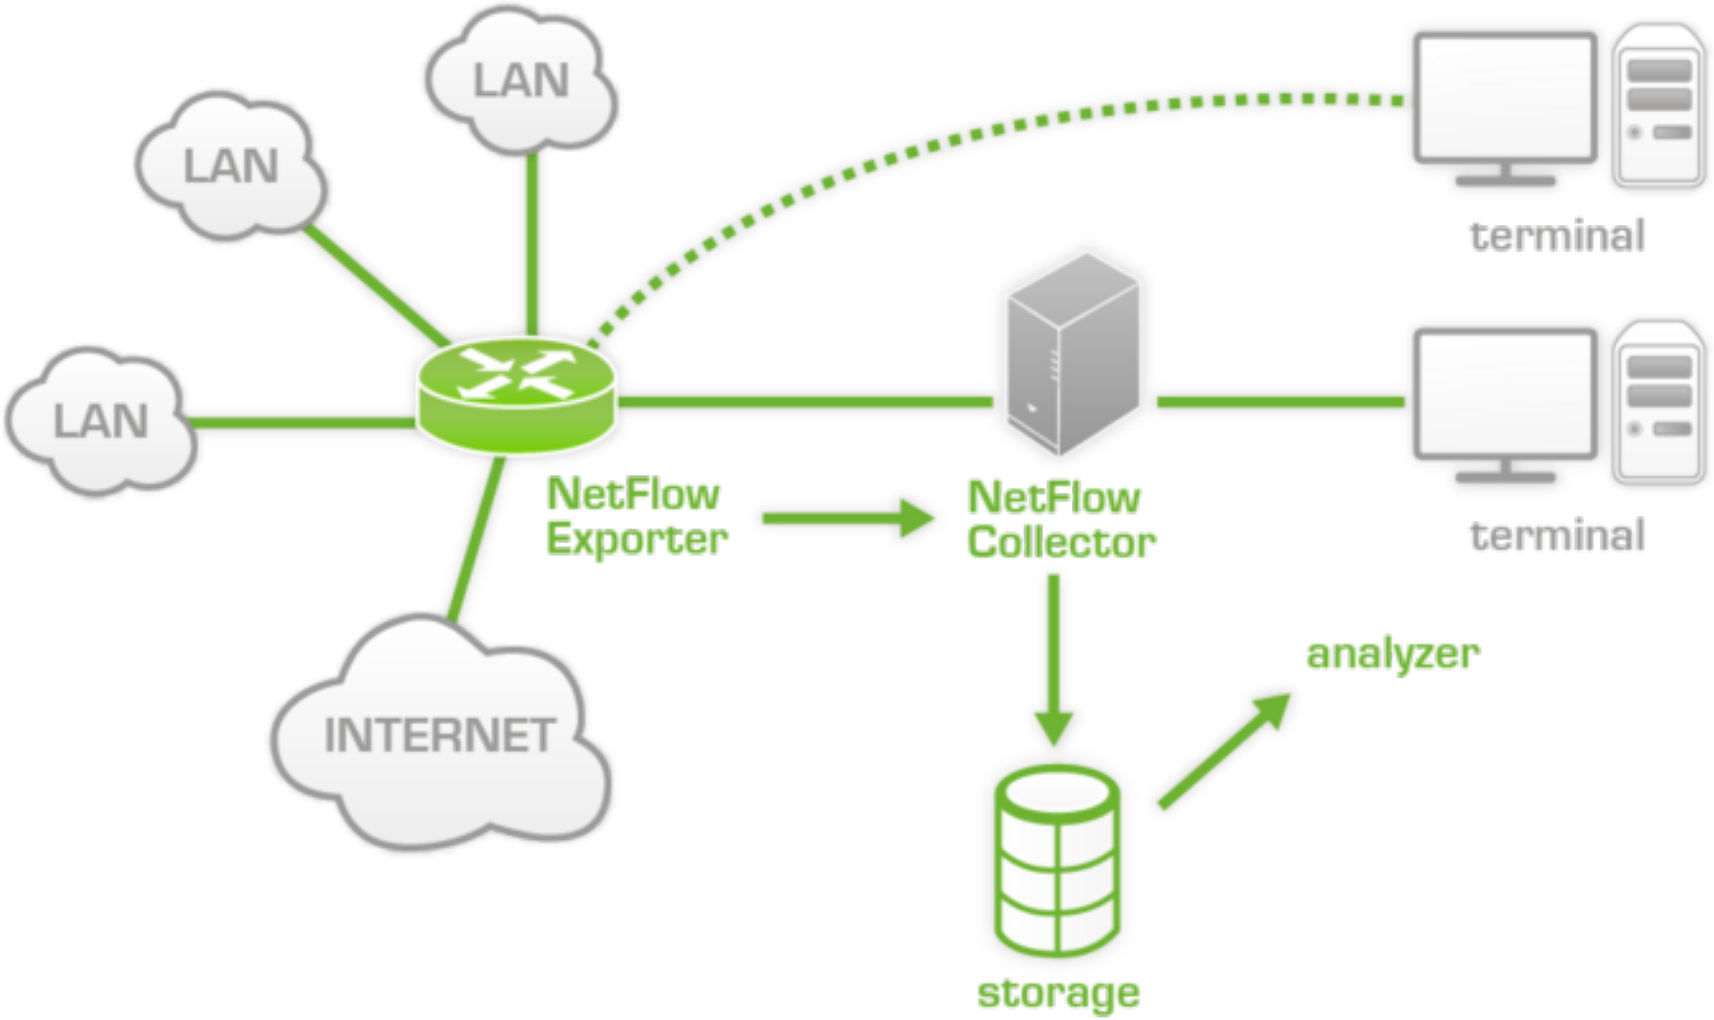

NetFlow Architecture (2)

- NetFlow Exporter
  - observes packet data and creates records from the monitored network traffic and transmits that data to the NetFlow collector.
- NetFlow Collector
  - collects the records sent from the exporter, stores them in a local database and forwards the records to an analyzer.
- NetFlow Analyzer
  - analyzes the NetFlow records for information of interest, which may include bandwidth usage, policy adherence, and forensic research.

NetFlow Records (1)

- The statistical information gathered from the network traffic is placed in a flow record.
- Each record is stored and managed in NetFlow cache
  - Once a flow has been created and placed in the cache, it remains active until it expires
  - After the flow expires, the record is added to a NetFlow Export datagram for transmission to the NetFlow collector

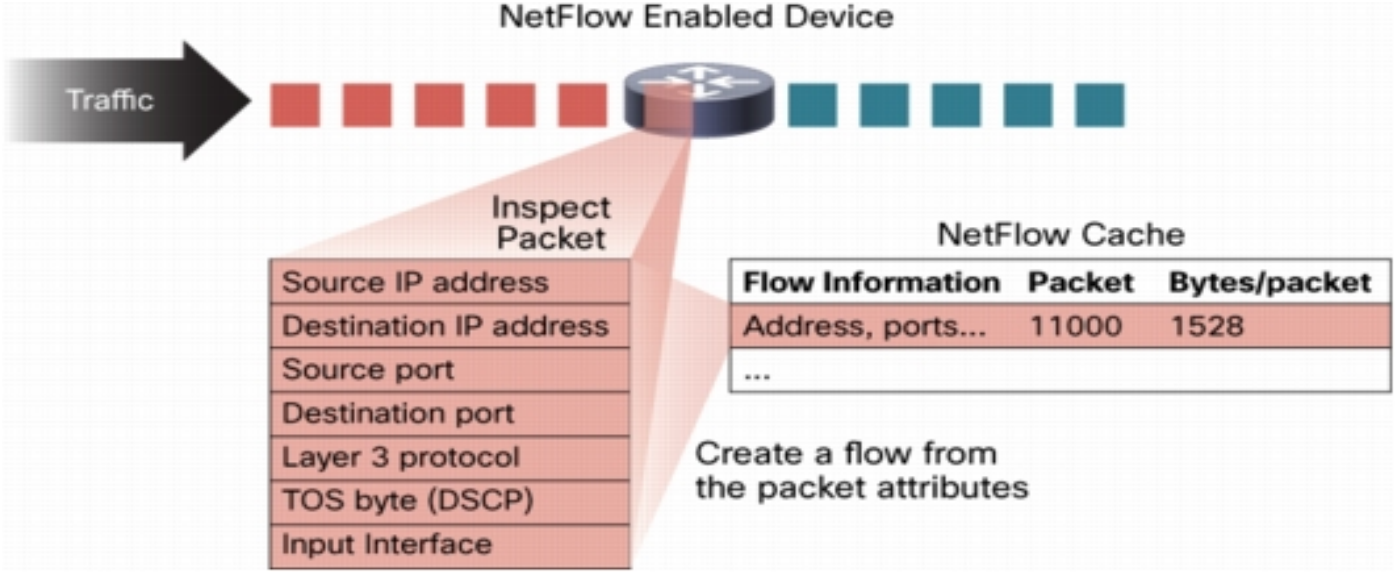

NetFlow Records (2)

- A NetFlow record may include many of or all of the following statistics:
  - NetFlow version
  - Flow Sequence (Identifier)
  - Input and output SNMP indices
  - Flow size in packets and bytes
  - Timestamp for flow start and stop times
  - Layer 3 header data (Source/Destination IP Addresses, IP protocol)
  - Port Numbers
  - Type of Service (ToS).
  - Layer 3 Routing information (IP address of the next-hop, Source and destination IP masks)
  - Multiprotocol Label Switching (MPLS) labels (version 9 only)
  - IPv6 addresses and ports (Netflow version 9 only)

Flow Tracking in NetFlow

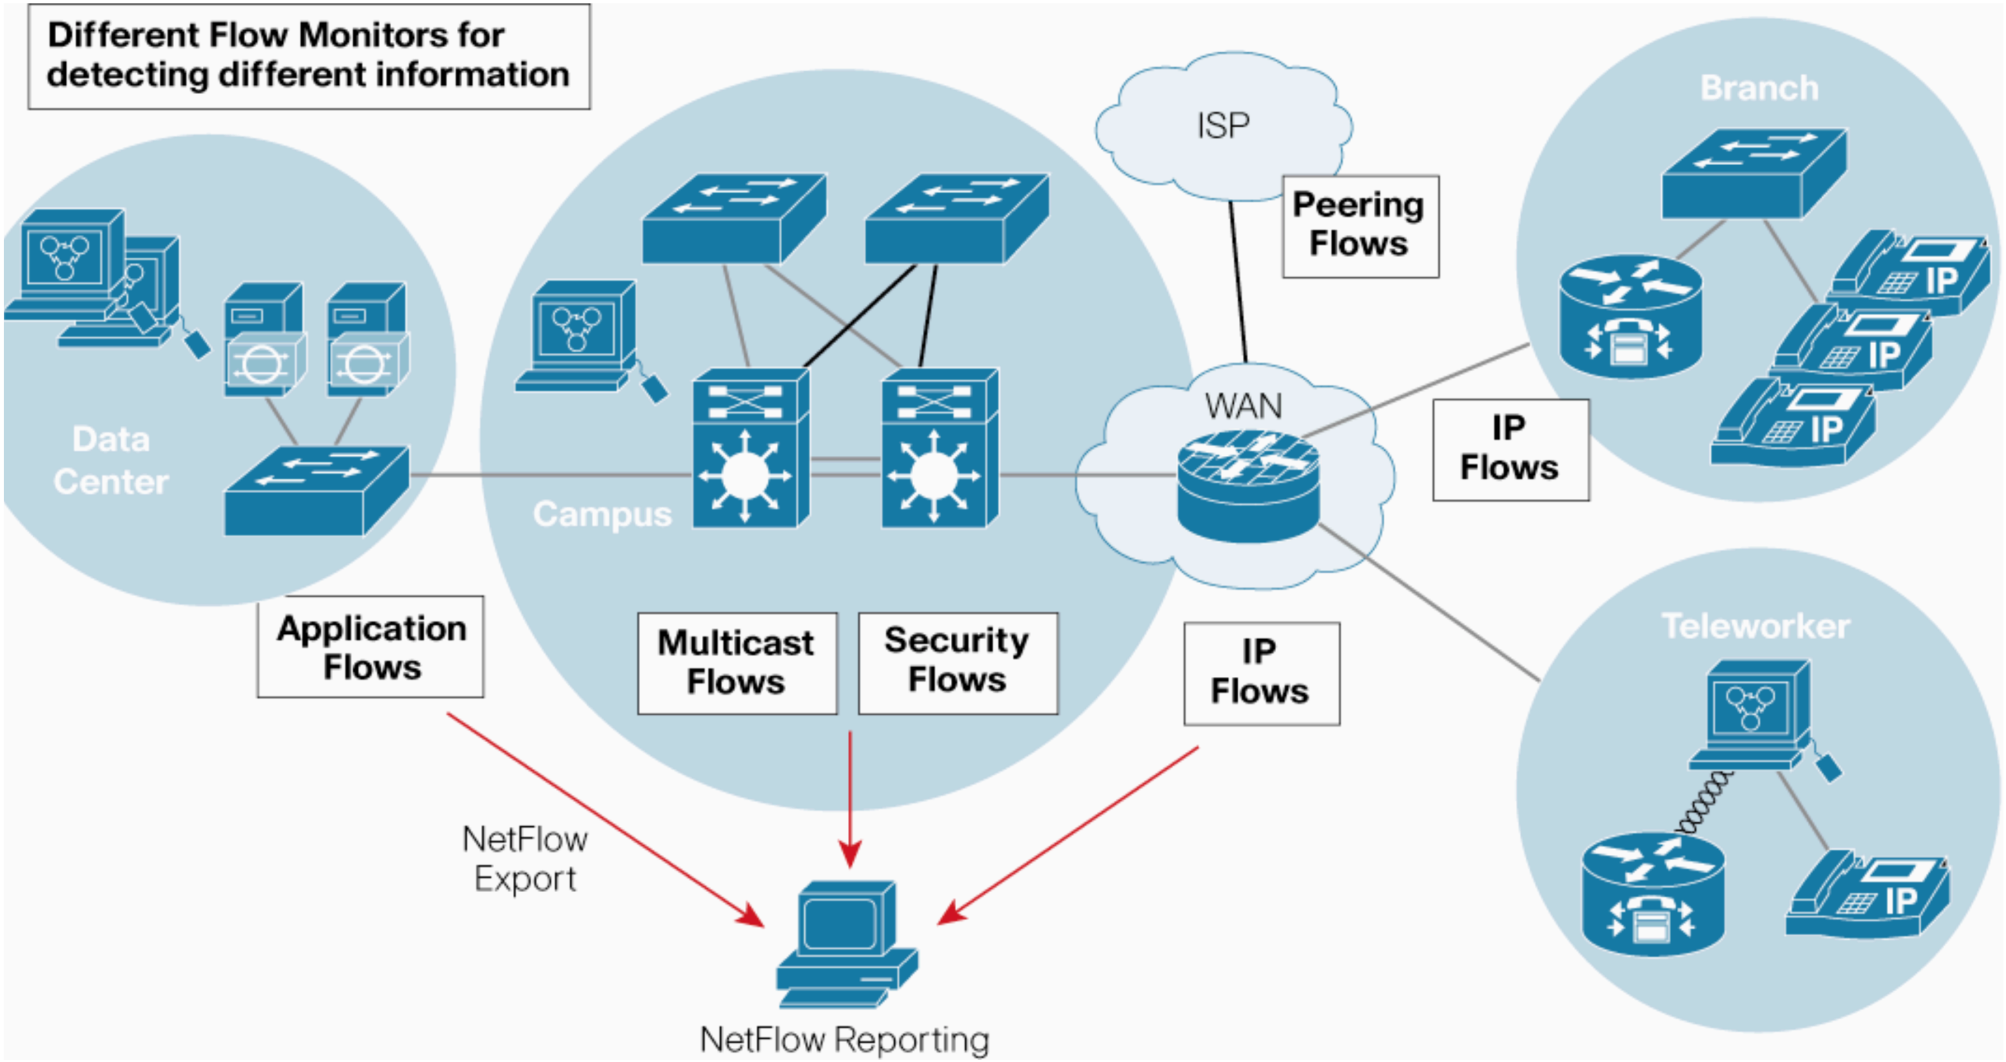

Summary

- What is network accounting?
- Tools: NetFlow