

Network Flows and Security

v1.01

Nicolas FISCHBACH

Senior Manager, Network Engineering Security, COLT Telecom
nico@securite.org - <http://www.securite.org/nico/>

Agenda

- The Enterprise Today
- Network Flows
- Netflow and NIDS
- Anomaly Detection
- Policy Violation Detection
- Peer-to-Peer
- Response and Forensics
- Conclusion

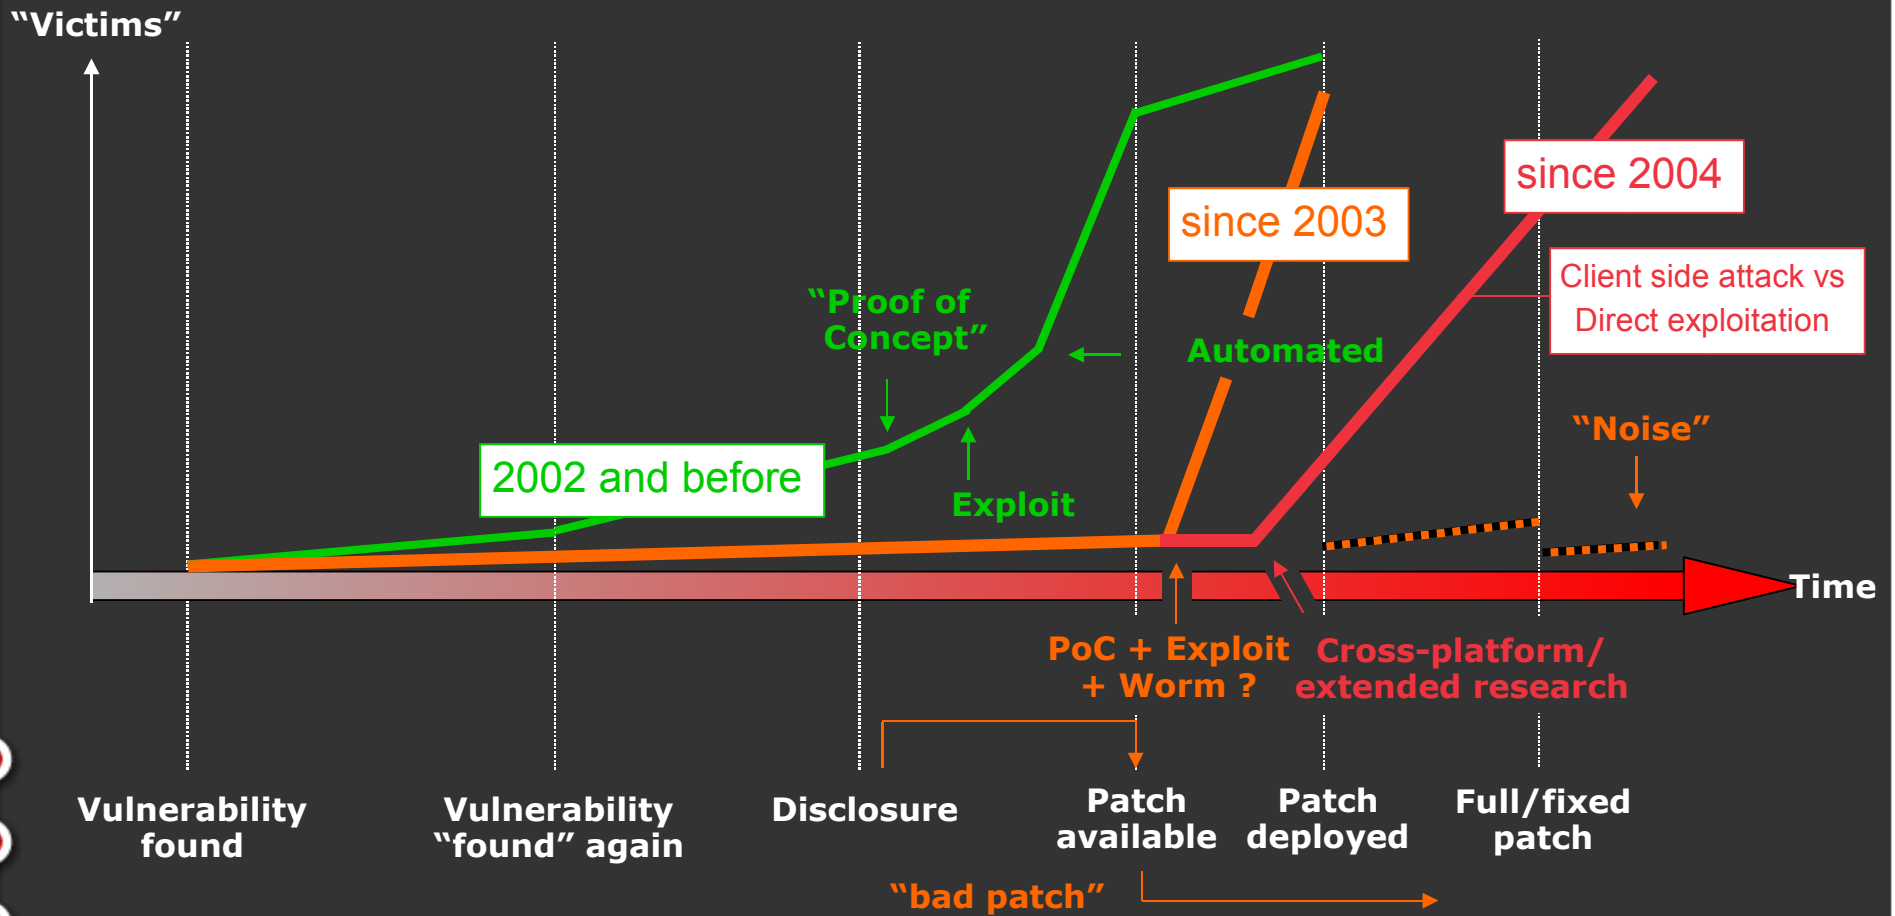
The Enterprise Today

- Where's my border ?
- WLANs, 3G devices, etc.
- Remote VPN/maintenance access: employees, partners, vendors and customers
- Client-side attacks
- Malware/spyware relying on covert channels
- Usually one "flat" undocumented network: no internal filtering, no dedicated clients/servers LANs, etc.
- More and more (wannabe) power users

The Enterprise Today

- Undocumented systems and applications
- Have you ever sniffed on a core switch's SPAN port ?
- Do you really need (expensive) NIDS to detect worms ?
- More and more communications are encrypted: SSH, SSL, IPsec, etc (even internally)

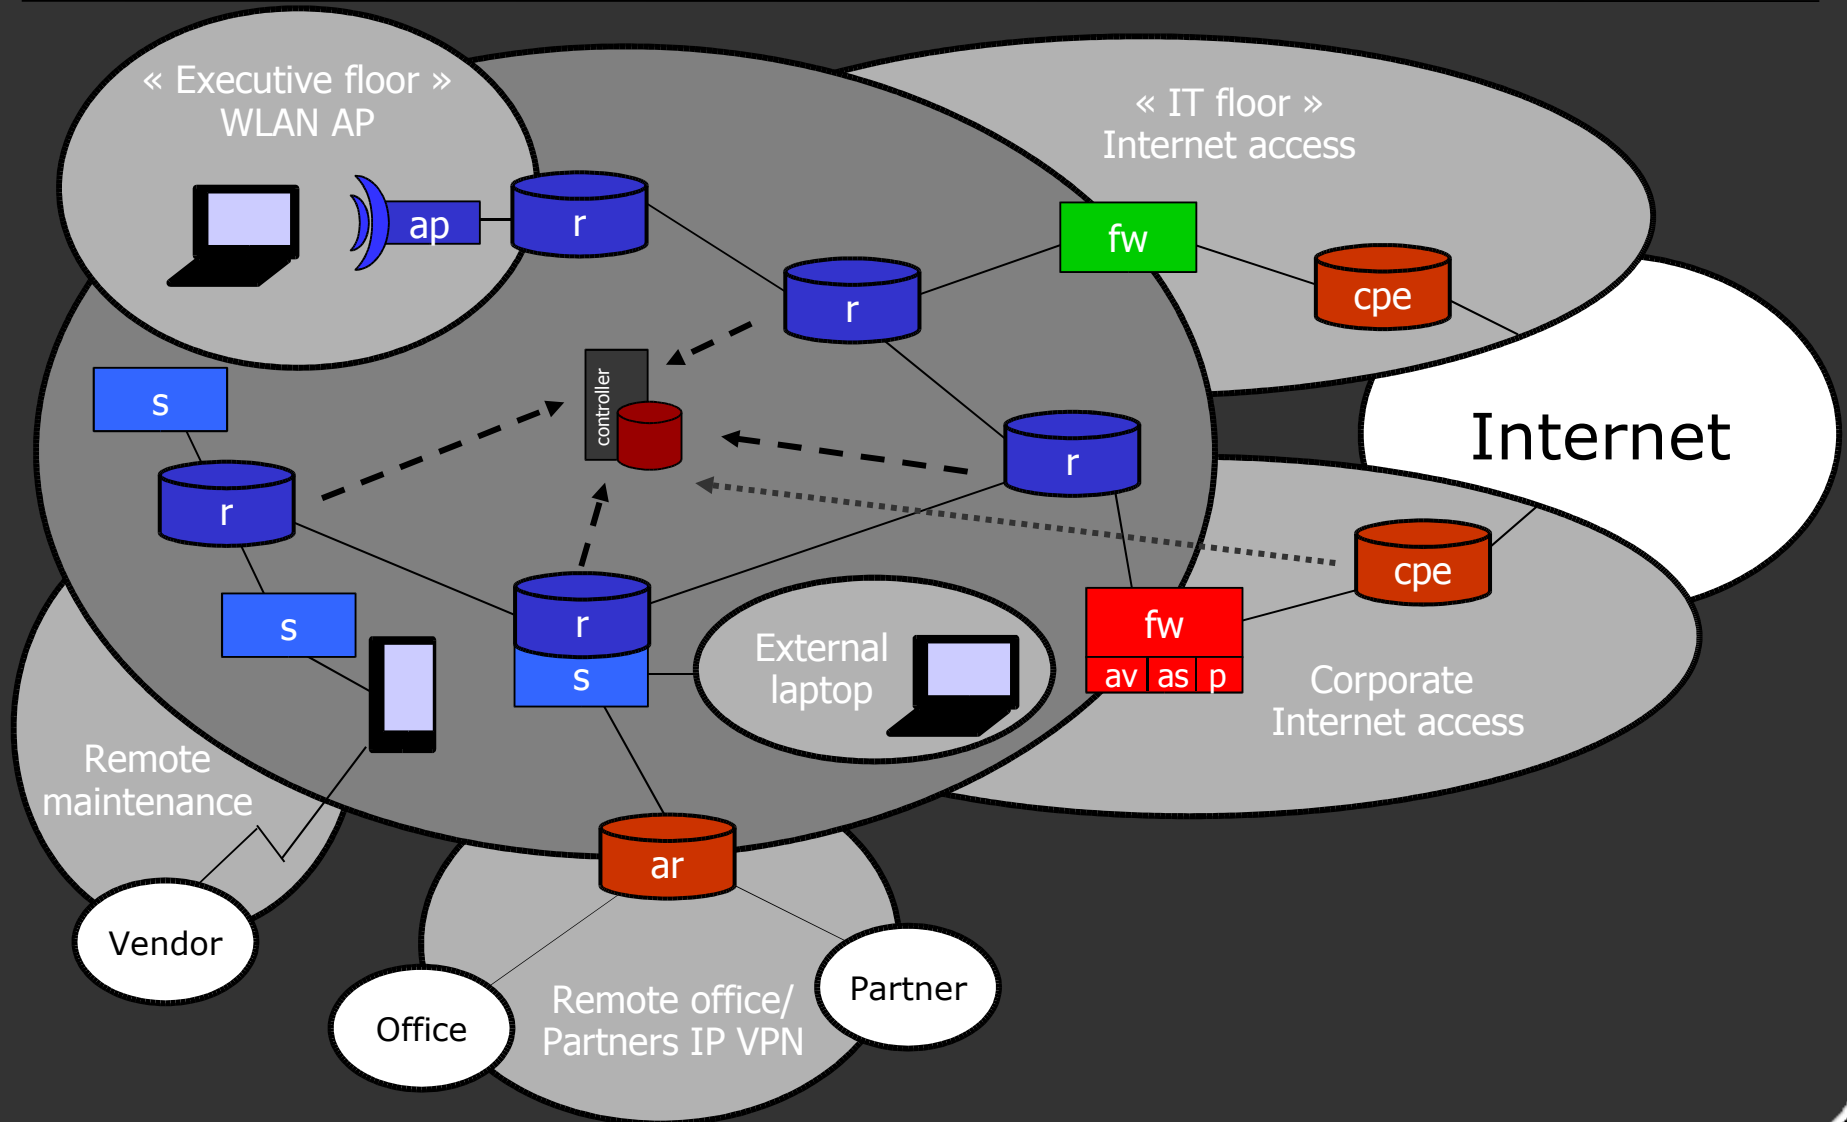
The Enterprise Today



Network Flows

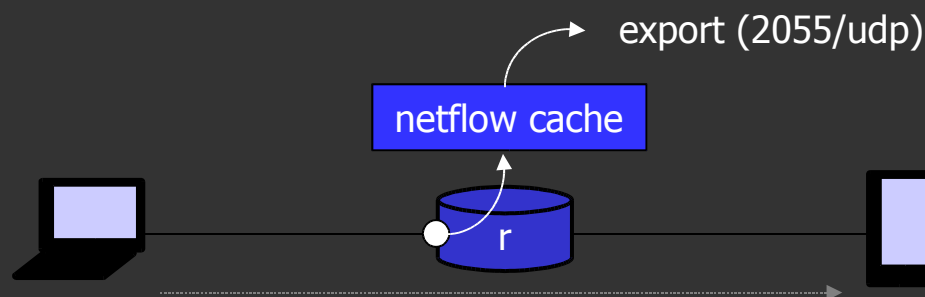
- What are network flows and why are they so interesting?
- Netflow (Cisco terminology) used to be a routing technology which became a traffic accounting solution
- Used since years by Service Providers to detect and traceback DDoS attacks and more recently for traffic engineering purposes
- In the enterprise network:
 - Network and application profiling, forensics, anomaly detection, policy violation, etc.
 - Netflow/NIDS: and/or ? Mix of macroscopic and microscopic views in high speed environments

The Connected Enterprise



Netflow

- A flow is a set of packets with common characteristics within a given time frame and a given direction
- The seven netflow keys:
 - Source and destination IP address
 - Source and destination port (code for ICMP)
 - Layer 3 protocol
 - Type of Service
 - Ingress interface (“one way”)



Netflow

- The following data are exported (Netflow v5)
 - The 7 key fields
 - Bytes and packets count
 - Start and end time
 - Egress interface and next-hop
 - TCP flags (except on some HW/SW combination on multilayer switches)
- And you may also see the AS number and other fields depending on version and configuration
- IPFIX is based on Netflow v9
- Egress Netflow and per class sampling in recent IOSes



Netflow

- The cache contains 64k entries (default)
- A flow expires:
 - After 15 seconds of inactivity (default)
 - After 30 minutes of activity (default)
 - When the RST or FIN flag is set
 - If the cache is full
- Counting issues: aggregation and duplicates (a flow may be counted by multiple routers and long lasting flows may be “duplicated” in the database)
- Security issues: clear text, no checksum, can be spoofed (UDP) and possible DoS (48 bytes per flow for a 32 bytes packet)

Netflow

- Sampling
 - By default, no sampling: each flow entry is exported
 - Sampled: percentage of flows only (deterministic)
 - Random Sampled: like sampled, but randomized (statistically better)
- “Full netflow” is supported on/by most of the HW/SW, sampled and random sampled only on a subset
- Sampling reduces load and export size but “losses” data:
 - OK: DDoS detection
 - NOK: Policy violation detection
- Avoid router-based aggregation

Netflow

- General configuration

```
router (config)# ip flow-export destination <serverIP> <port>  
router (config)# ip flow-export source loopback0  
router (config)# ip flow-export version 5
```

- Tuning

```
router (config)# ip flow-cache entries <1024-524288>  
router (config)# ip flow-cache timeout active <1-60>  
router (config)# ip flow-cache timeout inactive <10-600>
```

- Display the local cache

```
router# show ip cache flow
```



Netflow

- “Full”/unsampled

```
router (config)# interface x/y
router (config-if)# ip route-cache flow
```

- Sampled

```
router (config)# ip flow-sampling-mode packet-interval 100
router (config)# interface x/y
router (config-if)# ip route-cache flow sampled
```

- Random Sampled

```
router (config)# flow-sampler-map RSN
router (config-sampler)# mode random one-out-of 100
router (config)# interface x/y
router (config-if)# flow-sampler RSN
```

Netflow/NIDS

- Netflow is “header” only
 - Distributed and the network “speed” only has indirect impact
 - Often the header tells you enough: encrypted e-mails with the subject in clear text or who’s mailing whom =)
- NIDS may provide full packet dump
 - Centralized and performance linked to the network “speed”
 - Full dump or signature based dumps ?
 - PCAP-to-Netflow
 - May tell you the whole story (disk space requirements)

Netflow/NIDS

- Let's mix both: distributed routers sourcing Netflow and NIDS/sniffers in key locations!
- Decide how to configure your NIDS/sniffers:
 - PCAP-type packet sniffers
 - Standard ruleset
 - Very reduced and specific ruleset
 - How much data can you store and for how long ?
- Investigate ways of linking both solutions
- Storage (the older the less granular ?)
 - Flat files
 - Database

Anomaly Detection

- Discover your network
 - Enabling netflow will give you some insight on what your network actually carries :)
 - After the shock and the first clean up round:
 - Sniff traffic in specific locations
 - Introduce security driven network segmentation
 - Build a complete baseline
 - Update your network diagram

Anomaly Detection

- Distributed Denial of Service
 - Fairly easy to spot: massive increase of flows towards a destination (IP/port)
 - Depending on your environment the delta may be so large that you don't even require a baseline
 - You may also see some backscatter, even on an internal network
- Trojan horses
 - Well known or unexpected server ports (unless session re-use)
- Firewall policy validation
 - Unexpected inside/outside flow

Anomaly Detection

- Worms
 - Old ones are easy to spot: they wildly scan the same /8, /16 or /24 or easy to code discovery pattern
 - New ones are looking for specific ports
 - Each variant may have a specific payload size
 - May scan BOGON space
 - The payload may be downloaded from specific, AV identified, websites
 - The source address is spoofed (but that's less and less the case)

Anomaly Detection

- Covert channels / Tunnels
 - Long flows while short ones are expected (lookups)
 - Symmetric vs asymmetric traffic (web surfing)
 - Large payloads instead of small ones
 - Think ICMP, DNS, HTTP(s)
- Scans
 - Slow: single flows (bottomN)
 - Issue with bottomN: long tail
 - Normal/Fast: large sum of small flows from and/or to an IP
 - Return packets (RST for TCP and ICMP Port Unreachable for UDP)

Policy Violation Detection

- Workstation / server behaviour
 - Usually very “static” client/server communications
 - Who initiates the communication and to which destination ?
 - Office hours
 - New source/destination IPs/ports showing up
 - Tracking using DHCP logs, MAC address, physical switch port (SNMP)
 - Identify the “early” flows (auto-update and spyware)
 - After DHCP allocation or after login
 - Flows after the initial communication
 - Recurring flows (keyloggers) or flows towards the same destination but using various protocols (firewall piercing)

Peer to Peer (P2P)

- Legacy P2P protocols often use fixed ports or ranges
- Sometimes (like with FTP) the data port is the control port +/-1
- Recent P2P protocols have the session details in the payload: they can't be tracked using netflow but the flow size may give you a hint

Response

- Locate the source host
 - Requires the “netflow source” information (which router saw that flow)
 - Layer 3 and Layer 2 trace: identify the last layer 3 hop and then layer 2 trace or use previously SNMP polled MAC/port address
- Block the host
 - Port shutdown
 - ACLs
 - Blackhole route injection

Forensics

- Netflow and dumps storage need to be resolved first
- Clear post-mortem process
- Usual approach is to look for the flows and once identified extract the relevant dumps/logs
- In some environments only a couple of minutes/hours may be stored
- Legal/privacy issues
- Out-of-band network to push data and avoid multi-accounting

Tools

- argus (<http://www.qosient.com/argus/>)
- nfdump (<http://nfdump.sourceforge.net>) with nfsen (<http://nfsen.sourceforge.net/>)
- graphviz (<http://www.graphviz.org/>): human eye is good at catching things, but the graphs become really complex
- ntop (<http://www.ntop.org/>)
- Comprehensive list:
<http://www.switch.ch/tf-tant/floma/software.html>
- Commercial products

Conclusion

- Netflow: macroscopic view
- NIDS/sniffer: microscopic view
- Network switches: layer 0/1 view (MAC address/port)
- Mix them while controlling
 - CAPEX/OPEX
 - Storage
 - Search/detection capabilities
 - Avoid impact on the network
- Active response (quarantine/active defense) ?
- Q&A