



MaFIX: Using IPFIX for Scaling Threat Detection to High Data Rates

Gabriel Paradzik^{1,2}, Benjamin Steinert^{1,2}, Janik Steegmüller, Michael Menth¹

¹Chair of Communication Networks, ²Zentrum für Datenverarbeitung
University of Tübingen

<https://kn.inf.uni-tuebingen.de>





- ▶ Motivation
- ▶ Technical Background
- ▶ MalFIX Architecture and Implementation
- ▶ Performance Evaluation
- ▶ Conclusion



- ▶ Threat intelligence (TI) feeds provide information about indicators of compromise (IoC)
 - IoCs can be IP addresses, hostnames, signatures, etc.
 - Maintained by private companies, other network operators, or open-source projects
 - Examples: abuse.ch, AbuseIPDB

- ▶ Blocking all malicious IP addresses is unfeasible because of the large amount
 - Firewalls have limited amount of rules

- ▶ TI information can be used to identify bad actors on the network


- ▶ For networks with high volume, scanning every packet is not possible
 - Switching to flow-based scanning with IPFIX



- ▶ IPFIX protocol aggregates packets into flows
 - Flow represents communication between two endpoints

- ▶ IPFIX flow record consists of multiple Information Elements (IEs)
 - IE represents certain type data point
 - Packet payload is usually discarded

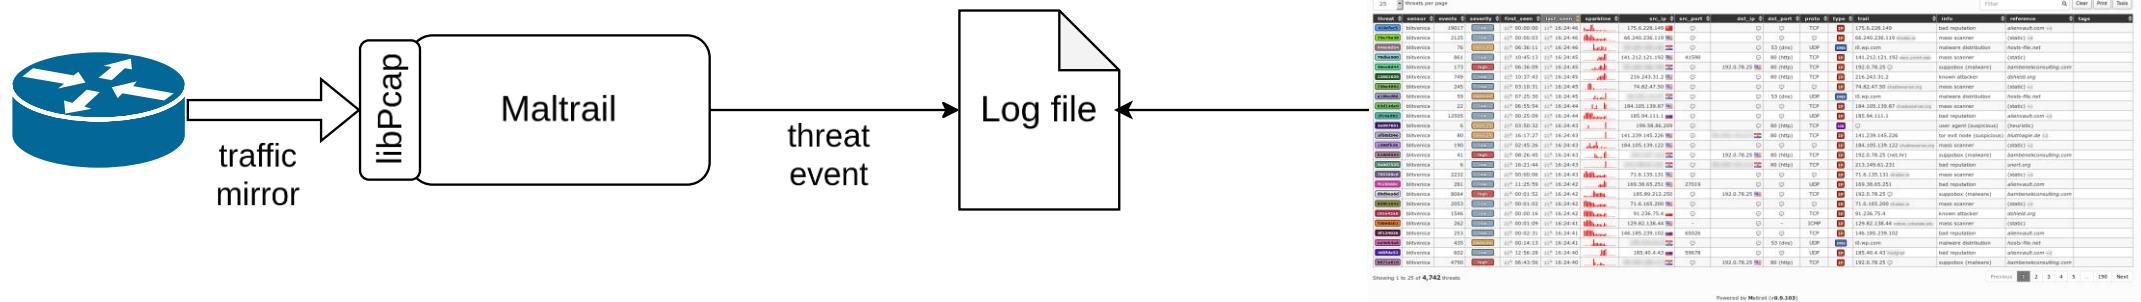
- ▶ IPFIX standard allows including arbitrary data via custom IEs
 - E.g., OS/application fingerprinting, observed TCP flags



Example flow record	
flowStart	2025-03-10 14:33:25.133
flowEnd	2025-03-10 14:33:29.021
sourceIP	1.2.3.4
destIP	6.7.8.9
srcPort	44276
destPort	443
protocol	TCP
octetCount	6345
packetCount	7
tcpFlagsUnion	SYN,ACK,FIN
flowEndReason	FIN
appLabel	HTTPS



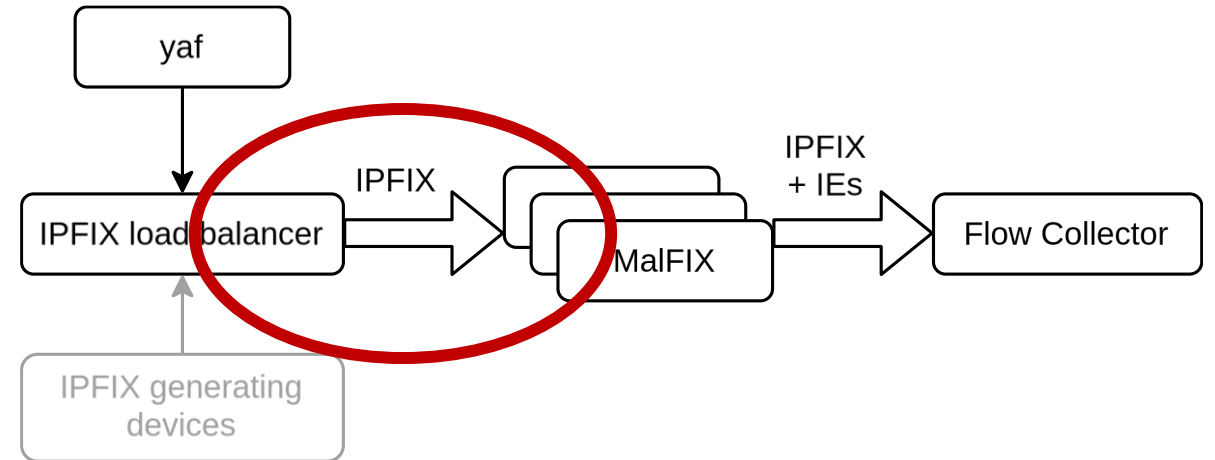
- ▶ Maltrail is an open-source all-in-one threat detection system written in Python
 - Actively maintained on GitHub
 - Utilizes a large number of TI feeds and static threat indicators



- Perfectly suited for small networks, but not performant enough for large networks with high traffic volumes
- ▶ Can we leverage Maltrail's up-to-date threat detection engine and use it for monitoring high traffic volumes?



- ▶ Maltrail was modified (“MaIFIX”) to allow high-performance threat monitoring
 - Changes are minimally invasive to allow easy merging with upstream
 - Input/Output capabilities were modified



▶ Input Adaptations

- Instead of raw packet captures, IPFIX is accepted
- Yaf generates IPFIX from traffic on an interface
 - High performance capturing library PF_RING™
- Run multiple instances of MaIFIX by employing IPFIX load balancer

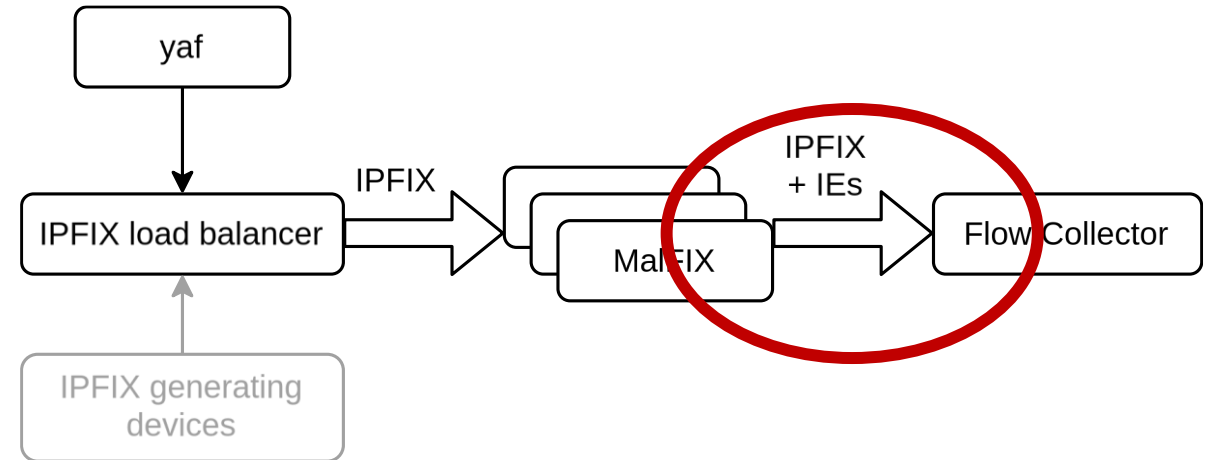


► Output Adaptations

- Use IPFIX custom IEs
- Detected threat information are attached via custom IEs
- Allows for subsequent processing with IPFIX-compatible tools

► MaIFIX fully IPFIX standard compatible

- Can be integrated into IPFIX pipelines






► Pipeline Mode

- All incoming flows to MaIFIX are exported
- Custom IEs are attached to malicious flows
- Useful for data enrichment scenarios



► Alert-Only Mode

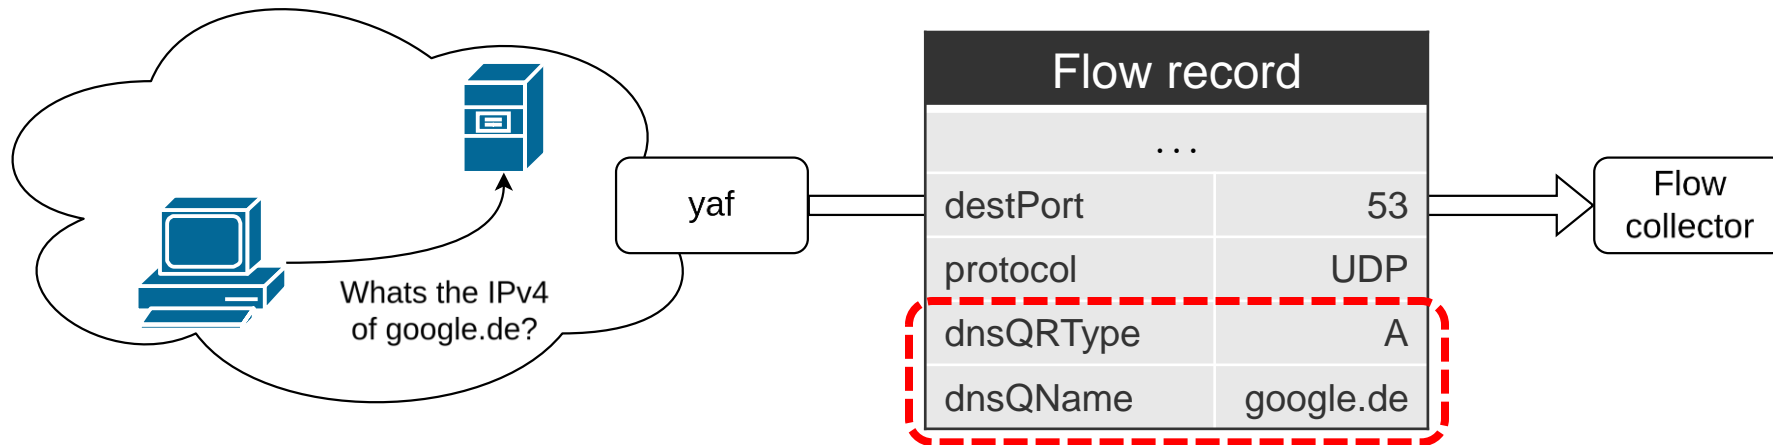
- Only malicious flows are exported
- Non malicious flows are dropped
- Useful for alerting





- ▶ By switching from packets to flows, we lose payload information
 - Payload information is lost in typical IPFIX setup

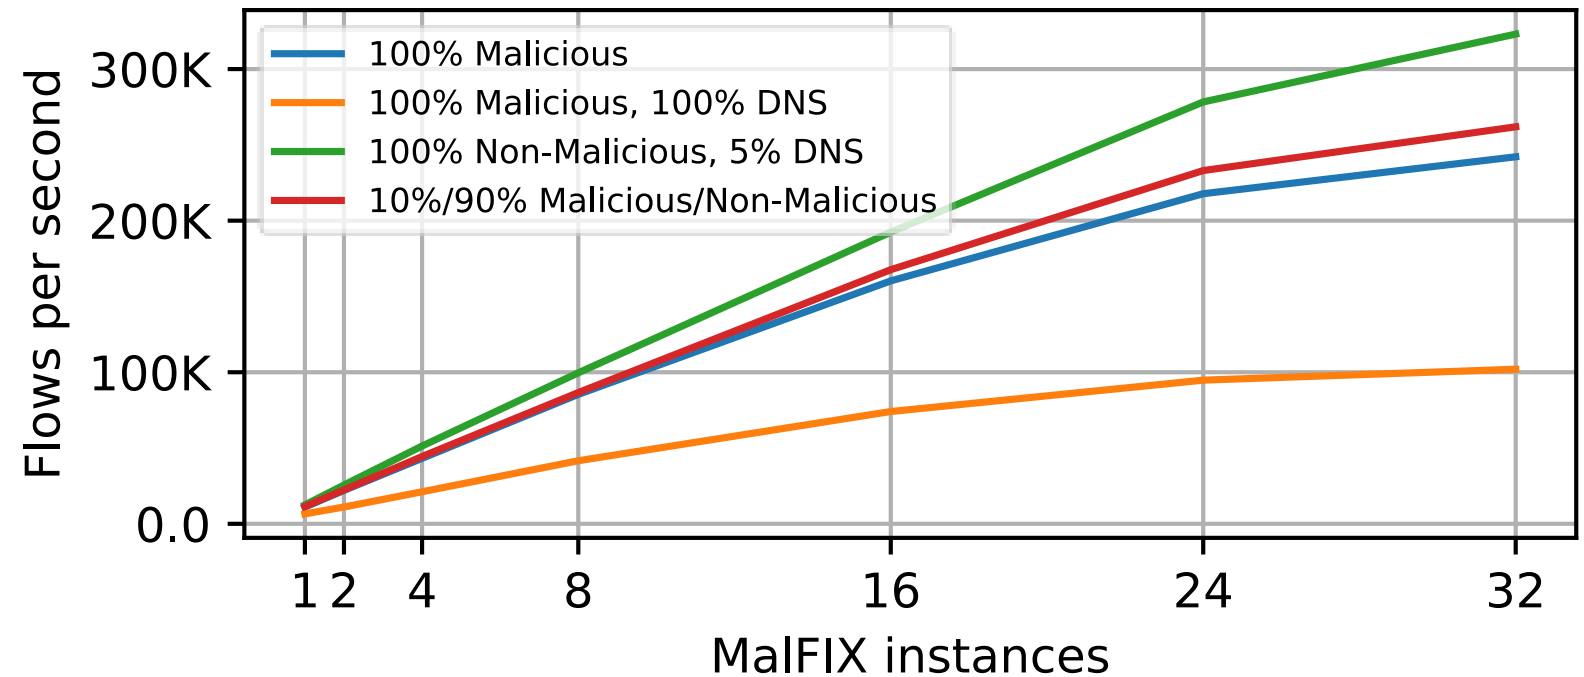
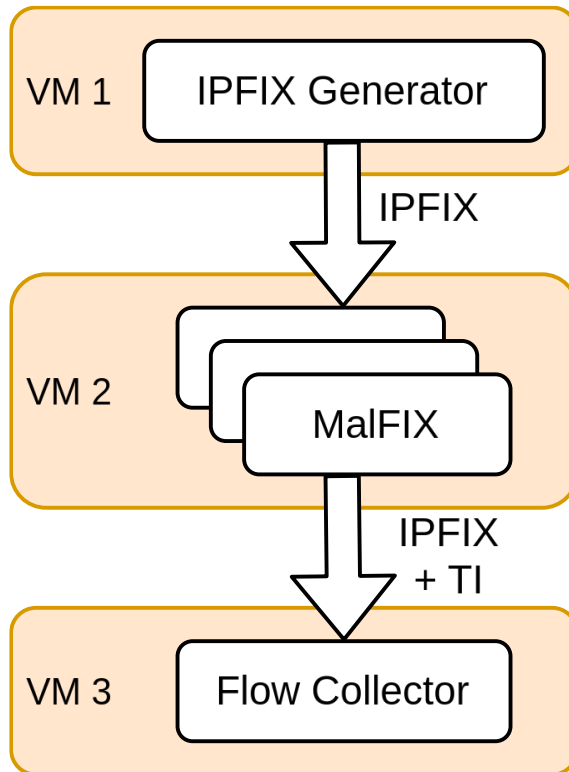
- ▶ Yaf has Deep Packet Inspection (DPI) capabilities
 - Search for payload information (DNS, HTTP, FTP, etc.)
 - Include results in custom IEs




- ▶ MalFIX also reads yaf's DNS DPI information
 - Domain names are checked with Maltrails internal threat detection engine



- ▶ Maximum flow processing speed was evaluated for Alert Only Mode
- ▶ Number of running MalFIX instances was varied
- ▶ Different traffic patterns were used





- ▶ Open-Source tool Maltrail was modified to fit a high-performance threat detection pipeline
 - By using standard conform IPFIX, MalFIX can be integrated with other data sources/sinks
 - MalFIX is Open-source:  <https://github.com/uni-tue-kn/MalFIX>

- ▶ Up to **300,000 flows/second** on 32 CPU cores can be scanned for threats
 - MalFIX can also be deployed across multiple machines

- ▶ MalFIX is deployed at the computation center of the University of Tübingen (ZDV)



Gabriel Paradzik^{1,2}, Benjamin Steinert^{1,2}, Janik Steegmüller, Michael Menth¹

¹Chair of Communication Networks, ²Zentrum für Datenverarbeitung
University of Tübingen

gabriel.paradzik@uni-tuebingen.de

THANK YOU!

QUESTIONS?

This work was supported by
the bwNET2.0 project funded
by the MWK

