Stream – Oriented Network Packet Capture

Hot topics in Security Research – the Red Book

Evangelos Markatos
FORTH-ICS

Two-part presentation

1. Stream oriented packet capture
   - For high-speed networks
     - appeared in IMC 2013 and JSAC 2014
2. Research Directions in Cyber Security
   - The Red Book experience

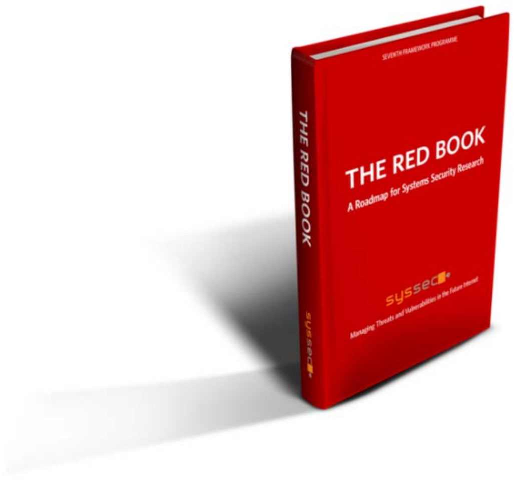

Scap: Stream-Oriented Network Traffic Capture and Analysis for High-Speed Networks

Antonis Papadogiannakis — *FORTH-ICS, Greece*

Michalis Polychronakis — *Columbia University, USA*

Evangelos P. Markatos — *FORTH-ICS, Greece*

Network Traffic Monitoring Systems

- Recent security applications depend on network traffic inspection
  - Network-level intrusion detection
  - Network traffic classification
  - Next-generation firewalls

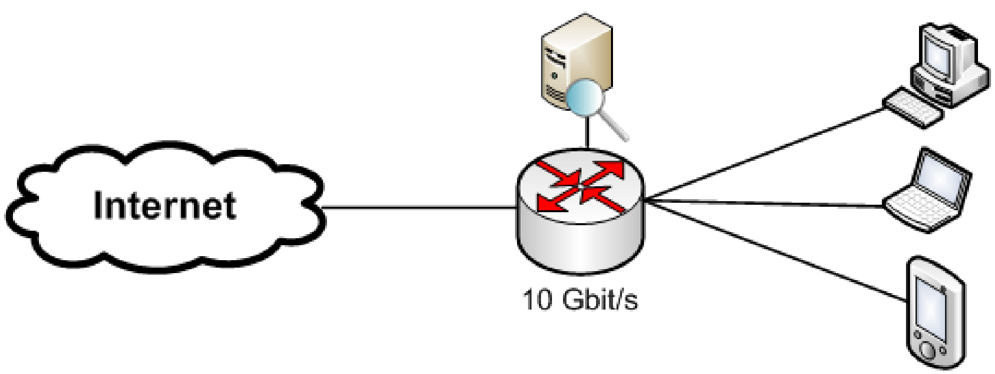

How do they do it?

- They capture IP packets
  - tcpdump/libpcap
  - PF_RING
  - PFQ
  - netmap
  - PacketShader I/O

- And they deliver **IP packets** to network monitoring applications

In this work we argue that
Packets are the wrong abstraction

- Connection-oriented analysis at L4 and beyond
- Attacks spanning multiple packets
- Protocol normalization to avoid evasion attempts

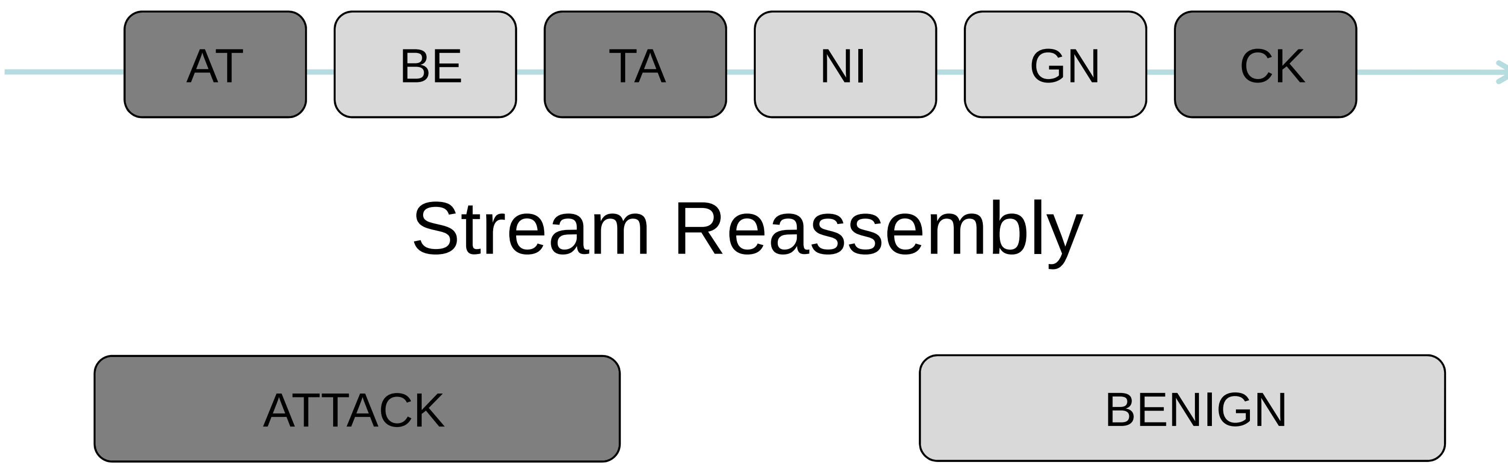

Stream Reassembly Libraries

- We are not the first to propose stream reassembly:
  - Libnids
  - Stream5
  - Bro's TCP stream reassembly
  - Custom stream reassembly implementations
- All such approaches are:
  - Implemented on top of packet capture libraries
  - Operate at **user-level**
  - Buffer packets or **copy** segments

Do we need another library?

Yes, to bridge a semantic gap

- *Monitoring applications* interested in analyzing traffic at **higher layers**
  - TCP flows, HTTP headers, SQL arguments, email messages
- *Traffic capture libraries* provide **raw packets**
  - Interleaved flows, out-of-order, duplicate, overlapping, uninteresting packets

Yes, to enable aggressive optimizations

- **Truncate** streams to handle large traffic volume
- Early discarding of **uninteresting** traffic
- Assign **priorities** to streams
- **Overload** control
- Utilize **multi-core** systems for stream processing

The Stream Capture Library (Scap)

- Captures and processes L4 **reassembled** streams
- Uses the abstraction of the **stream** (not packet)
- **Multi-core** support for stream capturing and stream processing
- Truncates streams with **subzero copy** at kernel or NIC
- Assigns **priorities** to streams and tolerate **overloads**

- Implements optimizations at the appropriate level (user level, kernel level, NIC)

Scap compared to other frameworks

*Support for multiple CPU cores and advanced NIC features (RSS, filters)

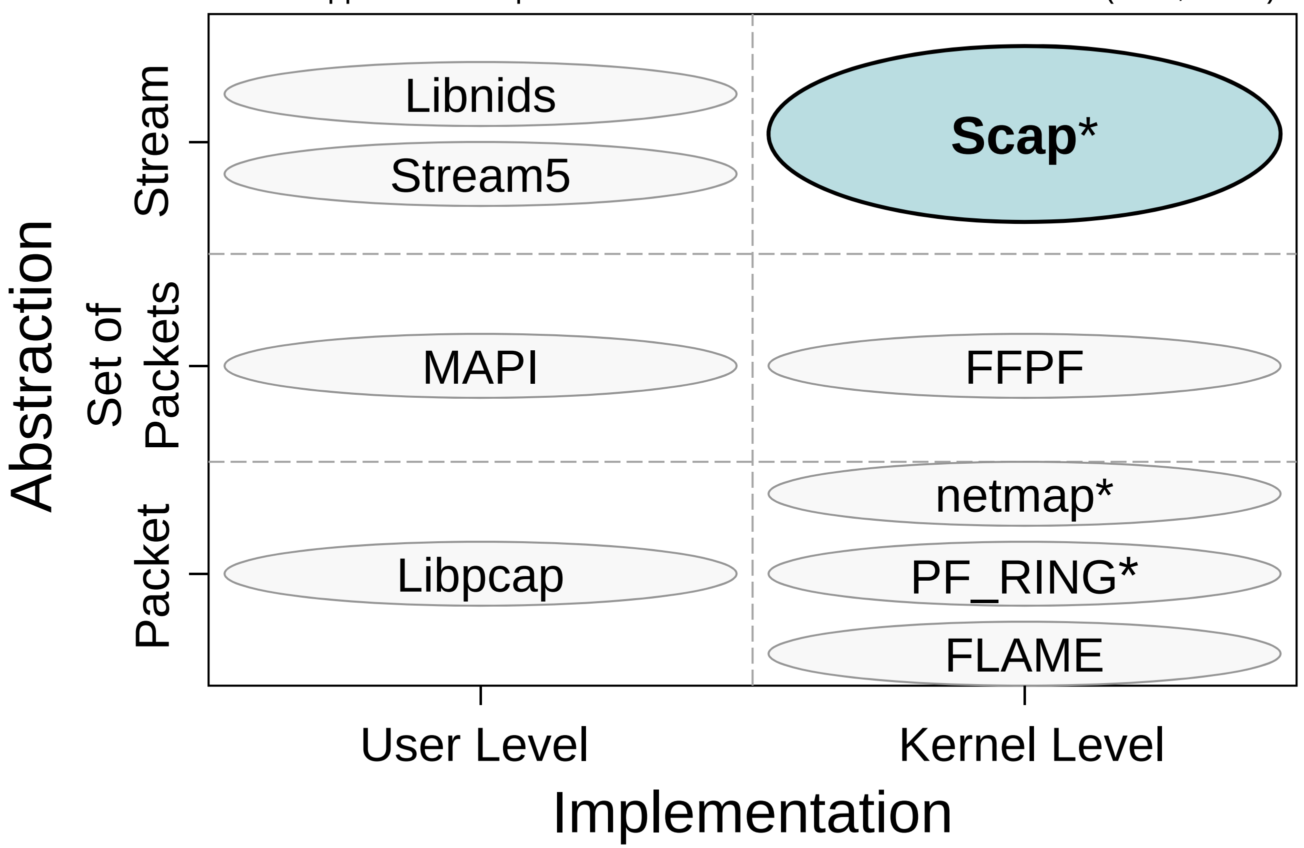

Stream Capture

- Store reassembled and normalized data segments to **stream-specific** memory locations
- Create events for stream creation, stream termination, data availability

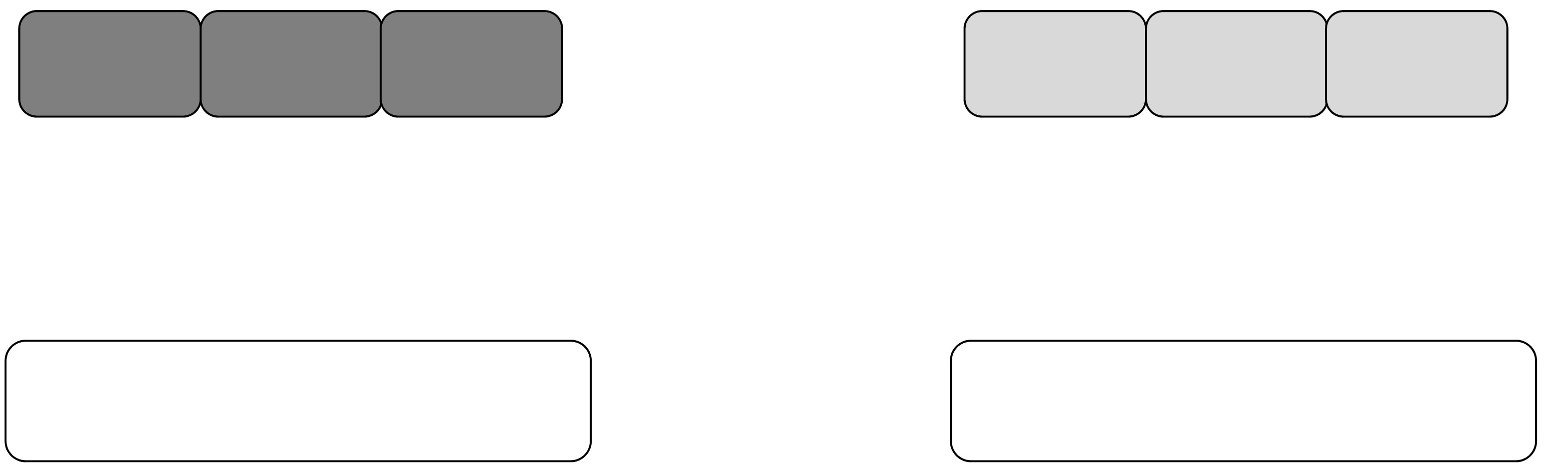

memory buffer for stream A   memory buffer for stream B

Stream Cutoff

- Analyze the **first N bytes** of each stream
- Cut the **long tails** if not interesting
- Use Scap and set **cutoff** size per each **stream**

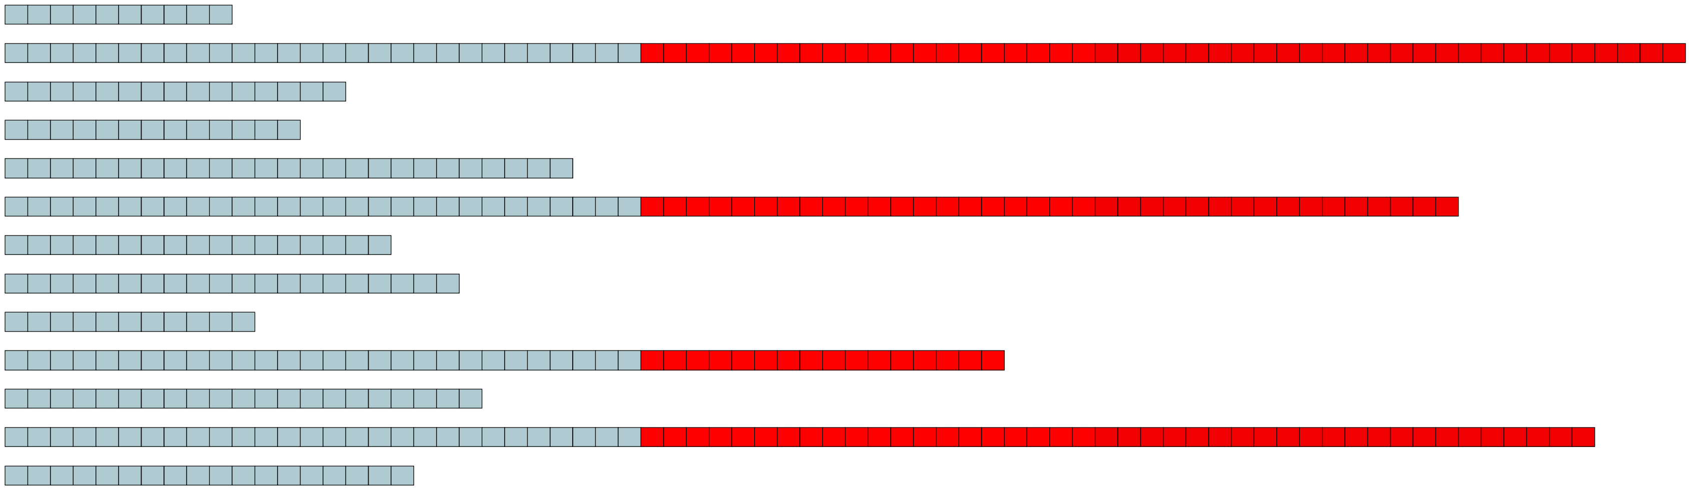

Subzero Packet Copy

- Early discarding of uninteresting traffic at the NIC
  - Utilize **NIC filters**
  - **No CPU cycles** spent for uninteresting packets
  - Receive TCP FIN/RST and use timeout for **stream termination** and accounting

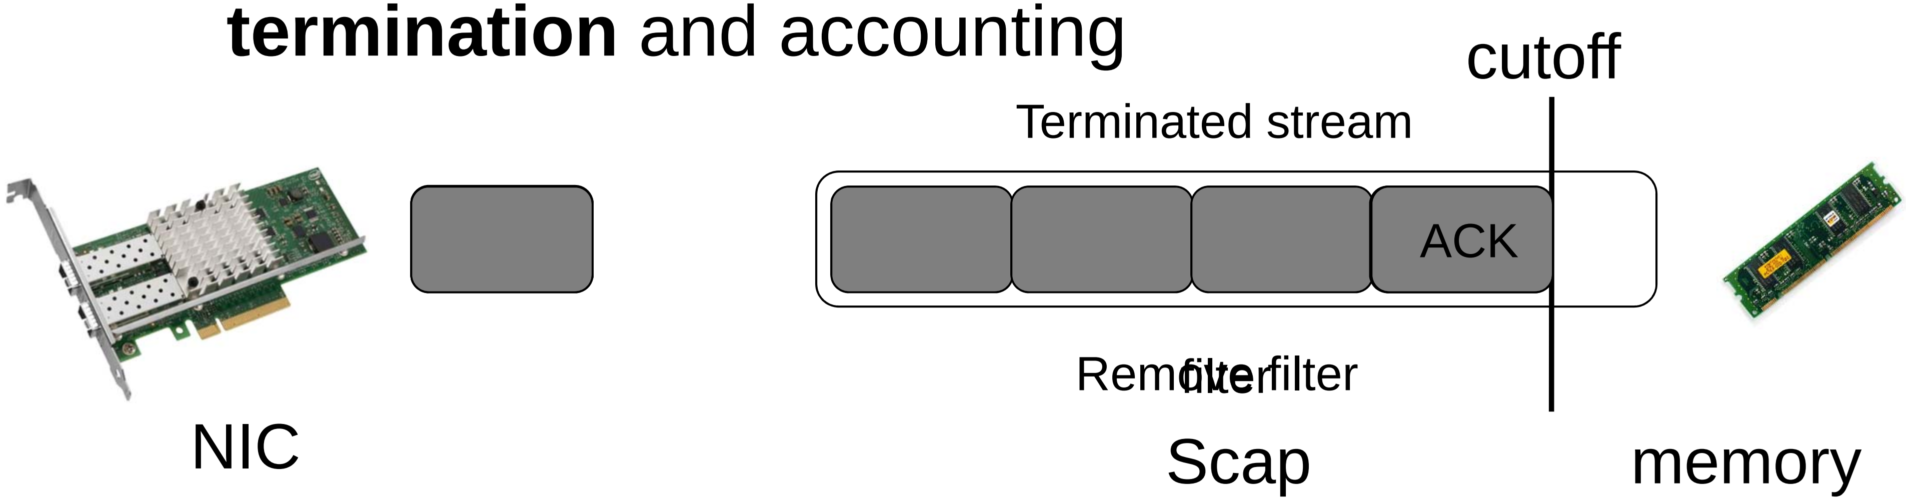

Prioritized Packet Loss

- In case of **overload** accommodate **high-priority** streams

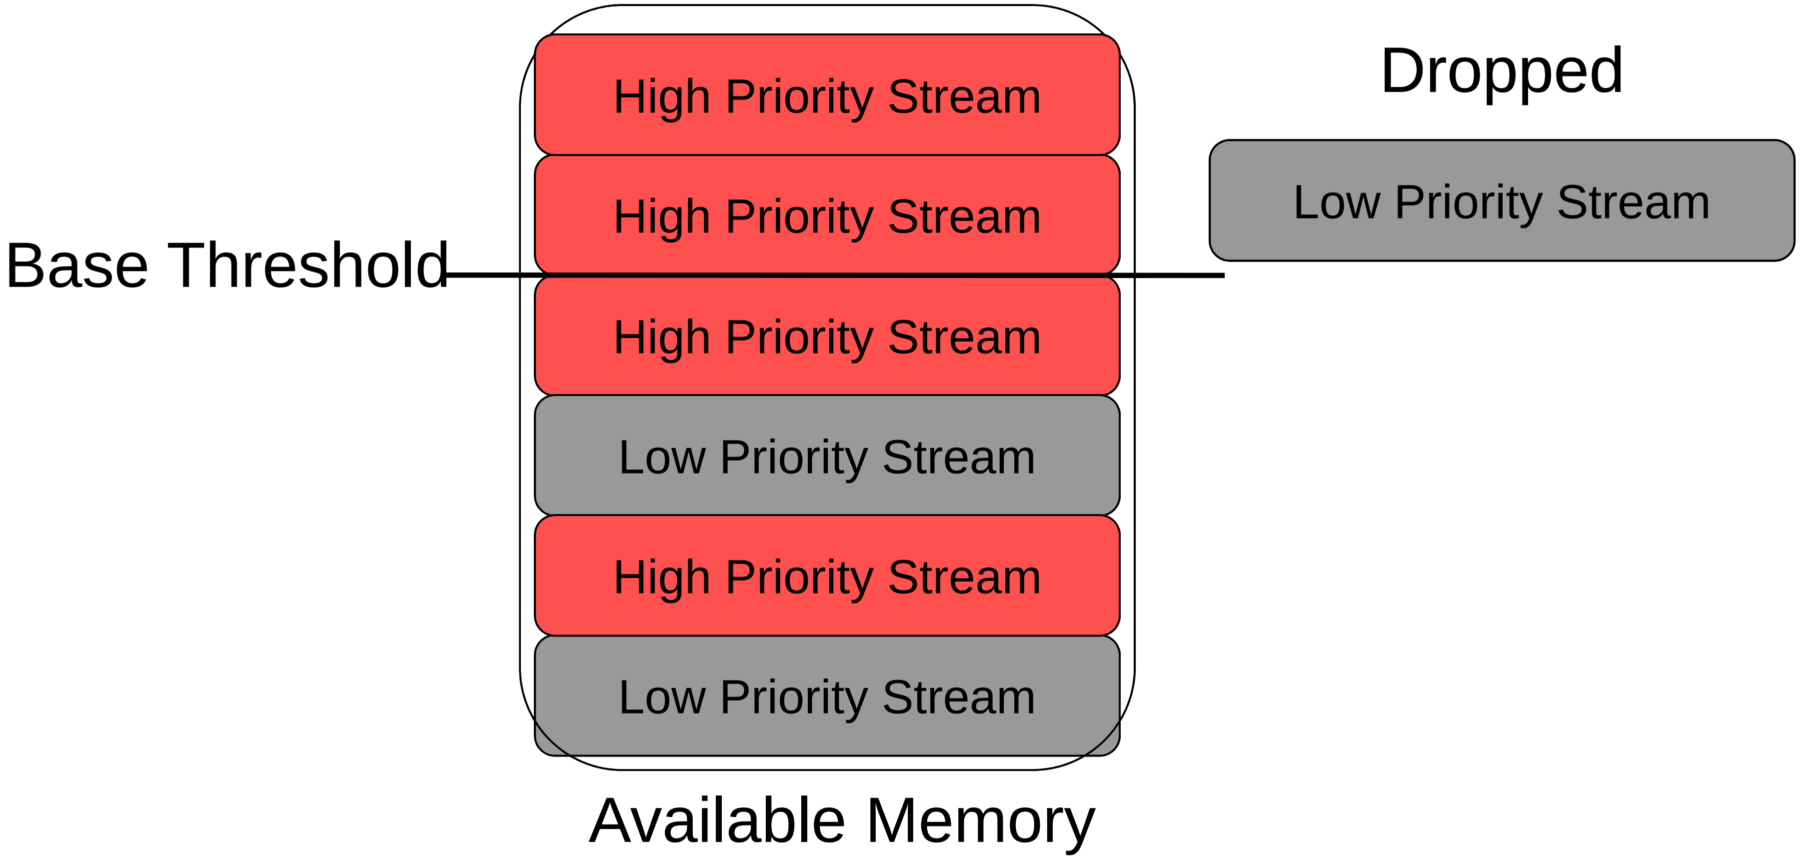

Flexible Stream Reasembly

- Stream reassembly at **user level** cannot control dropped packets
  - Missed control packets (TCP SYN, ACK, FIN, RST)
  - Captured streams with **holes**

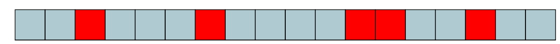

  - Extreme and pointless buffering in case of dropped packet
- Scap **controls** which packets are dropped
  - No control packet is missed
  - Captured **contiguous** streams

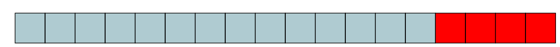

  - **Best effort** stream reassembly mode
- Different policy/OS implementation per stream

Other Features

- **Multiple applications** sharing the same reassembled streams
  - Best effort to satisfy all requirements
- **Packet delivery** per stream chunk if requested
  - Keep packet metadata and pointers to stream data
- **Parallel** stream processing using worker threads
- Improved **locality**
  - Keep related packets close
  - Capture and process a stream in the same core

The Scap Architecture

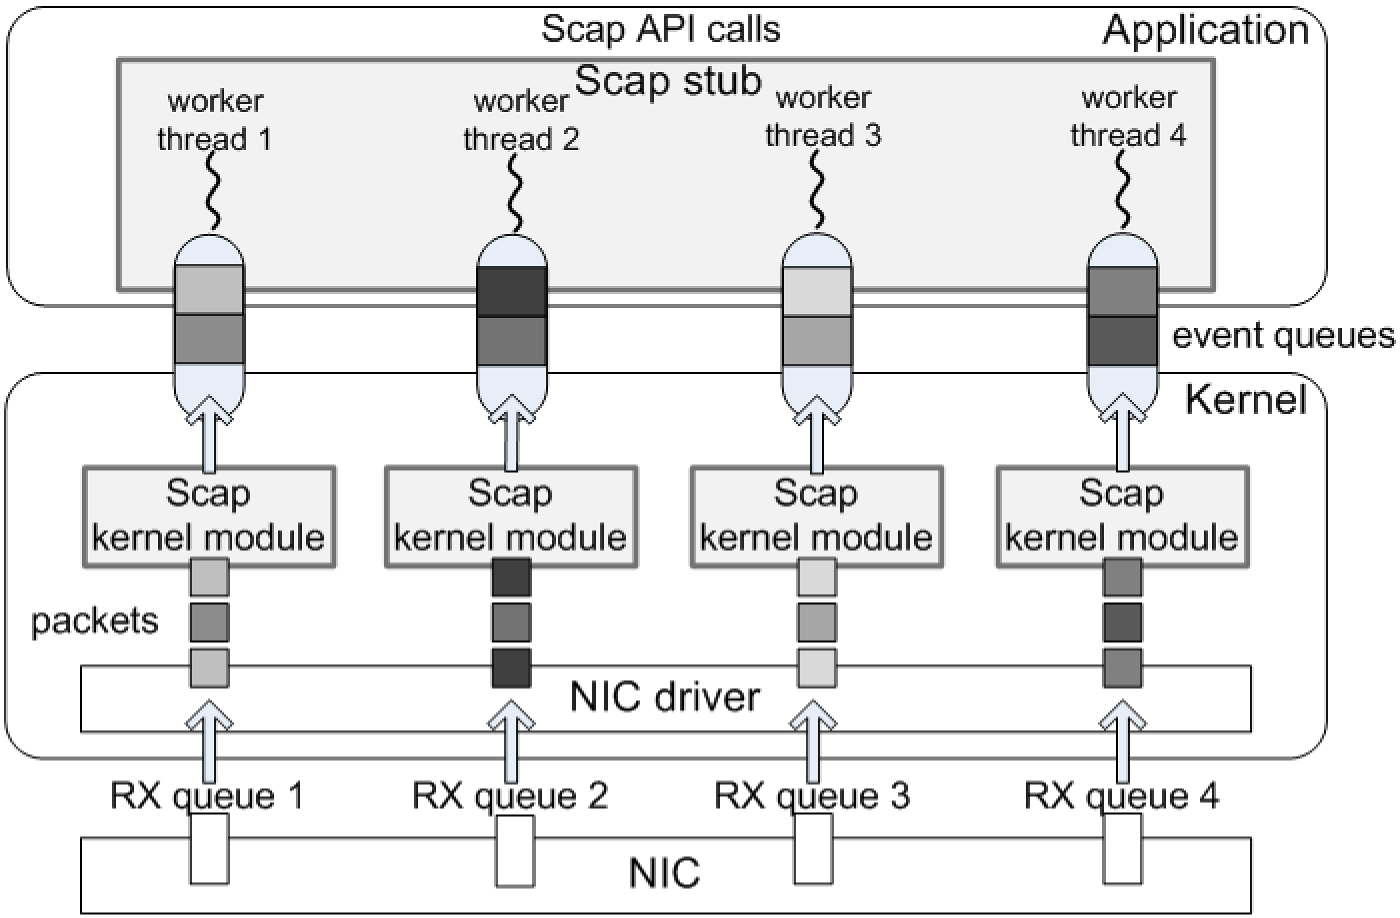

The Scap Kernel Module

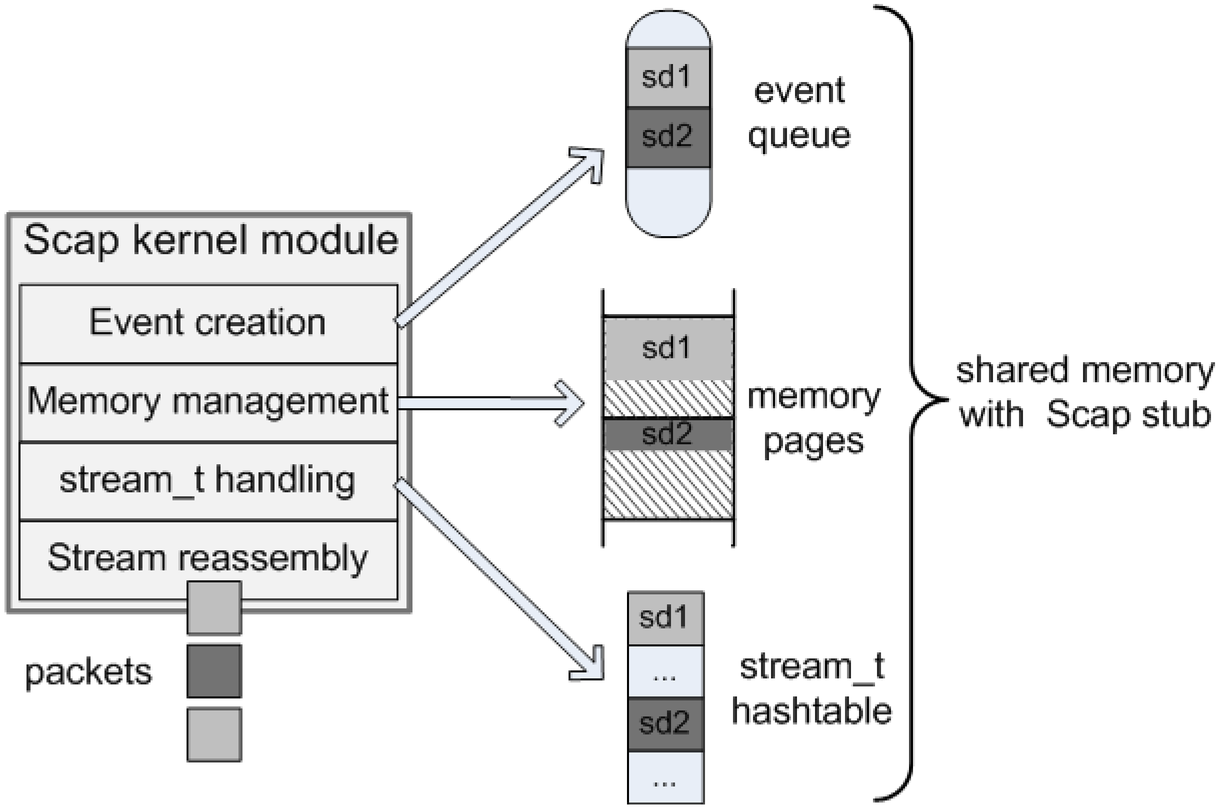

Experimental Evaluation

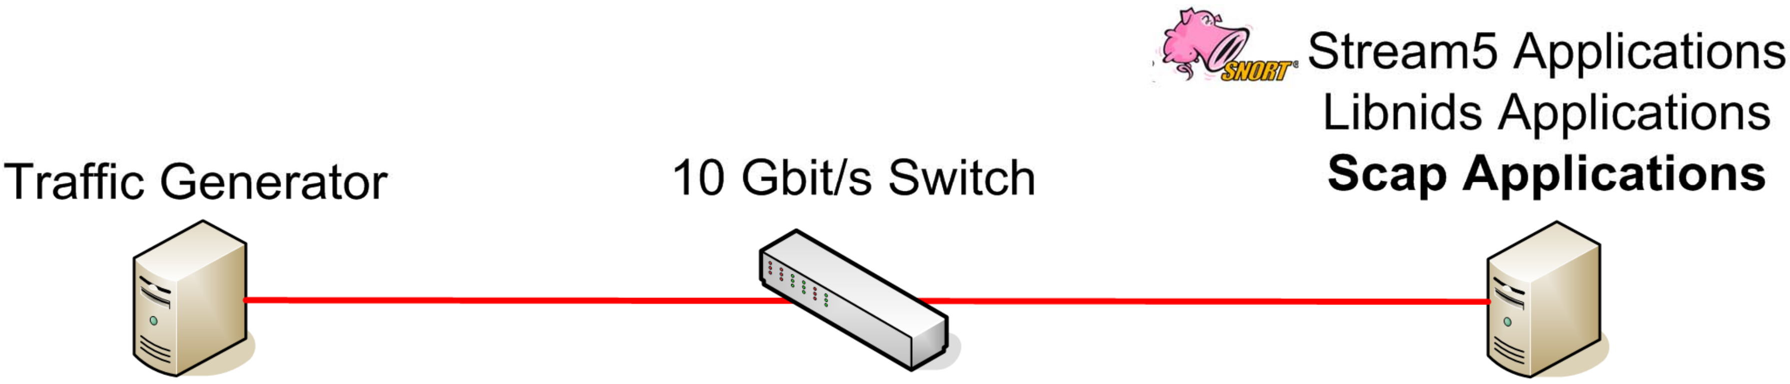

Flow Export

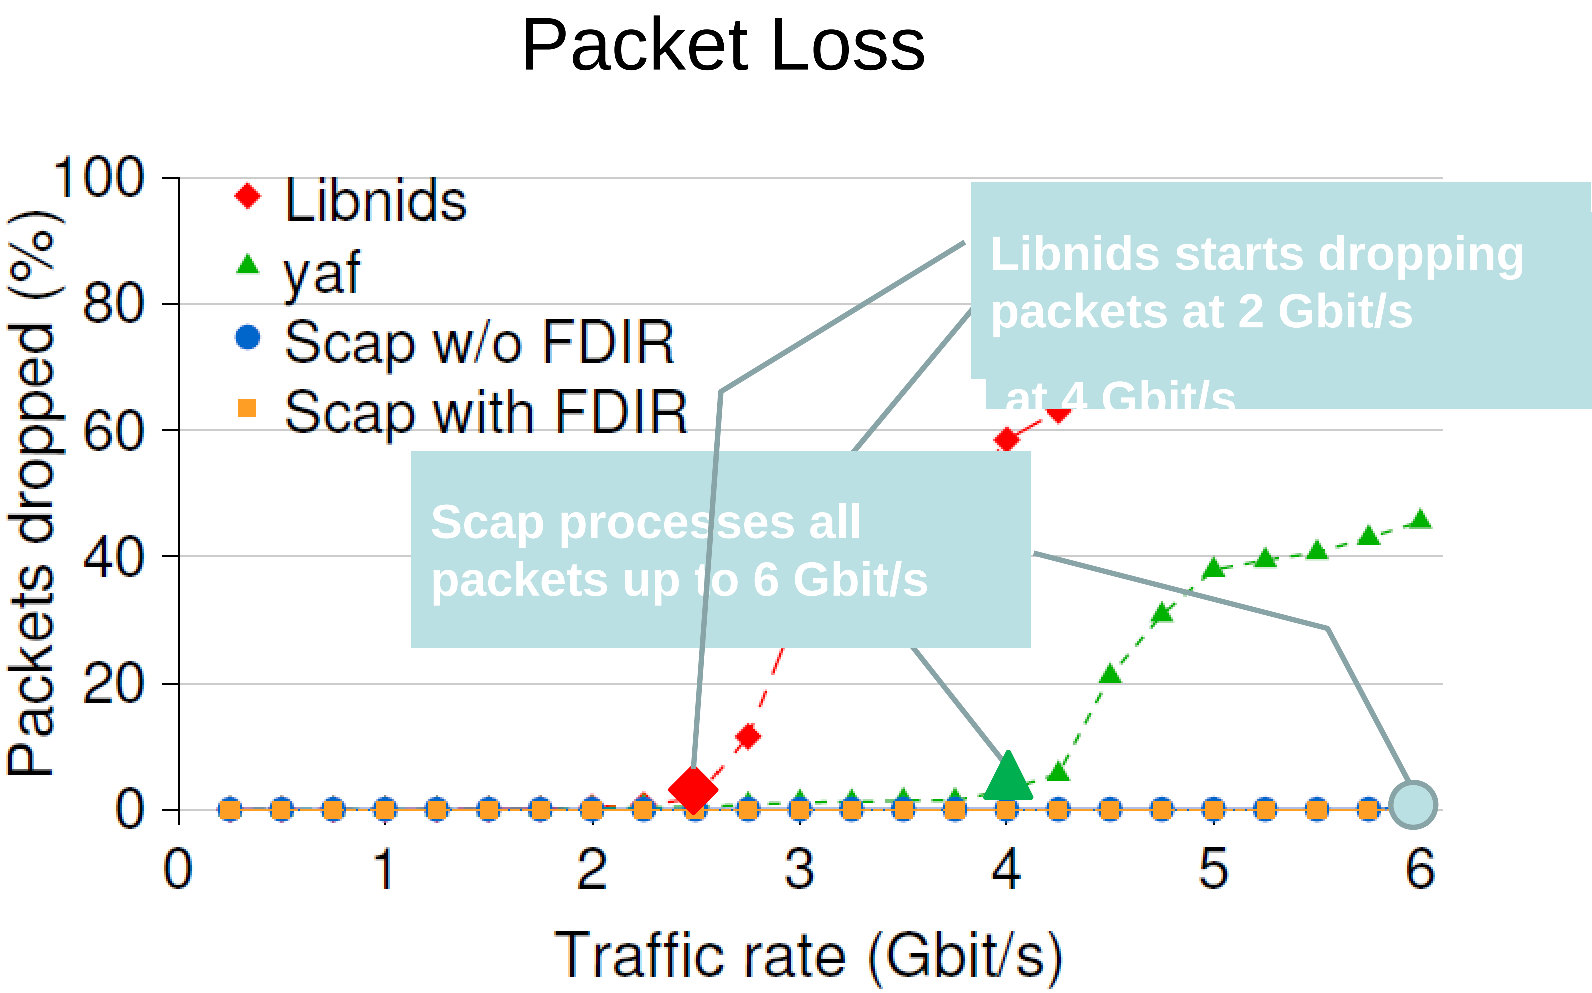

Packet Loss
Packets dropped (%)
100
80
60
40
20
0
Libnids
yaf
Scap w/o FDIR
Scap with FDIR
Libnids starts dropping packets at 2 Gbit/s
at 4 Gbit/s
Scap processes all packets up to 6 Gbit/s
0
1
2
3
4
5
6
Traffic rate (Gbit/s)

Stream Delivery

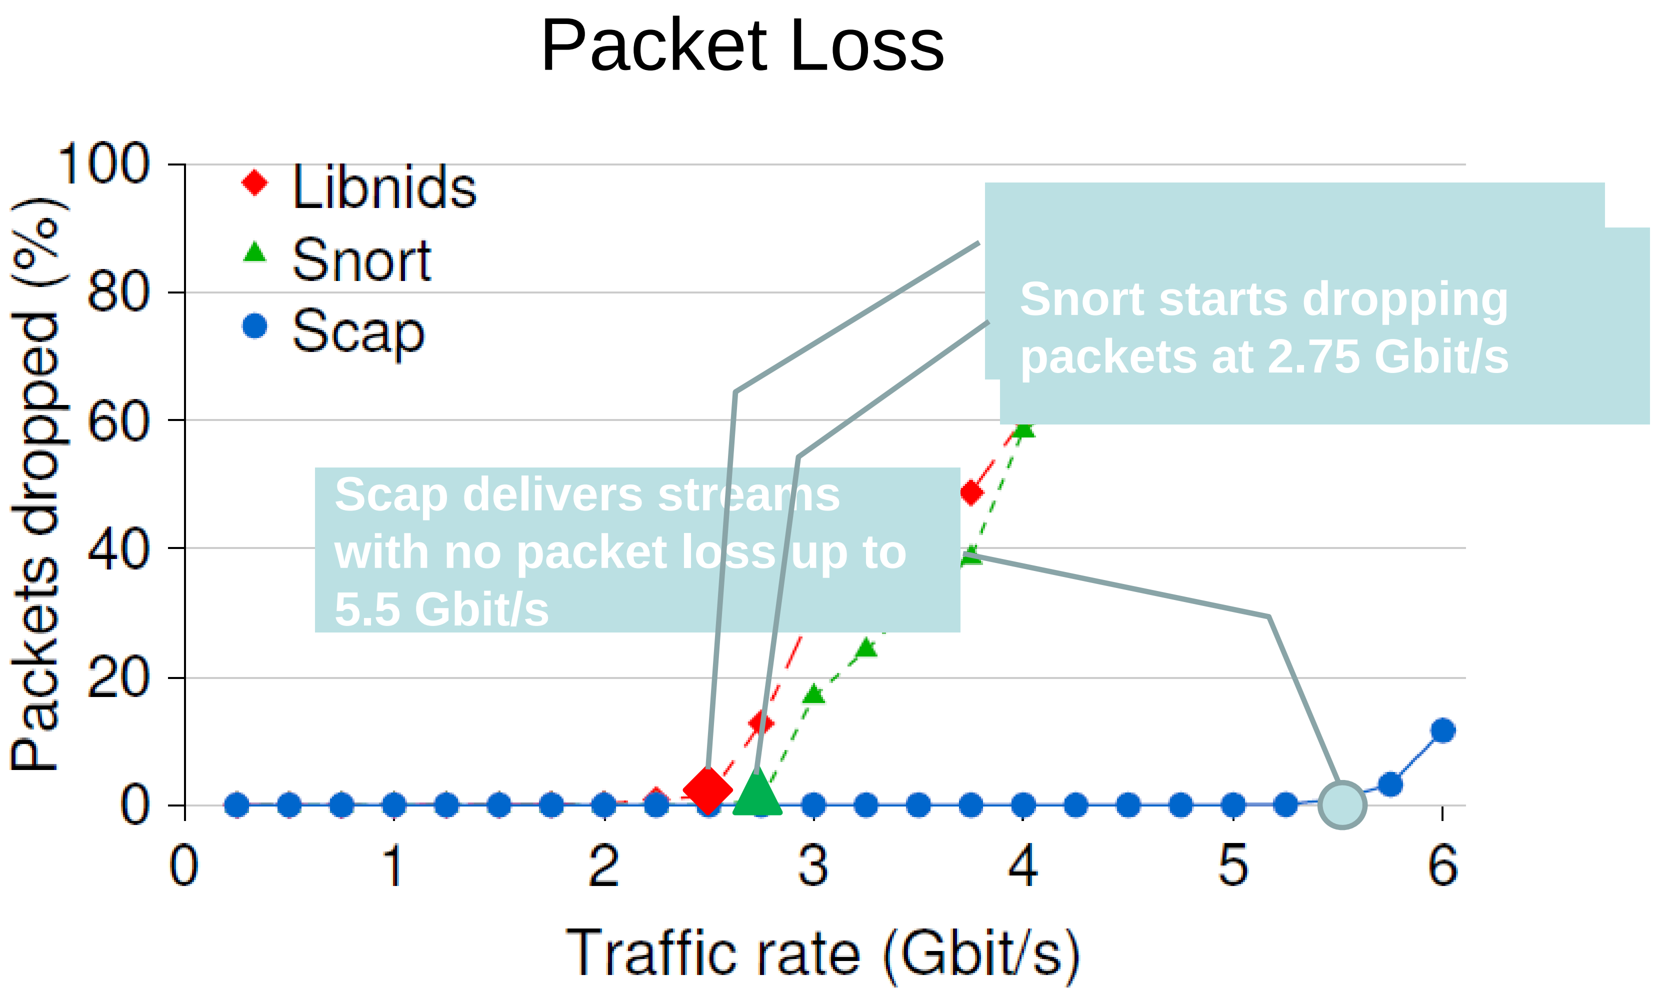

Cutoff Points

Packet Loss (Pattern matching, 4 Gbit/s)

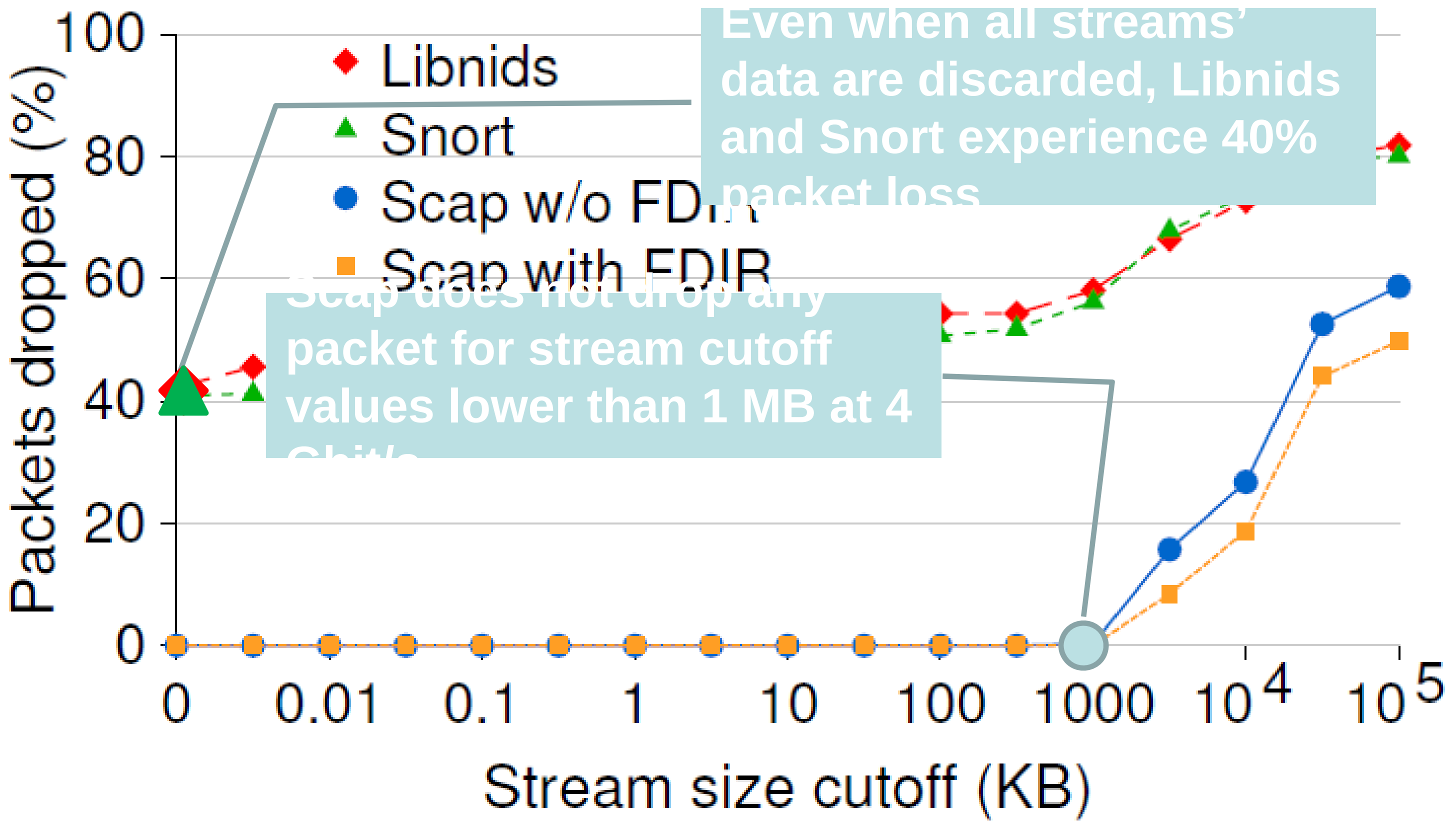

Even when all streams' data are discarded, Libnids and Snort experience 40% packet loss

Scap does not drop any packet for stream cutoff values lower than 1 MB at 4 Gbit/s

Conclusions

- Identified a **semantic gap**
  - Applications want high-level abstractions
  - Monitoring systems provide low-level packets
- Proposed the **stream** abstraction
- Better performance:
  - 2x throughput compared to Libnids and Stream5
- Truncate streams with **subzero copy** at the NIC
- Prioritized packet loss with **stream priorities**

Hot topics in Security Research – the Red Book

Evangelos Markatos
FORTH-ICS

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

Cyber Security is increasingly important

- The European Cyber Security Agenda:
  - 148,000 computers compromised daily
  - Symantec suggests that
    - Cybercrime victims lose 290 billion euros annually
  - 18% of users are less likely to buy goods online
  - 74% agreed that the risk of becoming a victim of cybercrime has gone up in the past year

Cyberattacks are getting more prelavent

- Hackers are getting more effective
- Users are getting more concerned
  - 12% of Internet users has experienced fraud
  - 8% have been victims of ID theft
    - » (src: Eurobarometer 390)

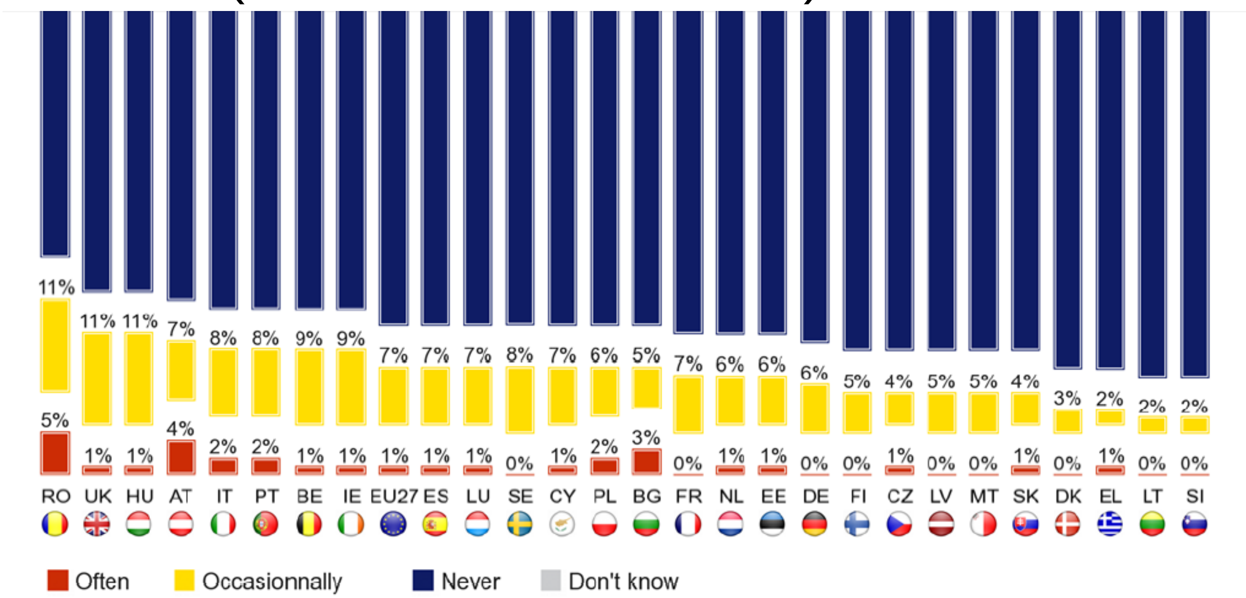

What is the impact of attacks?

"… potential (cyber)attacks against network infrastructures may have widespread and devastating consequences on our daily life: no more electricity or water at home, rail and plane accidents, hospitals out of service"

Viviane Reding

VP of the European Commission

European Cybersecurity Month

"in tomorrow's world if the internet isn't secured, nothing will be …"

Neelie Kroes

VP of the European Commission

How large is it?

- Cybercrime is *larger than*
  - *The global black market in marijuana, cocaine and heroin combined*

--Symantec

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

What shall we do?

- *Understand the important Research Issues*
- *Write them down in a book*
- *Circulate it widely*
  - *So that researchers can work on them*

- *The result:*
  - *The Red Book*
    - *in Cyber Security*

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

How did we do it?

- To build a winning team you need
  - Excellence,
  - Talent, and
  - Desire to work hard.

We assembled a Task Force of young European Researchers

Task Force

MEMBERS

Elias Athanasopoulos
Columbia University
Federico Maggi
Politecnico di Milano
Asia Slowinska
Vrije Universiteit

Lorenzo Cavallaro
Royal Holloway University of London
Michalis Polychronakis
Columbia University and FORTH
Iason Polakis
FORTH and University of Crete

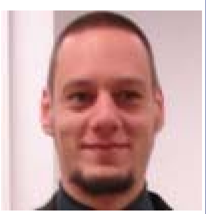



Contributors

Magnus Almgren
Chalmers
Sotiris Ioannidis
FORTH
Philippas Tsigas
Chalmers

Herbert Bos
Vrije Universiteit
Christian Platzer
TUV
Stefano Zanero
Politecnino di Milano

CONTRIBUTORS

Dennis Andriesse
Vrije Universiteit
Farnaz Moradi
Chalmers University
Simin Nadjm-Tehrani
Linköping University

Martina Lindorfer
TU Vienna
Zlatogor Minchev
Bulgarian Academy of Sciences
Christian Rossow
Vrije Universiteit

Chairs

SYSSEC TASK FORCE for the ROADMAP on SYSTEMS SECURITY RESEARCH

CO-CHAIRS

Evangelos Markatos
SysSec Project Manager
Foundation for Research and Technology - Hellas

Davide Balzarotti
SysSec WP4 Leader
Eurecom

The making of Red Book

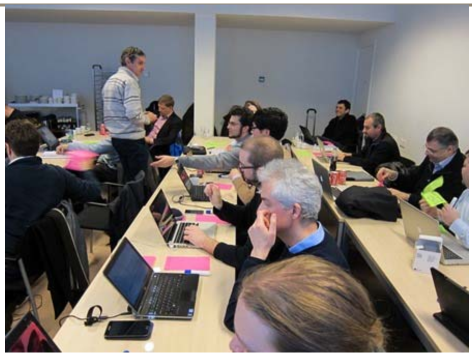

- "Rank the threats" workshop
  - Which are the important threats?
  - Rank them
- "What if" questions
- Grand Challenges

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- "What if" Questions
- The Threats
- The Grand Challenges
- Summary

“What if” Questions

- Examples from other disciplines
  - What if …
    - Antibiotics do not work anymore?
      – How would this impact medicine research?
    - There are no more fossil fuels to burn in 5 years?
      – How would this impact research in energy sources?
- “What if” questions
  - What if there is no more malware?
  - What if 50% of the computers are compromised?
  - What if there is no death? (for our data)
  - What if there is no Internet? (for a day or two)

“What if” Questions

- What if there is no more malware?
  - Will Security Research be over?
  - Will there be any security issues?
  - How about privacy issues?
- What if 50% of the computers are compromised?
  - How would you use them?
    - Why? When?
  - Would you do e-banking?
    - Under what circumstances?

“What if” Questions

- What if there is no death? (for our data)
  - Can we donate them?
  - Can we pass them on to our children?
- What if there is no Internet? (for a day or two)
  - What would work? What would not work?
  - Traffic? Air travel?
  - Will you be able to go home?
    - From work? from a business meeting?

Example “what if”

- What if there is no death? (for our data)
  - Will they be available after we pass away?
  - Can our children “inherit” our data?
    - Will they be able
      - to “own” our data?
      - to pass them on to the next generation?
        - » much like family photo albums?
  - Can we donate our data?
    - to Science?
    - Are there any security/privacy implications?
  - Can we incorporate all our data to an avatar?
    - Will the avatar be able to act on behalf of us?

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

The Threats

- “Rank the threats” workshop
  - Which are the important threats?
  - Rank them

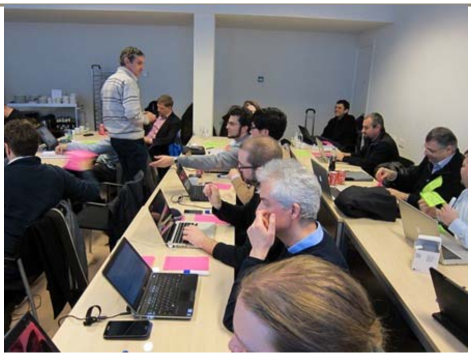


Cyber-security landscape

- Threat – Vulnerabilities
- Assets
- Domains
- Horizontal Research Areas

Threats - Vulnerabilities

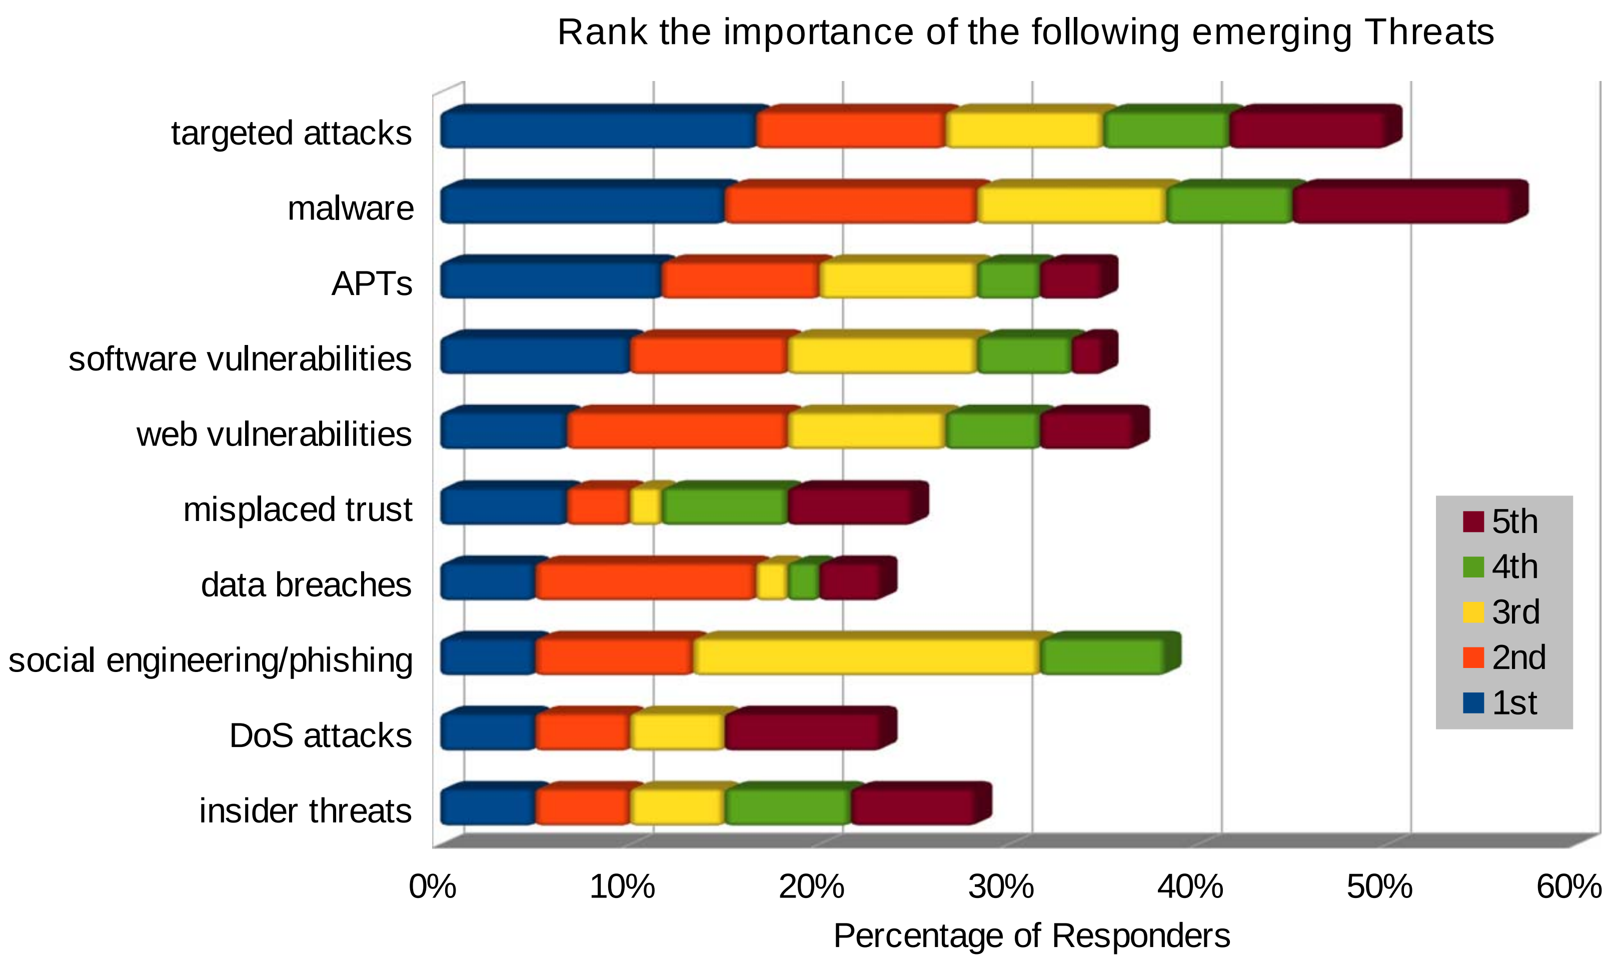

Assets

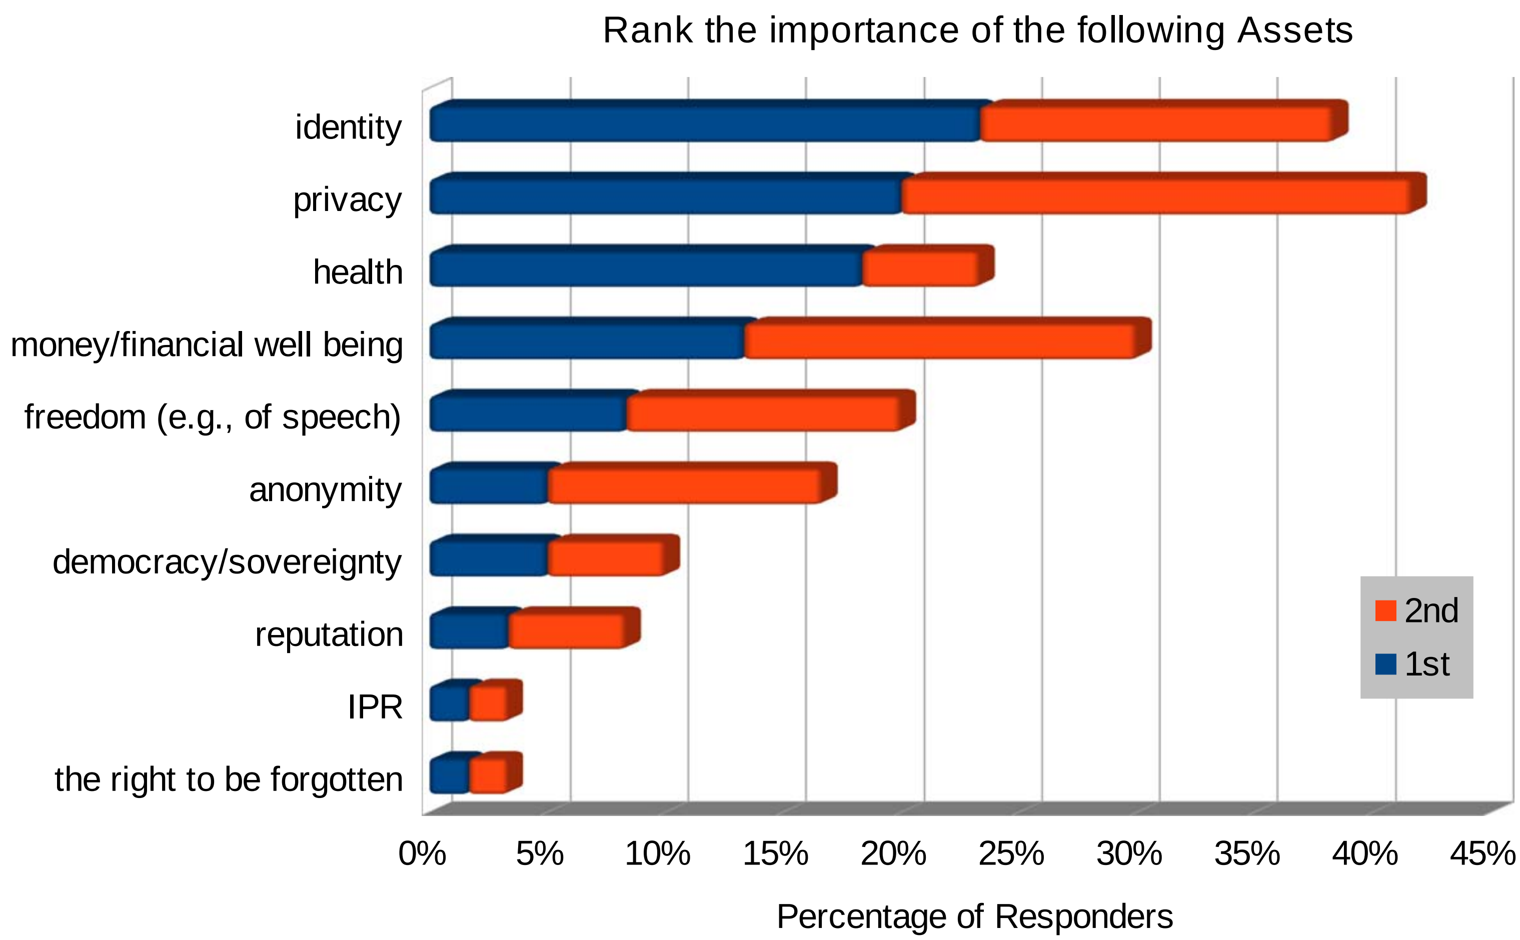
Rank the importance of the following Assets
identity
privacy
health
money/financial well being
freedom (e.g., of speech)
anonymity
democracy/sovereignty
reputation
IPR
the right to be forgotten
0%
5%
10%
15%
20%
25%
30%
35%
40%
45%
Percentage of Responders
2nd
1st

Domains

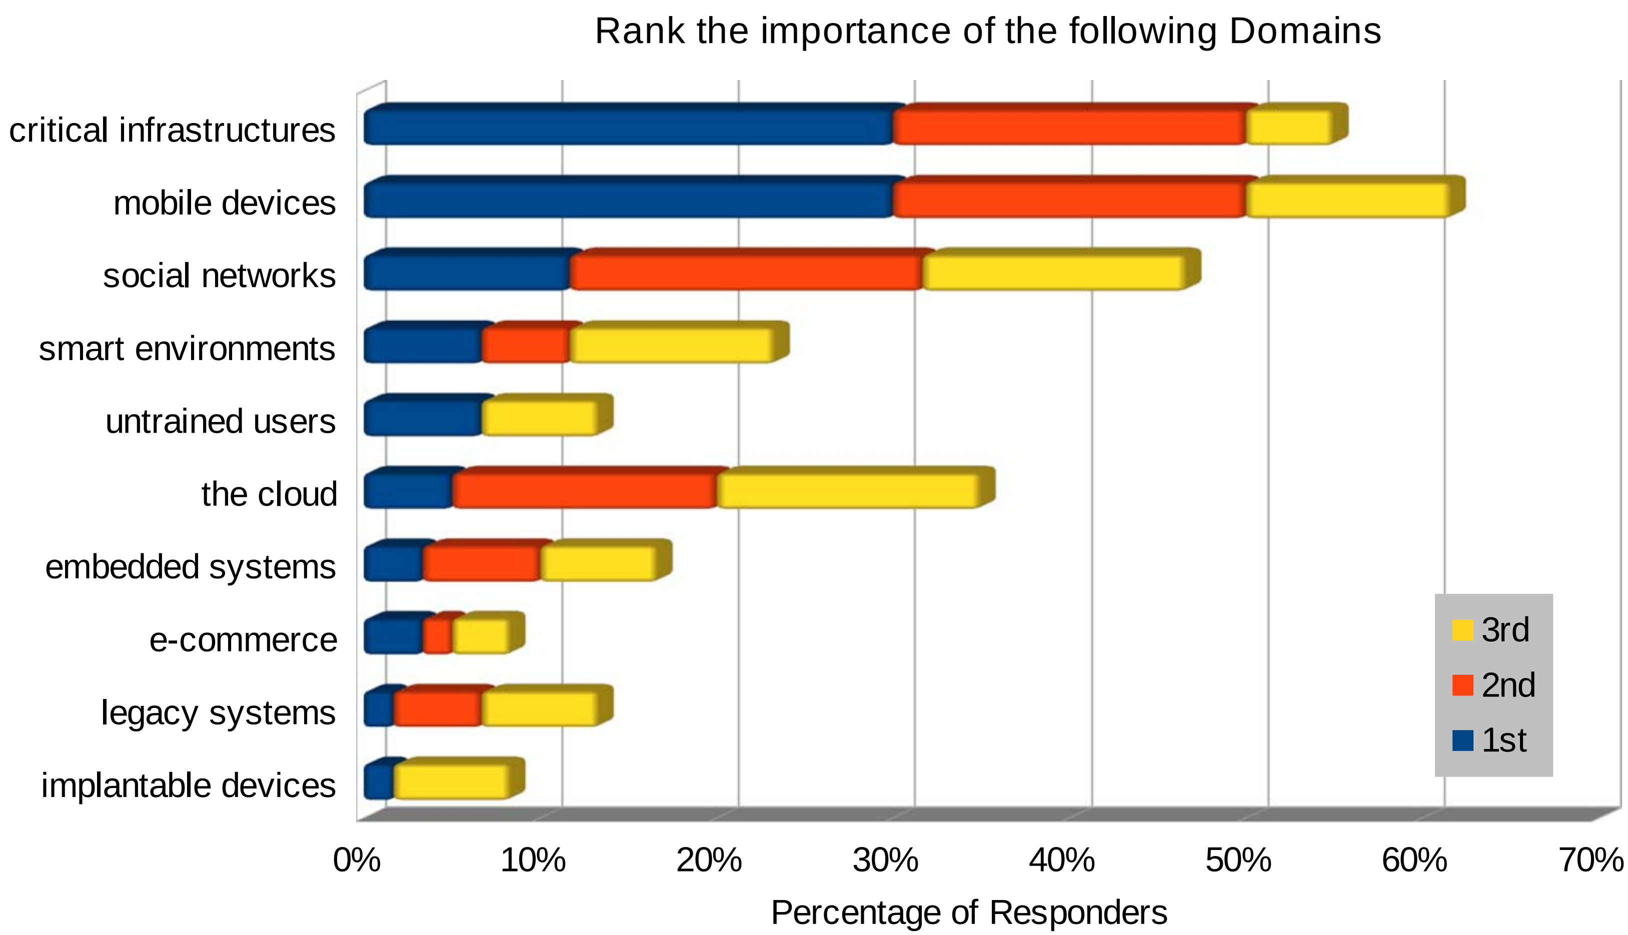

Rank the importance of the following Domains
critical infrastructures
mobile devices
social networks
smart environments
untrained users
the cloud
embedded systems
e-commerce
legacy systems
implantable devices
0%
10%
20%
30%
40%
50%
60%
70%
Percentage of Responders
3rd
2nd
1st

Most important threats

- Malware
- Targeted Attacks – Advanced Persistent Threats
- Social Engineering - Phishing

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

Grand challenges

- No device should be compromisable
- Give users control of their data
- Provide private moments in public places
- Develop compromise-tolerant systems

Example Grand Challenge

- Give users control over their data
- Users should be able to
  - know which data they have created
  - know which data they have given to which third parties
    - Cookies, accesses, IP addresssses, MAC addresses, etc.
  - Revoke all access to their data
  - Ask data to be deleted
    - if this is not prohibited by law

RoadMap of the talk

- Introduction
- The Red Book
- The making of the Red Book
- “What if” Questions
- The Threats
- The Grand Challenges
- Summary

Summary

- The Red Book:
  - Identify Research Directions in Systems Security
- The making of it:
  - Task Force of young excellent scientists
    - They drive the work
  - Workshop with the community
    - Everyone is engaged
- The result:
  - Threats, assets, priorities
  - Grand Challenges

Hot topics in Security Research – the **Red Book**

Evangelos Markatos
FORTH-ICS

The Red Book

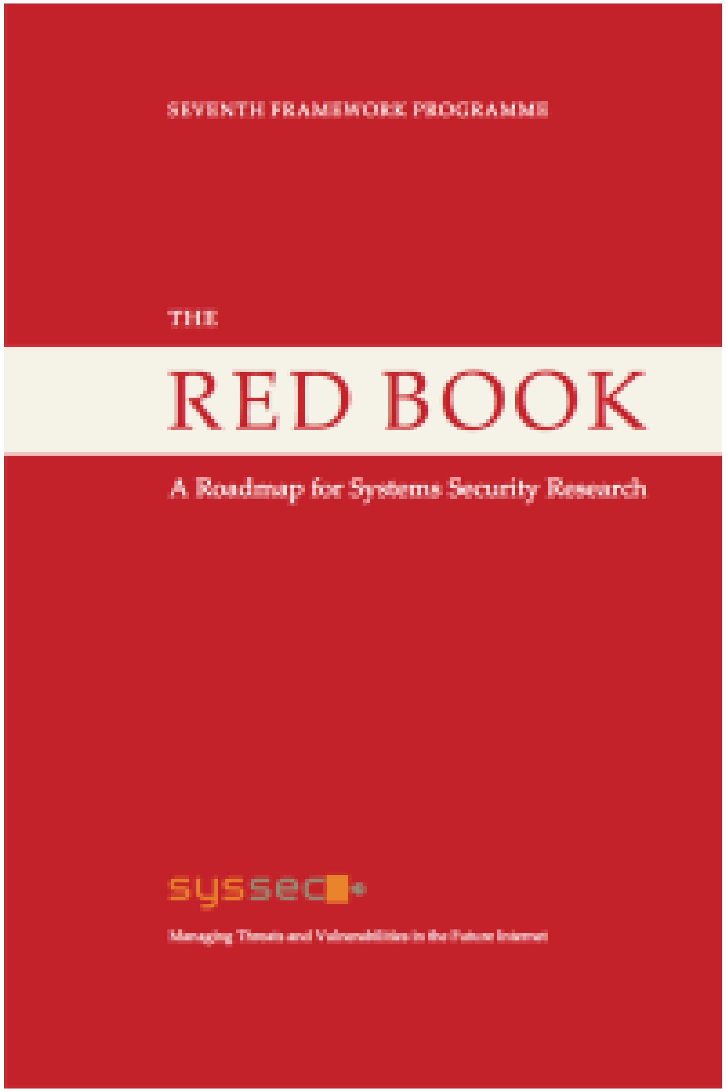
SEVENTH FRAMEWORK PROGRAMME
THE
RED BOOK
A Roadmap for Systems Security Research
syssec
Managing Threats and Vulnerabilities in the Future Internet

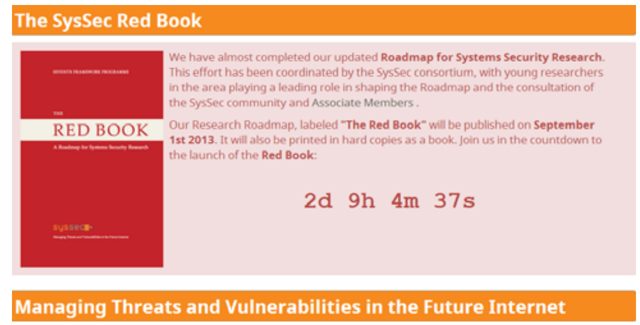
The SysSec Red Book
RED BOOK
We have almost completed our updated Roadmap for Systems Security Research. This effort has been coordinated by the SysSec consortium, with young researchers in the area playing a leading role in shaping the Roadmap and the consultation of the SysSec community and Associate Members .
Our Research Roadmap, labeled "The Red Book" will be published on September 1st 2013. It will also be printed in hard copies as a book. Join us in the countdown to the launch of the Red Book:
2d 9h 4m 37s
Managing Threats and Vulnerabilities in the Future Internet