Empirical study of the impact of Metasploit-related attacks in 4 years of attack traces

E. Ramirez-Silva and M. Dacier

Eurécom Institute - Sophia Antipolis, France

ASIAN'07

December 11, 2007- Doha, Qatar

Overview

- Introduction
  - The Leurré.com project
  - Experimental framework
  - Experimental results
- Conclusions

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Overall goal of the approach

- What can honeypots tell us about “script kiddies” related attacks?
  - How much impact do they have on these datasets?
  - Where do they come from?
  - When are we likely to see them?
  - Do they have a specific profile of activity?

To be or not to be a script kiddie

- Question:
  - Among all the attacks observed on a honeypot, how can we distinguish those likely due to script kiddies?
- Answer:
  - Define and detect on the honeypot the traces left by a specific tool, supposed to be used by script kiddies.

Our response

- Script kiddie tool:
  - We have decided to identify and study only instances of attacks likely due to metasploit plugins.
- Traces of the attack tool:
  - We have built an environment to run all attacks against a honeypot in a monitored environment.
  - The recorded traces are used to generate “network signatures” for each plugin.

To be or not to be at the right place

- Question:
  - What is the best “place” to look for attacks?
- Answer:
  - Everywhere in the world as previous work have shown that different blocks of addresses can be hit by different types of attacks.

Our response

- Our source of information:
  - We use attack traces collected by the low-interaction honeypots deployed within the Leurré.com project.
- Origin of the data:
  - This gives us access to 4 years of data collected in a large number of different environments, on the very same type of platform
    - 50 platforms in 30 different countries as of today
    - None in Qatar ... yet ...

Caveat

- We acknowledge the fact that, by focusing on Metasploit plugins only, we address a small fraction of the whole problem space
  - The experiments only derive lower bounds of the amount of attacks due to script kiddies.
- The lessons learned are, hopefully, of a much broader interest.

Overview

- Introduction
  - ✔ The Leurré.com project
  - ✔ Experimental framework
  - ✔ Experimental results
- Conclusions

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Leurré.com: a brief overview

- Ongoing effort since 2003:
  - Around 50 platforms running today in 30 different countries
  - All platforms have the very same configuration ; based on honeyd, each one implements 3 virtual machines
- Every day, tcpdump files are uploaded, enriched and stored into a centralized DB.
  - geographical location of the attackers, passive OS fingerprints of their machines, reverse name lookups, etc.

50 platforms in 30 different countries

In Europe


Norwegian Sea
ICELAND
NORWAY
RUSSIA
North Atlantic Ocean
North Sea
IRELAND
BELARUS
ROMANIA
BULGARIA
GREECE
TURKEY
Mediterranean Sea
MOROCCO
ALGERIA
MALTA

Win-Win Partnership

- Interested partner provides
  - An old PC (Pentium II, 256MB RAM, 233 MHz)
  - 4 unfiltered routable IP addresses
- The Project provides
  - Installation CD Rom containing OS + applications
  - Remote log collection and integrity checks
  - Access to the whole data set + wiki + various tools developed by the community (GUI, java applets, Matlab programs, alert ticketing system, etc.)

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Clusters of traces

- Among the various treatments, one important one aims at grouping together attack traces likely due to the same attack tool.
- This is done thanks to a simple clustering algorithm that group together attack sessions (traces of 1 IP against 1 platform) that share the same fingerprints
- Fingerprints are defined by means of 7 groups of attributes

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Attack fingerprints

1. Amount of targeted virtual machines,
2. Order in which they have been hit,
3. Amount of packets sent by the attacker to each virtual machine,
4. Sequence of ports,
5. Total amount of packets sent by the attacker,
6. Average IAT between packets received.
7. Duration of the attack.

Data used

- The experiments reported are based on the 4 years of collected data.
- They take advantage of the notion of clusters as defined and implemented by the project in the database available to all partners.

Overview

- Introduction
  - The Leurré.com project
  - Experimental framework
  - Experimental results
- Conclusions

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Sanity Check

- Question:
  - Have we ever observed a clear manifestation of a metasploit related cluster in the Leurré.com data set?
- Answer:
  - Yes, for instance, on May 15 2006, the one implementing an exploit against the *'RealVNC password authentication bypass vulnerability' (realvnc_41_bypass)*

Graphical Representation

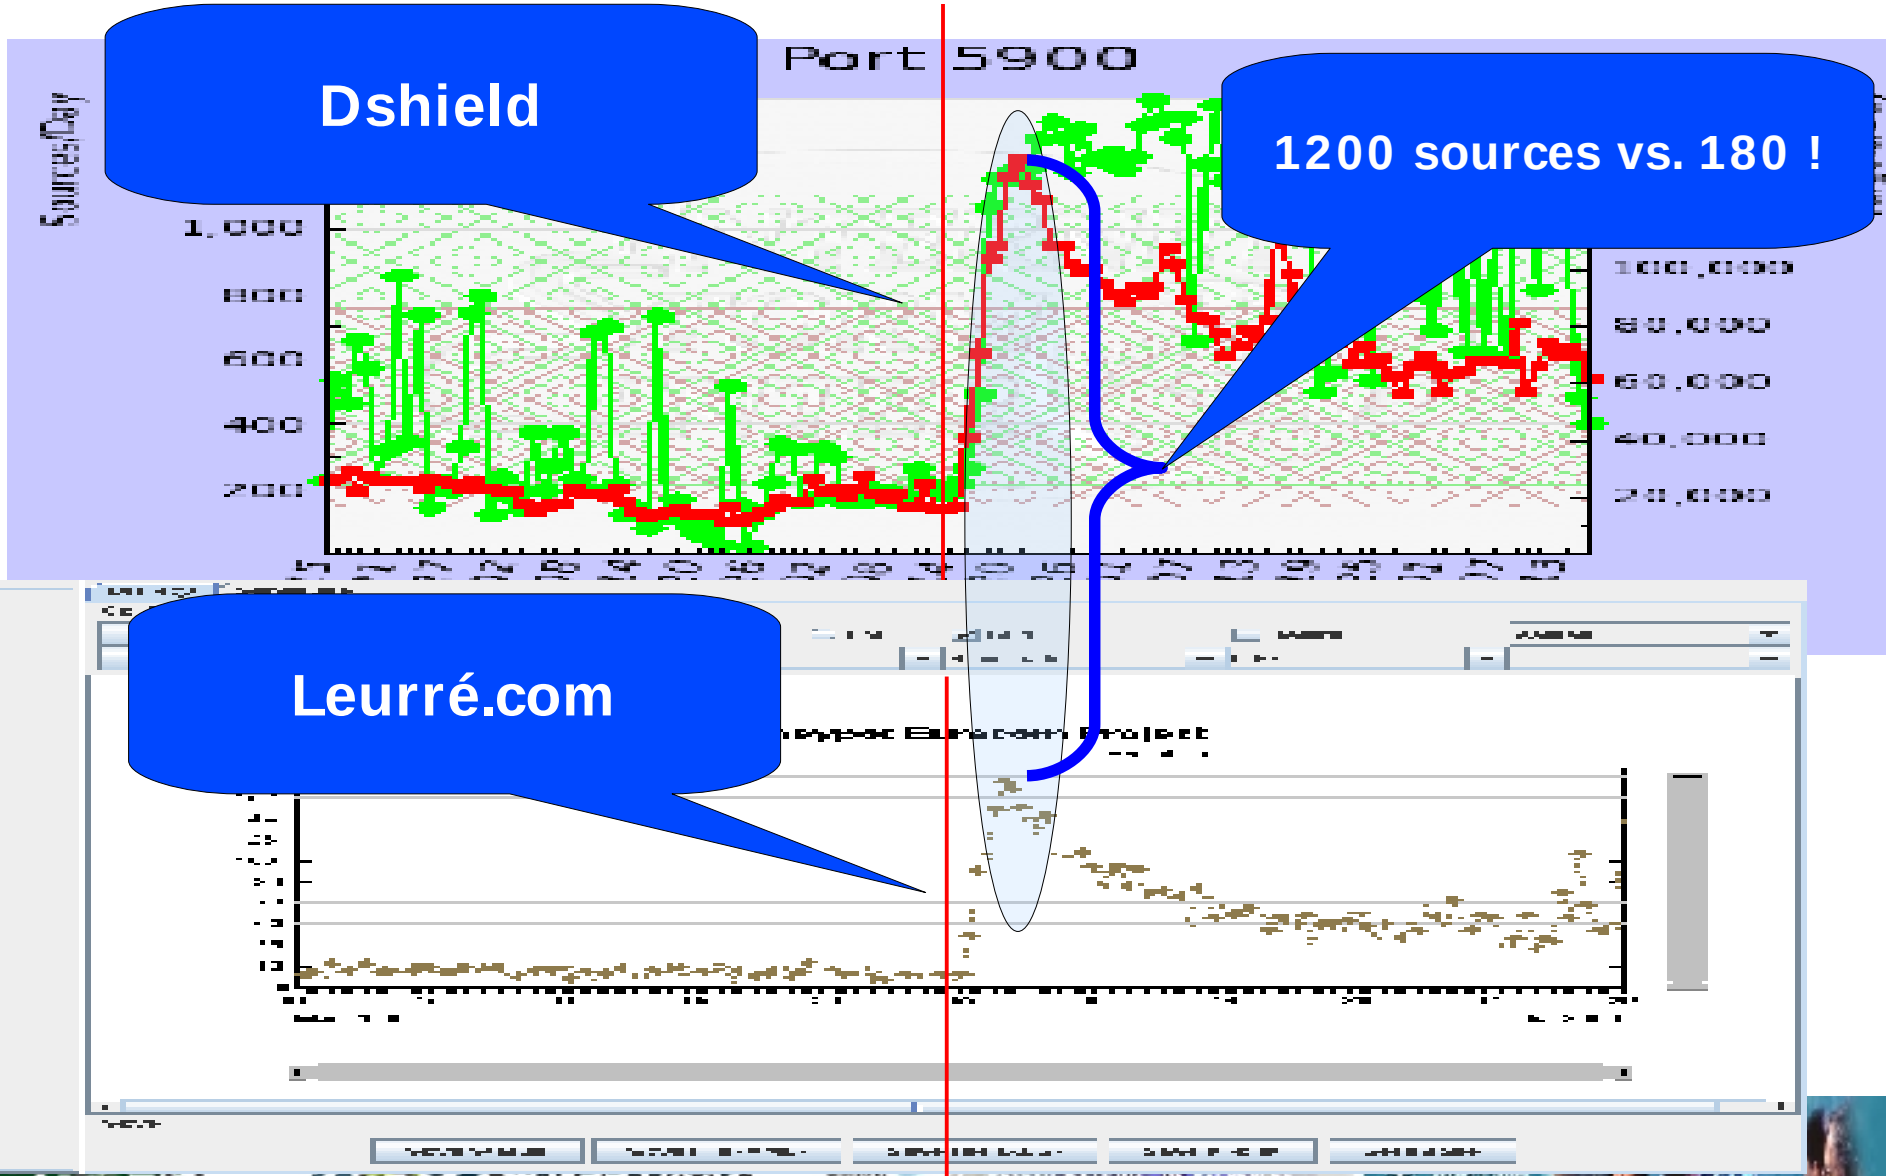

Metasploit framework

- It is often referred to as the most popular vulnerability exploitation tool
- Its ease of use makes it the ideal tool for script kiddies
- For practical reasons, we restrict ourselves to all versions of the Metasploit framework within the release 2 (2.0-2.7) to analyze their impacts on our dataset.

Method used

- We have run all attacks from all Metasploit releases, one by one, against one of our platforms, in a dedicated environment.
- Traces have been recorded and labels
- Cluster attributes have been derived from these traces
- Matching clusters have been retrieved from the DB for further analysis.

Metasploit signature generation

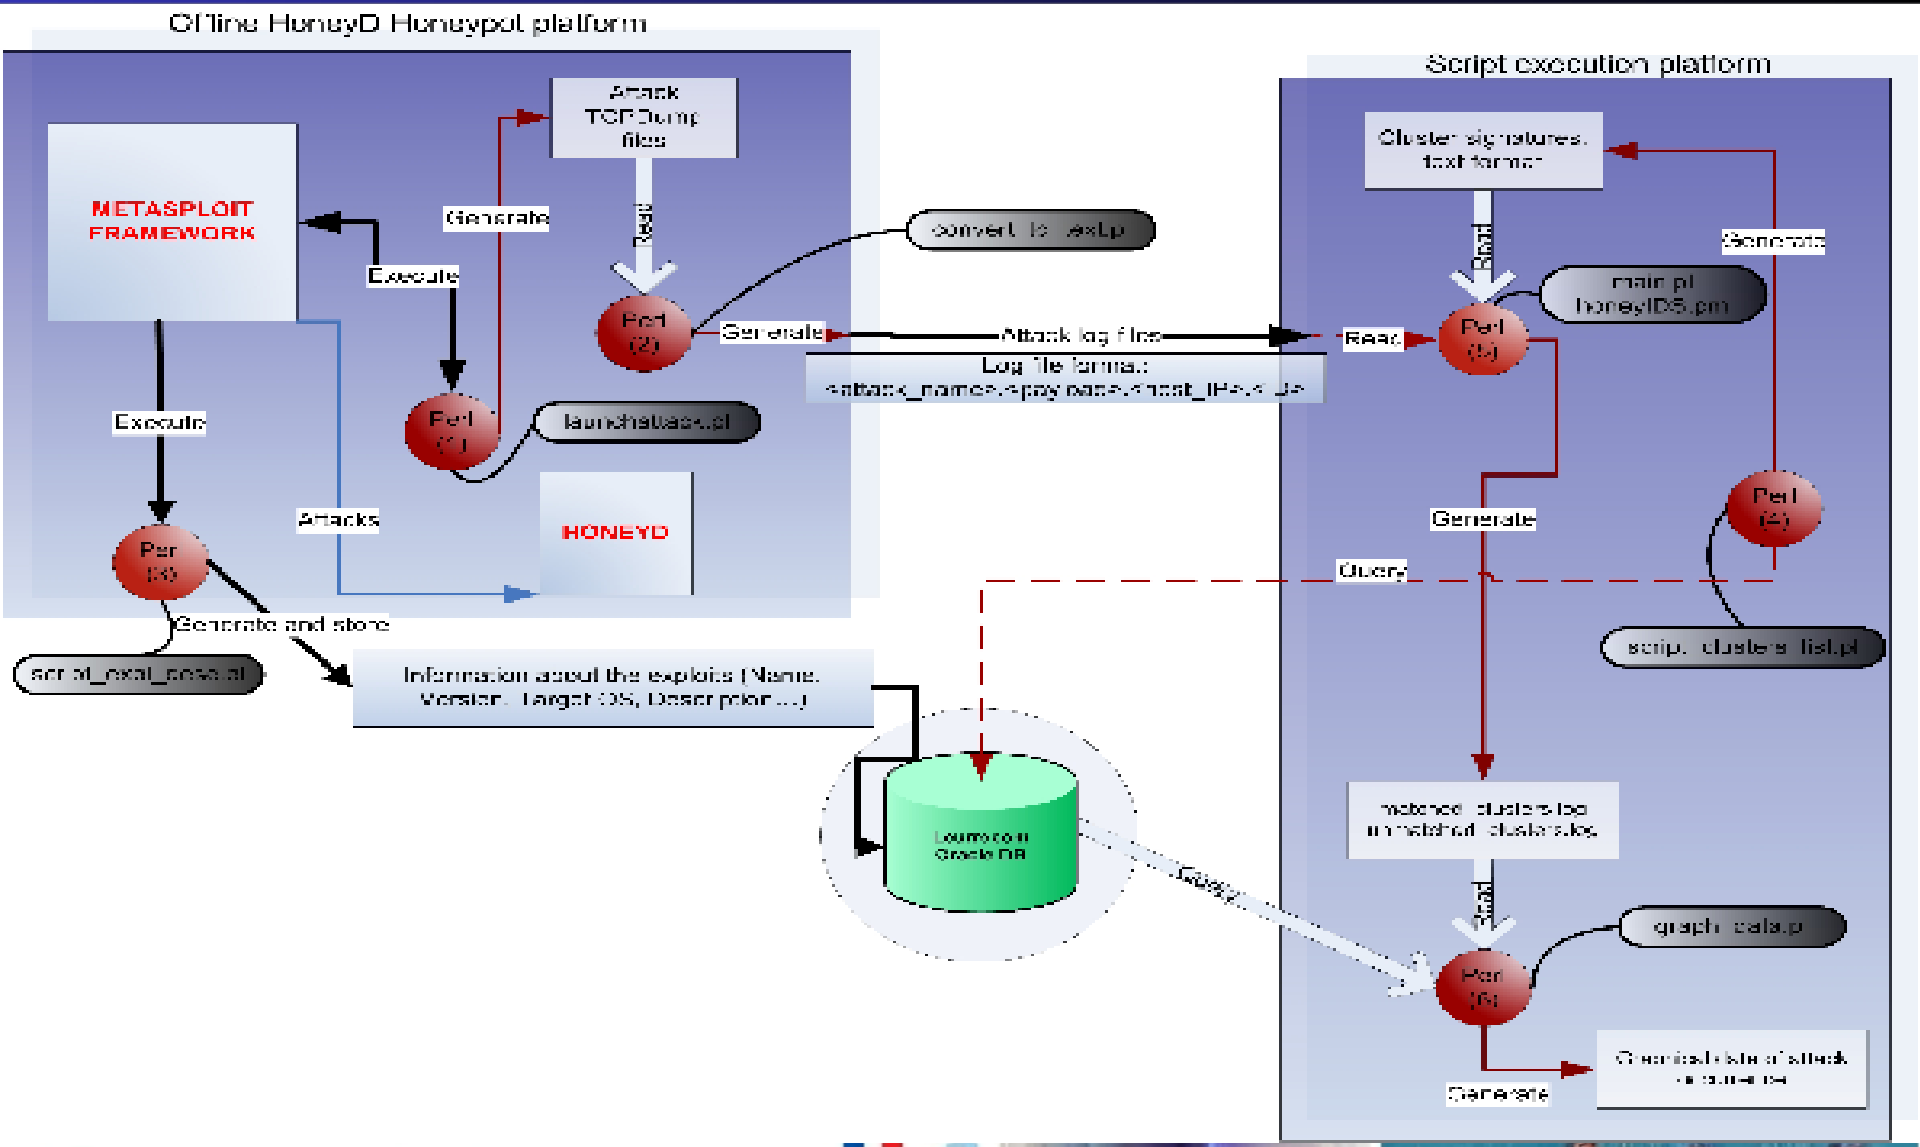

Metasploit signature generation

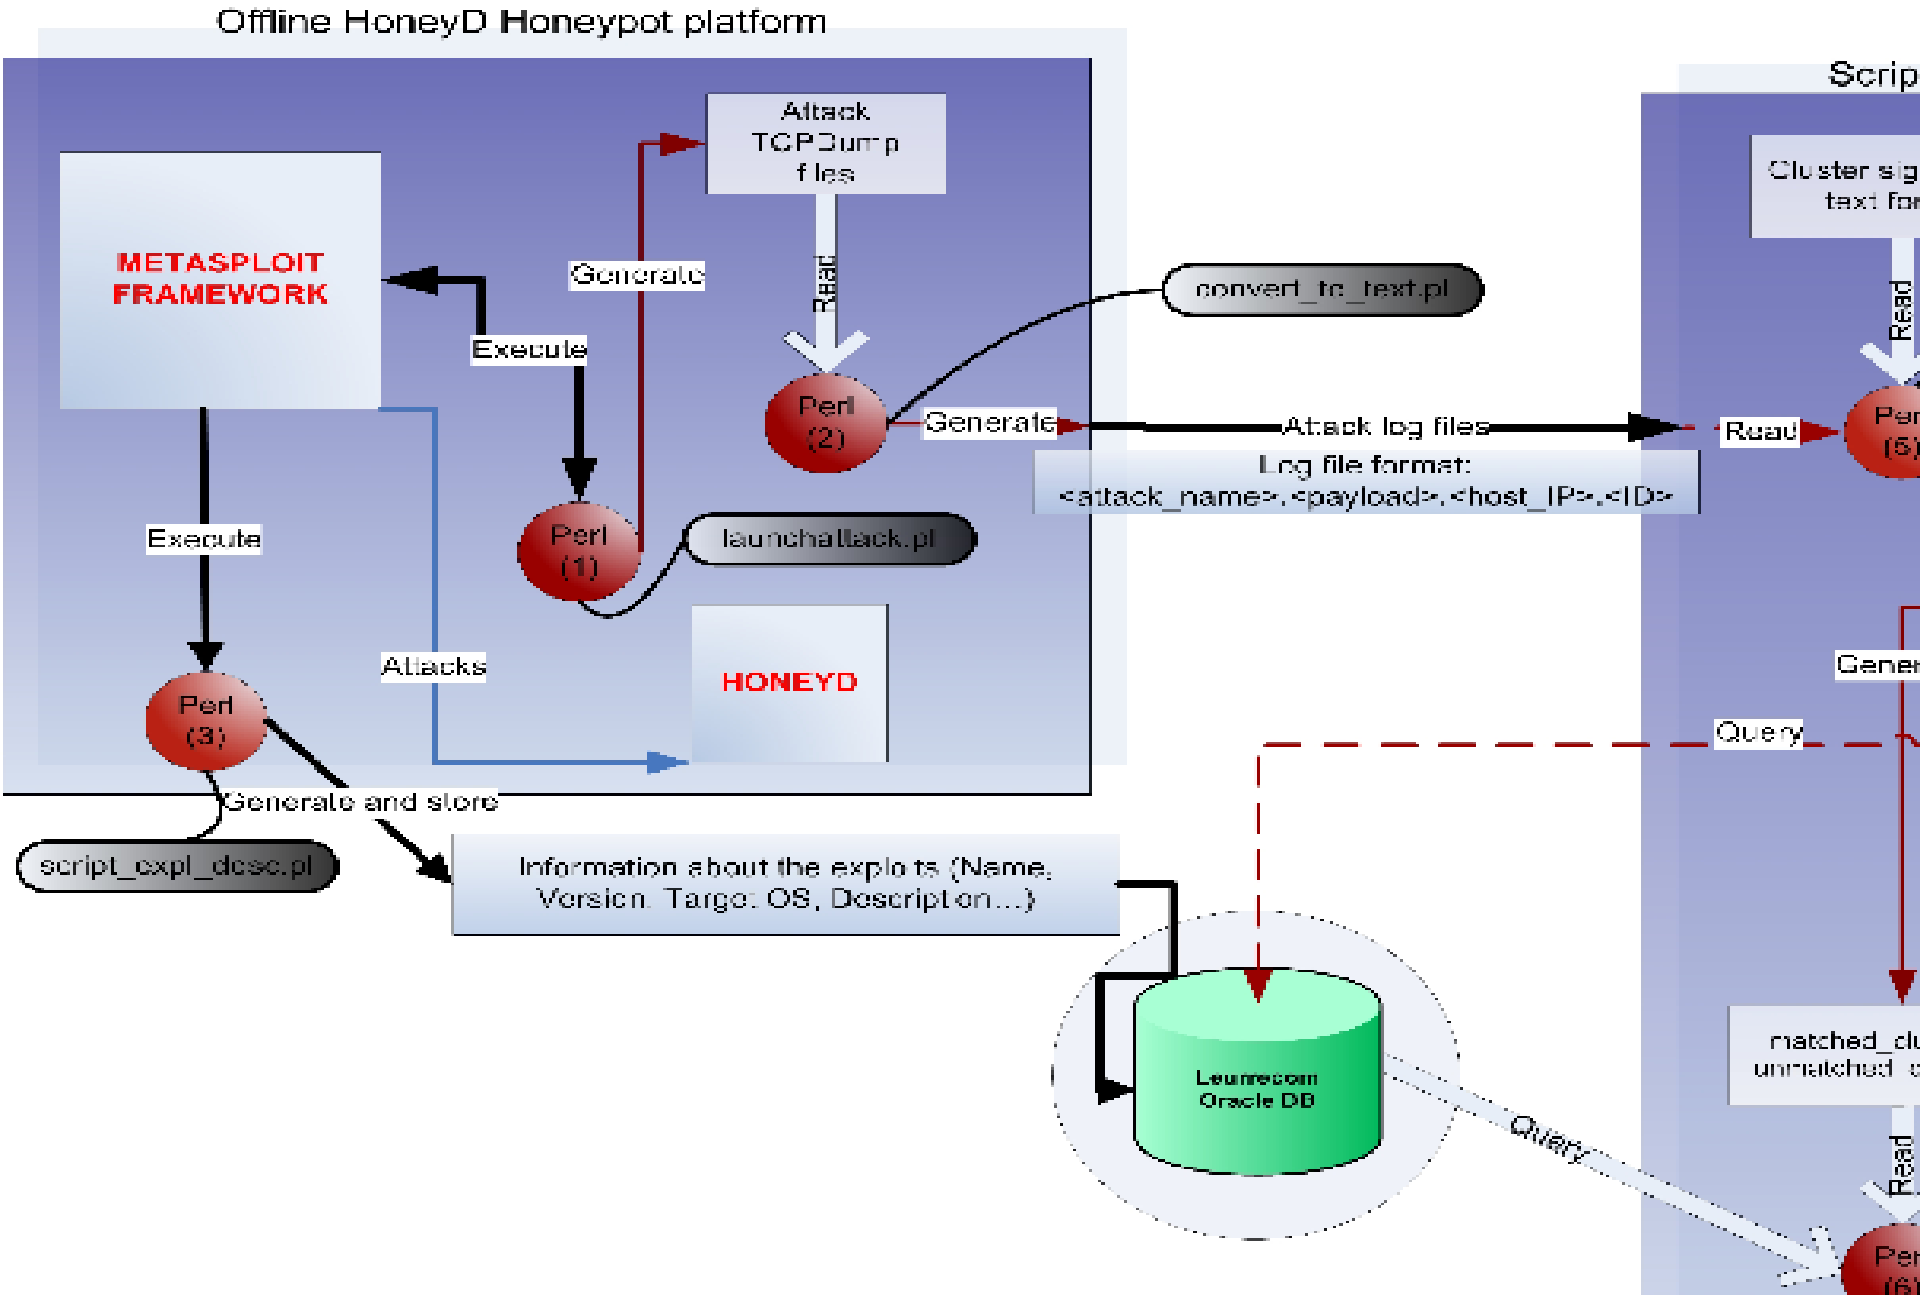

Metasploit signature generation

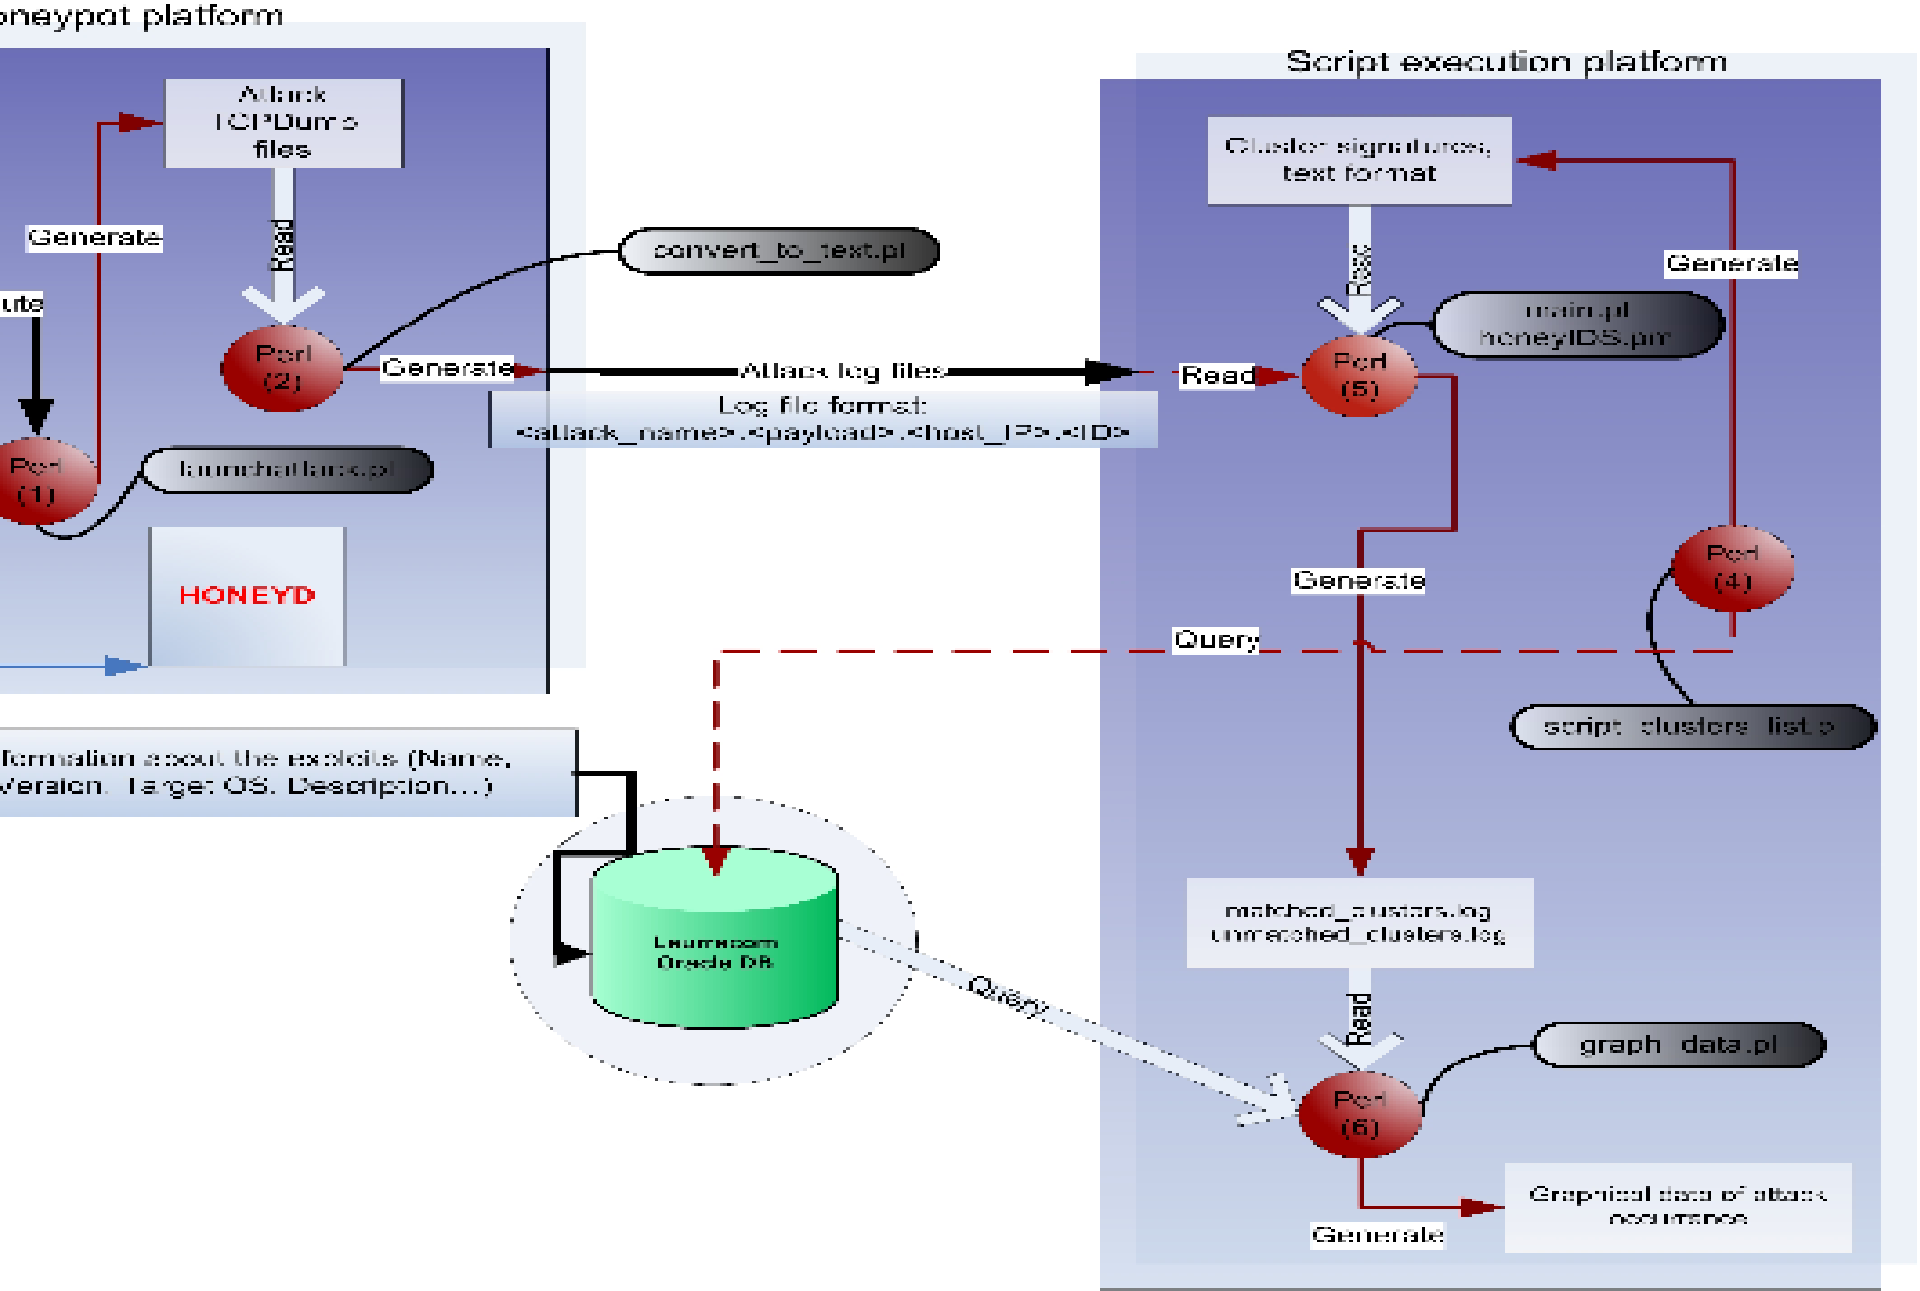

Overview

- Introduction
  - The Leurré.com project
  - Experimental framework
  - Experimental results
- Conclusions

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Initial Selection of Clusters

- 132 Metasploit modules used
- Running all of them in different ways, using various possible options, etc. led to 4000 distinct tpcdump files
- 19000 clusters (out of 150000) had their characteristics matching the ones of at least one of these files
- Clearly, we were selecting more than wanted!

Amount of exploits per cluster

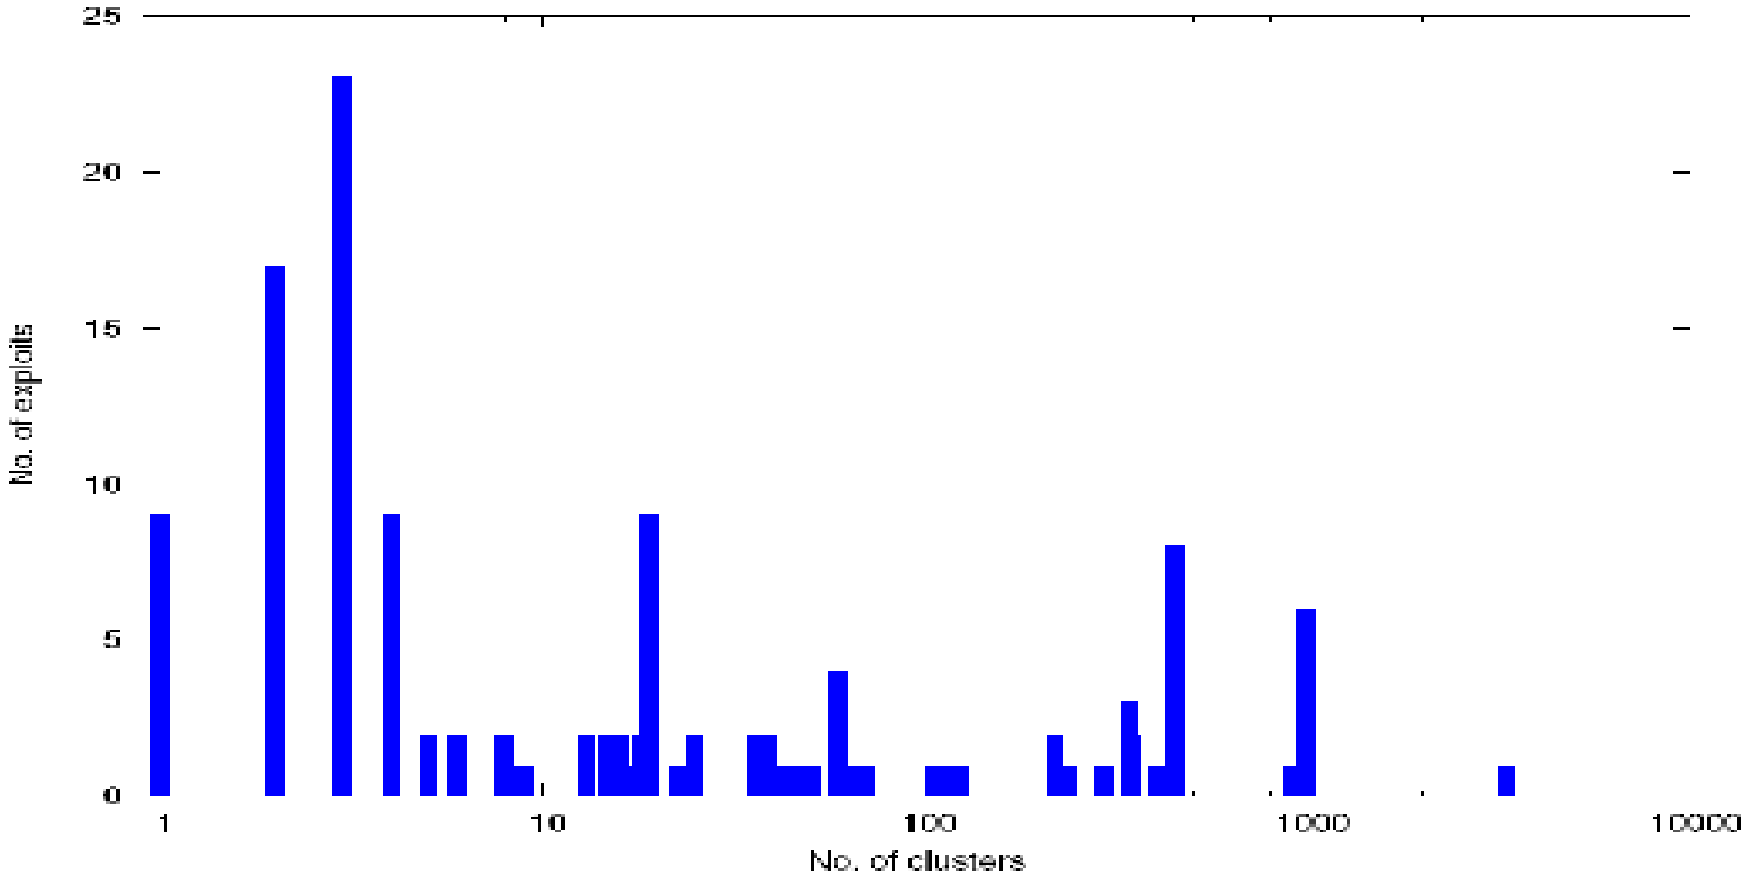

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Finding a few very “good” ones

- Question:
  - How to find which ones, among these 19000, are very likely to be related to a given Metasploit plugin?
- Answer:
  - Select only the clusters that have a substantial peak of activity very close to the plugin release date and no larger peak at any other point in time.

Algorithm 1

- For each of the 19000 selected clusters:
  - obtain the original plugin release date
  - compute the number of attacks, per day, observed for that cluster in the period ±30 days relative to the exploit release day
  - compute average (avg) and standard deviation (std) for the period ±30 days
  - If within a window of ±5 days centered at day 0, we have an activity larger than *avg + 2*std* then select the cluster as a good candidate
- For each candidate, search for its maximal number of attacks over its whole lifetime. Discard the candidate, if this value does not appear within the period ±5 days around day 0.
- Result: 700 clusters remain

Activities around day 0 of original release

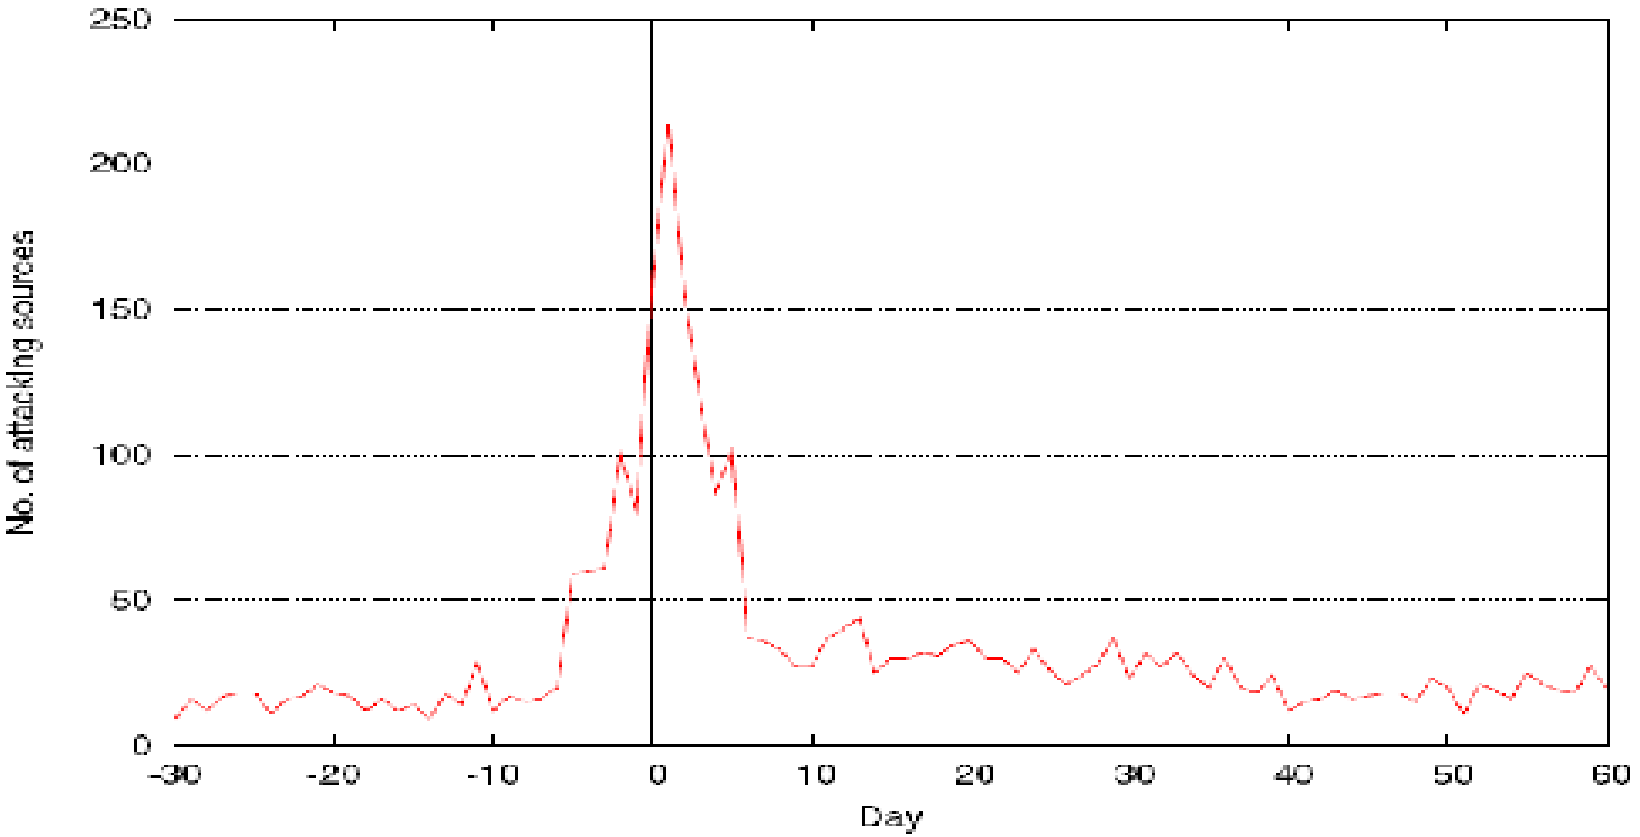

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Refinment

- Question:
  - How can we see if “old plugin” are reused when a new general release of the environment is made public?
- Answer:
  - Repeat the same experiment but consider each release date for all clusters now instead of the sole original plugin release date.
- Result:
  - This leads us to find ***1300 new*** matching clusters

Activities around day 0 of all releases

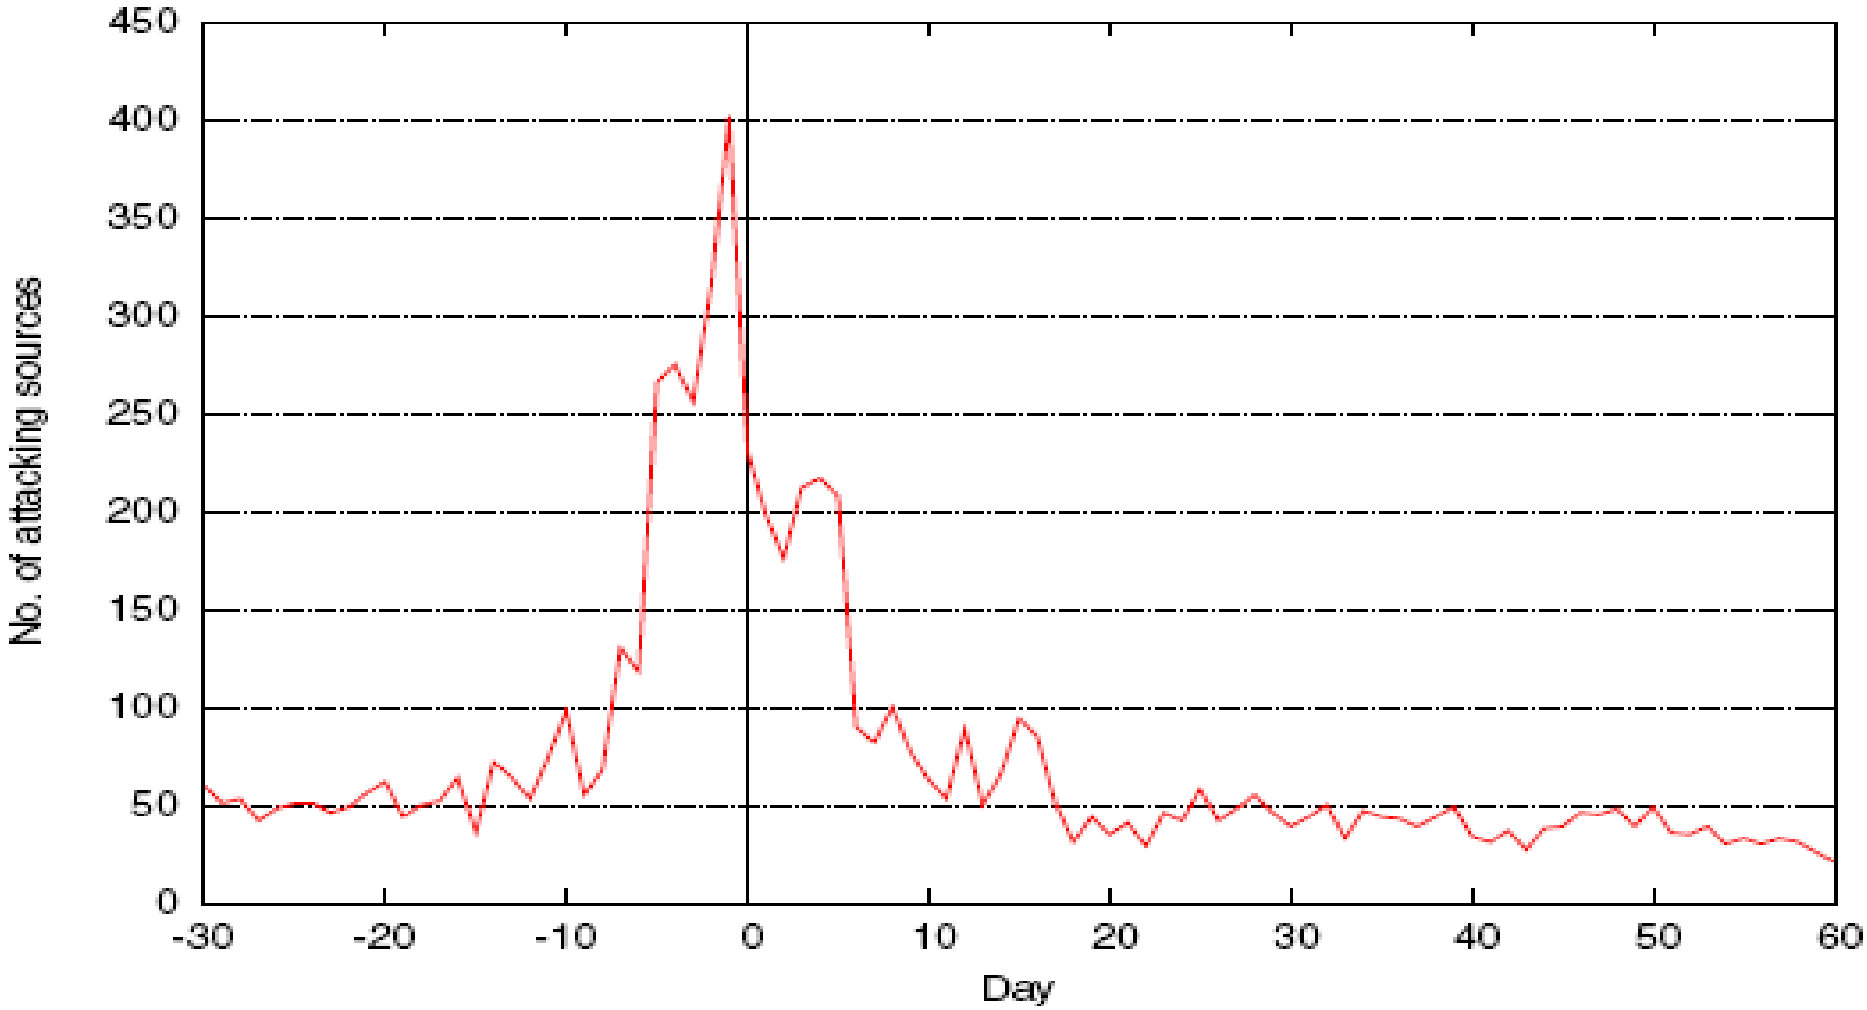

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Analysis of burst at day -2

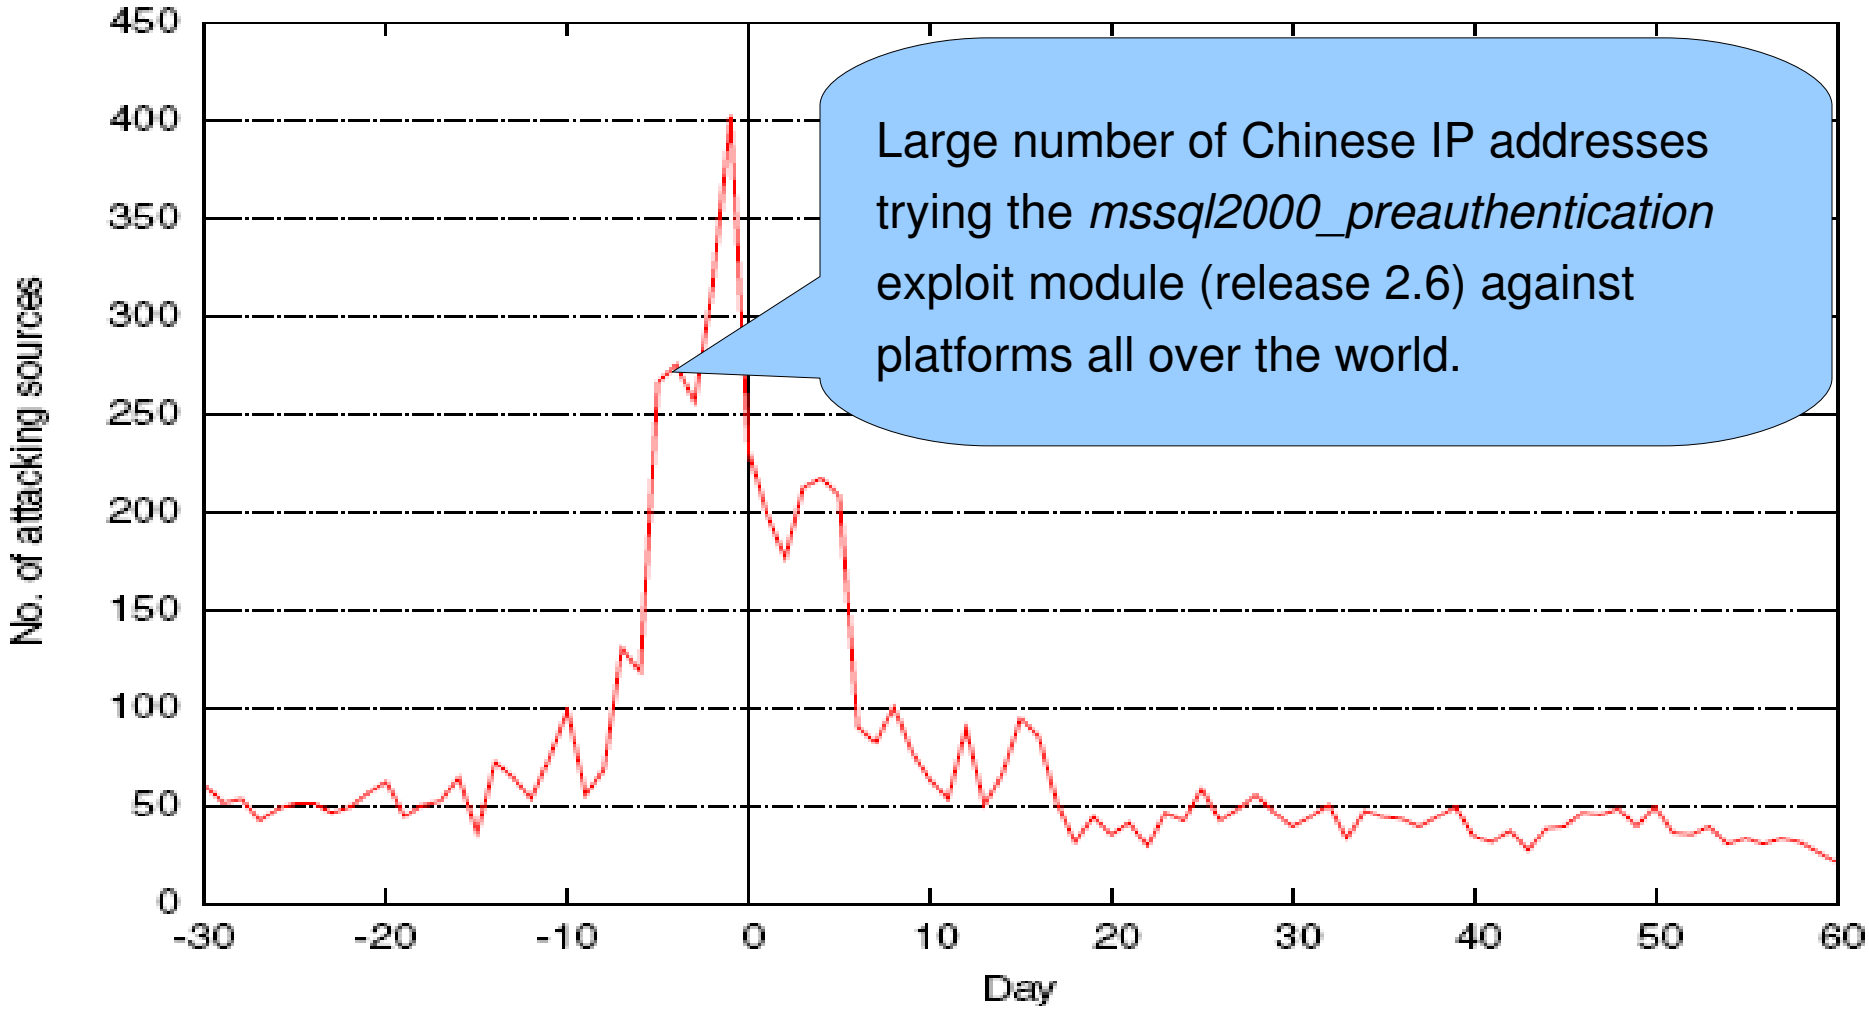

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Analysis of burst at day -1

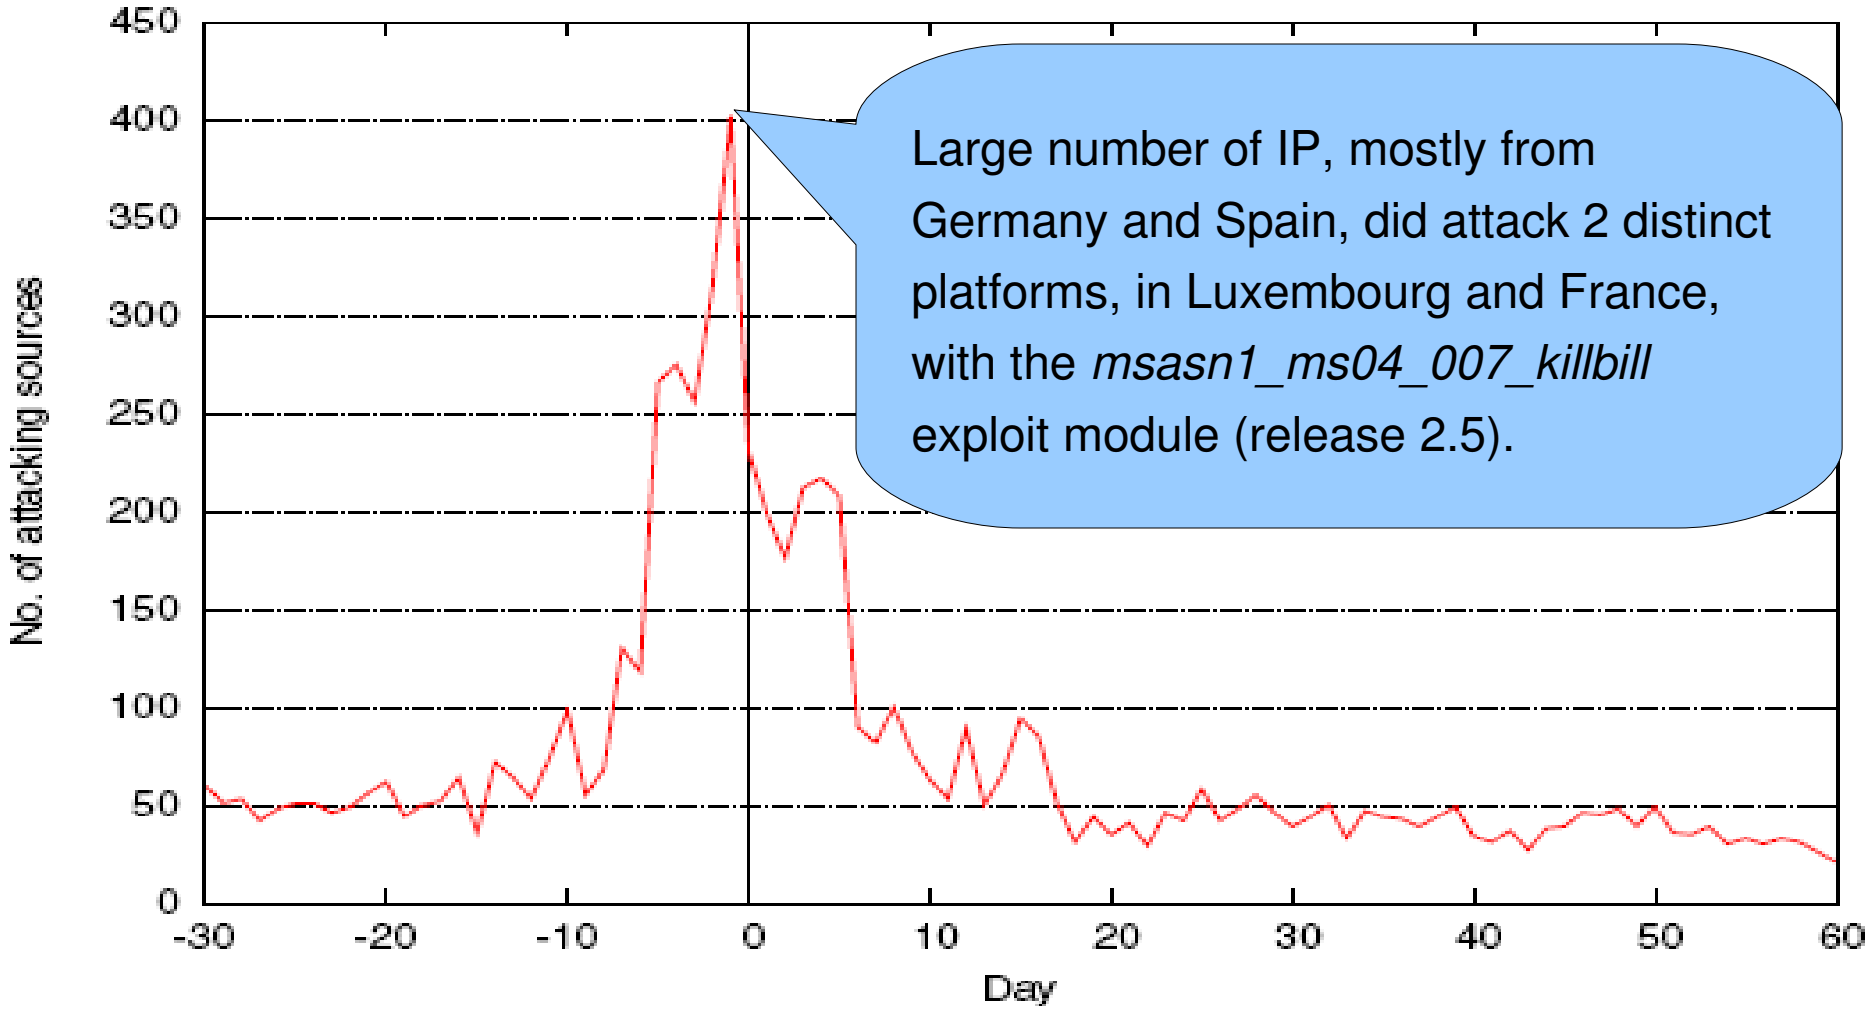

Large number of IP, mostly from Germany and Spain, did attack 2 distinct platforms, in Luxembourg and France, with the *msasn1_ms04_007_killbill* exploit module (release 2.5).

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Sanity Check

- Question:
  - How many good clusters did we lose because of the constraint regarding the maximal peak value around ±5 days?
- Answer:
  - select all clusters which very first manifestation was observed in a window of ±2 days around any of the 8 possible release dates.
- Result:
  - This leads us to find ***80 new*** matching clusters

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Activities of clusters unseen before day-2

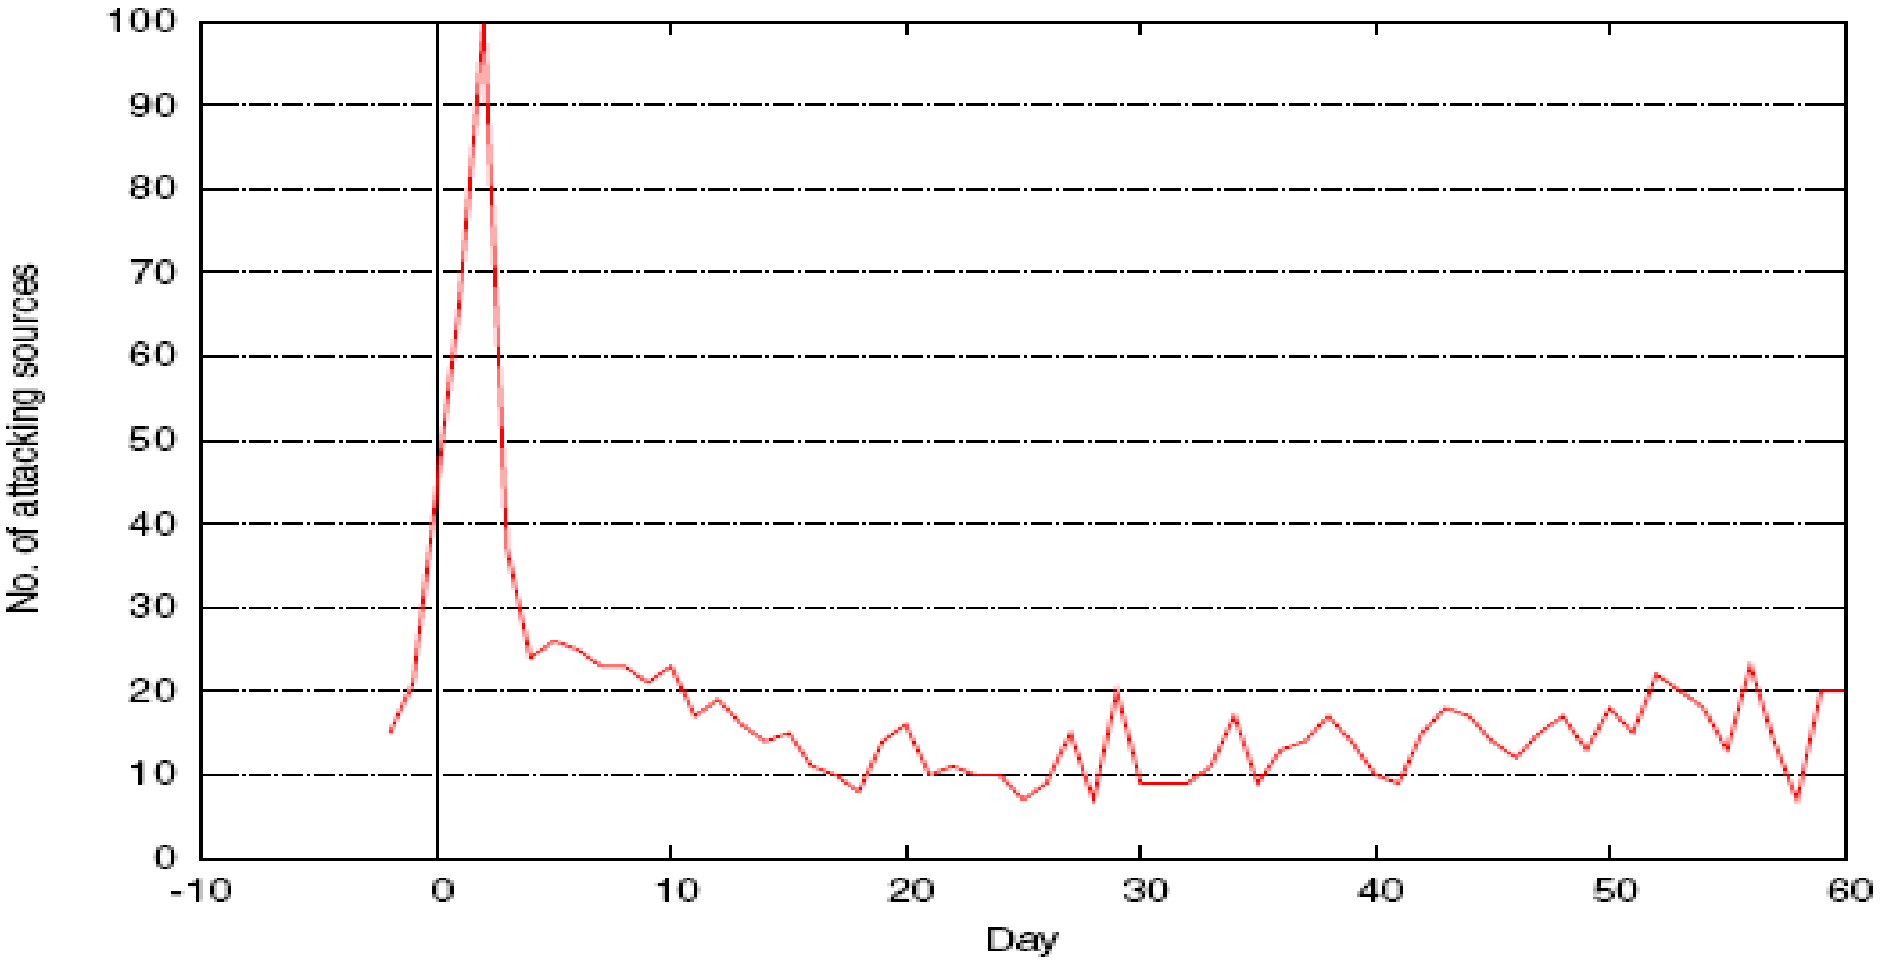

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Summing it up

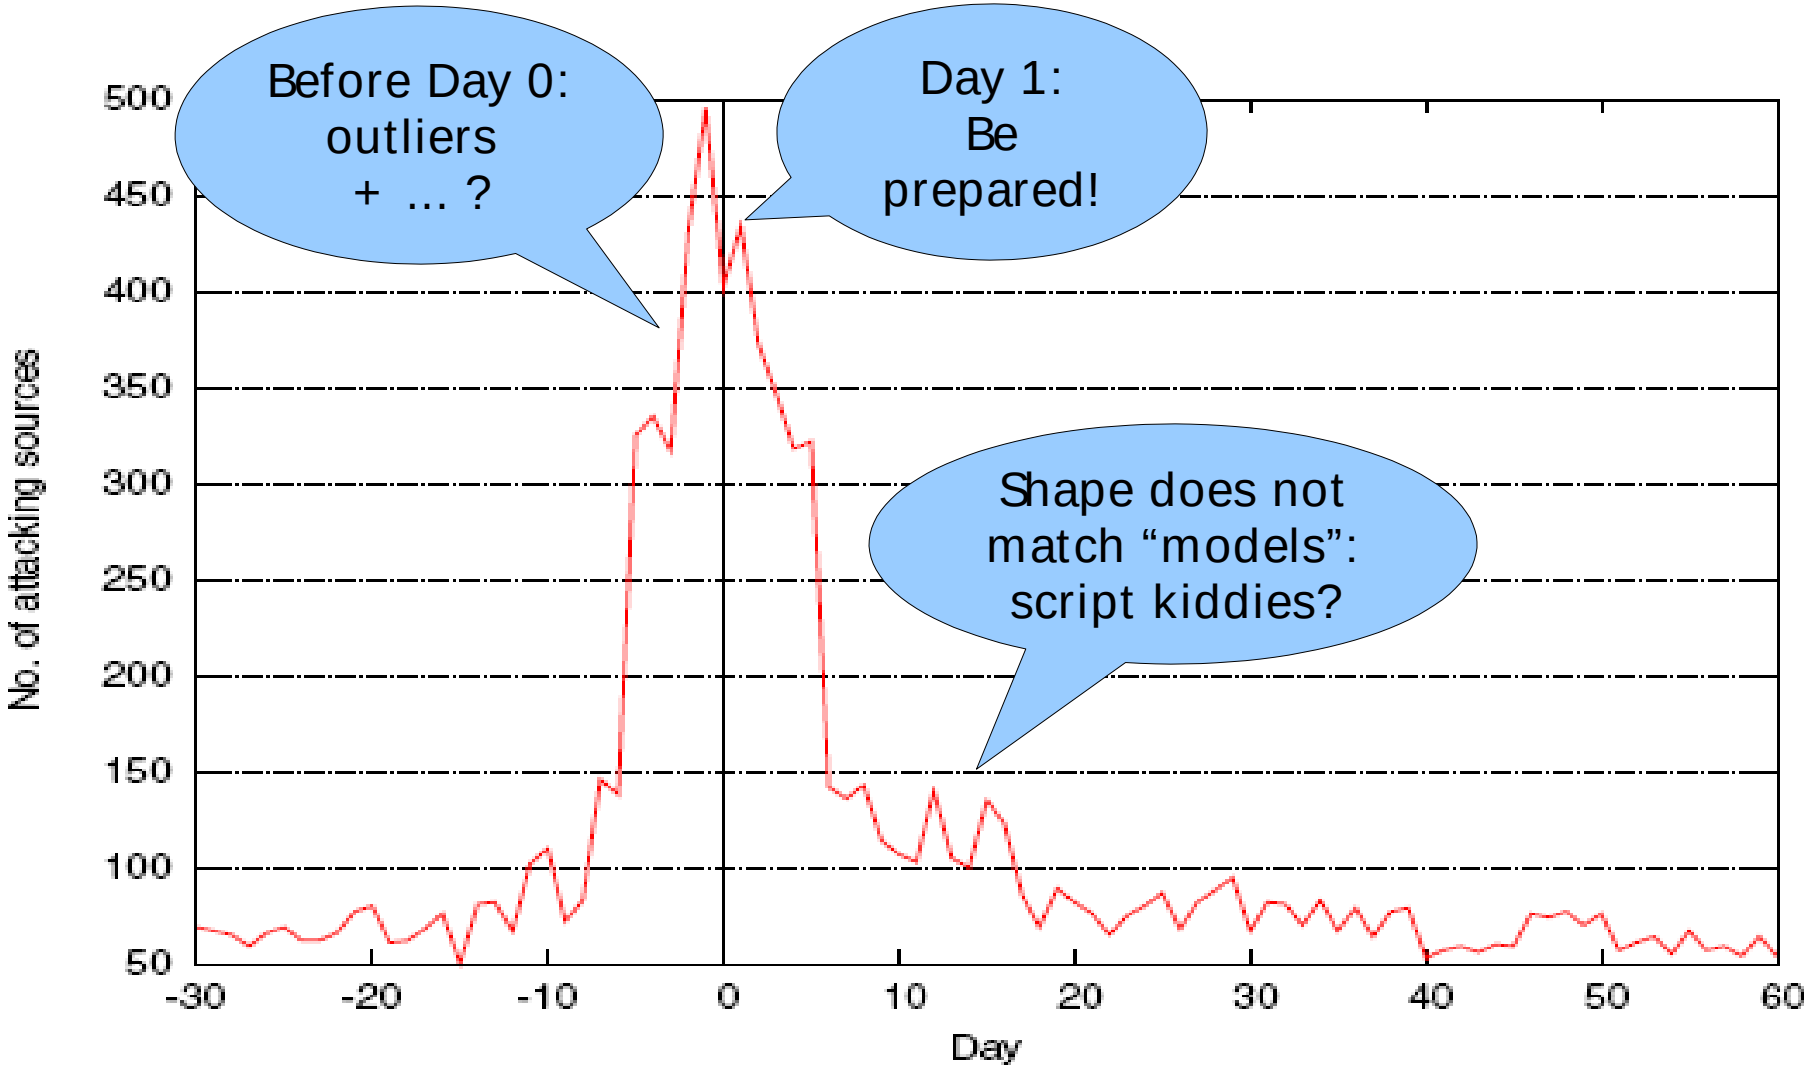

RESIST
ACI Sécurité & Informatique
RÉSEAU NATIONAL DE RECHERCHE EN TÉLÉCOMMUNICATIONS
International Science Linkages

Overview

- Introduction
  - The Leurré.com project
  - Experimental framework
  - Experimental results
- Conclusions

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

Conclusions

- Phenomena linked to Metasploit plugins releases have clearly been identified.
- Their amplitude is limited, as expected since we look at honeypots.
- Their mere existence as well as the shape of the curves tend to indicate that “script kiddies” tools are -also- used by well organized people.
- They are the ones predominantly observed in our dataset.

Further work

- Leurré.com V2.0 is about to be deployed:
  - based on Scriptgen (Eurecom, see ACSAC05, RAID06)
  - Enriched by *Argos (VU Amsterdam), Anubis (TUVienna), Nepenthes (Manheim), Virustotal (Hispasec).*
- It will offer much richer data under the same agreements.
- Downloads shellcode and malware and analyses them.
- You are welcome to participate.

ACI Sécurité & Informatique
RNRT Réseau National de Recherche en Télécommunications
International Science Linkages

WOMBAT

- A 3 years EC funded research project (STREP)
- Starts on January 1st 2008
- Involves 11 partners
  - 3 industrial partners
  - 1 CERT
  - 5 academic partners (VU Amsterdam, Eurécom, FORTH, Politecnico Milano, TU Vienna)
  - 2 non EC partners

WOMBAT: technical tasks

- Task 1: Federation of of malware collection techniques (existing and new ones such as Leurrecom honeyclient, wireless, etc..)
- Task
- Task

April 2008

by invitation Workshop for attack-related data producers/consumers.

Contact me if interested

dacier@eurecom.fr