Politecnico di Milano
Dip. Elettronica e Informazione
Milano, Italy

Observing the tidal waves of malware

Stefano Zanero, Ph.D.
Post-doc Researcher, Politecnico di Milano
CTO & Founder, Secure Network S.r.l.

Black Hat Briefings – Las Vegas NV, 01/08/2007

Presentation Outline

- A need for observing what is happening around us
  - Why do we need to do it
  - How do we need to do it
- Infrastructures we have
  - And their limitations
- Software we have
  - And their limitation
- The Great and Cunning Plan (TM)
  - Open to your critique and collaboration
- Conclusions and an awful lot of future work !

Caveat auditor

- Beware, listener, that this presentation includes forward looking statements that may be exaggerated, not quite correct or blatant lies. Additionally, it mostly deals with the presentation of a project which has yet to start, and may miserably fail before I even end speaking.
- Not really, but still most of what I will say is still in its infancy, not even under development. Any objections of “but this is a TODO presentation” will result in the phisical termination of the objector.
- Thanks to Jeff and Dominique for evaluating this talk positively even if I didn't know yet how much I could share of it; and for evaluating it though it was way late

Knowledge: granting success, since ~500 b.C.

- Knowing your enemy is the key to success
  - *"He will win who knows when to fight and when not to fight... He will win who, prepared himself, waits to take the enemy unprepared. Hence the saying: If you know the enemy and know yourself, you need not fear the result of a hundred battles. If you know yourself but not the enemy, for every victory gained you will also suffer a defeat. If you know neither the enemy nor yourself, you will succumb in every battle."* [Sun-Tsu]
- Perhaps the most often quoted, and less often practiced, sentence in history
- Understanding is the key to (re)acting *sensibly*, and we are failing in a lot of fields, notably anti-terrorism controls in the airports

Lies, Damn Lies, and Statistics – part 1

- "Asymmetric warfare potential of cyberspace will lead to an increase in electronic warfare and cyberterrorism". True or False ?
  - Repeated countless times, since 9/11/01 (at least)
  - "If we ever manage to get real-world terrorists to blow up computers instead of airplanes, it will be at our advantage, as computers have backups and humans don't" (R. Power, CSI)
  - No one has data to confirm or disconfirm cyberterrorism activities, also because there's no or little distinctive features of cyberterrorism from common cyberattacks
  - Someone says "there's data, but it's classified/top secret". My very humble opinion is that it's TS BS

Lies, Damn Lies, and Statistics – part 2

- FBI – CSI report: “croce e delizia”
- There is always a "rising wave of Internet crime"
- Reports of losses usually out of thin air
- Reports based on respondent's honesty and knowledge (“I have no intrusion detection process”, so how do you know?)
- Q: Why reported incident losses fall every year ?
- A: Because the numbers are not statistically solid
- From the CSI Alert Newsletter (quoted by A. Chuvakin)
- 5,000 members of CSI surveyed (they are not a representative set). Response rate 12% (616 of 5000). We do not know any statistics on these 12% and their dissimilarity to the others.

Lies, Damn Lies, and Statistics – part 3

- Prediction anonymized and mixed up to protect the innocent and clueless analysts out there
- "In July 2001, Code Red spread to $HUGE_INT systems within $SMALL_INT hours; the worldwide economic impact was estimated to be $INSANE_FIGURE billions. SQL Slammer was even faster.
- "We'll see an even greater increase in the speed and destructive capabilities of threats."
  - Warhol Worms, Flash worms, etc
  - Extremely good academic papers, but never incarnated

And by the way… where are the worms ?!

- We *all* thought that the Internet would get wormier
  - Don't try to deny it: I am sure you have AT LEAST one slide where you said that!
- The trend was clear:
  - 2001: Li0n, Code Red, Nimda
  - 2002: Slapper, Klez
  - 2003: SQL Slammer, Blaster, SoBig
  - 2004: Sober, MyDoom, Witty, Sasser
  - I have even an iDefense t-shirt with this list on it!
- Since then, silence on the wires. No new "major" worm outbreaks
  - Weaponizable vulns were there, we even collectively braced for impact a couple of times
  - Did we get *so better* at defending networks? I bet "not"

Rise of the Bots

- Bots, bots everywhere
  - When I was a youngster, bots were IRC warriors' stuff (~1999-2000)
  - We used to call remote control trojans “zombies”, and they were usually DDoS tools (2000-2)
- Today's bots are different
  - Intelligent, evolving, with complex C&C infrastructures, difficult to remove as well
  - Larger botnets (10k common, 1M+ seen)
  - Phishing, spamming and pharming bots... more difficult to track than DDoS events
- How do we track them? How do we analyze them?
  - Worm explosive propagation vs. bot slow and steady diffusion: there's no network telescope that can see them

Open wormy questions: example

- Why no worm has ever targeted the infrastructure?
  - (possible exception of Witty, targeting firewalls)
- Possible explanation: routers and the like are a difficult vector to exploit
  - Not really true anymore, see FX's and Michael Lynn's works
  - Can use a traditional worm for propagation + a specialized payload for infrastructure damage
  - Windows of opportunity were there:
    - June 2003: MS03-026, RPC-DCOM Vulnerability (Blaster) + Cisco IOS Interface Blocked by IPv4 Packets
    - April 2004: MS04-011, LSASS Vulnerability (Sasser) + TCP Vulnerabilities in Multiple IOS-Based Cisco Products (resets)
- So why, oh why, the /bin/ladens of the world were not there, grinning and reaping?

He who knows not the enemy, nor himself

- Summary of the worm rise and fall:
  - Most folks and consultants were clueless about worms in 2000 (lost preparing for the 2-digits-years cataclism)
  - Since 2004 lots of money and consultant-speak in the direction of fighting "the dreadful and impending Big One of the flash worms"
  - The era of the worms was actually almost over already
- The result
  - Not the disappearance of worms
  - Nor an improved resilience to them (infrastructure is just as exposed to a flash worm today as it was in 2004)
  - A mass distraction of resources from the real, impending threats (endpoint security and prevention of client-side attacks and botnets)
- "...every battle is a certain risk"

Observing attacks != Knowing attackers

- Various questions about the attackers
  - Attribution (tipically for law enforcement)
  - Characterization aka profiling
- Usually observation of attacks is not enough to answer such questions
  - In particular, characterization of attackers is still in its infancy
  - See www.ratingthehacker.net for an example of characterization based on the attacks
  - There are also various hacker profiling projects, but in most cases they are linked either to criminal case review or to dissemination of questionaires
  - The efficacy is highly debatable, to be honest

The need is felt also at political levels

- EU Commissioner Vivianne Reding recently stressed how difficult it is for decision-makers to create appropriate policies for fighting cybercrime without reliable data, models and theories on the root causes and the underlying generative processes of the tidal wave
- Testimonies in front of the House Committee on Homeland Security: Doug Maughan, Sami Saydjari, Daniel Geer: better sharing and analysis mechanisms needed
- DHS investments in Information Sharing & Analysis Centers (ISACs)
- National Strategy to Secure Cyberspace (NSSC) has 3 out of 8 action items related to log sharing

Today's observation points

- Efforts by vendors
  - ATLAS (Arbor)
  - DeepSight (Symantec, formerly SecurityFocus)
- Community and no-profit efforts
  - Dshield and the Internet Storm Center (SANS)
  - Network Telescope
  - The HoneyNet project
  - NoAH and Leurrecom projects

ATLAS

- Draws data from Arbor platforms which claim to monitor “70% of the Internet”
- Uses the unused address spaces as darknets
- The ATLAS portal is public: atlas.arbor.net
- Geolocation of attacks, top sources, top exploits etc.
- Data from multiple sources
  - Honeypot-captured payloads & malware samples, IDS logs, Scan logs, DoS logs, News & vulnerability reports
- ASERT analyzes data
- Alerts are pushed to customers and platforms
- Underlying technology and capabilities are proprietary and secret

DeepSight

- Symantec DeepSight Threat Management System consists of 40,000+ sensors in more than 180 countries
- Adds malicious code data along with spyware and adware reports from 120M+ client, server, and gateways
- Provides analysis capabilities to Symantec labs, and delivers reports and alerts to customers
- Commercial, therefore not (broadly) open to research community
- Underlying technology and capabilities are proprietary and secret

Other statistics are made (up) by vendors

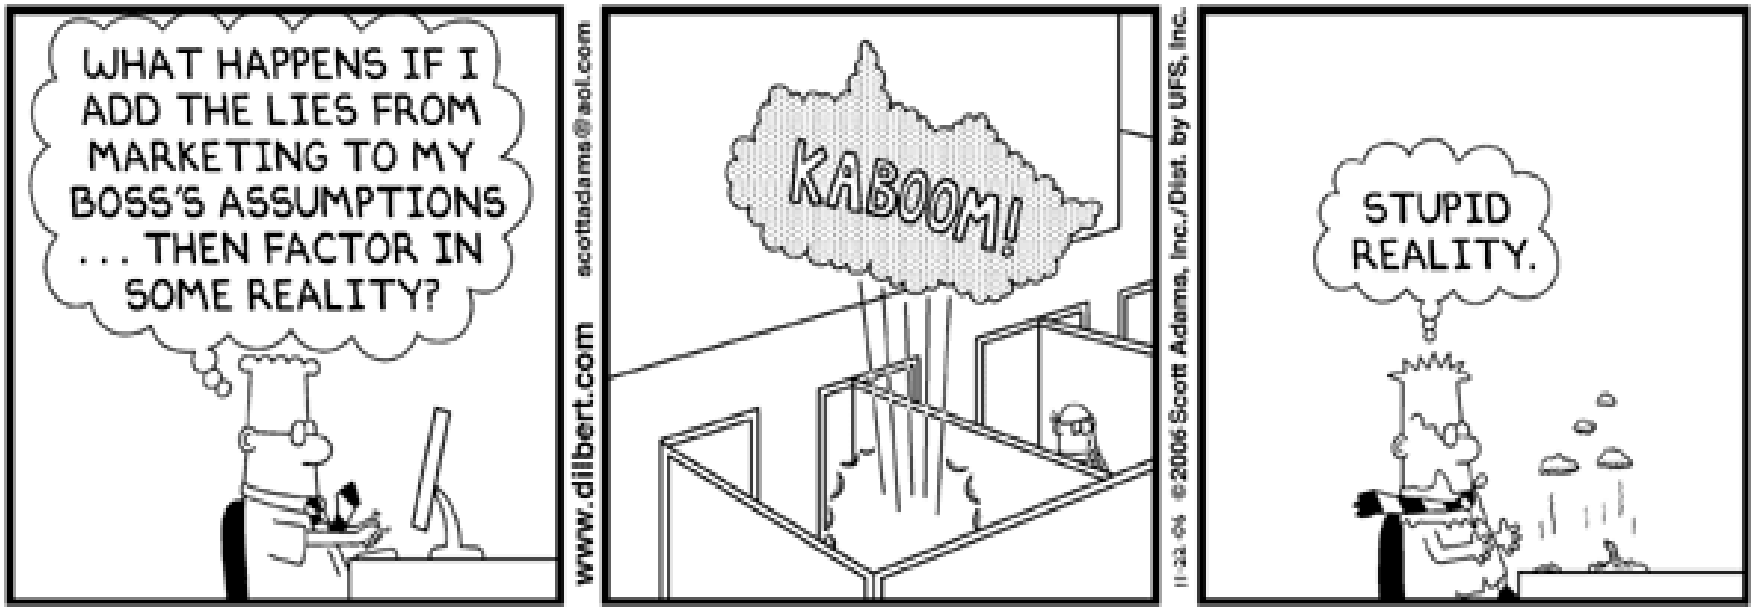

- "*** Report: Surge in Viruses and Worms Targeting Mobile Devices, Satellite Communications Anticipated in 2005"
- ... hell-looooooooo ? It's 2007... where are youuuuu ? :)

The Internet Storm Center

- Managed by the SANS institute
- Uses Dshield data
- Tens of millions of log entries received daily
- Volunteer incident handlers analyze detected problems and anomalies, then post a daily diary of analysis
- "Storm center": gathering data from thousands of small sources into a meaningful picture
- Raw TCP/UDP packets, dumps, IDS logs mean little by themselves, even if they are "a lot": the value here is the experience of the handlers (kudos)
- Arguably, the best experience of its kind
- Early warning potential

The ISC Process (as usually explained)

Data Collection

Analysis

Dissemination

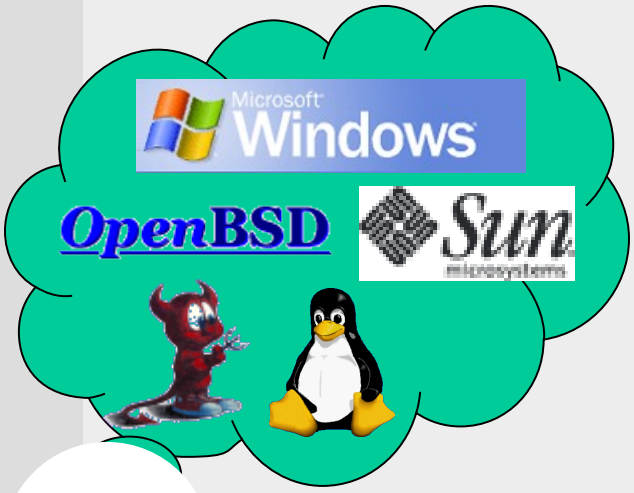

DShield Users

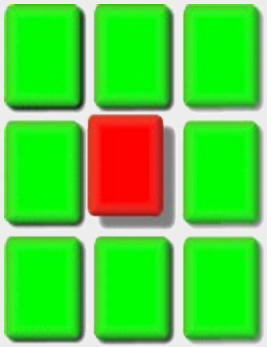

DShield.org

INTERNET STORM CENTER

Service Name	Port Number	Activity Past Month	Explanation
microsoft-ds	445		Win2k+ Server Message Block
epmap	135		DCE endpoint resolution
---	20525		
netbios-ssn	139		NETBIOS Session Service
---	1026		
icq	1027		icq instant messanger
---	1025		
www	80		World Wide Web HTTP
domain	53		Domain Name Server
netbios-ns	137		NETBIOS Name Service

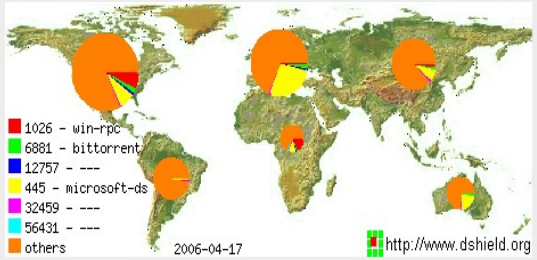

DShield PARTNERS

Telescope, blackholes and darknets

- Substantially similar
  - A telescope/blackhole is a large routed but unused address space
  - Darknets are unused address portions in an otherwise used network
- Traffic is the result of DdoS backscatter, worms, autorooters, mass scanners, or other banes
- A number of initiatives (all separated... :-( )
  - iSink,Team Cymru monitoring projects, CAIDA Telescope, IUCC/IDC Internet Telescope
- Internet Motion Sensor: a coordinated network of telescopes complemented with non-passive components (http://ims.eecs.umich.edu)
  - Initial /8 deployment in 2001. In 2005 60 address blocks at 18 networks on 3 continents

Limitations of ISC (and similar initiatives)

- Also because of privacy issues, raw data cannot be shared outside the handlers
- Just basic statistics about global current threats (e.g. hits per port, hits of specific malware as detected by an IDS, etc.).
- Uncontrolled sources: datasets contain also false positives, non-attacks, etc.
- Handlers are humans (exceptions in the direction of "demigod" may apply). While excellently skilled, this is a limitation for "early warning" capabilities
- A feeling that the collected data is just "not enough" for root cause analysis
  - How many times do we see the handlers manually asking for submission of some captures?

Example of data analysis gone awry (1)

- July 4 2007: some researchers (no url provided as no bashing intended) note a "deviation in global network traffic"
  - "Normally, global Internet traffic (as observed by the Internet Traffic Report) oscillates around 9% packet loss, with global response times of 138 ms. . . over the last 24 hours . . . packet loss has climbed to 11%, and the global response time to almost 150 ms. . . . When the figures are considered against the 7 day average, and the 30 day average, the deviation appears to be quite significant and seems to mark a distinct event or set of events"
  - They also note a geographical distribution of the deviation, and conclude that "either these regions are experiencing the first stages of a global event, or they contain networks that are under a sustained attack for some specific reason."

Example of data analysis gone awry (2)

- They also noticed that DShield was reporting a spike on Port 5901 (VNC)
  - An exploit supposedly targeting VNC was distributed earlier (actually it was against a VNC ActiveX control)
  - They concluded that VNC was probably the culprit
  - Post hoc ergo propter hoc
- ISC quickly downplayed the significance of the VNC spike
- Jose Nazario through ATLAS showed that most of the correlations sought between VNC attacks and loss of connectivity were just not there
- We don't know what happened, or if something happened, but definitely it wasn't VNC-related
  - What if we somehow reacted?

Other random examples

- July 24, Deborah Hale (ISC handler) observes a spike on port 57886 and asks readers for submissions
- On july 4, a spike is seen on port 1433 (MSSQL) and 5901, which is manually linked (by a reader) to the "ya bot" source code released one month before
- As a general rule, the diaries are much more effective at disseminating knowledge, raising attention to patches or disclosures, etc.

Project Honeynet

- One of the first and most successful "know-your-enemy" organized efforts
  - Kudos to Lance Spitzner and all the teams around the world
- Great insights gained through effort
  - In the form of books, so usually a recollection of forensic analysis
  - Scan of the month are a great teaching material for the academics among us :)
- Development of honeypot tools and tactics
  - Honeyd, sebek, web interfaces, etc.
- Not really tied together or usable for early warning
- Extremely dependent on the skills and the dedication of the volunteers running the honeypots

Today's (and tomorrow's) honeytools

- Honeyd (obviously !)
- ScriptGen
- Argos sensors
- Nepenthes
- MwCollect
- (there's a plethora of others, I won't have time to touch all of them)

Honeyd

- Simplest and most popular low-interaction honeypot
- Can monitor huge address spaces and create huge fake honeynets
  - up to 65k simulated hosts... in the real world!
  - Using arpd, darknets can be monitored
- Based on scripts that statefully emulate the various services listening to remote requests
  - Similar but stateless/high performance for ISP pipes: HoneyTank, iSink ActiveSink
- Writing a script = tedious task, impossible for undocumented proprietary protocol
  - For this reason, ScriptGen was invented

ScriptGen

- Autogenerate scripts that emulate a service
  - Impossible, a reverse engineer's wet dream :)
- Autogenerate scripts that emulate the answers of a service to a deterministic script (the exploit)
  - Far simpler
- Three steps approach
  - A real machine answers traffic, and a tcpdump is recorded
  - If the machine gets compromised, usual cleanup
  - Messages are analyzed and a state machine is derived, representing requests and replies
    - Using bioinformatics techniques from http://www.insidiae.com/PI
  - A honeyd script is produced from the state machine
- Similar effort: honeybee

MWCollect / Nepenthes

- (now the same thing) tool that collects malware
- Aka “medium interaction honeypot”
- Emulates vulnerable services, and analyzes malicious payloads to identify URLs
  - Provides a virtualized filesystem and a virtualized shell to allow the exploit to run harmlessly
  - Emulates **specific** vulnerabilities, in modules
  - Does not need to **look** for the payload, it knows where it is
- Downloads and stores the malicious software
- MwCollectAlliance for deploying nepenthes and collecting the results
- Honeytrap: similar concept with FTP/TFTP clients as well

Argos high-interaction honeypots

- Argos: HIH that extends Qemu to detect exploits via taint analysis
- Core idea: identify when code that came from the network is executed
  - Untrusted data is tagged and an alert is generated (only) if and when it is executed
  - Can tag zero-days!
  - Used for IPS already (Minos: hw-oriented, cannot track back to the exploit; Vigilante: sw-oriented, per-process, does not work on kernel exploits)
- Argos supports multiple guest operating systems including Linux, Windows 2000 and Windows XP
- Also automagically extracts exploit signatures which are then refined globally with SweetBait
  - Honeycomb signatures can be refined as well

Leurré.com

- www.leurrecom.org, project operated by Institut Eurécom (Sophia-Antipolis, France)
- Broad network of honeypots covering more than 30 countries
- Architecture of distributed **low-interaction** honeypots and a central server, using ScriptGen
- All traces captured on each platform are uploaded on a daily basis into a centralized relational database
- All project partners can access the whole database. Simple queries are open also to the outside

Sample results

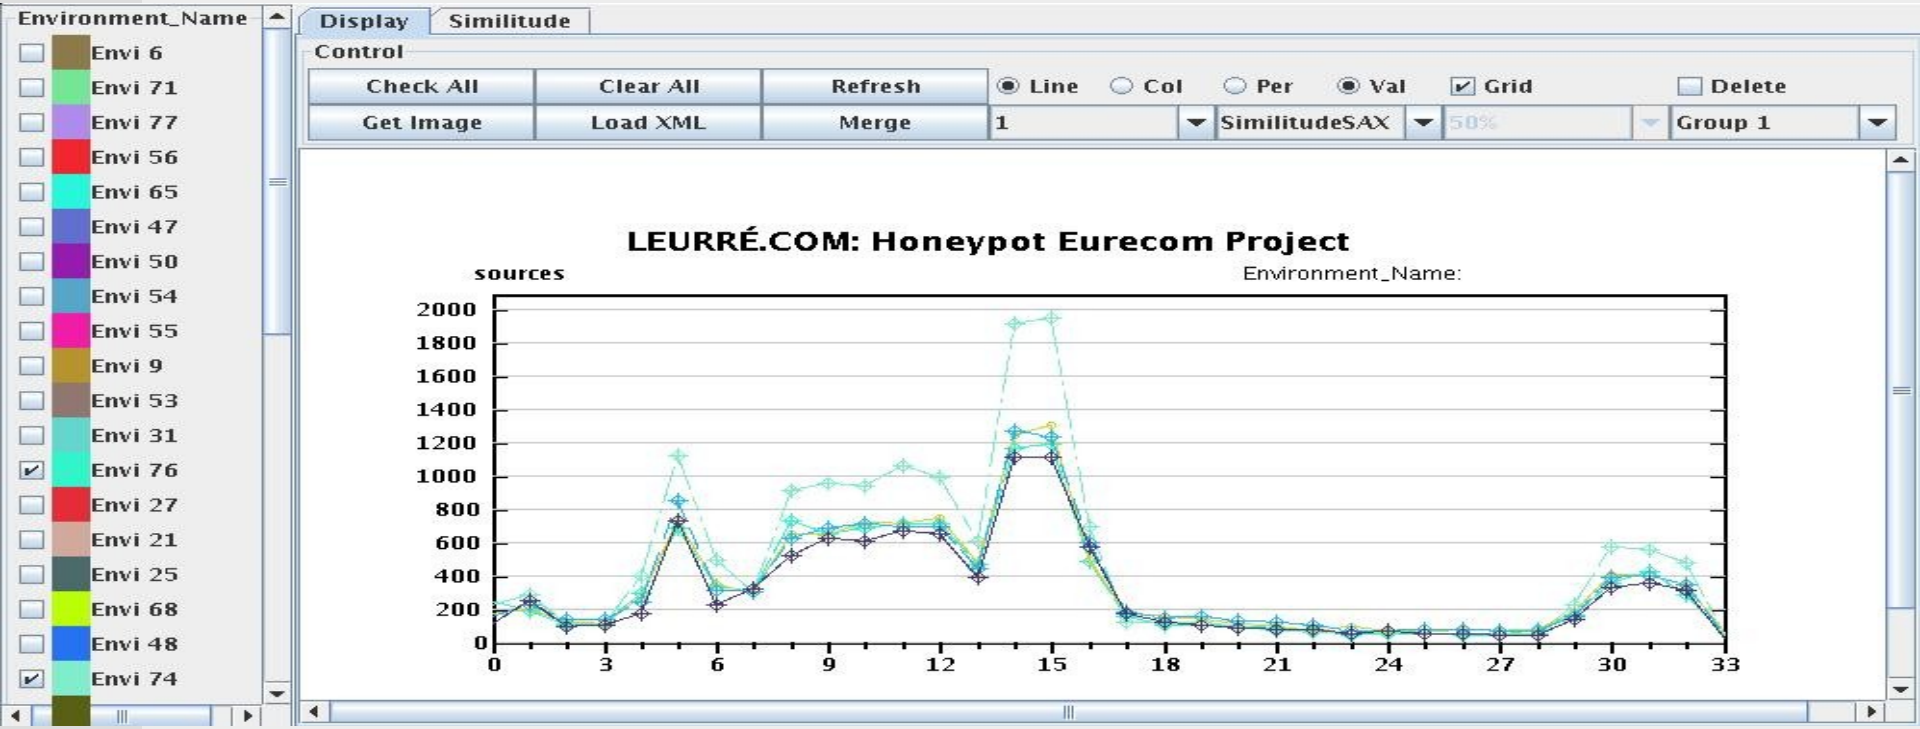

- Groups of platforms sharing the same attack profile
- Algorithm which discovers these cliques automatically

Sample results

- Data:
  - Feb 1, 2005 until Feb 1, 2007
  - Backscatters only
  - Grouped
    - by Country of origin
    - by Platform
- Small influence but
- Viewpoint matters !

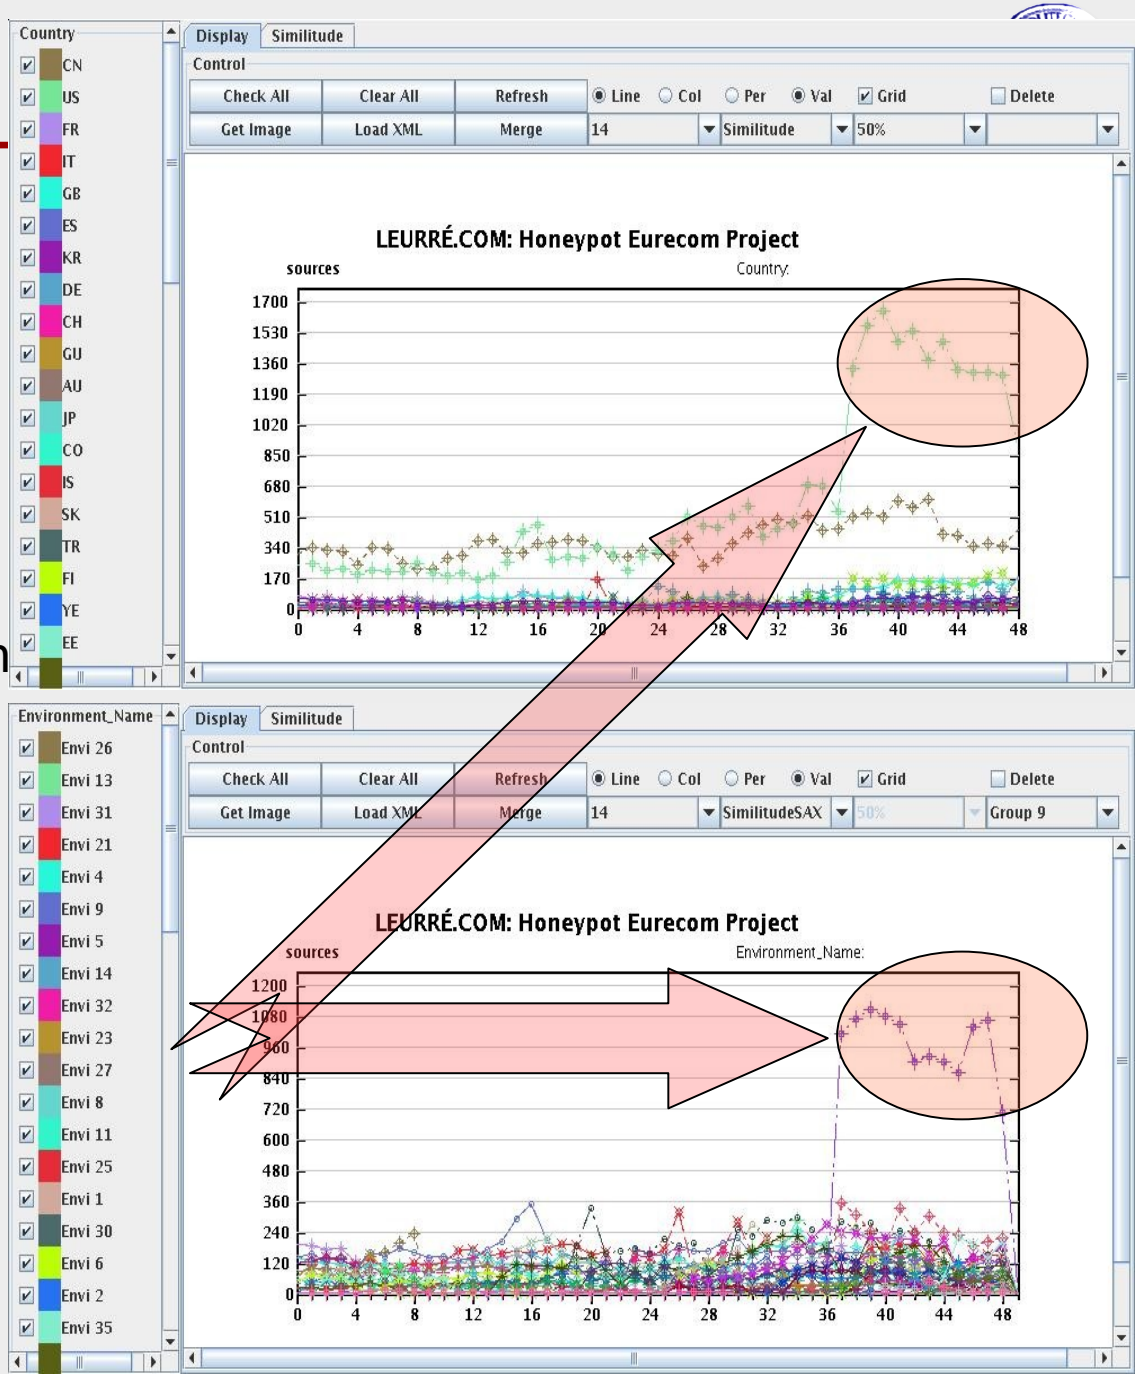

Sample results

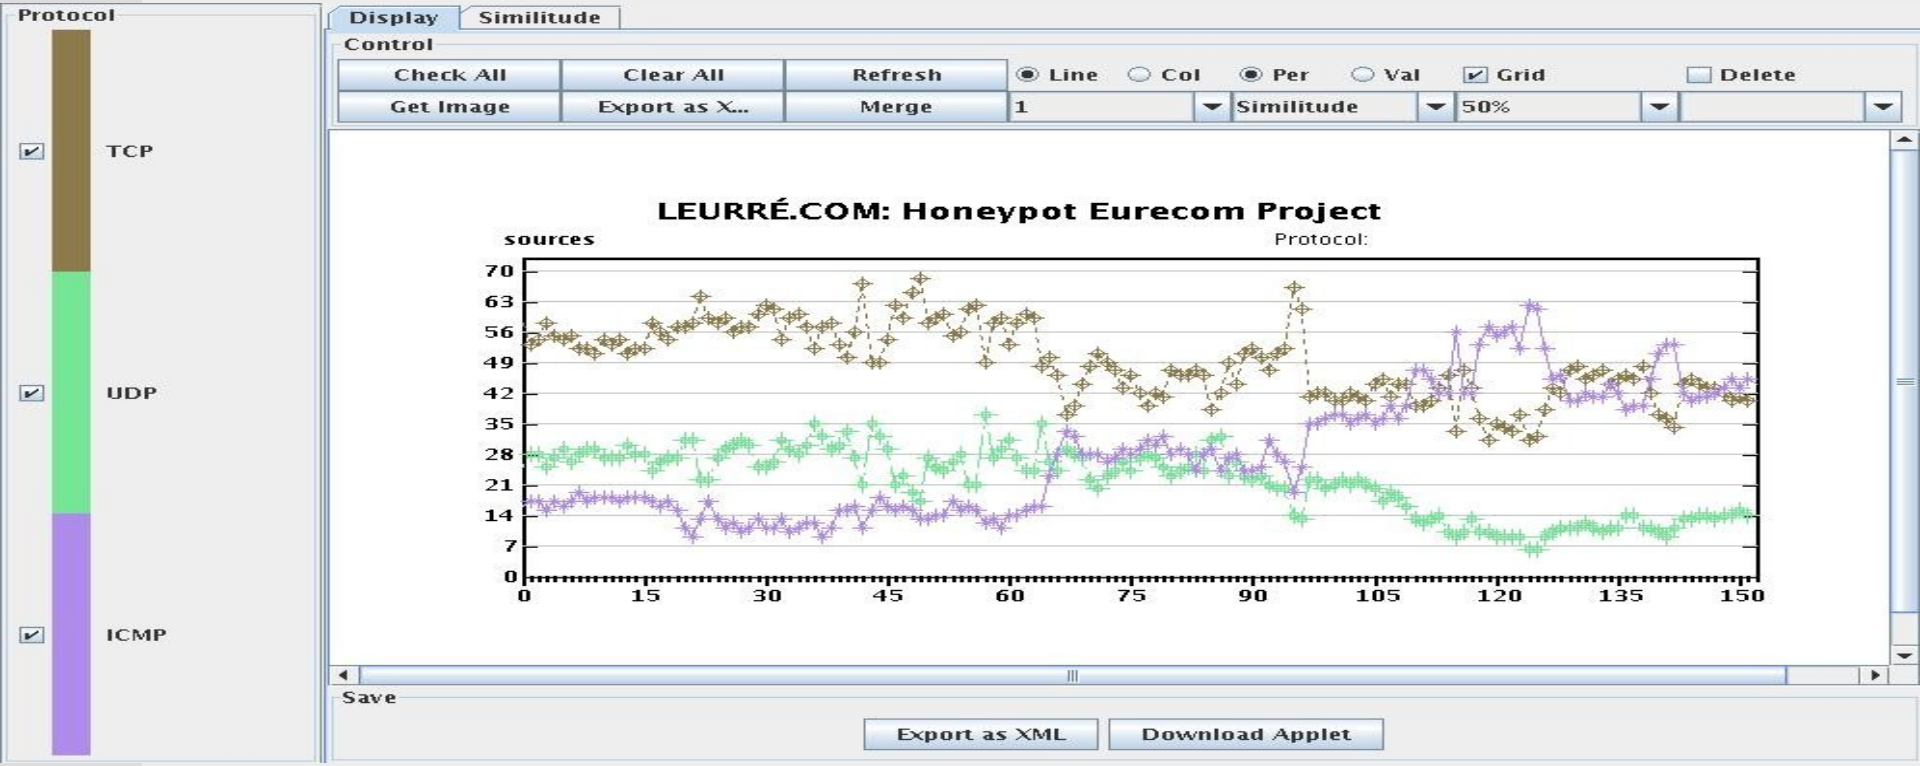

- Still some things are unexplicable from this data alone
- Sudden change in ICMP ratio (Sep 06 through Jan 07) around Decembe

Similar scalable architectures

- NoAH (Network of Advanced Honeypots)
  - FP6 project, designed a network of LIH and HIH using Argos sensors
- Collapsar (Purdue University)
  - centralized network of HIH + traffic redirectors
  - Redirector implemented as a UML virtual machine, honeypots are VMware or UML machines
- Potemkin honeyfarm infrastructure
  - large number of virtual HIH on top of Xen VM
  - uses cloning, recycling and mempage sharing techniques to run as many VMs as possible on a single machine
  - Outgoing traffic produced by honeypots redirected to another honeypot of the honeyfarm
- Bailey et al: hybrid scalable honeypot architecture where LIH hand off to HIH filtering out traffic

Mixed other projects worth a mention

- Billy Goat
  - IBM's own LIH with focus on worm detection, very similar to honeyd+arpd
- MyNetWatchman
  - similar to Dshield but focused on automatic notification in order to clean up hacked machines
- Surfnet IDS
  - A distributed IDS project
- Protected Repository for the Defense of Infrastructure Against Cyber Threats (PREDICT)
  - So protected that no one has access to date, and that no one outside the US will ever have access afterwards
  - Seemingly won't aim to be global and comprehensive, but to create datasets for (vetted) (US) researchers

Worldwide Observatory of Malicious Behavior and Attack Tools

Basic facts on WOMBAT

- A project which will be funded by the EU (and partner countries) and several partner institutions in the Seventh Framework Programme of European research
- 5.2MEUR budget over 3 years (3MEUR contribution by the EU), more than 40 collective m/y, starting at the beginning of 2008
- Participants:
  - Academics (T.U. Vienna; Vrije Universiteit Amsterdam; Politecnico di Milano; Queensland Univ. of Technology)
  - Research Institutes (Institut Eurecom; FORTH; Institute for Infocomm Research - Singapore)
  - CERTs (NASK)
  - Corporations (France Telecom R&D;Hispasec; a leading vendor of security solutions which we cannot name yet)

External liaisons

- Internet Motion Sensor (IMS)
- NICTER (Network Incident Analysis Center for Tactical Emergency Response), a Japanese project which shares some of our objectives
- CCIED (Collaborative Center for Internet Epidemiology and Defenses), a joint effort of UCSD and the International Computer Science Institute's Center for Internet Research
- MAAWG (Messaging Anti-Abuse Working Group), a global organization focusing on preserving electronic messaging from abuse
- TERENA (Trans-European Research and Education Networking Association)
- Clearstream, leading European supplier of post-trading services
- HP Labs, Trusted Systems Laboratory

Three core areas

- Data Acquisition
- Data Enrichment
- Threat Analysis

Data Acquisition

- Need to foster international collaboration
  - Ideally: creation of a standard and an infrastructure for data sharing
  - Look out for announcements on this, or get in touch with me if interested to participate
- Creation of an infrastructure for storage, access and analysis
- Development of new/improved types of sensors
  - client-based honeypots and their integration into monitoring systems
  - Wireless and Bluetooth honeypots
- Building upon NoAH and Leurré.com know-how, build a scalable network of LIH, MIH and HIH

Data Enrichment

- Commonly acquired data have proven not to be sufficient to reveal root cause(s)
  - Collecting thousands of malware: easy
  - Identify and classify them automagically: more difficult
  - Figuring out who's developed them and why: priceless
- Examples of the types of analysis we are studying to integrate:
  - code behavior characterization;
  - structure of the malicious code and philogeny
  - attack contextual information (how it was performed; scanning activities; type of deployed payload; subsequent actions)
- Experiences from the NoAH and Nepenthes projects will be invaluable

Threat analysis

- Final goal:
  - Find out the root causes of the observed attacks
  - Build upon this acquired knowledge in order to better predict upcoming threats.
- Tools
  - Data and metadata correlation (very different from correlating alerts for intrusion detection purposes)
  - Statistical analysis
- Delivered results:
  - Early warning capabilities
  - Security investments and policy making decisions support

Milestones

- Infrastructural
  - Early 2008: invitation workshop for setting up cooperation and gathering requirements (open workshops will follow in 2009 and 2010)
  - Late 2008: infrastructure design and integration of existing sensors
  - 2009: development and deployment of new sensors
- Characterization
  - End of 2008: code behavior analysis specifications
  - 2009: automated behavior and structure analysis tools
  - End of 2009-Early 2010: finalization of gathering and analysis of contextual informations
- The early warning prototype and root cause analysis are expected somewhen in 2010

Conclusions & Future Work

- Conclusions:
  - We need to be able to observe, understand and infer
  - We are currently partially able to observe, to understand (but generally late), and not to infer
  - We need to improve collection (a little bit), data analysis and enrichment (a lot), and to devise automatic inference mechanisms for root cause analysis
- WOMBAT:
  - Everything is a future work ;)
  - Funded global initiative for studying attacks and threats
  - Trying to make good use of the excellent work that has already been done in this area
  - Aiming to coordinate, rather than compete, with other large initiatives

Thank you!

Any question?

I would greatly appreciate your feedback !

Stefano Zanero
zanero@elet.polimi.it
www.elet.polimi.it/upload/zanero/eng