

Privacy-Preserving Sharing and Analysis

Patrick Lincoln
SRI International

Outline

- Context

- Information security, internet-scale defense

- Requirements

- Preserve privacy
- Enable aggregation and correlation
- Practical performance

- Approach

- Encryption, hashing, anonymization

Main Point

- Enable interesting interfaces to data
 - Protected by mathematics (cryptography)
 - Not just human trust (though that's very useful too)

General Goal

Enable enterprises to share information for sharing, regulatory and business purposes while maintaining a high level of security and organizational privacy.

Privacy-Preserving Sharing and Correlation of Information Security Alerts

Patrick Lincoln

Phil Porras and Vitaly Shmatikov

SRI International

Context: Information Security

- Networked information assets
 - Huge and growing value of information on networks
 - JV2020 predicates national security on information superiority
- Defense against
 - Bots
 - Email viruses
 - Worms
- The scale of the problem is large and growing rapidly
 - Individual attacks taking over 100,000's machines
 - More and more successful attacks

Examples of Past Threats

- Code Red, Code Red II, Nimda, SQL Slammer, MBlaster, MyDoom, Sasser
- Attacks of growing sophistication
 - Control, speed, infection success %
- Already serious consequences
 - Though fortunately, generally benign payloads

Name	Date	OS	Service	Infected Machines	Time
Lion	March 2001	Linux	BIND	10,000?	Days
Code Red	July 2001	Windows	IIS	200-400k	Days
Nimda	Oct 2001	Windows	IIS	100-200k	Hours
SQL Slammer	January 2003	Windows	IIS	100-200k	Minutes
MSBlaster	August 2003	Windows	IIS	300k?	Hours

Potential Future Threats

- Wormhol worms
 - Famous for 15 minutes
- Flash worms
 - Even shorter timespans of visible attack
- These are too fast for human-oriented response mechanisms to address

Why Are Attackers More Successful?

- More systems being created
 - More features being added
 - More vulnerabilities
 - More instances
 - More networks
-
- Practice of sharing and building on others work

Out-Share or Lose

- Attackers and their machines share vulnerability information (rapidly, efficiently)
 - Virus writers, worm writers, crackers, script kiddies, etc.
 - Email, chat rooms, IRC, web sites, p2p, IM, etc.
 - Network effect works to their advantage
 - Build on others work
- Currently many defenders (network administrators) do not share detailed information automatically with others outside their organization
 - Precious little. Ponderously slow publication cycle

Attackers Share Openly

- Open-source code sharing
- Large email networks
- Chat rooms

- Example: SDbot
 - Toolkit for building IRC-controlled bot armies
 - Rapid community uptake, extension of functionality, exploitation

Defenders Do Not Share As Openly

- Network defense viewed as local responsibility
- Individual sites defend themselves (only)
- Works OK against low or moderate levels of attack

- Internet-scale threats not well addressed in this mode

Defenders Do Not Share As Openly

- Sharing security alert information compromises organizational privacy and other business interests
 - Network topology
 - Network defenses
 - Which IDS, configured how, on which subnets?
 - Which Firewall, configured how?
 - Network monitoring capabilities
 - Customer information
 - Attacks come through 'normal' channels of communication
- Sharing takes time and energy

This is Complex Problem

- At least partially sociological
- Perhaps there are technological solutions that can play a role
- If it was easier and safer to share, more good guys would share more data more often
- Need to give defenders tools
 - Theories
 - Implementations
 - Accepted practicesthat make their life easier


Starting to Get Technical

- What can the theoretical information security community do to help?
- How can we enable security alert sharing without compromising privacy interests?

Requirements

- Security alerts gathered from thousands of security sensors
 - Firewalls, Network IDS systems, Host-based IDS systems, Antivirus filters
- Network defenders (the good guys) need to perform coarse groupings of alerts
 - Geographic region, industry group, OS type, port
- Network defenders need to perform fine-grained correlation of certain alerts
 - Patterns of alerts highlights new attacks, vulnerabilities, and potential defenses

Example: DShield.org



[Member Login](#)
[Signup](#)

Getting There

[What's New](#)
[Introduction](#)
[Internet Primer](#)
[How to submit your logs/reports](#)
[Client Programs](#)
[Web Interface](#)
[DShield Reports](#)
[Using DShield Feeds](#)
[FightBack](#)
[Mail Lists](#)
[Links](#)
[About us](#)
[Privacy Statement](#)

Distributed Intrusion Detection System

DShield.org

Records Added

Last Month
294,679,239

Last Week
116,077,068

Today
15,034,132

As of Tue May 25 10:36:56 2004 UTC

Internet Storm Center Status

green Akamai Problems, New Angle(r) On An Old Phish

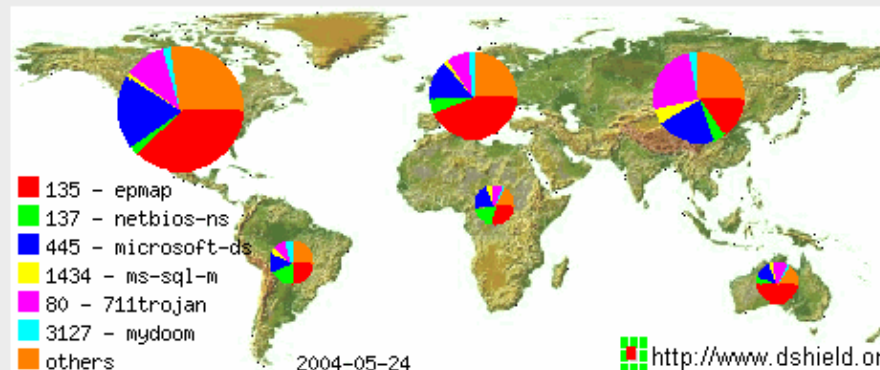
stop | start ticker

🔴 1027 | 🔴 445 | 🔴 80 | 🔴 1026 | 🔴 3128 | 🔴 25 | 🔴 113 | 🟢 4672 | 🟢 139 | 🟡

(ISC Daily Trends Page)


Top Attacker: 211.157.101.25

Most Attacked Port: 135



Geographic Distribution of attack sources. Last days

DShield.org



DShield.org

Distributed Intrusion Detection
System

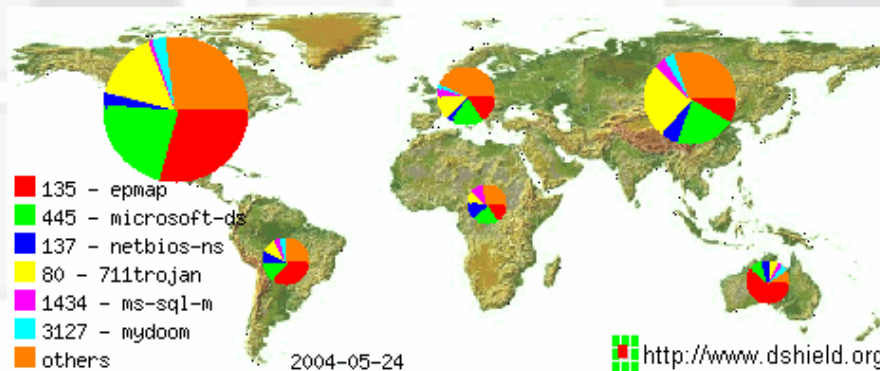
Google Search



ZoneLog Analyser

*The ZoneAlarm
Log Analyser*

Sorts the Threats from the Noise



Previous

Start

Stop

Next

[[Home](#) | [Login](#) | [What's New](#) | [Intro](#) | [Submit](#) | [Clients](#) | [Web Submission](#) | [All Reports](#) | [Mail Lists](#) | [Links](#) | [About](#) | [Privacy](#)]

Contact info@dshield.org for more information.

Internet Traffic Report

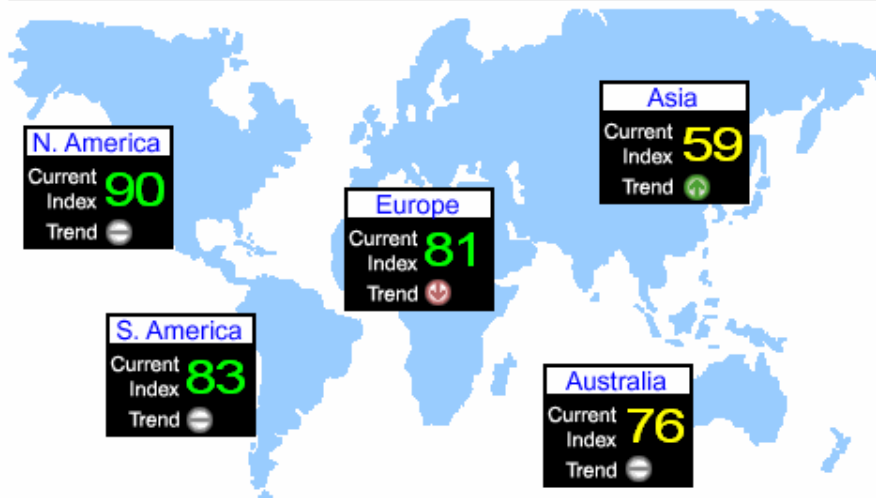


The Internet Traffic Report monitors the flow of data around the world. It then displays a value between zero and 100. Higher values indicate faster and more reliable connections.

Last update (MST):
5/24/2004 09:05
Global Index **84**
Trend



[Home](#) | [FAQ](#) | [Events](#) | [Contact](#) | [Links](#)



[View Graphs](#) or [Click a Continent](#) to view more detailed information.

The Internet Traffic Report (ITR) wants to continue to provide useful information about networks from around the world. We want to make this information as accurate as possible!
[More Information.](#)

Want to add a live statistics display to your website?
[Click here](#) to select your graphic.

Got questions? We've got answers!
Check out the [ITR FAQ](#)

Continent	Current Index	Avg. Response Time (ms)	Avg. Packet Loss (%)
Asia	59	393	17 %
Australia	76	230	0 %
Europe	81	174	1 %
North America	90	84	1 %

Symantec DeepSight

Help | Login | Documentation

symantec. DeepSight™ Threat Management System

symantec ThreatCon

5/24/2004 4:09:19 PM

instantaneous measurement of global threat exposure

[ThreatCon Definitions](#)

Login

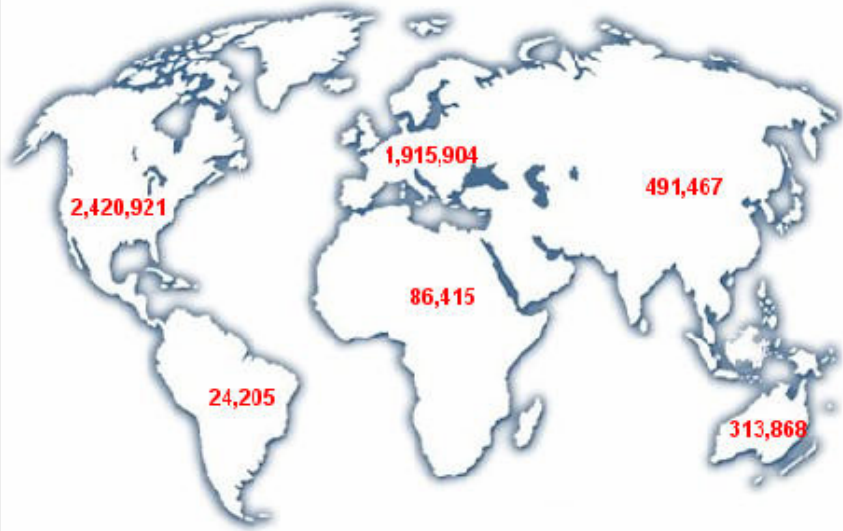
Member ID:

Password:

Save Username?

Login

Events Over Last 24 Hours



Region	Count
North America	2,420,921
South America	24,205
Europe	1,915,904
Africa	86,415
Asia	491,467
Oceania	313,868

Information

[About DeepSight.](#)
[Frequent Questions.](#)
[Privacy Policy.](#)

Events

Today	3,161,354
7 Days	62,109,116
Total	9,542,251,775

Attacking IPs

Today	185,638
Last 7 Days	2,703,839
Total	87,143,089

Symantec Early Warning Solutions

The Symantec DeepSight Threat Management System is the industry's most effective enterprise security threat management solution, providing early warning of attacks and bulletproof countermeasures to prevent attacks before they affect your enterprise.

Over 19,000 organizations in over 180 countries have registered to upload incident information to the Symantec Event Database. The expert team of Symantec Threat Analysts examines this global data, as well as hundreds of primary and secondary public and confidential sources, identifying imminent attacks and delivering comprehensive, detailed analysis based on your specific network configuration.

With the DeepSight Threat Management System, you can focus your security resources on proactively deploying critical countermeasures to mitigate the impact of attacks, rather than spending hours searching dozens of Web sites or hundreds of emails frantically trying to gather information on an attack and how to react and respond to it.

Deepsight (from their website)

- Monitors vulnerabilities in more than 18,000 technologies, operating systems, and application product versions from 2,200 vendors,
- Vulnerabilities monitored 24x7
- Enables secure, Web-based queries to an industry-leading vulnerability database
- Prioritized alerts
- Current and historical alerting and response reports
- Administrative user status gives control over subordinate users in order to share information, collaborate for early mitigation, and increase accountability
- Alert status tracking streamlines task assignment and reporting by providing status and documenting resolutions

SANS internet storm center

Port Lookup: [graph](#) [details](#)

+ Port Graph

- Port History

- Today's Diary

+ Papers and Analysis

- Diary Archive

+ Trend History

- ISTS News

- World Map

+ Top Rising

+ Internet Traffic

Join our discussion lists and contribute to analysis and information exchange. [click](#)

Interested in participating? [Click here to find out how to do so.](#)

Today's Diary

[Previous](#)

Handlers Diary May 23rd 2004

Updated May 24th 2004 15:21 UTC (Handler: Mike Poor)

Akamai problems. Quiet, well kinda quiet, day on the Internet

Update (Mon. May 24th 9 am EST, 13:00 UTC, 15:00 CEST)

It appears that websites that use Akamai's distribution system are currently not reachable. Security related web sites effected are symantec.com and trendmicro.com. Virus updates may fail as a result. Further details are currently not available and updates will be posted here as they become available. Thanks to Vidar Wilkens for alerting us of this problem.

According to a post to NANOG, the outage may be the result of a DDOS attack. At this point, Akamai has not ETA for a resolution.

Update 09:45 EST: Looks like some of the Akamai hosted sites start to come back.

Akamai posted this statement: " Due to a peering problem between ATT and UUNet, a subset of UUNet users may have experienced problems accessing Akamai delivered sites between 8-10pm EDT on Saturday May 22, 2004. The problem has been fully resolved. "

Port History

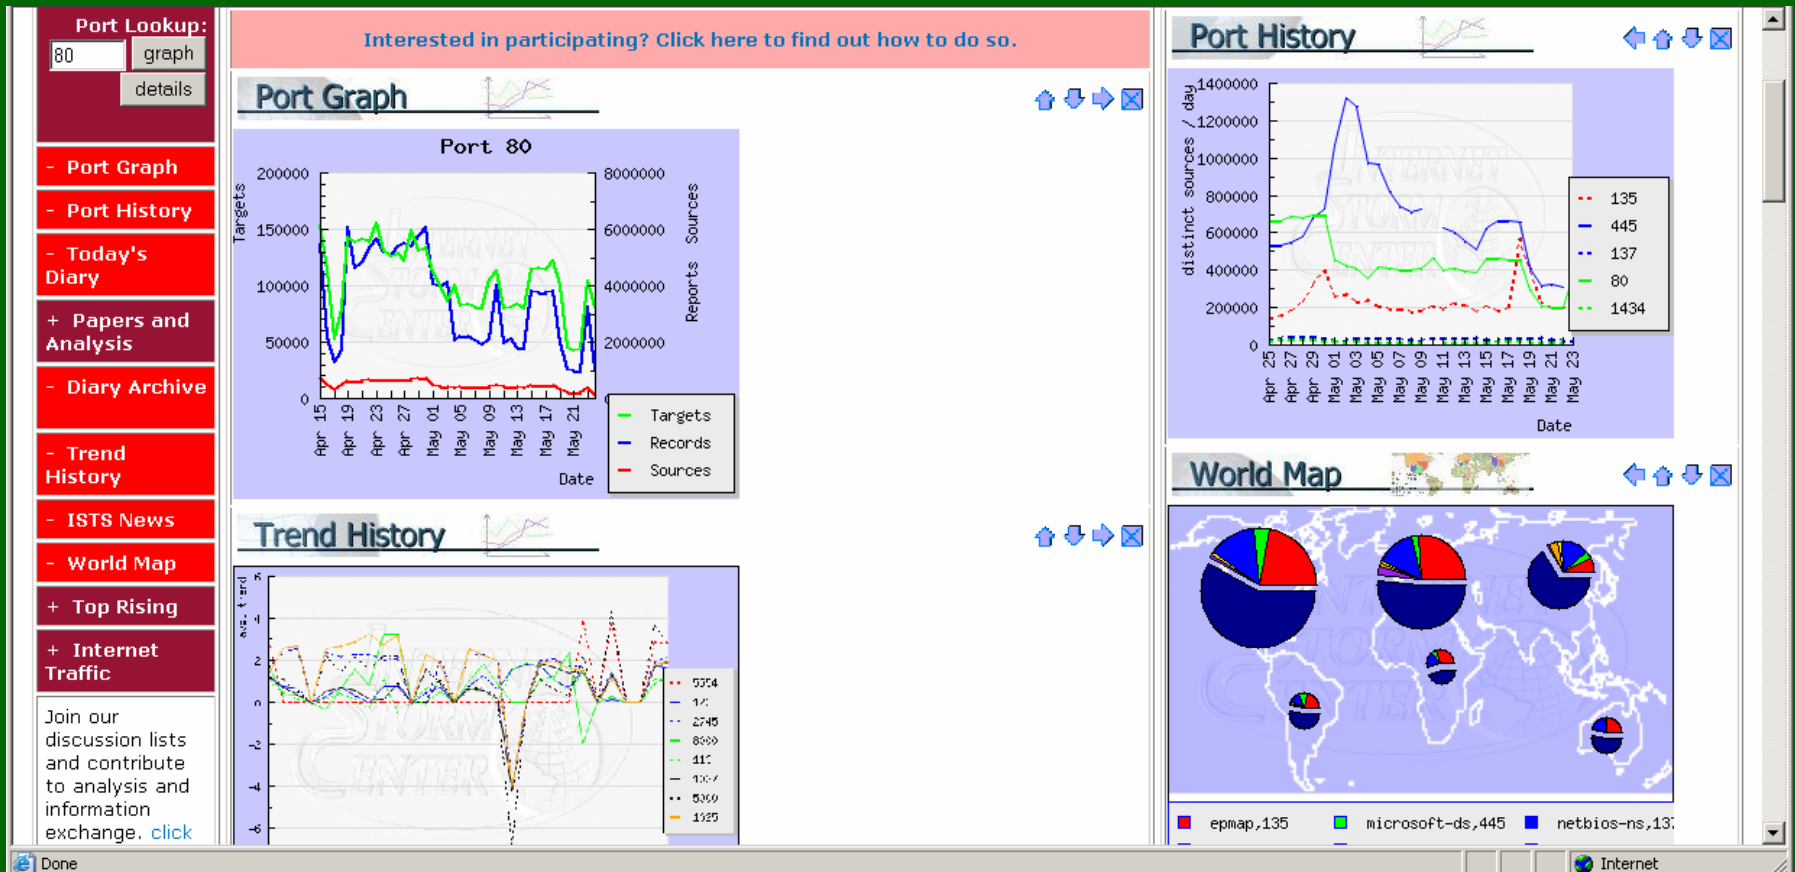
Date	135	445	137	80	1434
Apr 25	100000	500000	100000	500000	100000
Apr 27	150000	600000	100000	500000	100000
Apr 29	200000	700000	100000	500000	100000
May 01	300000	800000	100000	500000	100000
May 03	250000	1000000	100000	500000	100000
May 05	200000	900000	100000	500000	100000
May 07	200000	700000	100000	500000	100000
May 09	200000	600000	100000	500000	100000
May 11	200000	500000	100000	500000	100000
May 13	200000	400000	100000	500000	100000
May 15	200000	300000	100000	500000	100000
May 17	200000	200000	100000	500000	100000
May 19	200000	100000	100000	500000	100000
May 21	200000	100000	100000	500000	100000
May 23	200000	100000	100000	500000	100000

World Map

Legend: epmap,135 (red), microsoft-ds,445 (green), netbios-ns,137 (blue)

<http://isc.incidents.org/top10.php?isc=5b5a2847462b8f9f9577145a4544db2c>

SANS internet storm center



Microsoft Security

The screenshot shows the Microsoft Security website with a blue header and a white content area. On the left is a navigation menu. The main content area features a large banner for the Sasser worm, a 'Security Bulletins' section with a list of links, a 'Virus Alerts' section with a list of updates, and a 'Free Support' box. On the right side, there are three vertical boxes: 'Stay secure' with an email subscription link, 'Enterprise Security Guidance' with a link to a strategy document, and 'Update Your Software' with links to Windows and Office updates. At the bottom right, there is a 'Communities and Events' section with a link to a webcast.

Microsoft
Microsoft Security

Security & Privacy Home
Trustworthy Computing ▶
Home Users
IT Professionals (TechNet)
Developers (MSDN)
Businesses
Partners ▶
Microsoft Privacy Policies
Worldwide Security Sites

Free Support
Call **1-866-PCSAFETY** for free virus-related support (U.S. and Canada only)
Please call your local Microsoft subsidiary. Find your subsidiary

Sasser worm: Important information
What to do to protect against or remove the worm.
Actions you can take. ➔

Stay secure
Get e-mail about new security updates

Enterprise Security Guidance
Get Tools, Training, and Guidance to Plan and Manage a Strategy for Your Organization

Update Your Software
Windows Update
Office Update

Communities and Events
Webcast: Technical Update on Sasser
Webcast: Info About

Security Bulletins

- Action: Install the Windows Security Update
- More Security Bulletins ...

Virus Alerts

- Update: Sasser Worm May 11, 2004
- Update: Mydoom Worm March 11, 2004
- Bagle Virus Recommendation: Be Cautious About Opening E-Mail Attachments
- Specific Actions for the Blaster Worm
- More Antivirus Information...

Learn More
Microsoft Progress Report

Technical Information
Patch Management Security

Is The Problem Solved? Not Completely

- In order to use these services, one has to trust these organizations
 - Must trust them not to reveal private information
 - Must trust their sysadmins, janitorial staff, etc.
- Their privacy policies may not match yours
- They do not make data globally available for analysis by researchers and other services
 - Because their business model doesn't allow that
 - Because that would compromise privacy of submitted alerts

Privacy Requirements

- No public release of information about vulnerabilities
 - May invite and enable more attacks
- No public release of information about defenses
 - May enable attackers to adapt and avoid defenses
- No public admission of successful attack
 - May compromise public confidence
- No public release of information even about network topology

In Addition

- Don't require (too much) trust in a repository
- Don't require unreasonably large computational or communications overhead
 - Shared computation
 - Secure multiparty computation
 - Threshold cryptography

Are These Requirements in Conflict?

- At least partially...
- Any information gleaned from exposed reports may be viewed as a privacy breach
- Need to agree on acceptable disclosures

Specific Question

- Can we partially resolve the tension between sharing information security alert information and the maintenance of organizational privacy?

Proposed Approach

- Cleanse alerts
- Share anonymously
- Protect alert repositories

Proposed Approach

- Protect privacy
 - Encrypt certain fields of alerts
 - Anonymize reports where appropriate
 - Route alerts anonymously
- Enable correlation of alerts
 - Exploit "weaknesses" in cryptographic and hashing algorithms
 - Malleability, etc
 - Make certain properties of alerts checkable

Example Security Alert Content

Source_IP	FW,ID	Typically refers to the source IP address of the machine that initiated the session or transferred the transaction that caused the alert to fire. In IDS alerts, this field may represent the victim, not the attacker, since some systems alert upon an attack reply rather than request.
Source_Port	FW,ID	Source TCP or UDP port of the machine that initiated the session or transferred the transaction that caused the alert to fire.
Dest_IP	FW,ID,AV	Typically refers to the destination IP address of the machine that initiated the session or transferred the transaction that caused the alert to fire. In AV systems, Dest_IP can identify the machine in which the infection is discovered.
Dest_Port	FW,ID	Destination TCP or UDP port of the machine that initiated the session or transferred the transaction that caused the alert to fire.
Protocol	FW,ID	Protocol type (<i>e.g.</i> , UDP, TCP, ICMP).
Timestamp	FW,ID,AV	May incorporate incident start time, end time, incident report time.
Sensor_ID	FW,ID,AV	May incorporate the brand and model of the sensor and a unique identifier for the individual instantiation of the sensor.
Count	FW,ID,AV	Often used to represent some notion of repeated activity, either at the alert or event (<i>e.g.</i> , packet) level.
Event_ID	FW,ID,AV	Uniquely defines the alert type for the given sensor.
Outcome	FW,ID,AV	Reports the status or disposition of the reported activity. For firewalls, it may report whether the log entry was associated with an allow or deny rule. For AV, it may indicate infection disposition (<i>e.g.</i> , Symantec's AV indicates whether the infected file is cleaned or quarantined). Outcome fields for IDS tools are highly vendor-specific.
Captured_Data	ID	Some IDS sensors have the ability to report part or all of the data content in which the alert was applied.
Infected_File	AV	Antivirus logs include the identity of the file that was infected.

Basic Privacy Protection

- Scrubbing sensitive fields
 - Infected file
 - Outcome
 - Captured data
 - In future, hope to enable analysis, for now, scrub
- Hiding IP addresses
 - Source_IP
 - Dest_IP

Hiding Information

- Encrypting IP addresses under private key is unacceptable
 - Prevents correlation of attacks
- Hashing under SHA-1 or MD5 has issues
 - Enables dictionary attacks
 - Attacker could precompute hash values of subset of network, monitor repository for hits
 - Only 65536 or 256 addresses (or smaller)
- Solution: balance privacy and utility
 - Own network use keyed hash
 - External networks use standard hash function

Abstractly

- Want to enable (approximate) equality tests under hash or encryption, but not enable simple dictionary attacks, etc
- Decreased functionality is acceptable
 - Matching Source_IP in one attack to Dest_IP address in another may not be simple

Other Privacy Protection

- Re-keying by repository
- Randomized hot list thresholds
 - For collaborative detection of high-volume alerts, it is sufficient for the repository to publish only the hot list of reported alerts that have something in common where the number exceeds some threshold
 - Vulnerable to flooding attacks
 - Defense includes random thresholds that the attacker cannot predict
- Controlled delay alert publication
 - For reports on historical (hours) trends, delayed alert publication provides feasible defense against probe-response attacks

Example Firewall Alert Sanitization

Field ID	Raw firewall alert	Sanitized firewall alert
Source-IP	172.16.30.2	0x16e9368f
Source-Port	1147	1147
Dest-IP	173.19.33.1	0x78a65237
Dest-Port	135	135
Protocol	6	6
Timestamp	09032003:01:03:10	09032003:01:03:00
Sensor	PIX-4-10060231	PIX
Count	1	1
Event-ID	Deny	Deny
Outcome	none	none
Capture-Data	none	none
Infected-File	none	none

Example IDS Alert Sanitization

Field ID	Raw IDS alert	Sanitized IDS alert
Source-IP	172.16.30.49	0xb09956c2
Source-Port	1299	1299
Dest-IP	176.20.22.43	0xd6e79b79
Dest-Port	80	80
Protocol	6	6
Timestamp	10132003:11:41:09	10132003:11:41:00
Sensor	EM-HTTP-90209321	EM-HTTP
Count	1	1
Event-ID	CGI-ATTACK	CGI-ATTACK
Outcome	NO-REPLY	none
Capture-Data	/scripts/.%255c%255c./winnt/system 32/cmd.exe?/c+dir	none
Infected-File	none	none

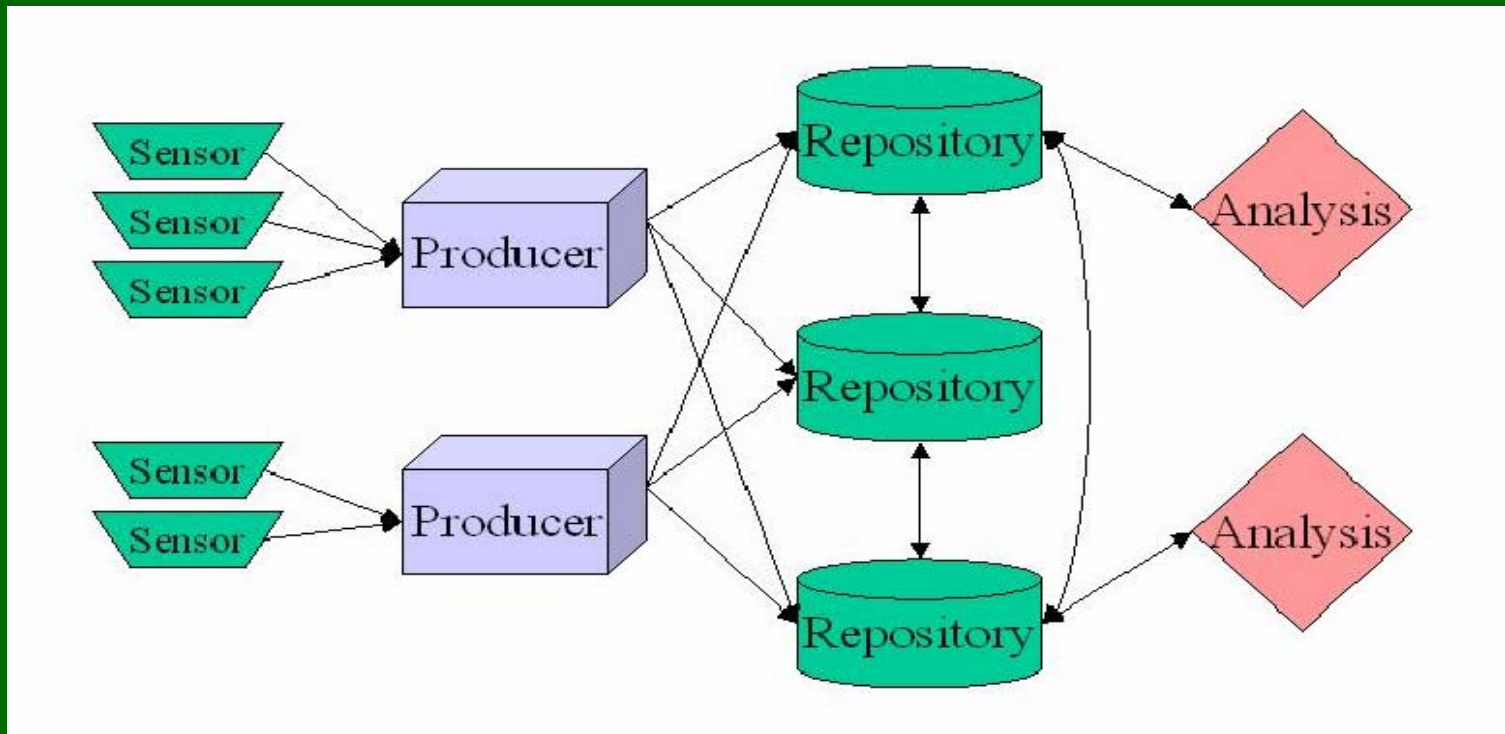
Example Antivirus Alert Sanitization

Field ID	Raw AV Alert	Sanitized AV alert
Source-IP	none	none
Source-Port	none	none
Dest-IP	176.30.22.11	0xb4ddc807
Dest-Port	none	none
Protocol	none	none
Timestamp	11172003:09:39:00	11172003:09:39:00
Sensor	NORTON-AV-02209302	NORTON-AV
Count	1	1
Event-ID	W32.Sobig.F.Dam	W32.Sobig.F.Dam
Outcome	Left alone	none
Capture-Data	none	none
Infected-File	A0014566.pdf	none

Further Privacy Protection

- Multiple alert repositories
 - Sensor information can be shared directly with multiple repositories
 - Spreading alerts reduces opportunities for collaborative pattern finding, but still enables real-time monitoring
- Randomized alert routing
 - Overlay network for p2p routing of alerts hides sources
 - Mix networks, Crowds, or Onion routing

Alert Sharing Infrastructure



Anonymity Estimates

- With n routers, if c are controlled by attacker, the probability a route contains a node controlled by the attacker is

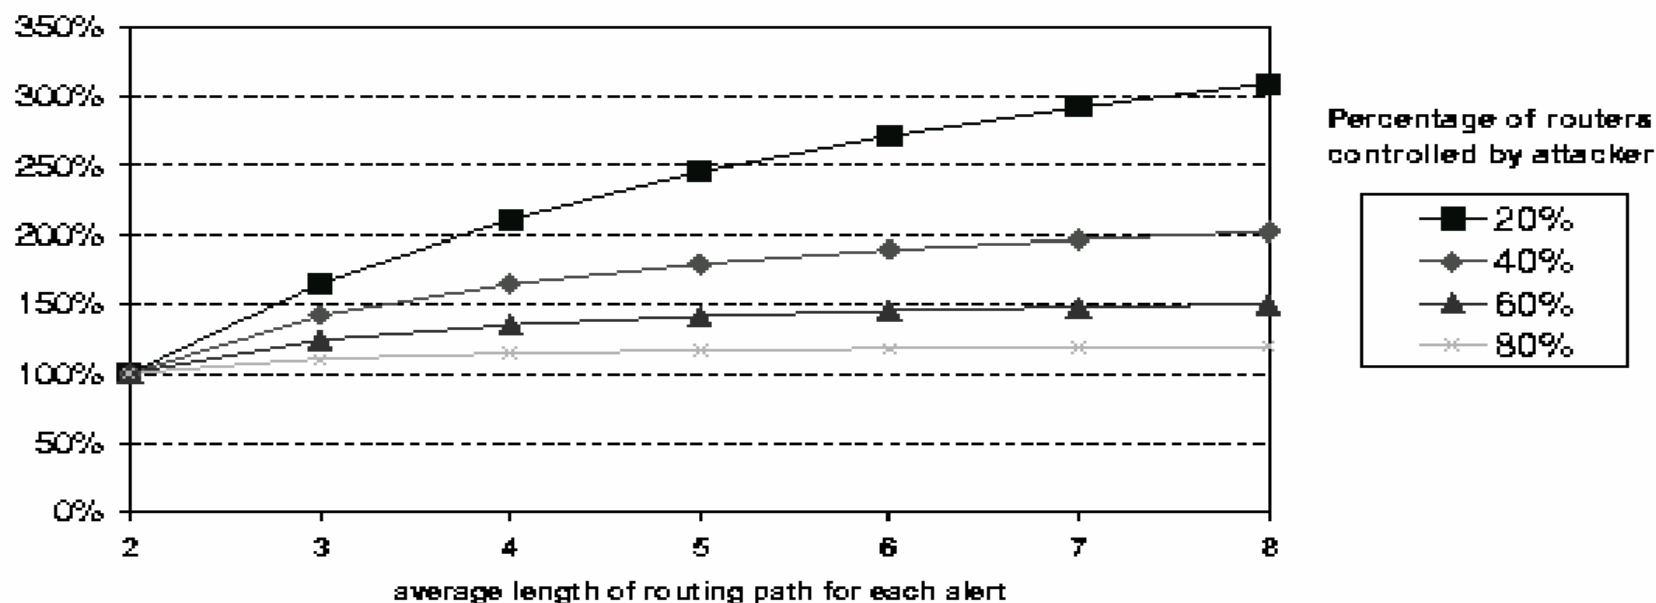
$$\frac{c(n - np + cp + p)}{n^2 - np(n - c)}$$

- This is close to $\frac{c}{n}$
- So 1- that many alerts are not observed by the attacker at all.
- For each of the alerts that is observed by the attacker, the probability that its apparent source is the actual source is

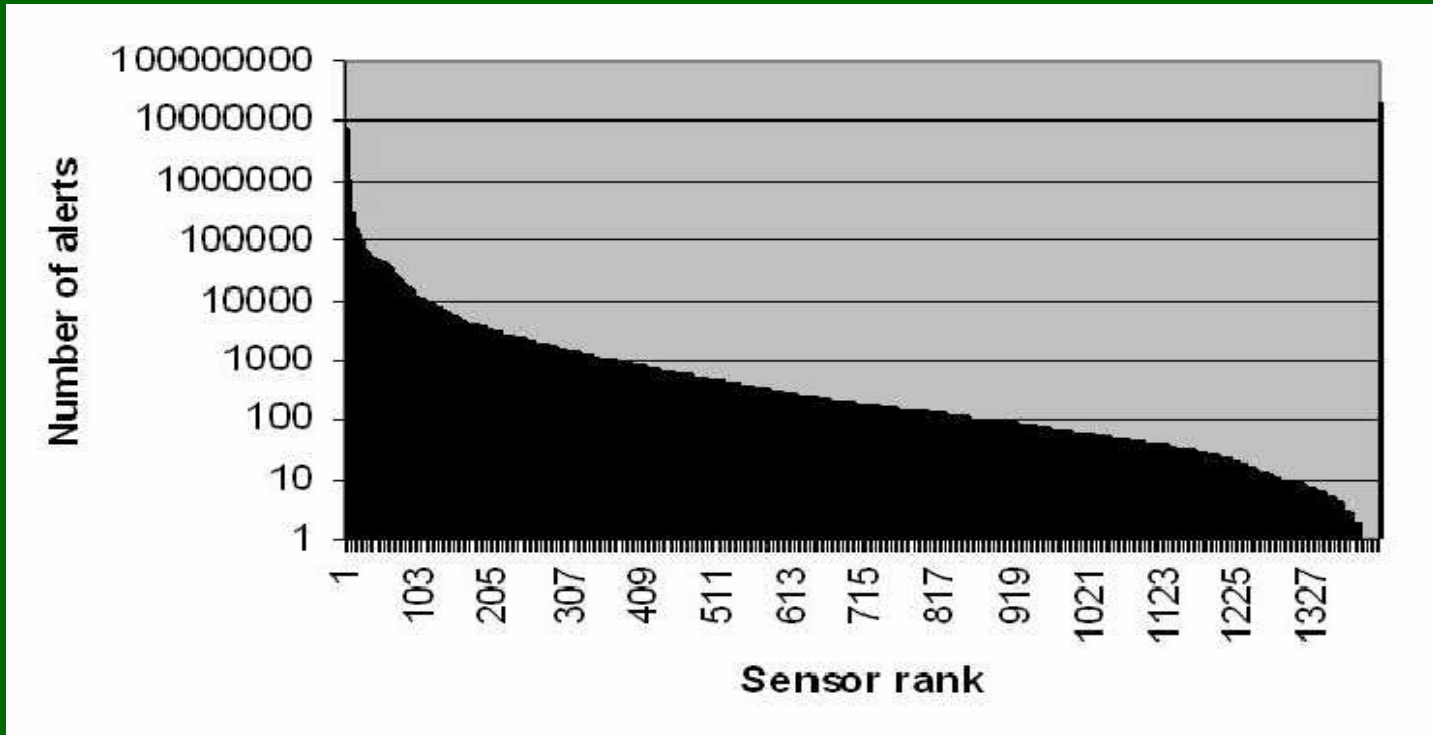
$$\frac{n - p(n - c - 1)}{n}$$

Attacks on Anonymous Sources

Attacker workload to determine origin of observed alerts



A Few Loud Sensors



Thus we have to be careful how we anonymize sensor IDs

Implementation

- Reasonable-seeming performance costs
 - CPU at the sensor
 - Outbound networking
- 1.5GHz P3, Mark Shellor's SHA and HMAC
- SRI firewall during Kuang2 outbreak
 - 4M alerts
- 1 day of DShield records
 - 19M alerts

Implementation

- Costs are modest

	baseline	hashed	delta	cached-8	delta
DShield.org	29.81	64.16	34.35	56.84	27.02
Laboratory	75.80	110.34	34.54	106.20	30.40

Table 5: CPU Impact of IP Hashing (seconds per 1 million alerts).

Analysis Mode

- After anonymous routing, analysis is needed
- Enables key EMERALD alert aggregation and correlation
 - EMERALD is an intrusion-detection and hierarchical correlation infrastructure built at SRI
 - Similar tools using similar APIs into alert repositories could also use this data

Next Steps

- Expand types of analysis can we enable on alert records?
- Connect with Cornell work on multiparty computation
- Handle rare alert events better
- Implement Snort plugin for wide deployment
- Connect to actual onion routing infrastructure
- Convince people to use it