010xxx011xx101010011001110101xx110100001

WORMINATOR

COLLABORATIVE INTRUSION DETECTION

http://worminator.cs.columbia.edu

Objective

Solve the problem of ranking IDS alerts to focus on the most sophisticated and dangerous attacks

- Difficult to differentiate legitimate versus truly dangerous, illegitimate traffic from just one point on a network
- The best IDSes do not see slow, stealthy activities spread out over time and space
- IDS noise makes it difficult to track zero-day worm attacks

Motivation

With an increase in both hitlist and zero-day worms, need to rapidly identify attacks from a variety of globally distributed sources

- If a malicious scan attempt is detected by one IDS, it can mean anything
- If similar malicious scan attempts from the same source are detected by IDSes at other sites, we have more confidence it's not just noise or "a coincidence"
- If a scan attempt is detected at some sites but not others, it's less likely a worm drone and instead a targeted scan

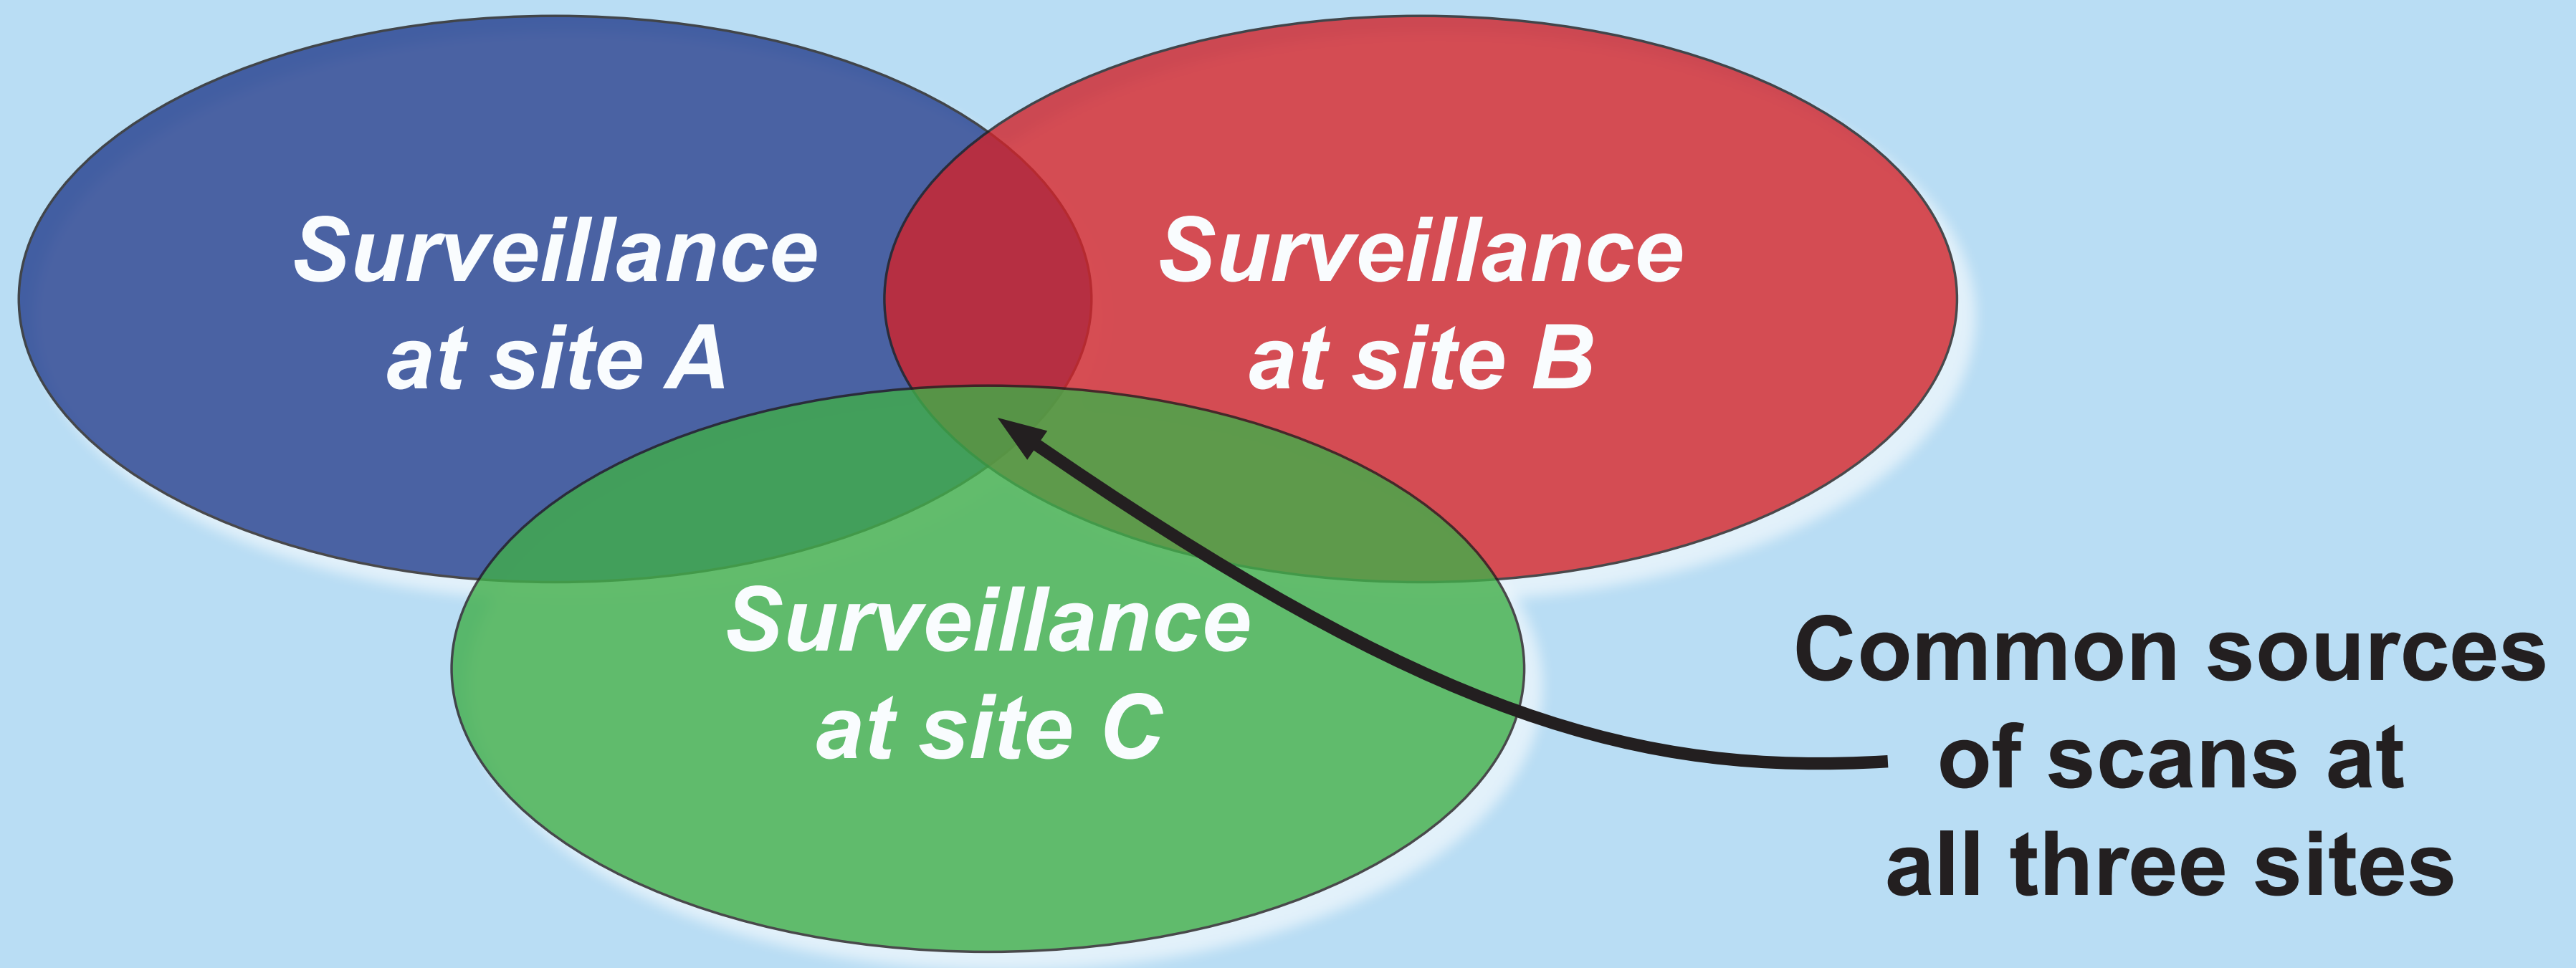

Architecture

- IP watchlists: transfer natively, or use a one-way hash data structure called a Bloom filter to encode IPs and ports into a privacy-preserving, compact data structure
- Signatures: generate using payload anomaly detection algorithms; they may be exchanged natively, as a Z-string, or in a Bloom filter as well

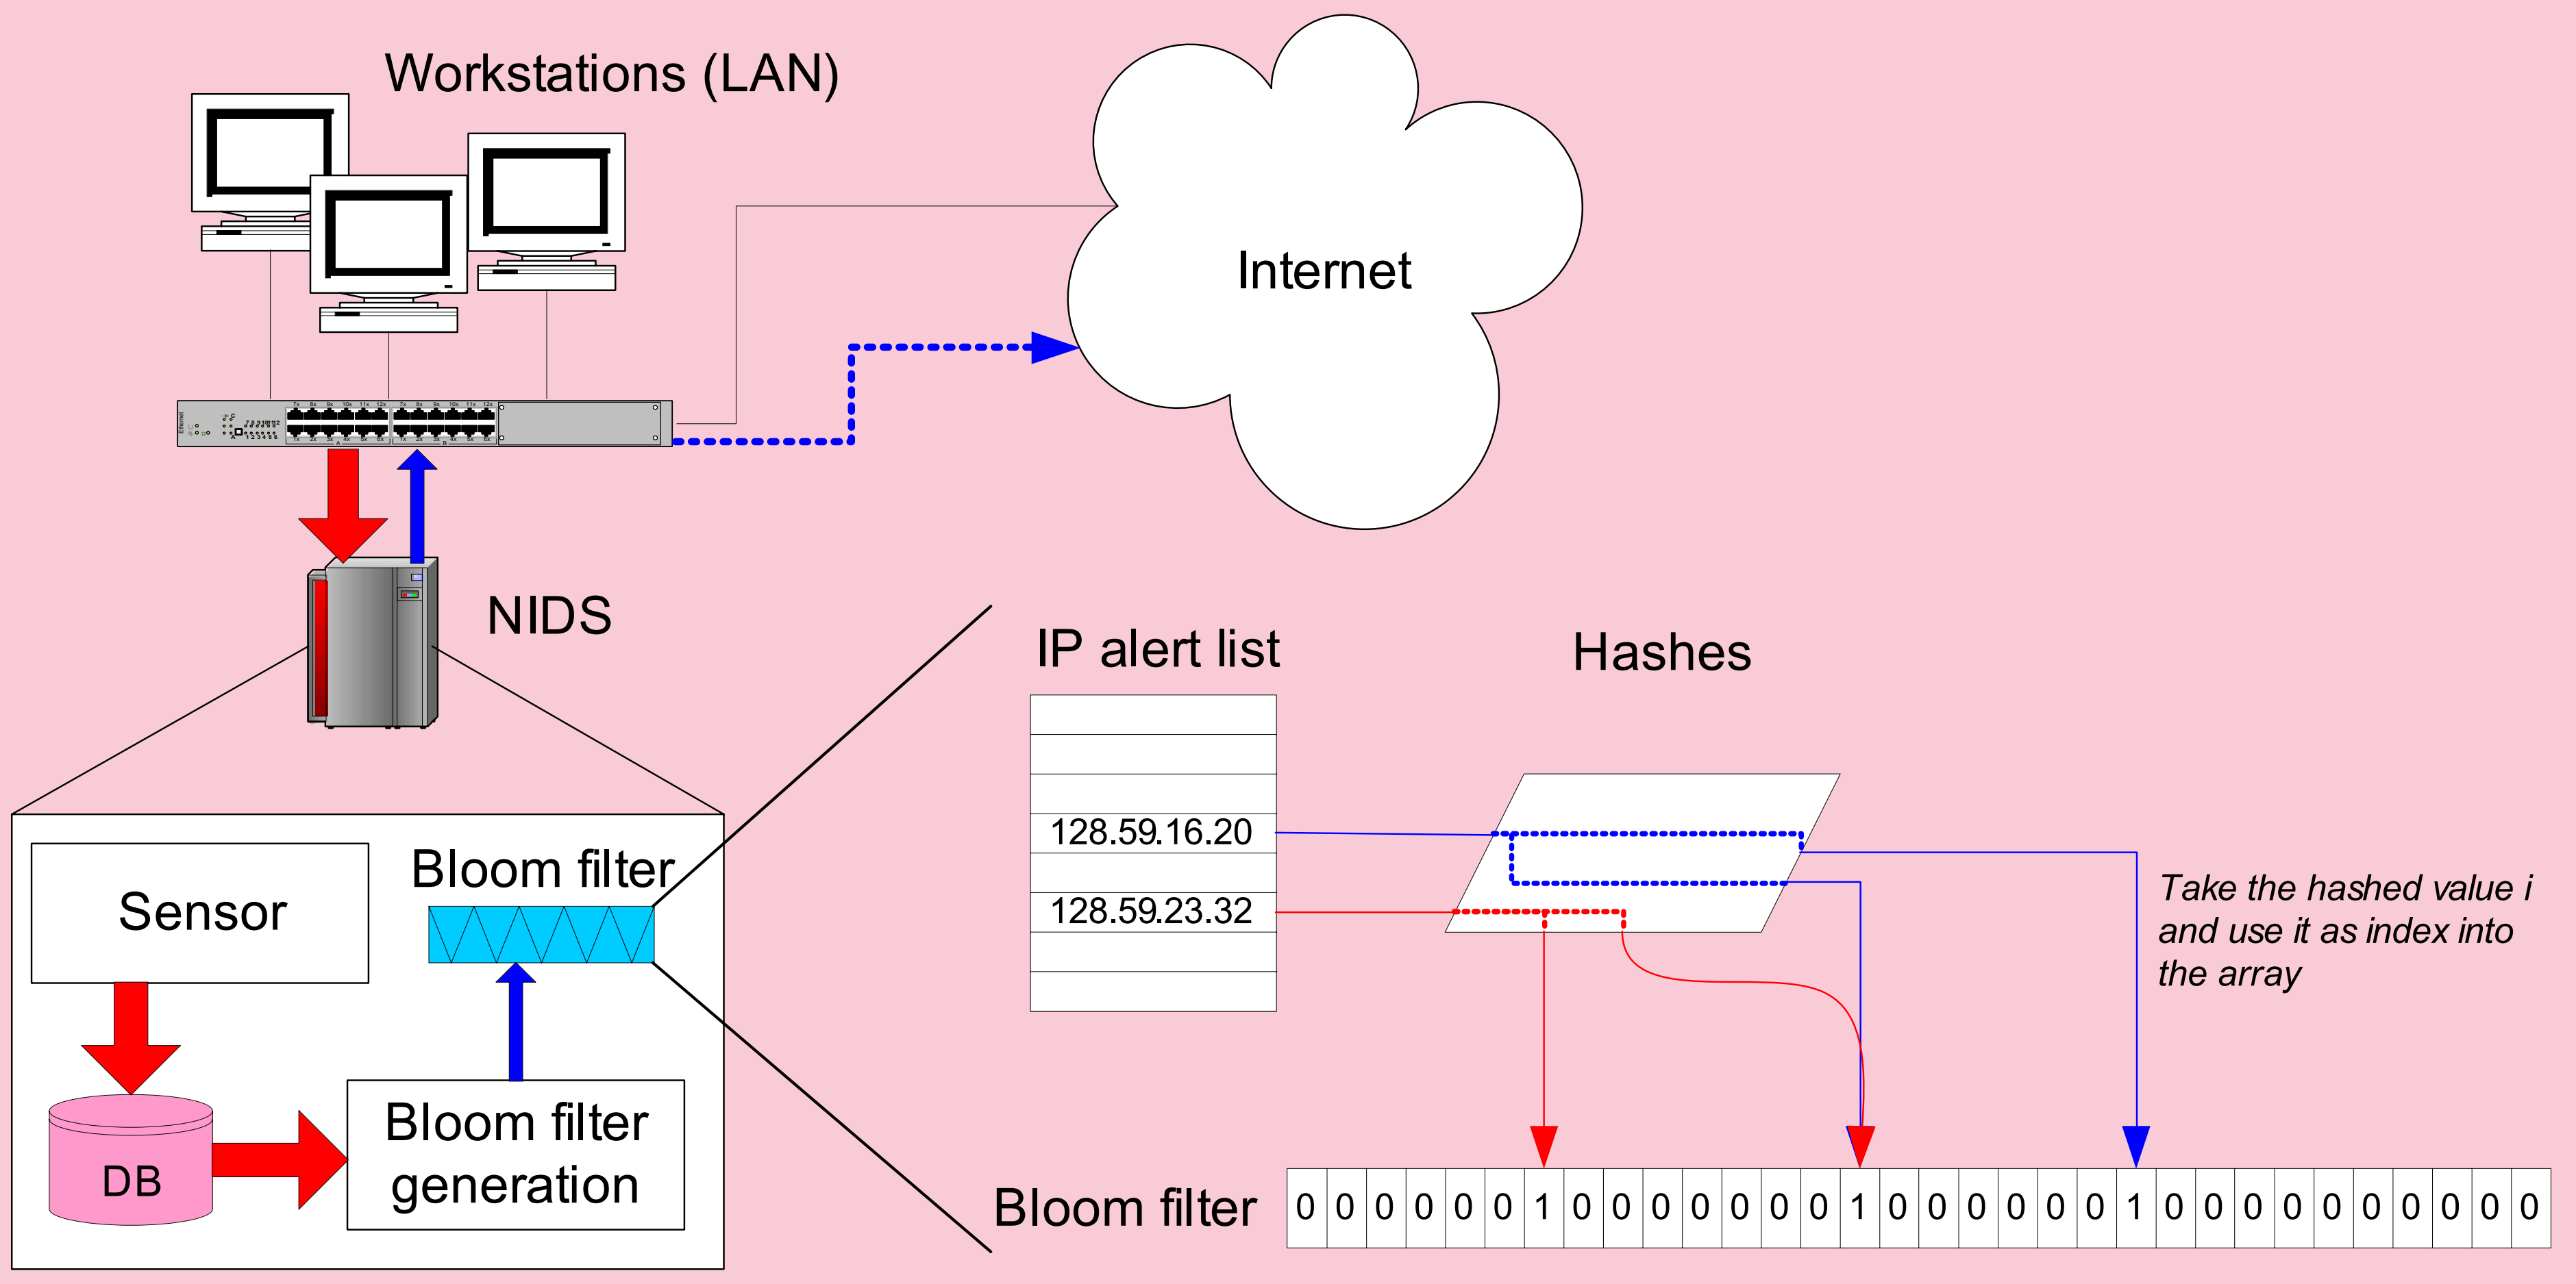

- Distribution: centralized server (e.g., broadcast, publish/subscribe), hierarchical, or P2P approach. In the latter case, the problem of *network scheduling* becomes significant (e.g., how to distribute the data with as few transmissions as possible given a large set of nodes)

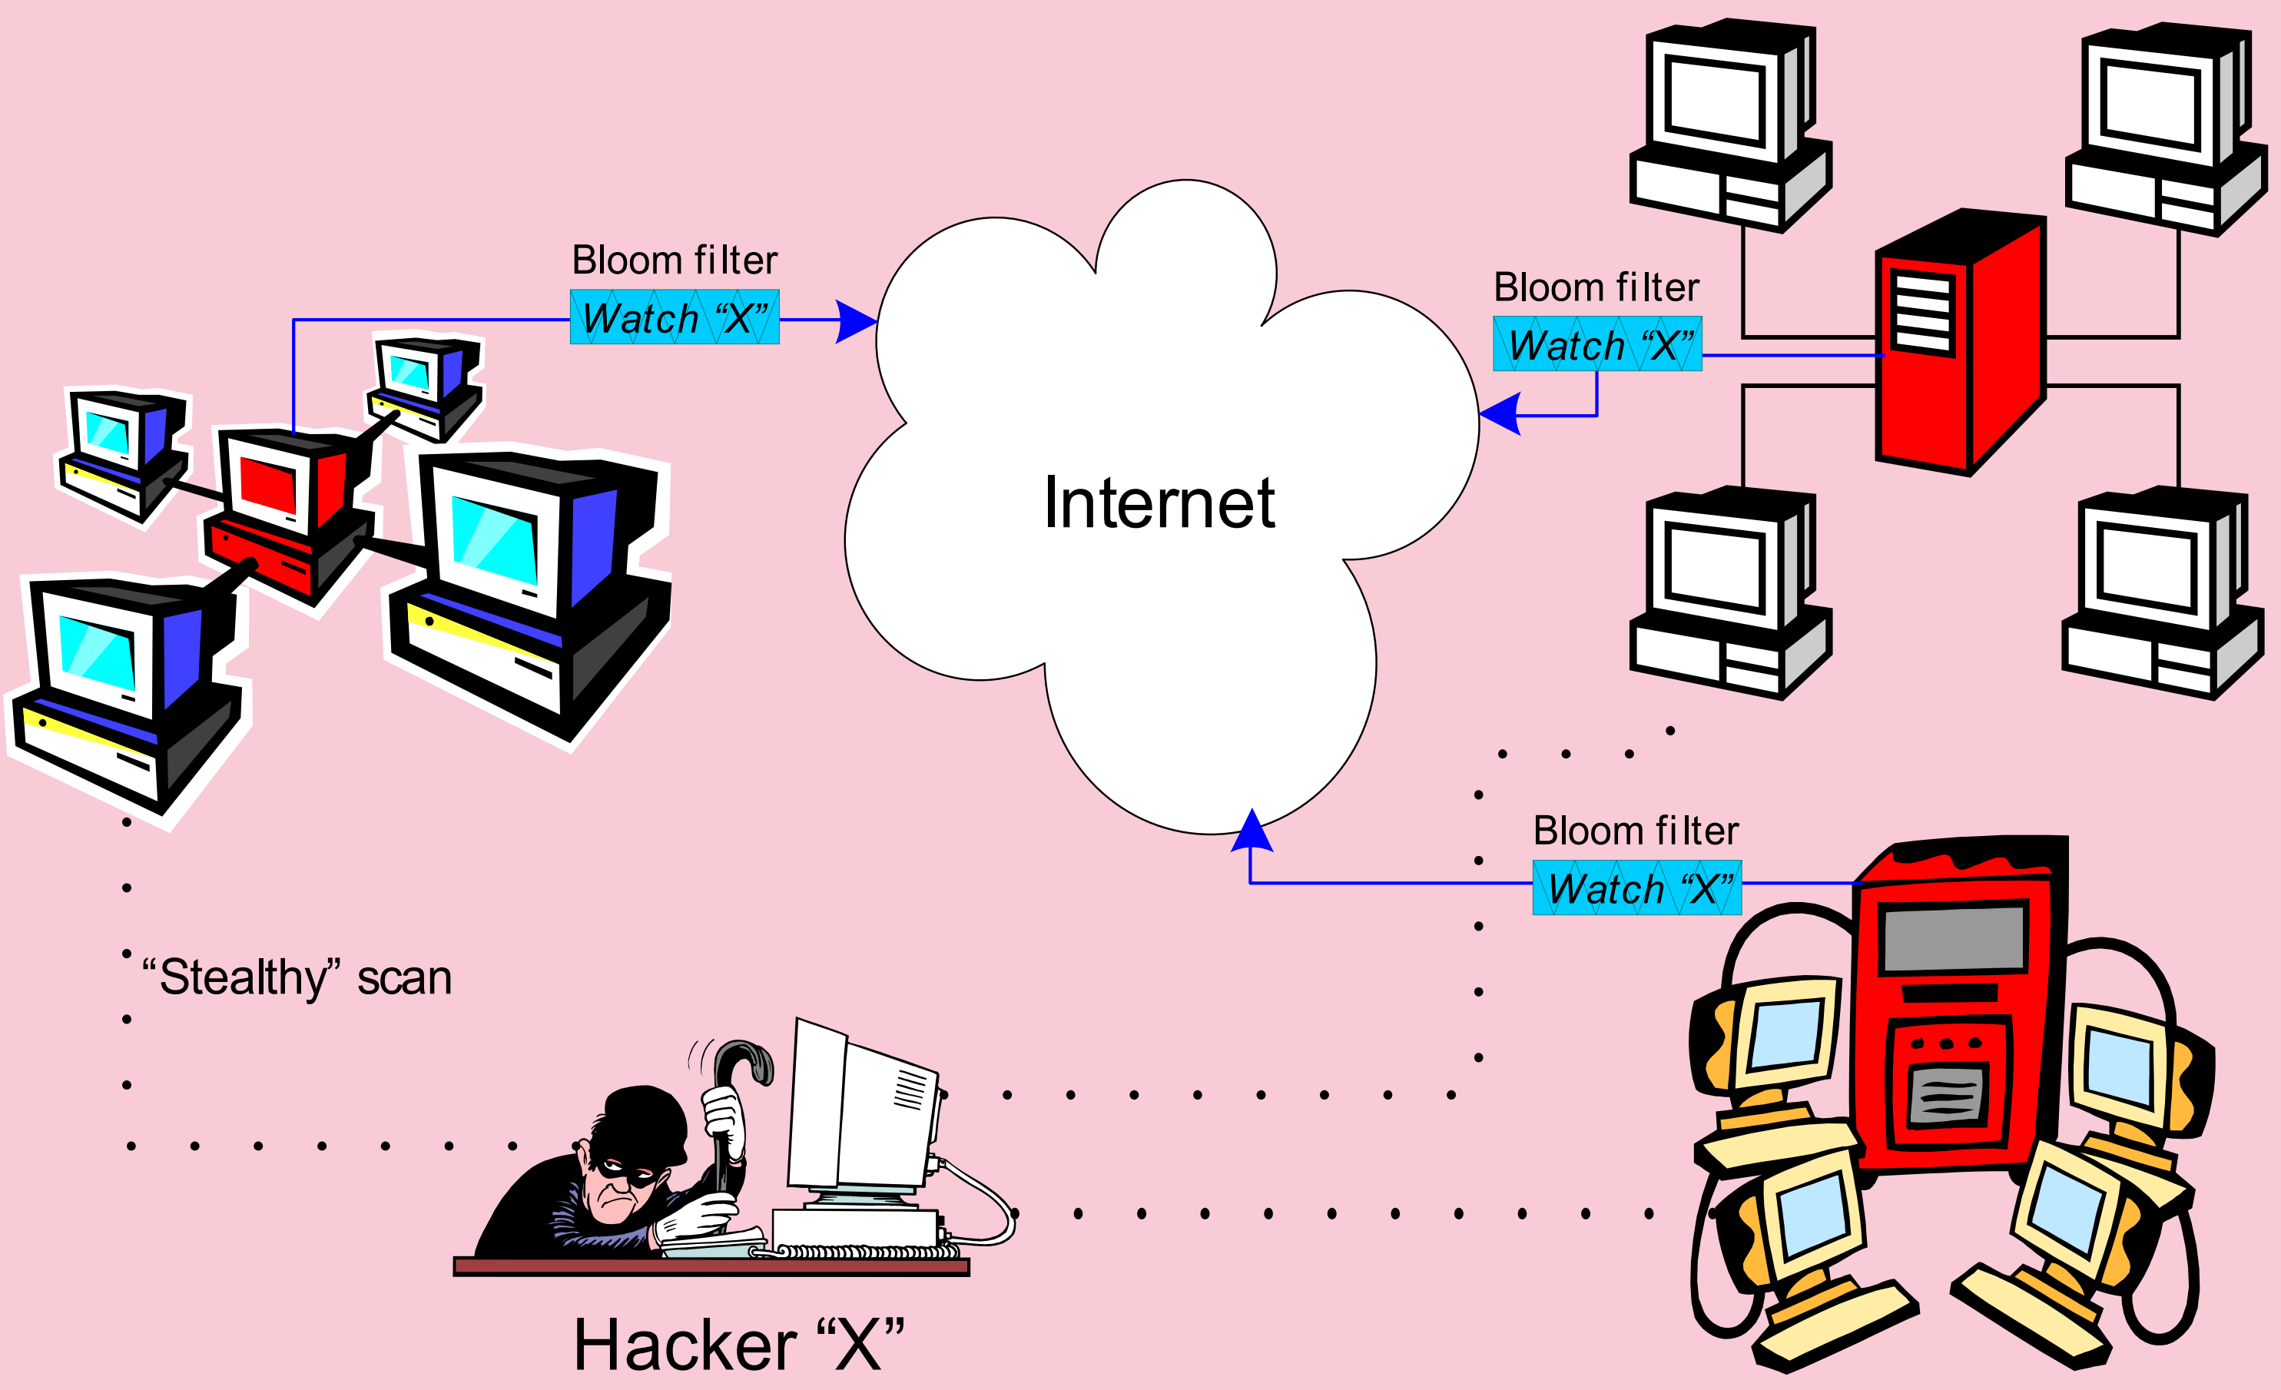

NIDS systems are in RED; they pick up the stealthy long-term scans and start exchanging Bloom watchlists.

Project status

- Proof-of-concept architecture developed using Java 5; web interface uses JSP 2.0/Servlet 2.4 platform
- Uses JMS publish/subscribe infrastructure for rapid alert exchange (~ 1s latency under normal congestion)
- Works with off-the-shelf Counterstorm AntiWorm-1 product, based on Columbia IDS technology
- Currently deployed in 5 networks; more forthcoming shortly

Columbia University IDS Lab
Prof. Sal Stolfo, PI

Screenshots

Main reporting screen, showing participating sites (anonymized)

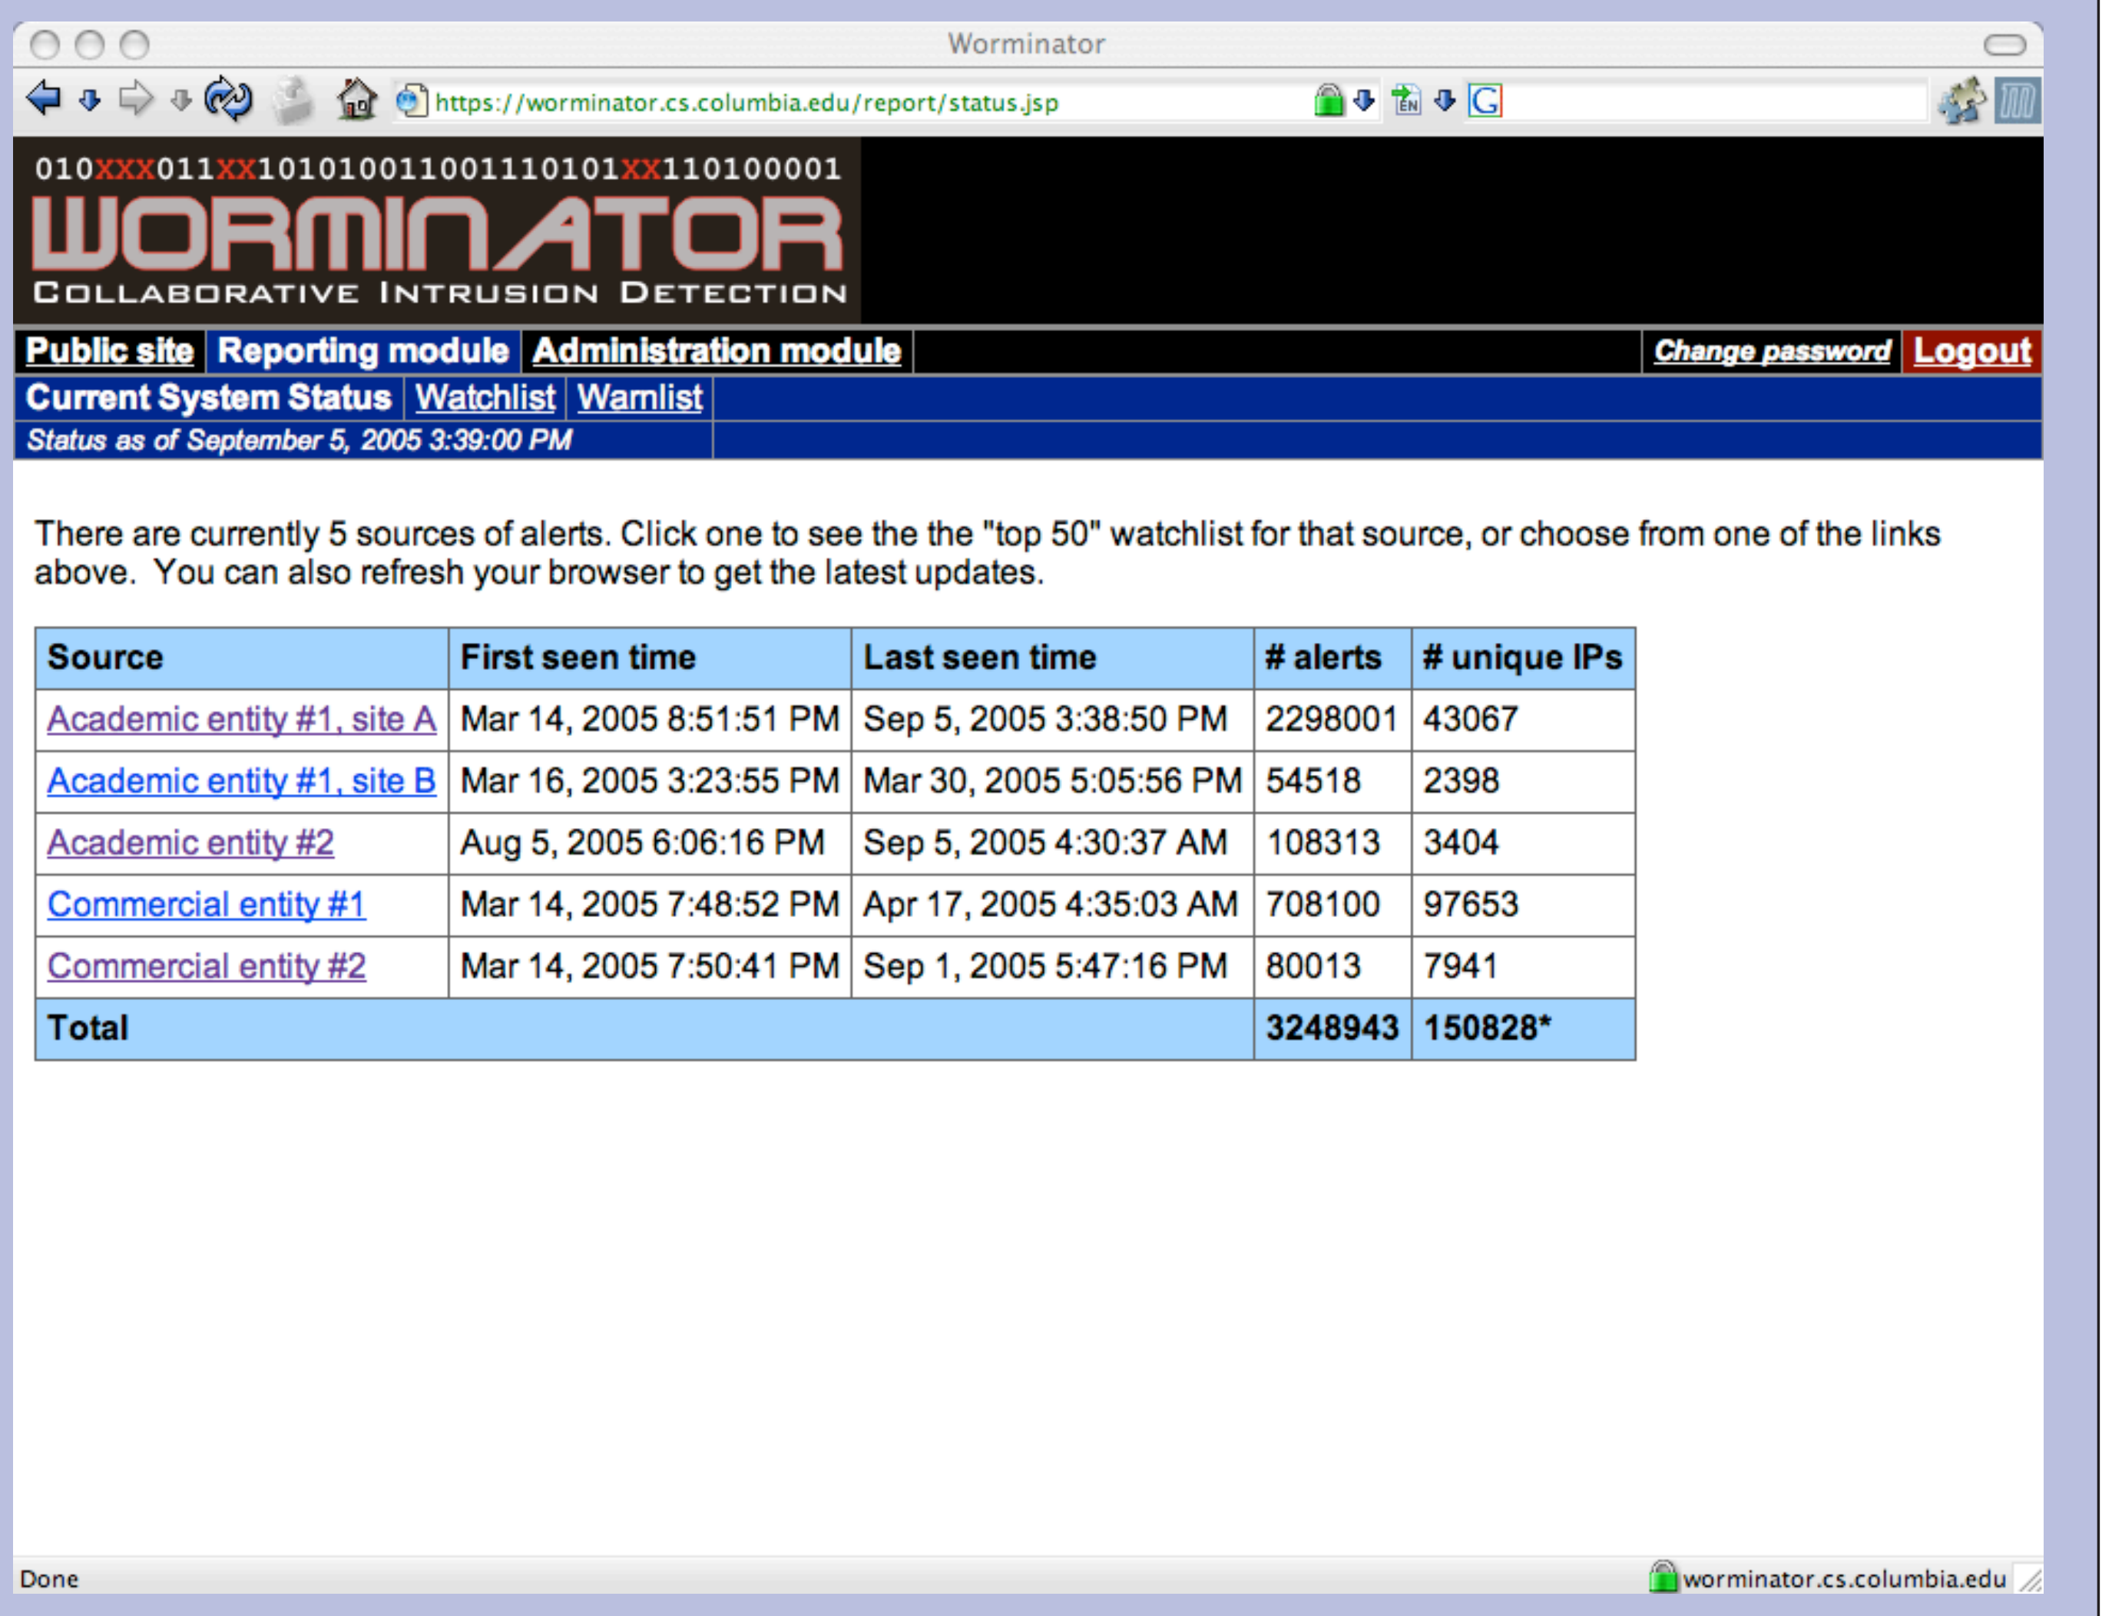
Worminator
https://worminator.cs.columbia.edu/report/status.jsp
010xxx011xx1010100110011101010xx110100001
WORMINATOR
Collaborative Intrusion Detection
Public site
Reporting module
Administration module
Change password
Logout
Current System Status
Watchlist
Warnlist
Done
worminator.cs.columbia.edu

Status as of September 5, 2005 3:39:00 PM

There are currently 5 sources of alerts. Click one to see the the "top 50" watchlist for that source, or choose from one of the links above. You can also refresh your browser to get the latest updates.

Source	First seen time	Last seen time	# alerts	# unique IPs
Academic entity #1, site A	Mar 14, 2005 8:51:51 PM	Sep 5, 2005 3:38:50 PM	2298001	43067
Academic entity #1, site B	Mar 16, 2005 3:23:55 PM	Mar 30, 2005 5:05:56 PM	54518	2398
Academic entity #2	Aug 5, 2005 6:06:16 PM	Sep 5, 2005 4:30:37 AM	108313	3404
Commercial entity #1	Mar 14, 2005 7:48:52 PM	Apr 17, 2005 4:35:03 AM	708100	97653
Commercial entity #2	Mar 14, 2005 7:50:41 PM	Sep 1, 2005 5:47:16 PM	80013	7941
Total			**3248943**	**150828***

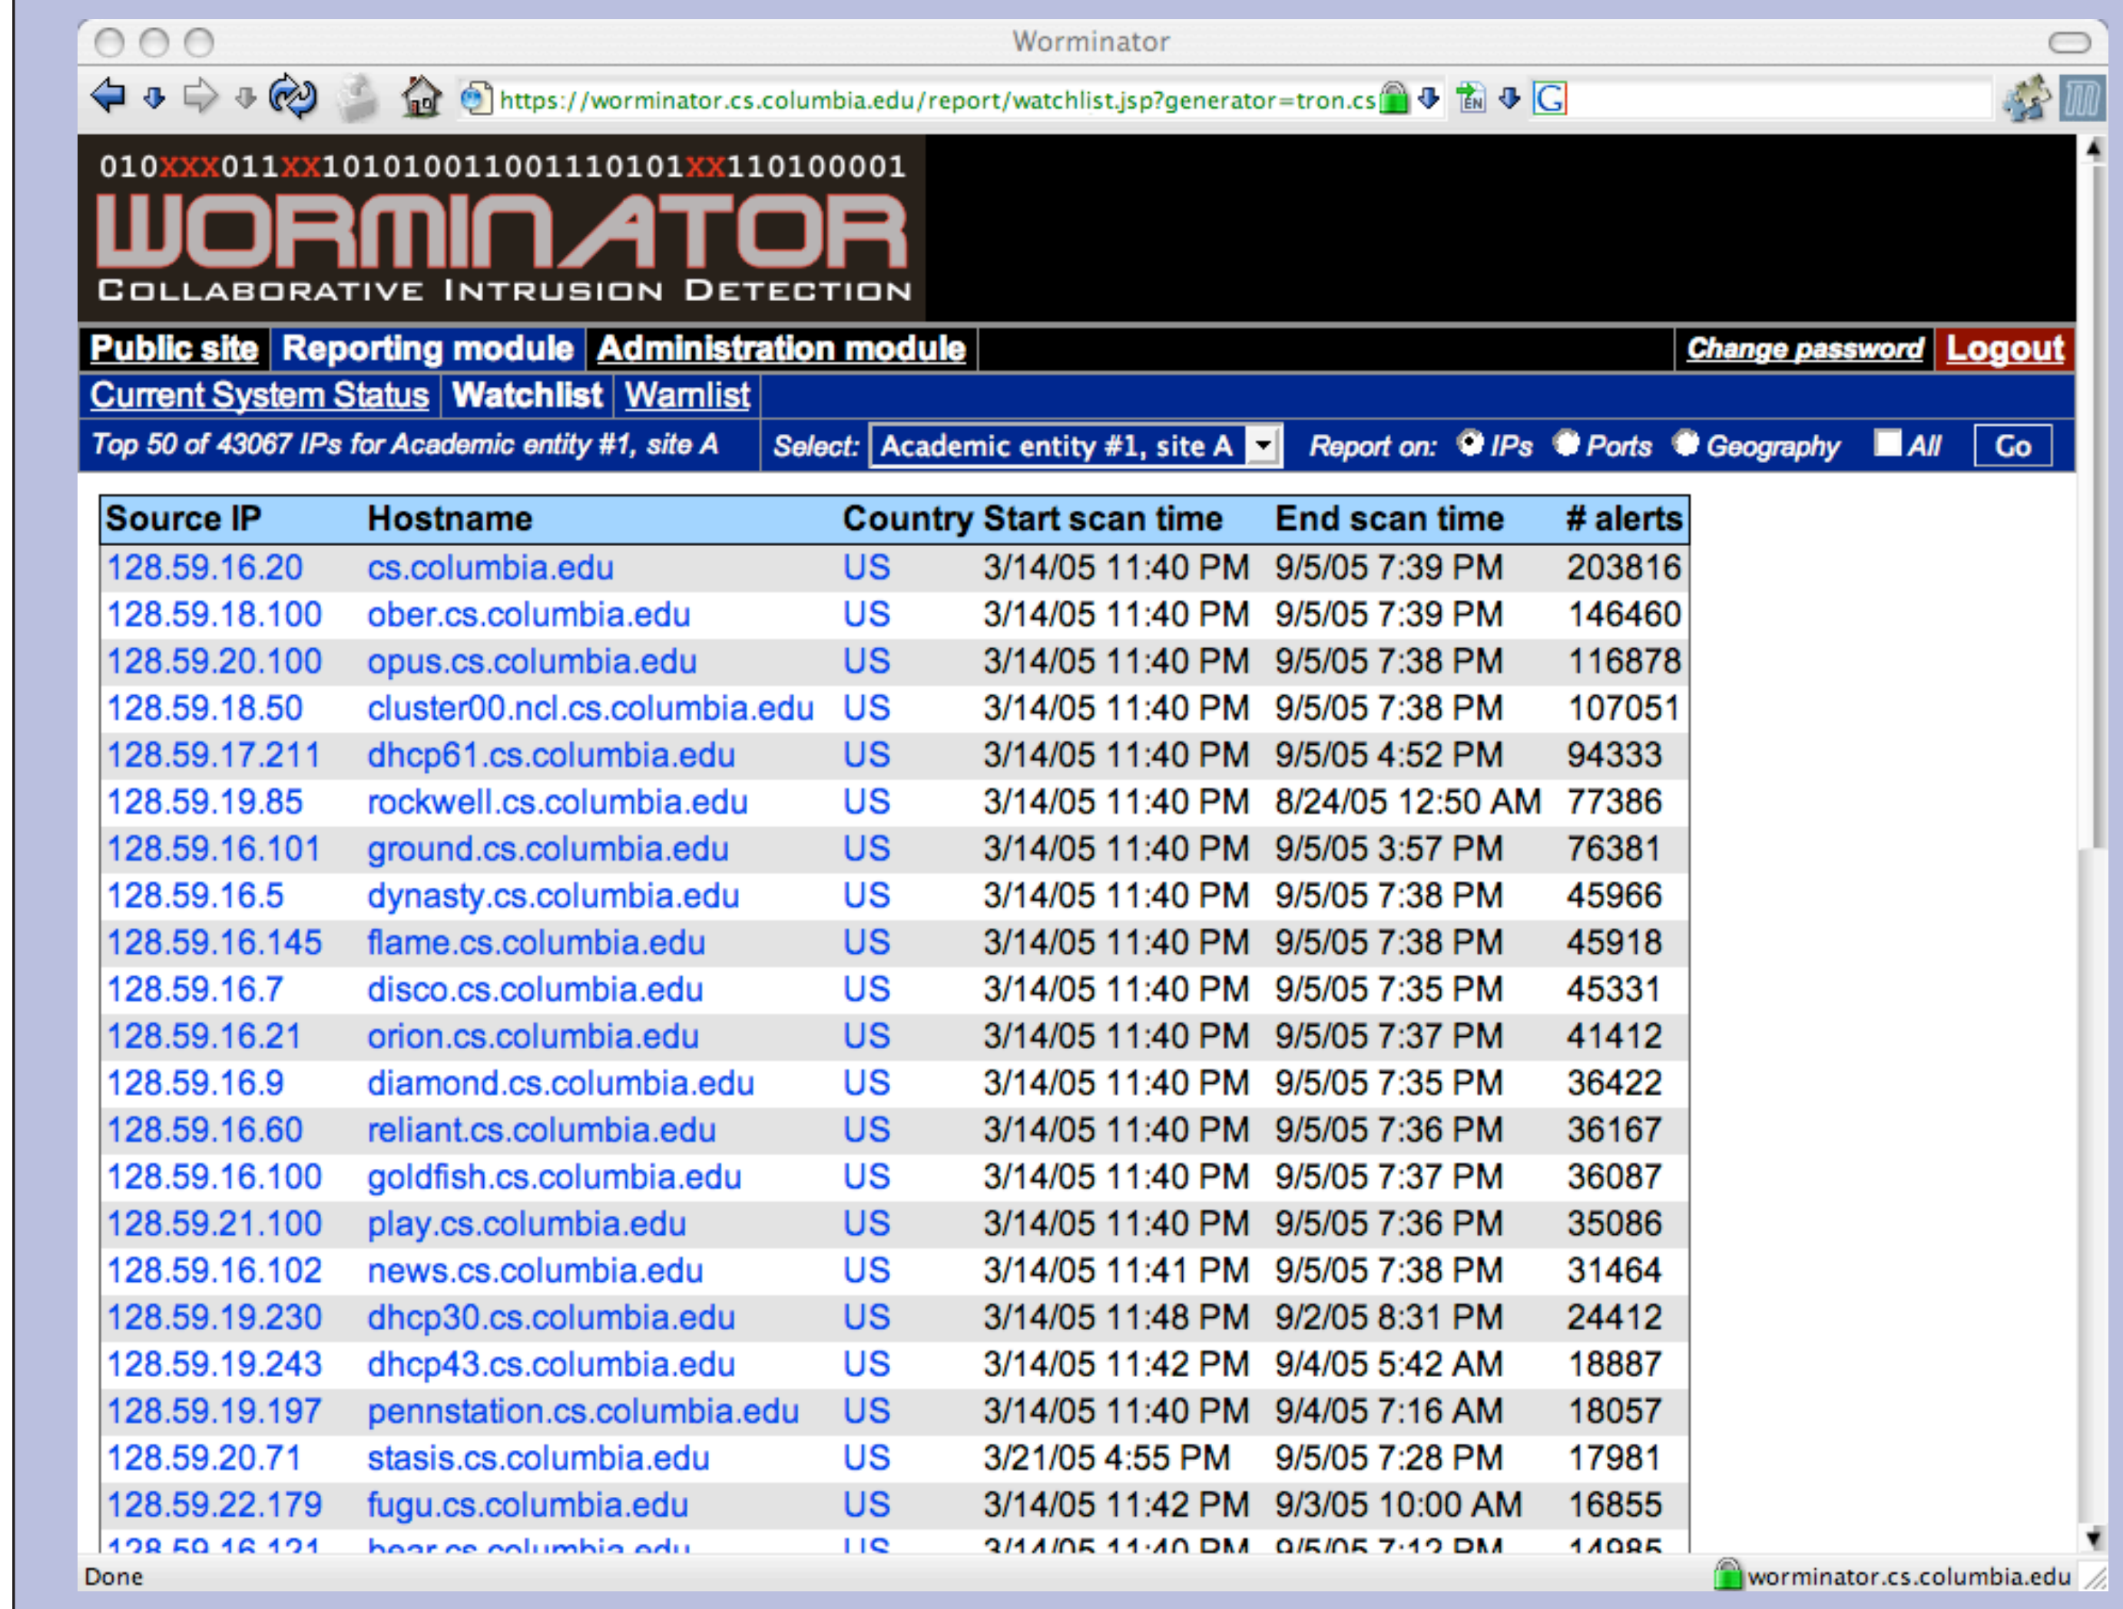
Worminator
https://worminator.cs.columbia.edu/report/watchlist.jsp?generator=tron.cs
010xxx011xx1010100110011101010xx110100001
WORMINATOR
Collaborative Intrusion Detection
Public site
Reporting module
Administration module
Change password
Logout
Current System Status
Watchlist
Warnlist
Go
Done
worminator.cs.columbia.edu

Top 50 of 43067 IPs for Academic entity #1, site A — *Select:* Academic entity #1, site A — *Report on:* ● *IPs* ○ *Ports* ○ *Geography* □ *All*

Source IP	Hostname	Country	Start scan time	End scan time	# alerts
128.59.16.20	cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:39 PM	203816
128.59.18.100	ober.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:39 PM	146460
128.59.20.100	opus.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:38 PM	116878
128.59.18.50	cluster00.ncl.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:38 PM	107051
128.59.17.211	dhcp61.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 4:52 PM	94333
128.59.19.85	rockwell.cs.columbia.edu	US	3/14/05 11:40 PM	8/24/05 12:50 AM	77386
128.59.16.101	ground.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 3:57 PM	76381
128.59.16.5	dynasty.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:38 PM	45966
128.59.16.145	flame.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:38 PM	45918
128.59.16.7	disco.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:35 PM	45331
128.59.16.21	orion.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:37 PM	41412
128.59.16.9	diamond.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:35 PM	36422
128.59.16.60	reliant.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:36 PM	36167
128.59.16.100	goldfish.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:37 PM	36087
128.59.21.100	play.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:36 PM	35086
128.59.16.102	news.cs.columbia.edu	US	3/14/05 11:41 PM	9/5/05 7:38 PM	31464
128.59.19.230	dhcp30.cs.columbia.edu	US	3/14/05 11:48 PM	9/2/05 8:31 PM	24412
128.59.19.243	dhcp43.cs.columbia.edu	US	3/14/05 11:42 PM	9/4/05 5:42 AM	18887
128.59.19.197	pennstation.cs.columbia.edu	US	3/14/05 11:40 PM	9/4/05 7:16 AM	18057
128.59.20.71	stasis.cs.columbia.edu	US	3/21/05 4:55 PM	9/5/05 7:28 PM	17981
128.59.22.179	fugu.cs.columbia.edu	US	3/14/05 11:42 PM	9/3/05 10:00 AM	16855
128.59.16.121	bear.cs.columbia.edu	US	3/14/05 11:40 PM	9/5/05 7:12 PM	14985

Single-site alert report

(Contains noise as CUCS does not have a firewall)

Multiple-site warnlist correlating alerts

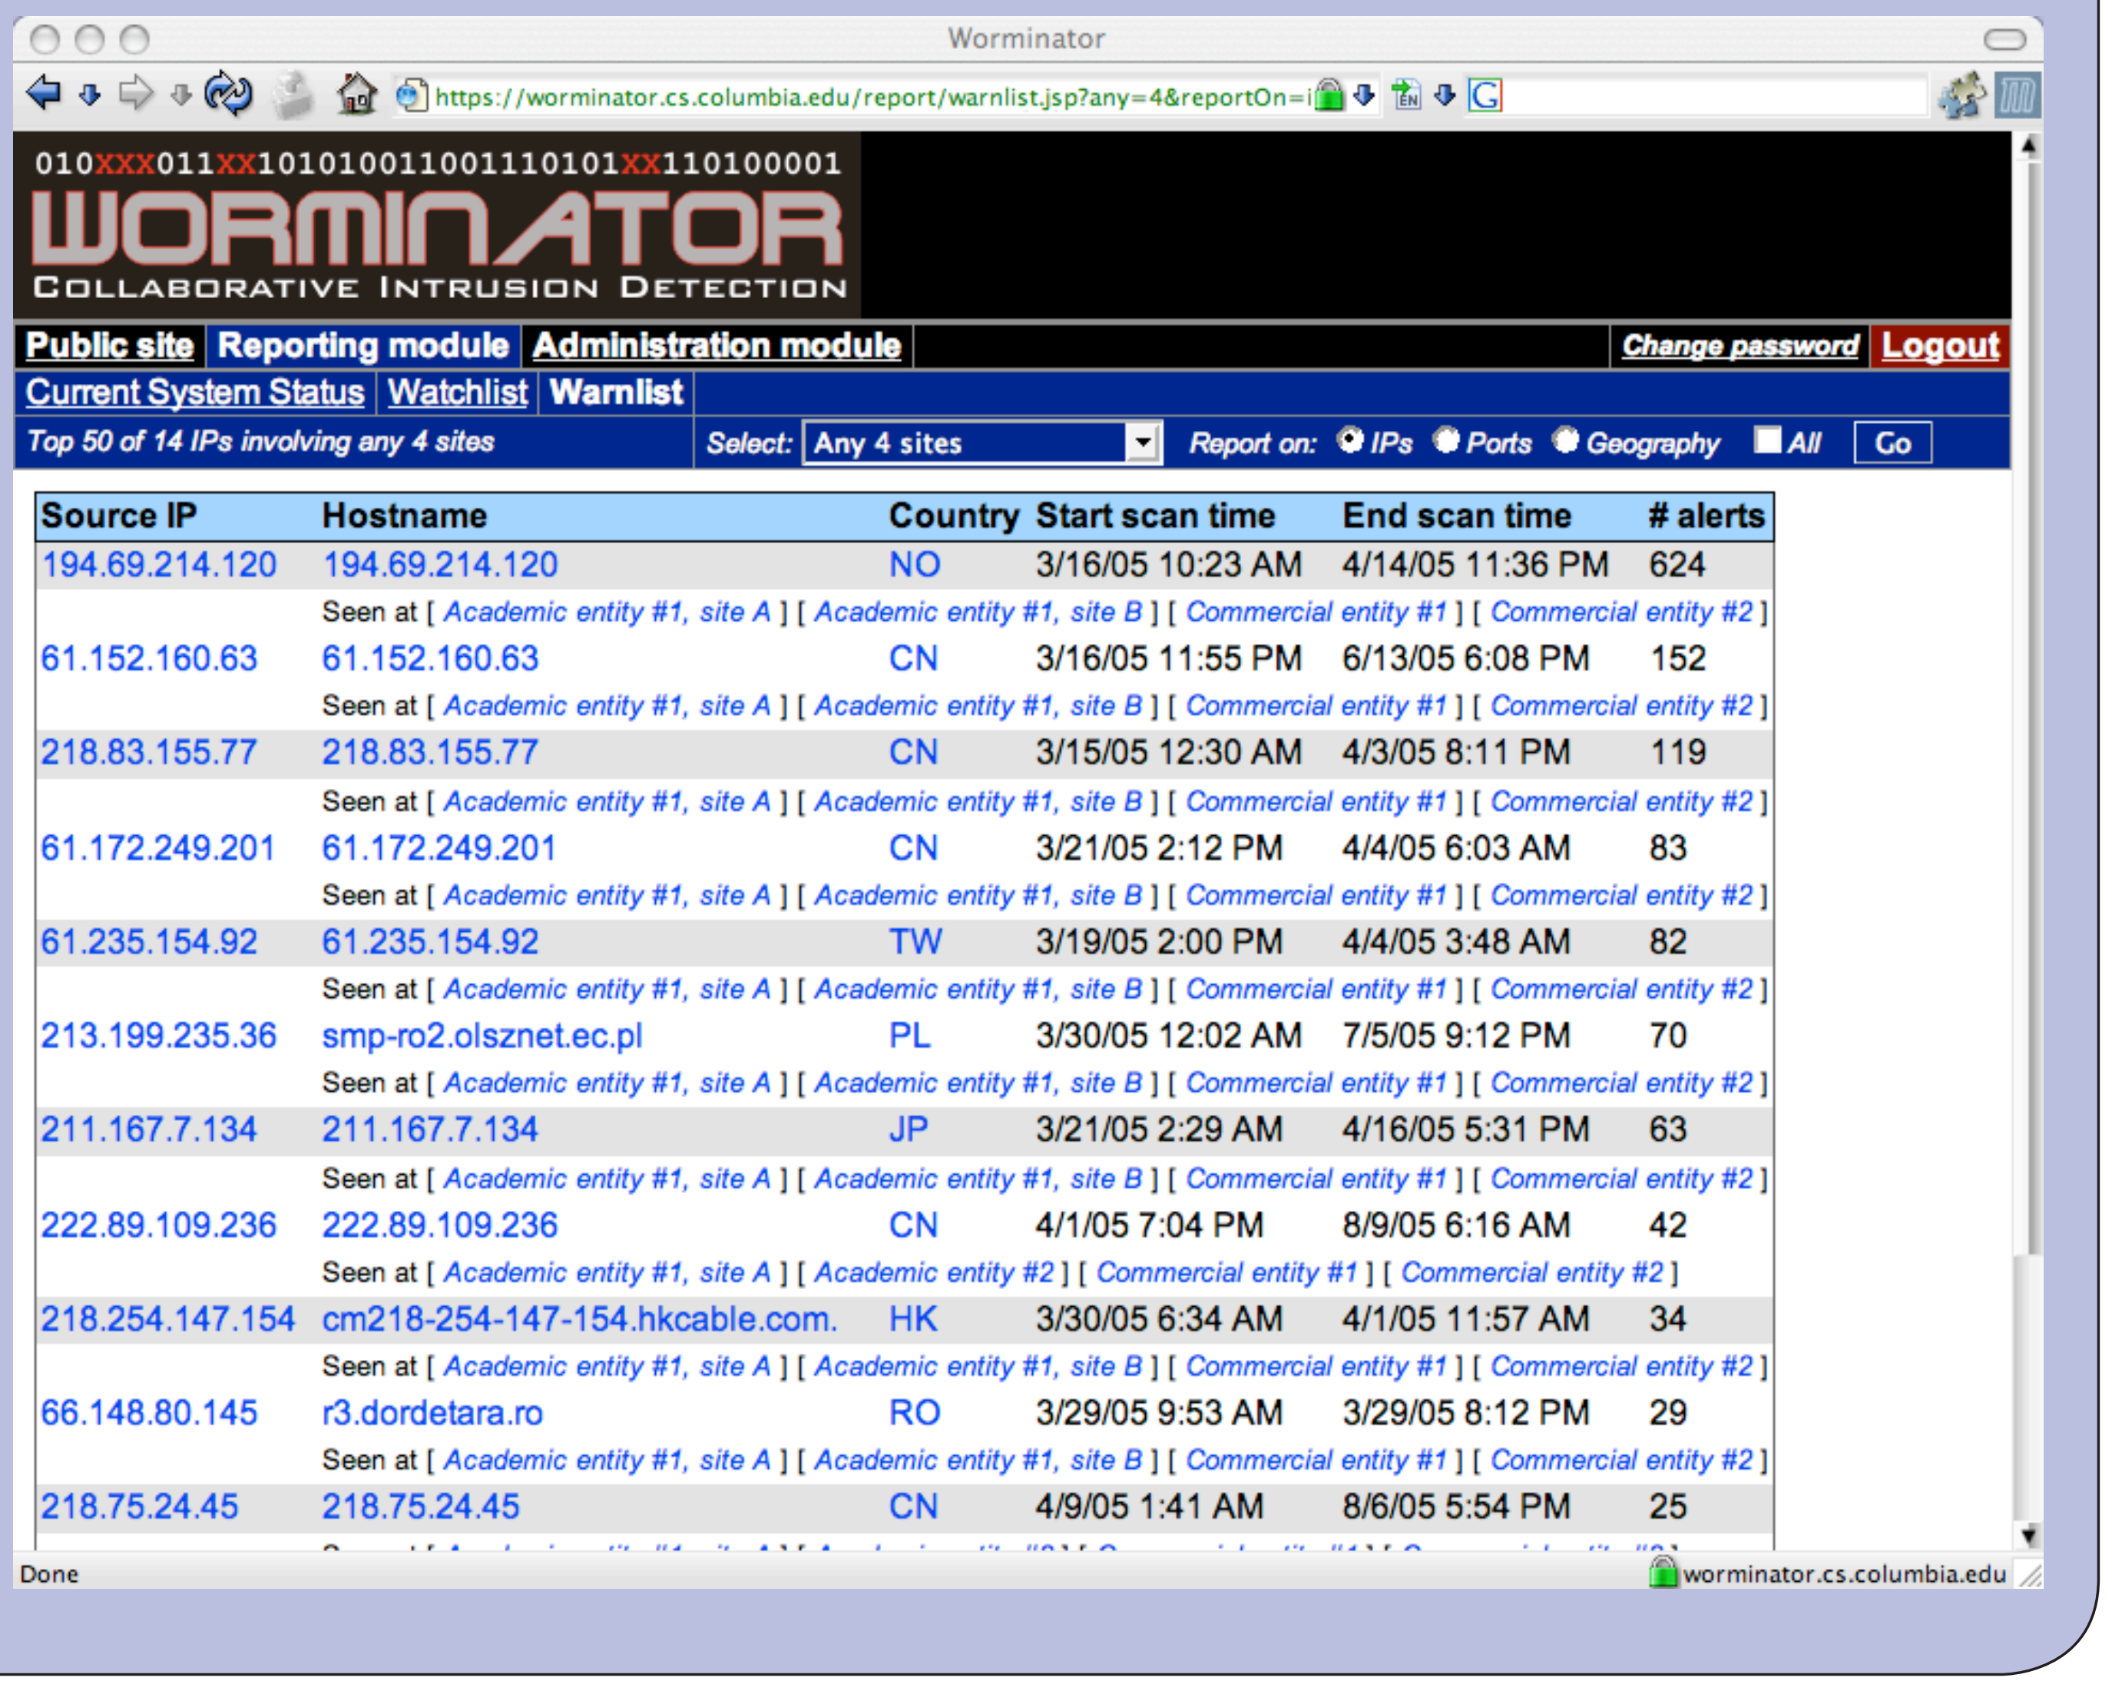
Worminator
https://worminator.cs.columbia.edu/report/warnlist.jsp?any=4&reportOn=i
010xxx011xx1010100110011101010xx110100001
WORMINATOR
Collaborative Intrusion Detection
Public site
Reporting module
Administration module
Change password
Logout
Current System Status
Watchlist
Warnlist
Go
Done
worminator.cs.columbia.edu

Top 50 of 14 IPs involving any 4 sites — *Select:* Any 4 sites — *Report on:* ● *IPs* ○ *Ports* ○ *Geography* □ *All*

Source IP	Hostname	Country	Start scan time	End scan time	# alerts
194.69.214.120	194.69.214.120	NO	3/16/05 10:23 AM	4/14/05 11:36 PM	624
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
61.152.160.63	61.152.160.63	CN	3/16/05 11:55 PM	6/13/05 6:08 PM	152
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
218.83.155.77	218.83.155.77	CN	3/15/05 12:30 AM	4/3/05 8:11 PM	119
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
61.172.249.201	61.172.249.201	CN	3/21/05 2:12 PM	4/4/05 6:03 AM	83
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
61.235.154.92	61.235.154.92	TW	3/19/05 2:00 PM	4/4/05 3:48 AM	82
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
213.199.235.36	smp-ro2.olsznet.ec.pl	PL	3/30/05 12:02 AM	7/5/05 9:12 PM	70
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
211.167.7.134	211.167.7.134	JP	3/21/05 2:29 AM	4/16/05 5:31 PM	63
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
222.89.109.236	222.89.109.236	CN	4/1/05 7:04 PM	8/9/05 6:16 AM	42
	Seen at [*Academic entity #1, site A*] [*Academic entity #2*] [*Commercial entity #1*] [*Commercial entity #2*]				
218.254.147.154	cm218-254-147-154.hkcable.com.	HK	3/30/05 6:34 AM	4/1/05 11:57 AM	34
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
66.148.80.145	r3.dordetara.ro	RO	3/29/05 9:53 AM	3/29/05 8:12 PM	29
	Seen at [*Academic entity #1, site A*] [*Academic entity #1, site B*] [*Commercial entity #1*] [*Commercial entity #2*]				
218.75.24.45	218.75.24.45	CN	4/9/05 1:41 AM	8/6/05 5:54 PM	25
	Seen at [*Academic entity #1, site A*] [*Academic entity #2*] [*Commercial entity #1*] [*Commercial entity #2*]				

Results and Experiences

- *Significant reduction in the number of alerts (orders of magnitude), enabling more aggressive response*

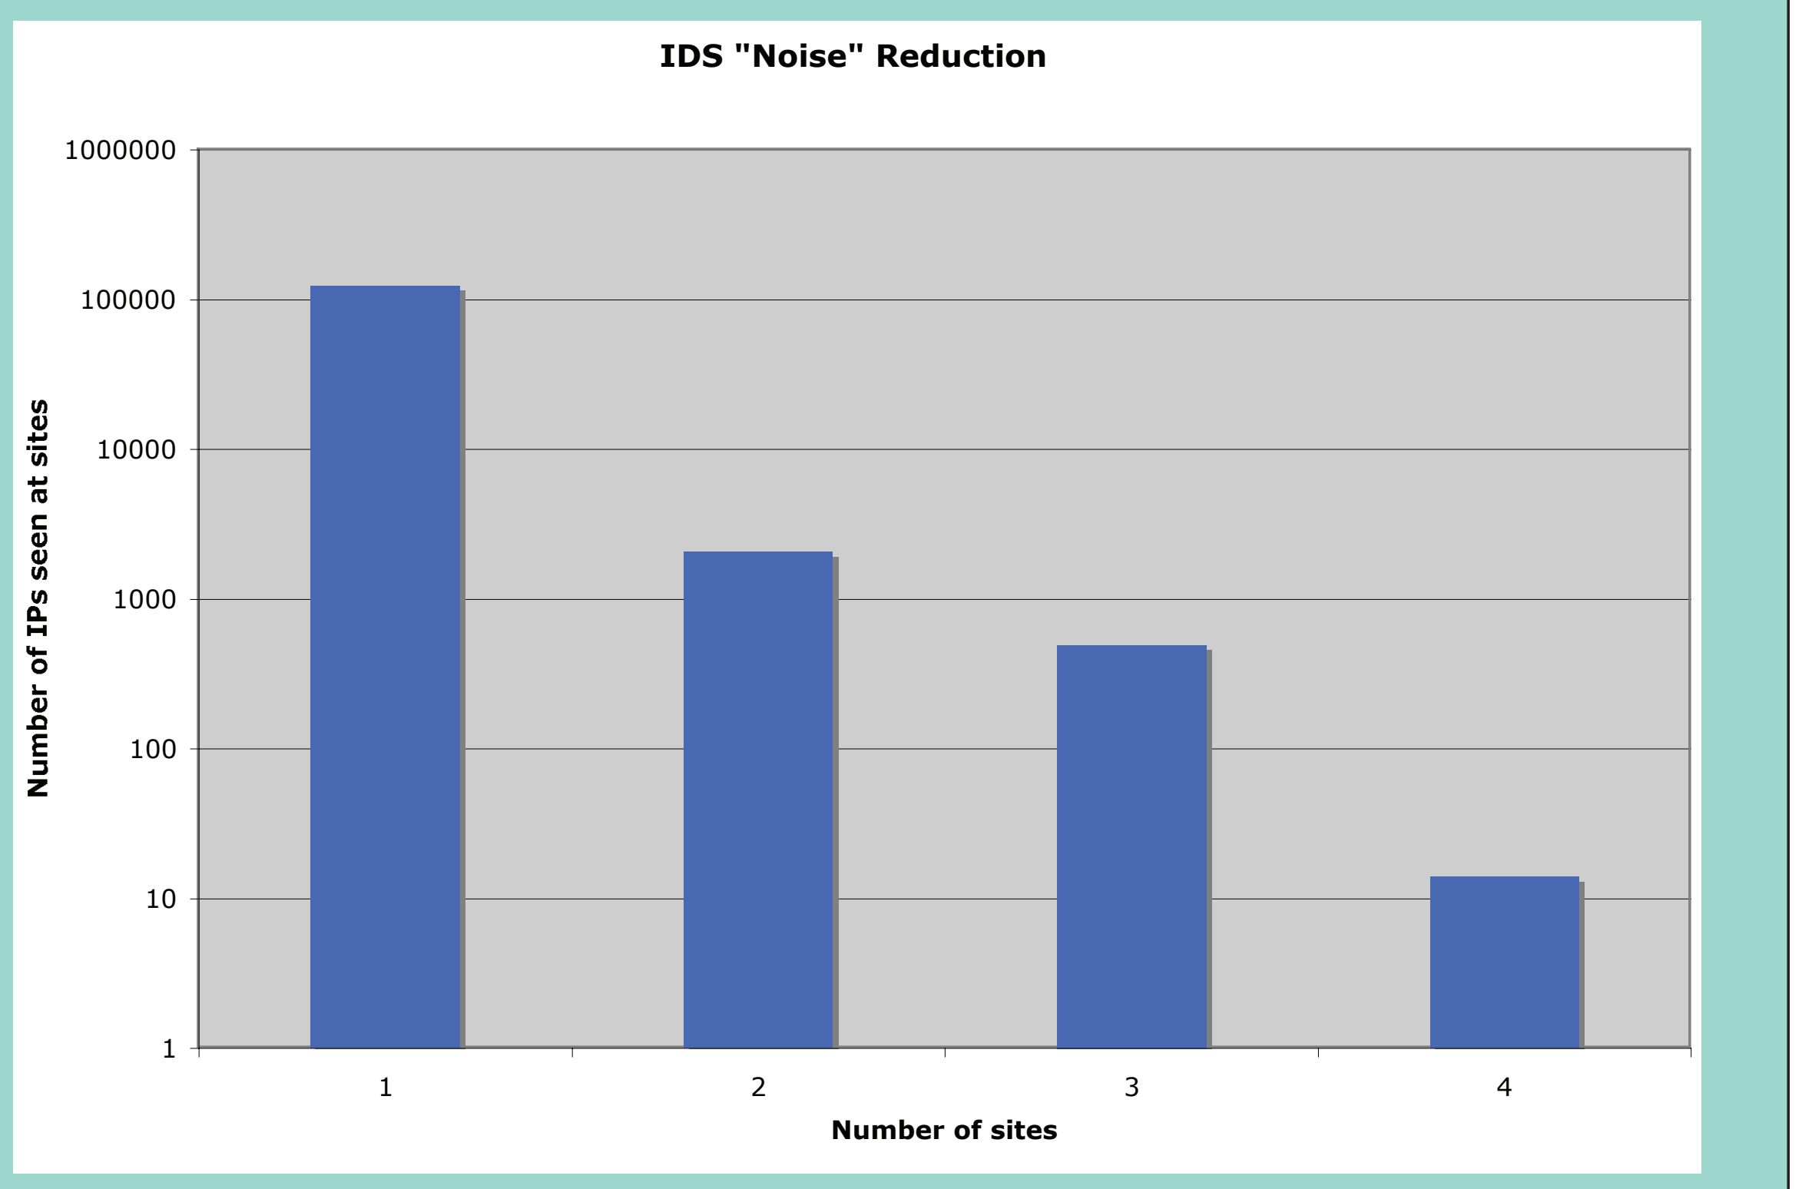
IDS "Noise" Reduction

- *Scanning behavior varies by source; some observe many sites rapidly; others tend to spread out their activities*

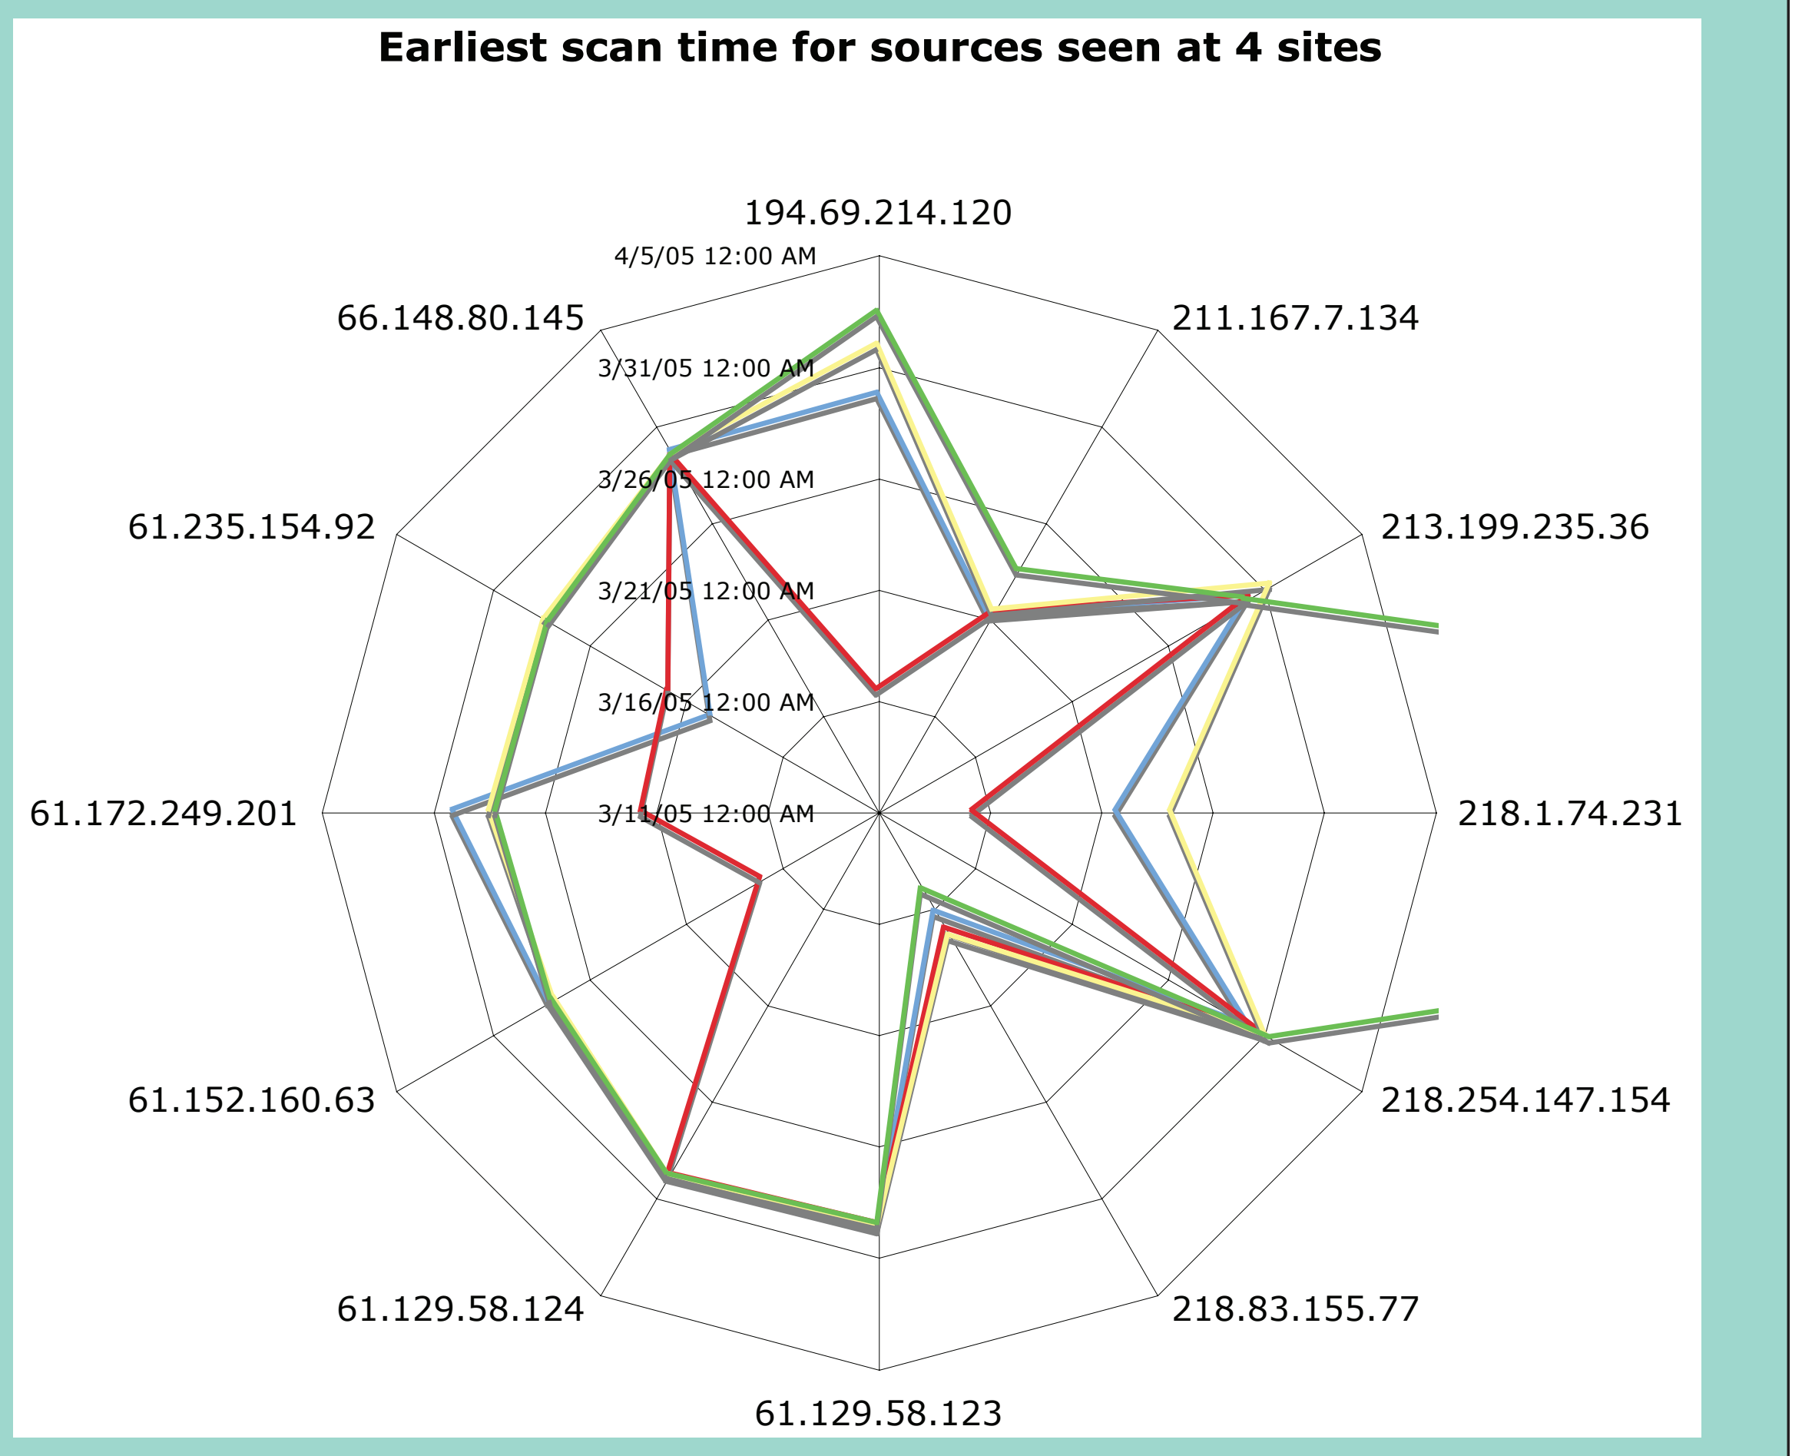
Earliest scan time for sources seen at 4 sites

- *Most popular ports targeted are Windows services or backdoors installed by others (worms, etc.)*

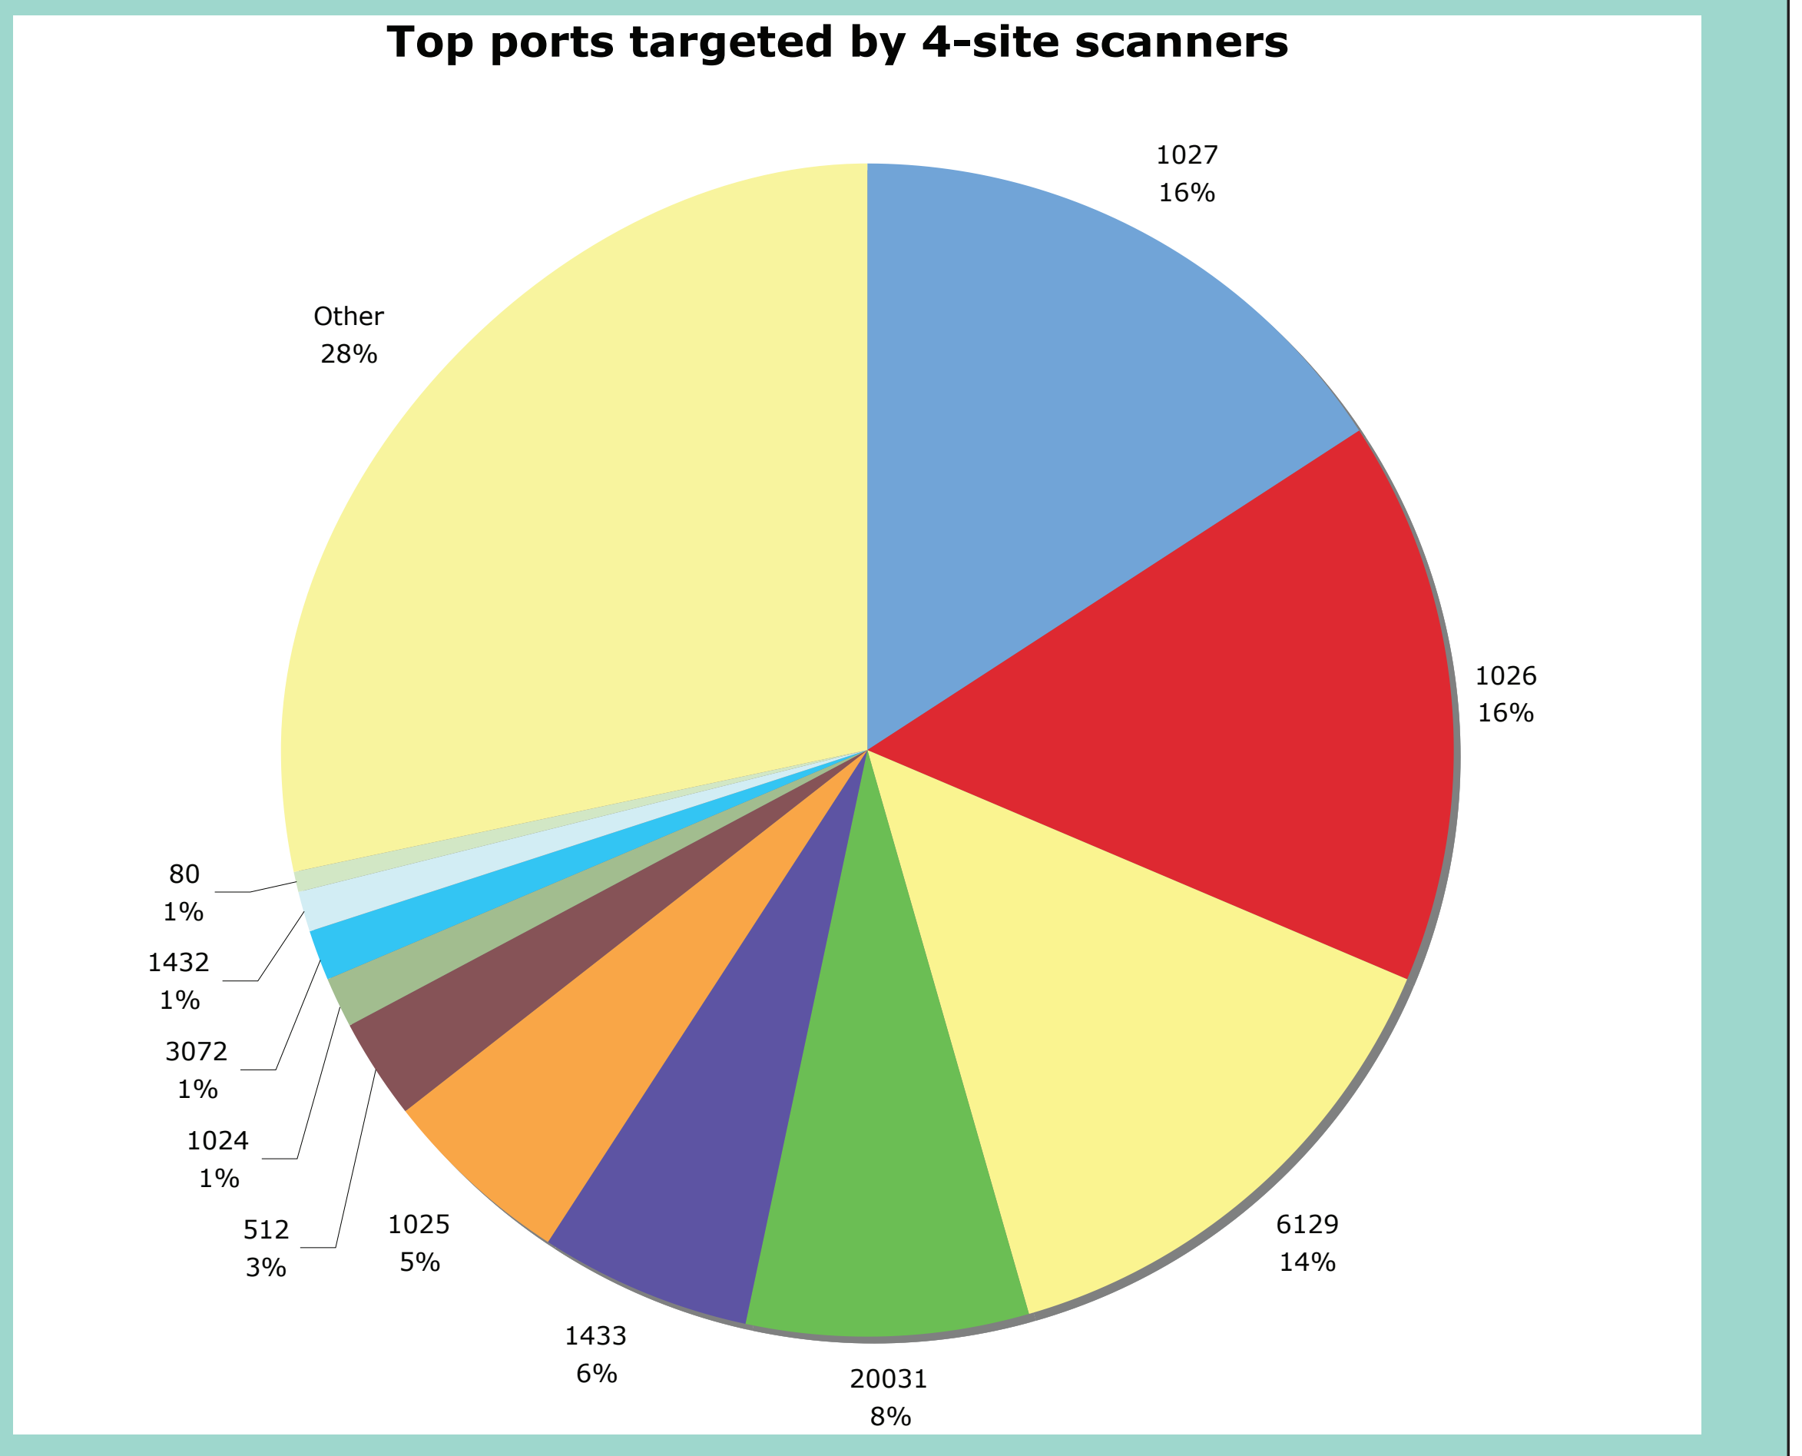
Top ports targeted by 4-site scanners

- Biggest challenge is getting organizations to participate - not for technical reasons, but rather due to organizational, legal, or political issues
- Supporting privacy-preserving mechanisms makes a big difference, especially when non-academic sites are involved

Future work

- Additional site deployment
- Longitudinal study on incoming data
- Integrate Whirlpool network scheduling model
- Integrate support for PAYL anomaly detector to automate content signature and model exchange
- Research into using Worminator in more diverse network environments, e.g., MANETs