A Traffic Analysis of a Small Private Network Compromised by an On-line Gaming Host

Ron McLeod, BCSc, MCSc.
Director - Corporate Development Telecom Applications
Research Alliance
Doctoral Student, Faculty of Computer Science, Dalhousie
University

Abstract

(Abridged)

In the early months of 2006 a small private network (the Network) suffered a Noticeable degrading of its network performance. A network traffic capture and analysis was conducted and used to investigate the network performance issues. This paper presents partial results of that analysis. During the first analysis of the captured data it was discovered that the Network contained a host that had been compromised at some time in the past and was currently being used to support the on-line gaming activity of over 174,000 distinct player source addresses around the globe. The initial finding was the result of a manual investigation of unusual time and volume traffic spikes from arbitrarily chosen time slices. Subsequent work was conducted on searching for a traffic signature which could be representative of the presence of the Game such that future discovery of Game activity could be automated. Gaming traffic is predominantly UDP traffic of high byte volumes, typically targeted at a given range of destination ports. This analysis also searches for a specific TCP traffic pattern that is suggestive of a Game signature. Network traffic patterns that emerge after access to the compromised host has been closed are labeled as SCAR traffic, for **S**evered **C**onnection **A**nomalous **R**ecords

Presentation Outline

- Summary of the event
- A UDP Profile of the Infection
- The Search for a TCP Signature
- The Search for Residual Traffic (SCAR)
- Concluding Remarks

Event Chronology

- A Traffic Capture was initiated on February 3.
- On February 11 the first slice of data was extracted for analysis.
- On February 13 a Game Server was discovered on a compromised host.
- For the next 30 days this server supported the on-line Game playing of over 174,000 unique Source Addresses.
- During this time the traffic to and from the game server accounted for greater than 50% of the traffic byte volume and 34% of the network flows.

Network Description

- A Multi-tenant Network consisting of:
  - ~ 40 user assigned hosts, actual number subject to minor fluctuations over time.
  - ~40 special hosts not assigned to individual users. These hosts form parts of various temporary development and experimental environments.
  - Users were apprised that Network flow data was now being captured for experimental and management reasons.
  - Payload data was neither collected nor examined.
  - Analysts did not have access to the content of specific hosts for further investigation.
  - For confidentiality reasons the identity of the Network is not specified in this Presentation.

First Capture

- On February 11 the first sample of network traffic (Slice) was extracted for analysis.
- The time period from midnight to 7:00 AM local time on February 8 was chosen for the first slice.
- This was partially a random choice and partially due to the fact that the author expected minimal traffic volumes during this time. The institution which houses the Network is closed during these hours.
- Only Non-port 80 and Non-Null traffic was initially examined.

The First Capture Image

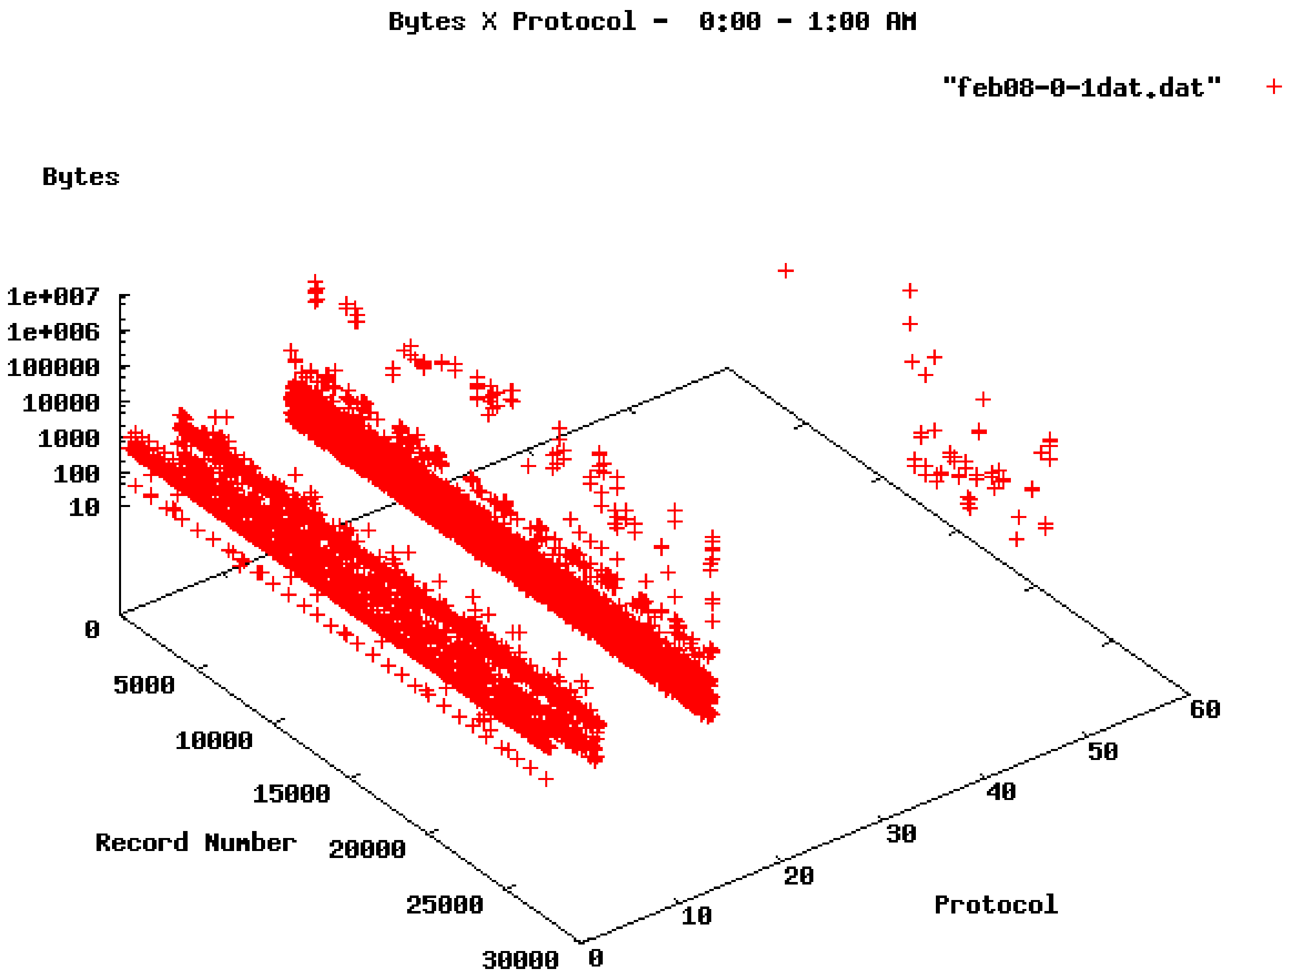

First Traffic Capture Observations

- Protocols are as one would expect (1,6,17,50*,53*).
- Size raised suspicion: 27,000+ records per hour *seemed* large for a network with no active users.

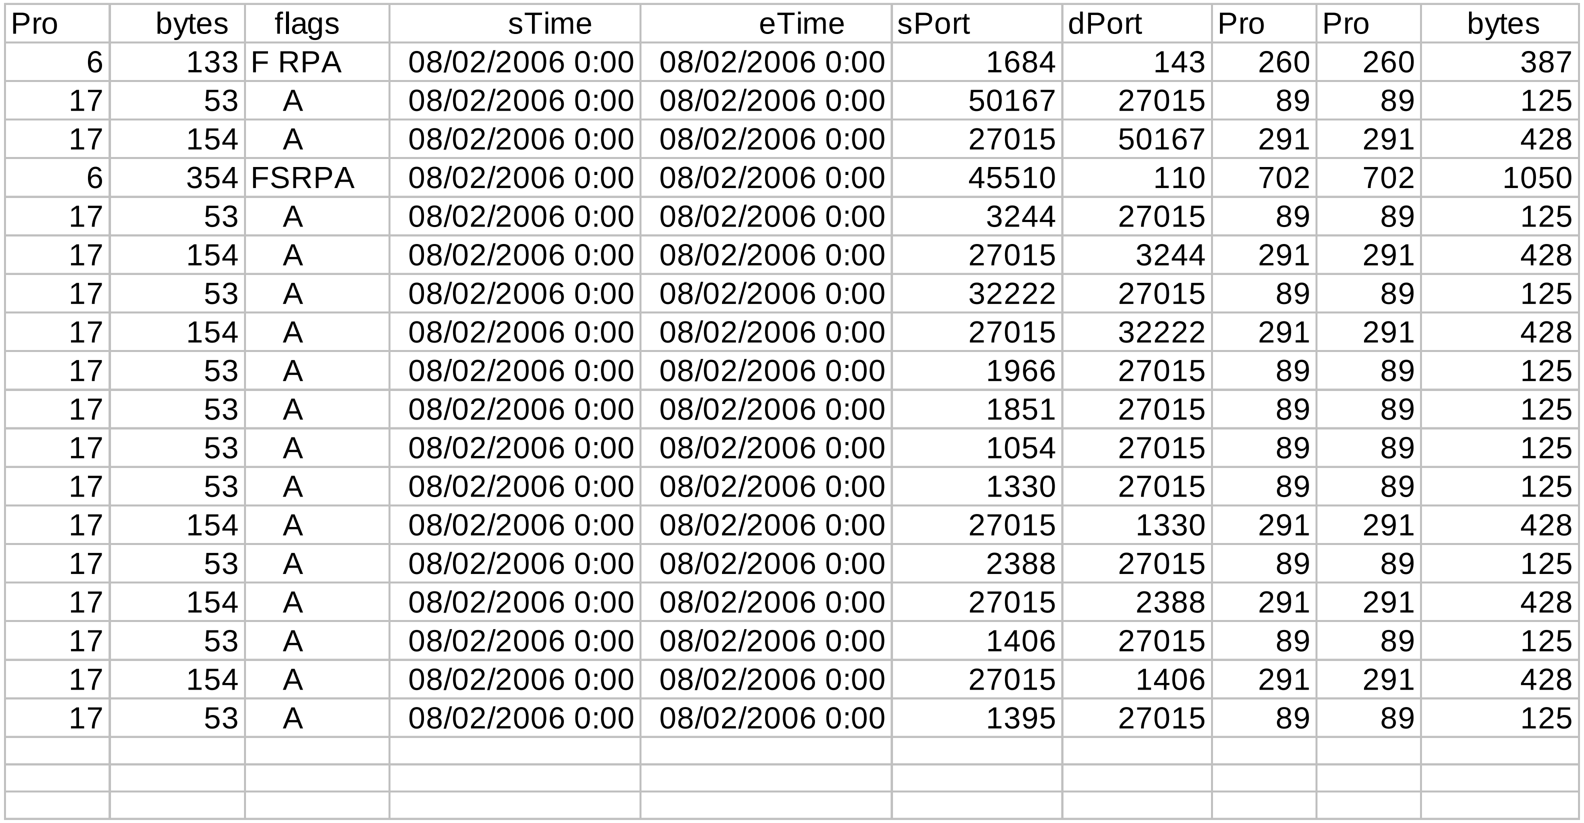

Pro	bytes	flags	sTime	eTime	sPort	dPort	Pro	Pro	bytes
6	133	F RPA	08/02/2006 0:00	08/02/2006 0:00	1684	143	260	260	387
17	53	A	08/02/2006 0:00	08/02/2006 0:00	50167	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	50167	291	291	428
6	354	FSRPA	08/02/2006 0:00	08/02/2006 0:00	45510	110	702	702	1050
17	53	A	08/02/2006 0:00	08/02/2006 0:00	3244	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	3244	291	291	428
17	53	A	08/02/2006 0:00	08/02/2006 0:00	32222	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	32222	291	291	428
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1966	27015	89	89	125
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1851	27015	89	89	125
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1054	27015	89	89	125
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1330	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	1330	291	291	428
17	53	A	08/02/2006 0:00	08/02/2006 0:00	2388	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	2388	291	291	428
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1406	27015	89	89	125
17	154	A	08/02/2006 0:00	08/02/2006 0:00	27015	1406	291	291	428
17	53	A	08/02/2006 0:00	08/02/2006 0:00	1395	27015	89	89	125

First Traffic Capture Observations

- Records were then ordered by byte size

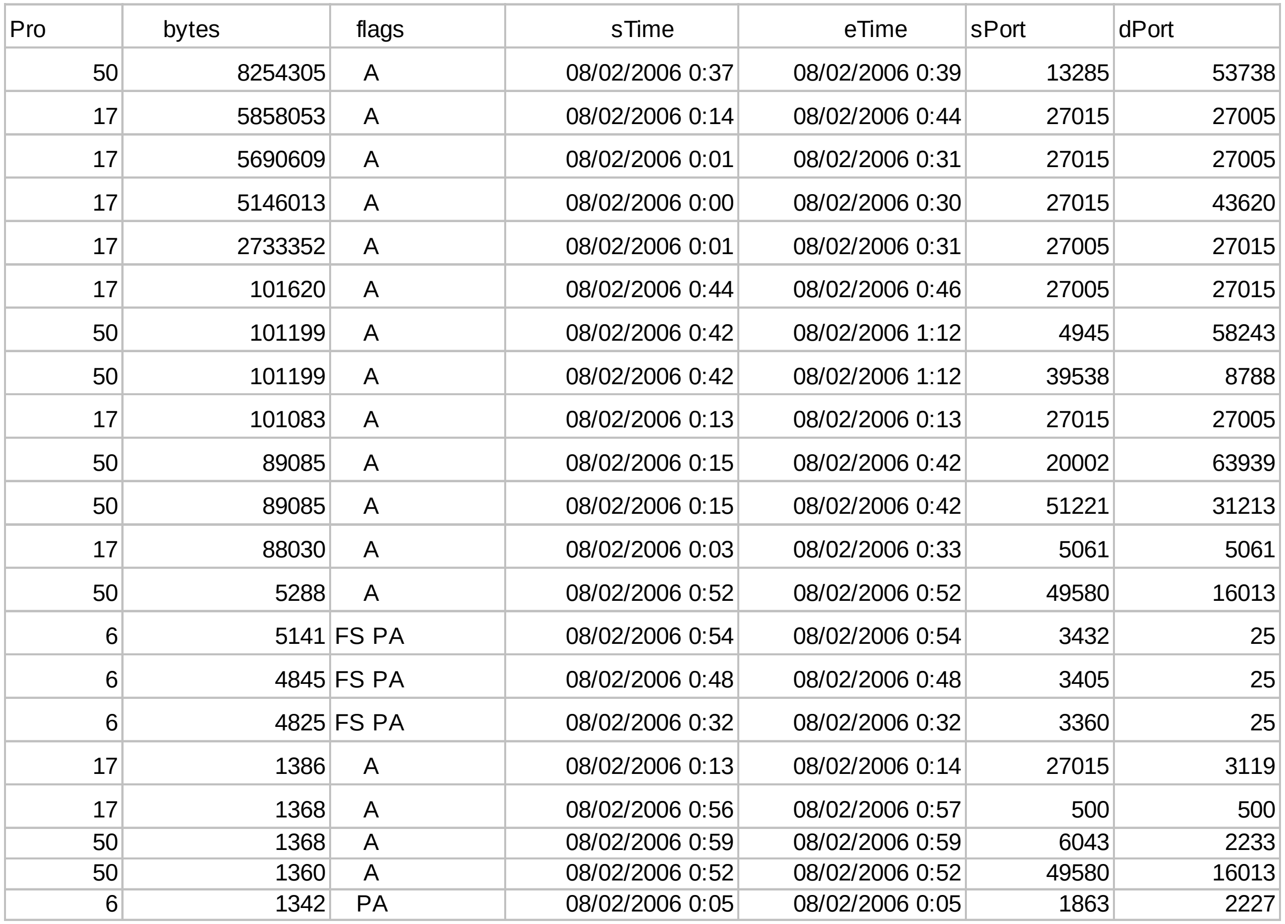

Pro	bytes	flags	sTime	eTime	sPort	dPort
50	8254305	A	08/02/2006 0:37	08/02/2006 0:39	13285	53738
17	5858053	A	08/02/2006 0:14	08/02/2006 0:44	27015	27005
17	5690609	A	08/02/2006 0:01	08/02/2006 0:31	27015	27005
17	5146013	A	08/02/2006 0:00	08/02/2006 0:30	27015	43620
17	2733352	A	08/02/2006 0:01	08/02/2006 0:31	27005	27015
17	101620	A	08/02/2006 0:44	08/02/2006 0:46	27005	27015
50	101199	A	08/02/2006 0:42	08/02/2006 1:12	4945	58243
50	101199	A	08/02/2006 0:42	08/02/2006 1:12	39538	8788
17	101083	A	08/02/2006 0:13	08/02/2006 0:13	27015	27005
50	89085	A	08/02/2006 0:15	08/02/2006 0:42	20002	63939
50	89085	A	08/02/2006 0:15	08/02/2006 0:42	51221	31213
17	88030	A	08/02/2006 0:03	08/02/2006 0:33	5061	5061
50	5288	A	08/02/2006 0:52	08/02/2006 0:52	49580	16013
6	5141	FS PA	08/02/2006 0:54	08/02/2006 0:54	3432	25
6	4845	FS PA	08/02/2006 0:48	08/02/2006 0:48	3405	25
6	4825	FS PA	08/02/2006 0:32	08/02/2006 0:32	3360	25
17	1386	A	08/02/2006 0:13	08/02/2006 0:14	27015	3119
17	1368	A	08/02/2006 0:56	08/02/2006 0:57	500	500
50	1368	A	08/02/2006 0:59	08/02/2006 0:59	6043	2233
50	1360	A	08/02/2006 0:52	08/02/2006 0:52	49580	16013
6	1342	PA	08/02/2006 0:05	08/02/2006 0:05	1863	2227

First Traffic Capture Observations

From this sorting it was discovered that a small group of SourceIPs using protocol 17 appeared to be responsible for a large portion of the traffic bytes.

However, given the size of the database it was not immediately apparent if There was a subset of these hosts that were unusually heavier than the rest.

The traffic was then sorted by Source IP and the total bytes over all flow Records were accumulated for each SourceIP.

One SourceIp, labeled Suspicious Host, accounted for more than 56% of the traffic volume in bytes

Total Bytes for 12:00 – 1:00AM	142,129,799
Total Byte Volume for Suspicious Host 12:00 – 1:00AM	79,865,126

Feb 08 UDP Traffic

Bytes

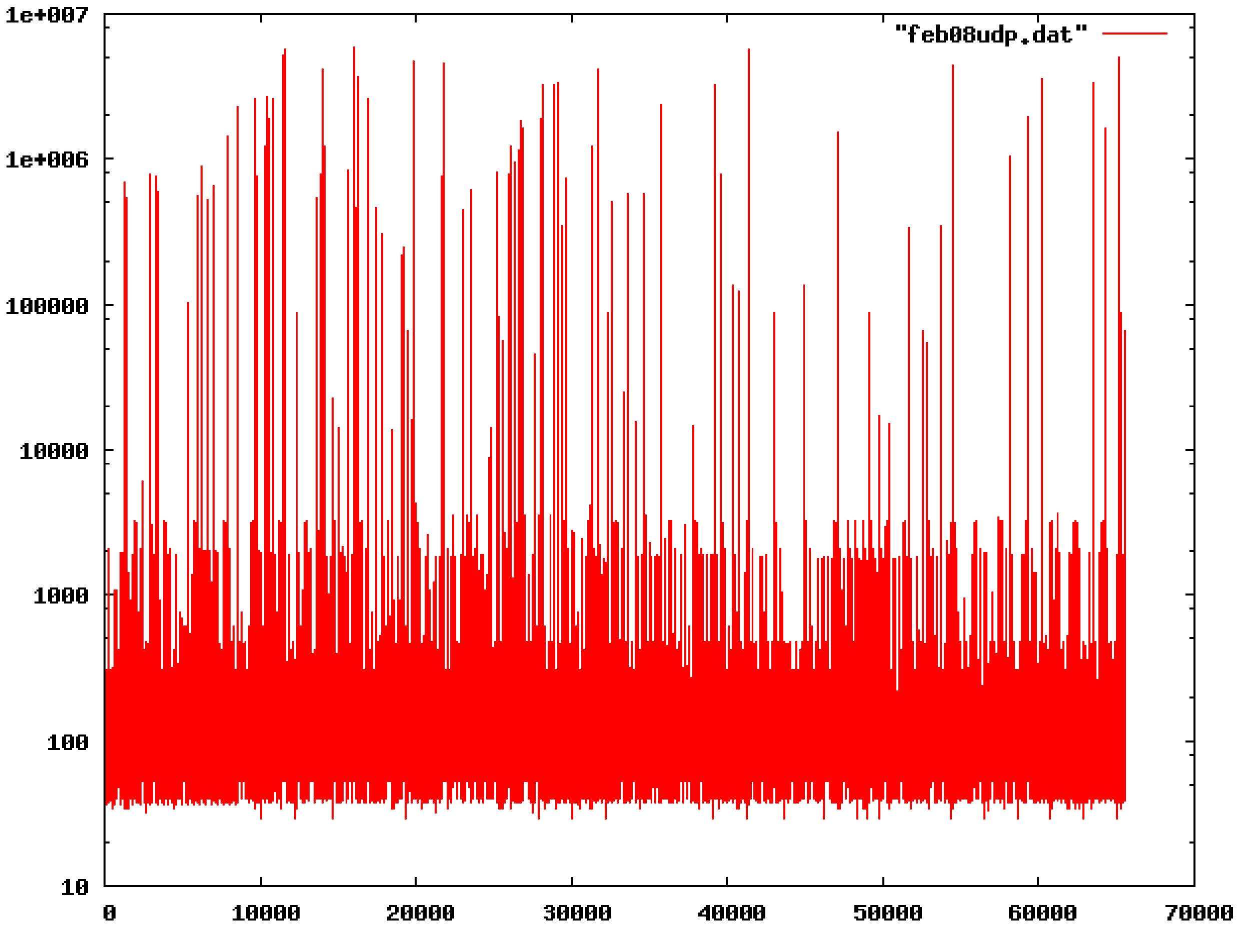

12:00 – 2:00 AM

First Traffic Capture Observations

The Next Step was to examine the use of UDP Ports.

This was done by creating Port Bags and reporting on Key counts greater than 10,000.

Port Bag For Key Counts > 10,000

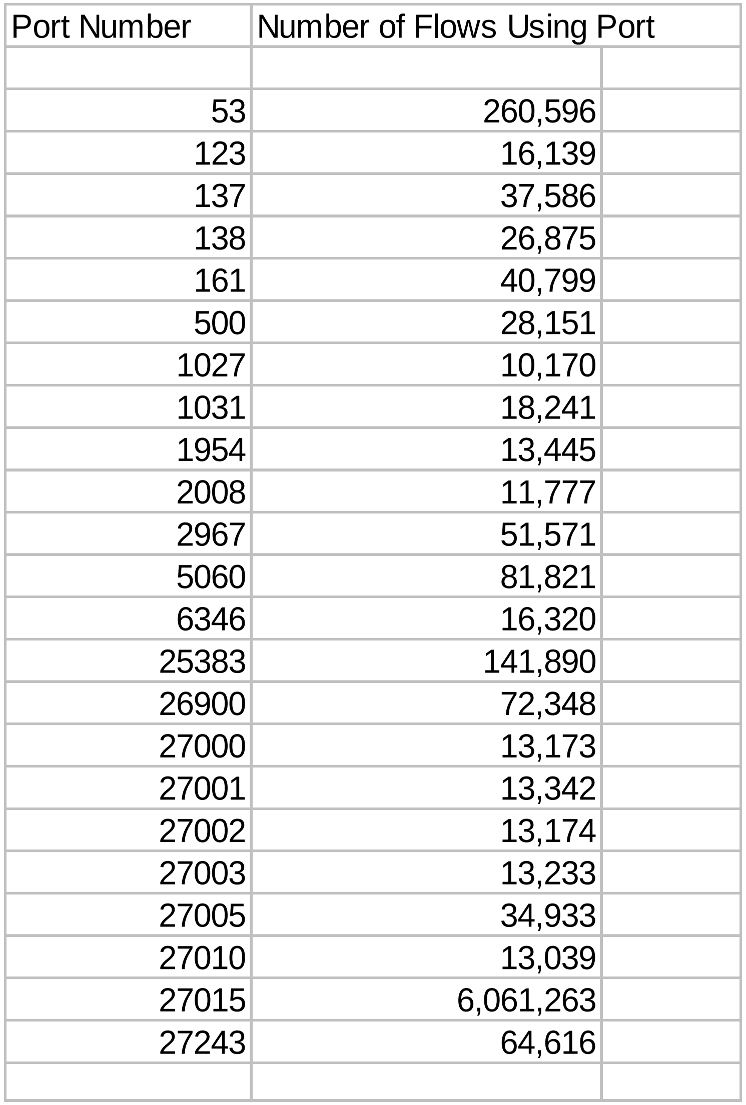

Port Number	Number of Flows Using Port	
53	260,596	
123	16,139	
137	37,586	
138	26,875	
161	40,799	
500	28,151	
1027	10,170	
1031	18,241	
1954	13,445	
2008	11,777	
2967	51,571	
5060	81,821	
6346	16,320	
25383	141,890	
26900	72,348	
27000	13,173	
27001	13,342	
27002	13,174	
27003	13,233	
27005	34,933	
27010	13,039	
27015	6,061,263	
27243	64,616	

Port Bag For Key Counts > 10,000

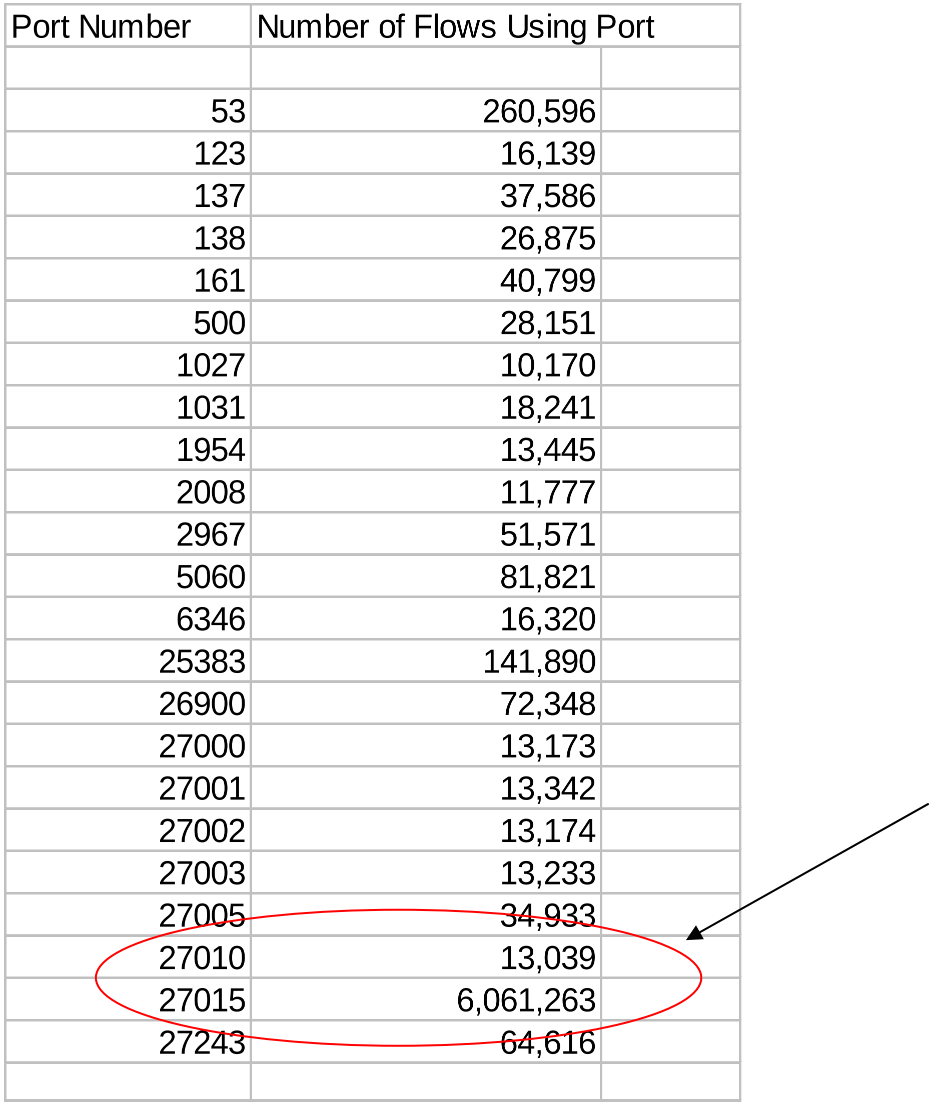

Port Number	Number of Flows Using Port
53	260,596
123	16,139
137	37,586
138	26,875
161	40,799
500	28,151
1027	10,170
1031	18,241
1954	13,445
2008	11,777
2967	51,571
5060	81,821
6346	16,320
25383	141,890
26900	72,348
27000	13,173
27001	13,342
27002	13,174
27003	13,233
27005	34,933
27010	13,039
27015	6,061,263
27243	64,616

Ah hah!

Port Bag For Key Counts > 10,000

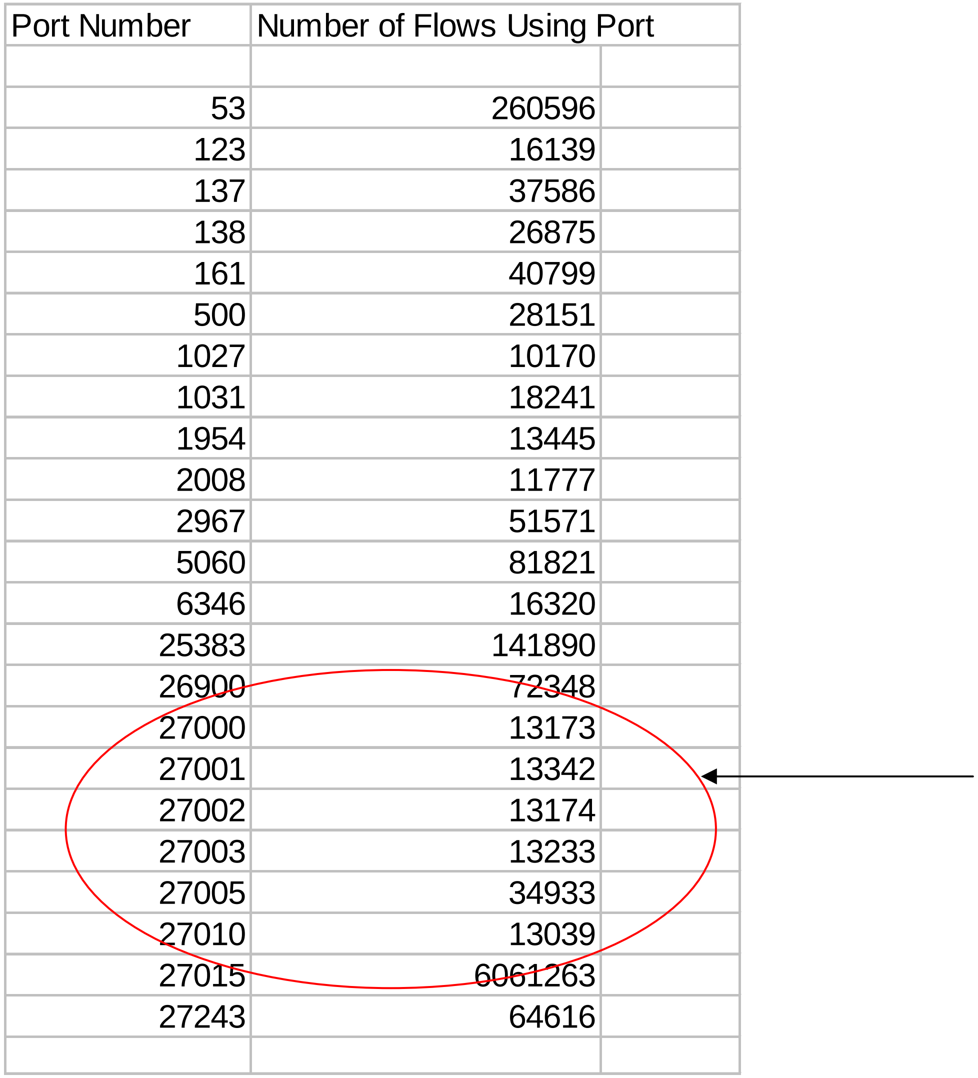

Port Number	Number of Flows Using Port	
53	260596	
123	16139	
137	37586	
138	26875	
161	40799	
500	28151	
1027	10170	
1031	18241	
1954	13445	
2008	11777	
2967	51571	
5060	81821	
6346	16320	
25383	141890	
26900	72348	
27000	13173	
27001	13342	
27002	13174	
27003	13233	
27005	34933	
27010	13039	
27015	6061263	
27243	64616	

Also note for future reference

First Traffic Capture Observations

- Next we look at the pattern of traffic accessing 27015

UDP Traffic Any Port = 27015

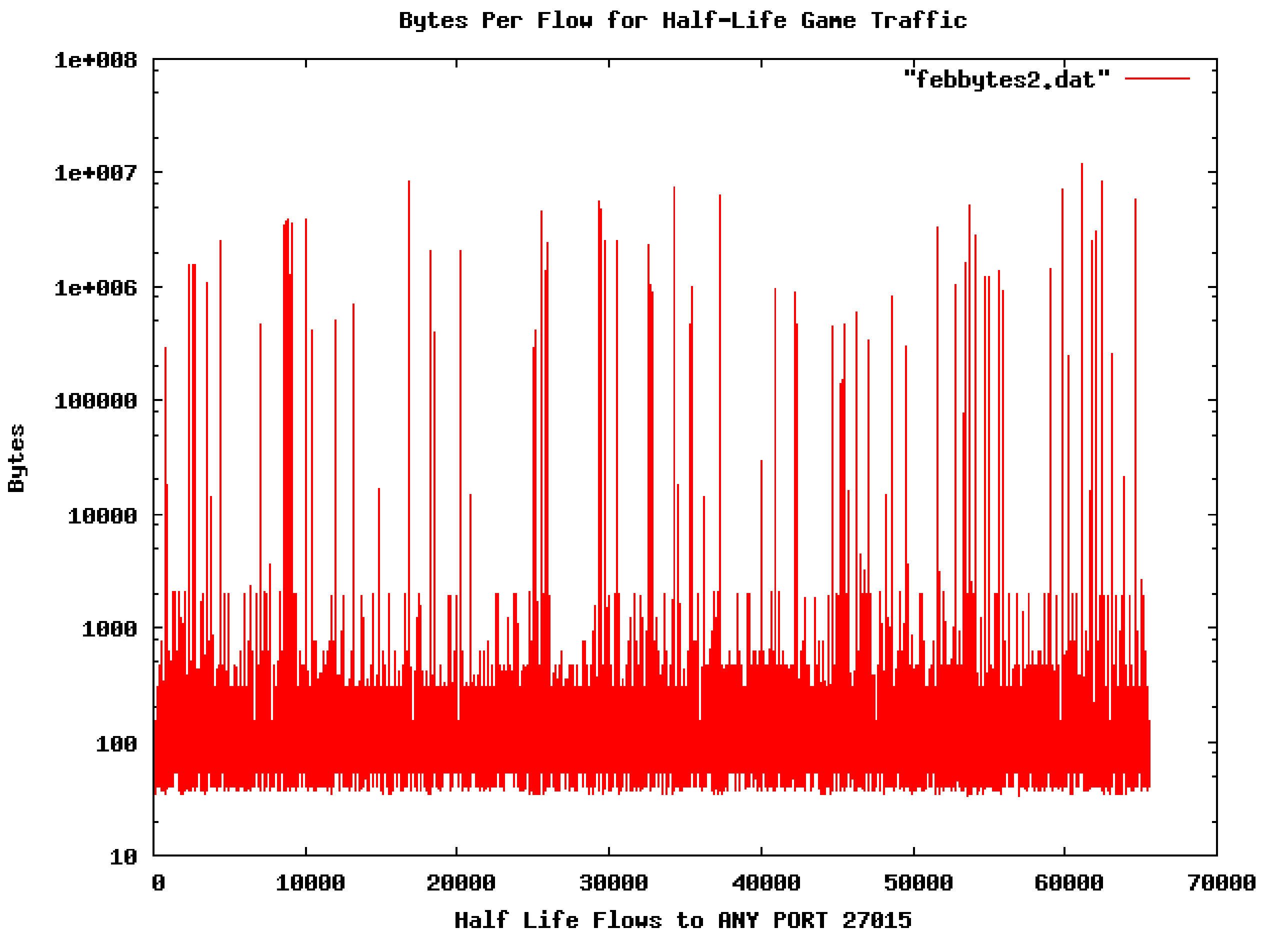

First Traffic Capture Observations

- Then we do a side by side Comparison to the behavioural pattern with the total UDP traffic.

Influence of Port 27015 on All UDP

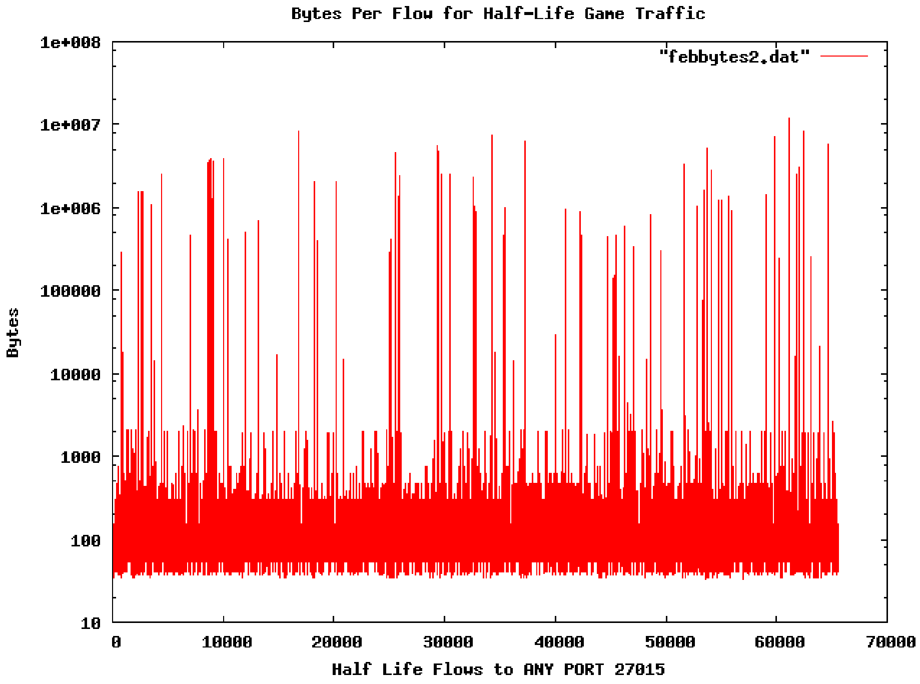

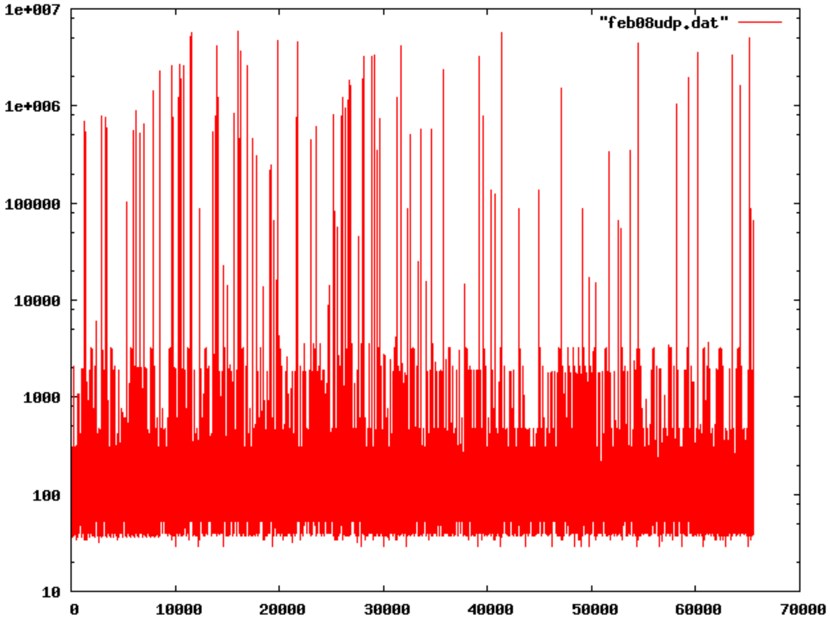

This dominating behavioural pattern was assumed to represent a single application's protocol.

First Traffic Capture Observations

- Additional information on the characteristics of the Suspicious Host:
  - Suspicious Host was the 34% of all flow records during the hour tested.
  - Suspicious Host communicated with 5,987 separate DestinationIP's during the hour.
  - Almost all traffic from SourceIP's that targeted the Suspicious Host as the DestinationIP was using protocol 17 and destination port 27015
  - A significant amount of the traffic to and from Suspicious Host was directed at a university campuses in the United States and consumer ISP's around the world.
  - Finally, the Suspicious Host was identified as an experimental development machine that had been part of a development and testing project in the previous year. Although it was still connected to the network it was not supposed to have any active users.

Summary Characteristics of Suspicious Host

- Responsible for 56% of Network Non-Port 80 Byte Volume
- Responsible for 34% of Network Non-Port 80 Flow Volume
- Constant Communication with Thousands of Hosts around the World
- Some Preference for University Campuses and Consumer ISP's
- Primarily uses UDP Port 27015
- Should Have Little or No Traffic
- WHAT AM I?

Half-Life^{tm}

- An on-line First Person Shooter Game produced by Valve Software
- Based on earlier versions of on-line game engines (Quake) and exists in many variations.

IMPORTANT DISCLAIMER

It is important to point out that Valve Software, the maker of Half-Life^{tm} is a legitimate company that would never knowingly allow its products to be part of an unauthorized network compromise. Indeed, in these circumstances, companies such as Valve Software are as much a victim as the owner of the compromised network.

FURTHER DISCLAIMER

It is important to point out that since the experimenter had no access to the actual machine or payload data this conclusion is simply conjecture.

Game Characteristics

- Clients communicate with Servers on destination port 27015.
- Game Servers may be initiated by players.
- Meta or Master Servers track available game servers.
- Game servers communicate with Meta servers on UDP port 27010.
- Some TCP Traffic associated with game network management.

Recall the Presence of Uniform Access in the 27,000 – 27,010 Port Range

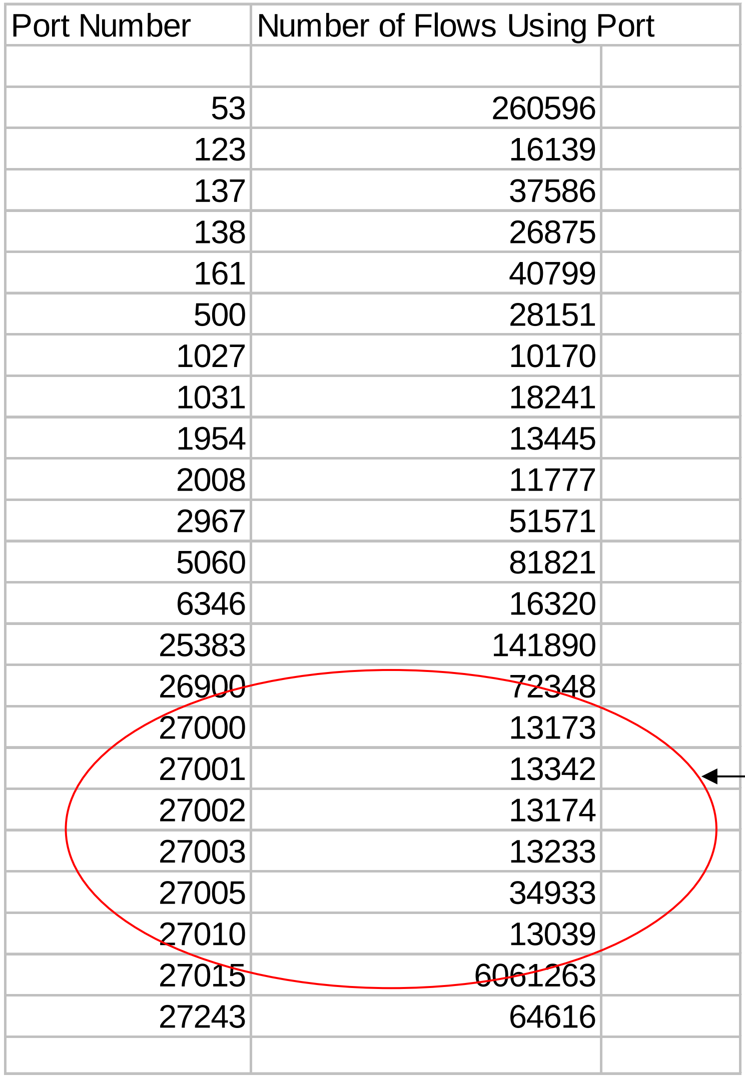

Port Number	Number of Flows Using Port
53	260596
123	16139
137	37586
138	26875
161	40799
500	28151
1027	10170
1031	18241
1954	13445
2008	11777
2967	51571
5060	81821
6346	16320
25383	141890
26900	72348
27000	13173
27001	13342
27002	13174
27003	13233
27005	34933
27010	13039
27015	6061263
27243	64616

Also note for future reference

Signature is By No Means Unique

- UDP port can be chosen by any application.
- Large byte volume is a relative term
- User demographic (Consumer ISP’s, Campus networks) is determined by looking.
- Would like to find a TCP management signature

Strategy To Isolate TCP signature

- We know that one exist’s from on-line developer discussions.
- Build a set of Game SIP’s.
- Slice out all TCP traffic.
- Isolate the TCP traffic associated with the Player SIP’s.

Strategy To Isolate TCP signature

- Build a set of Game SIP's.
  - Create a game host file:
    - rwfilter - -aport=27,005,27,014,27015 - -pass=hltraffic.f out*
  - Create a set of unique IP's for Game Hosts
    - rwset - -sip-file=hlsipfile.set hltraffic.f
- Slice out all TCP traffic.
    - rwfilter - -proto=6 - -pass=tcptraffic.f out*
  - Create a set of Unique IP's for the TCP traffic
    - rwset - - sip-file=tcpsipfile.set tcptraffic.f
- Intersect the sets to get the Game hosts using TCP
    - setintersect - -add-set=hlsipfile.set - -add-set=tcptraffic.f - -set-file=hltcp.set

TCP Game Traffic

- Upon completion hltcp.set contained only four unique SourceIP’s, one of which was the Suspicious Host. The other three were not within the address space of the Network.
- Removing these SIP’s from the complete file of TCP traffic created an artificial normal TCP traffic slice
- Comparing the *Artificial Normal Data* to the actual data revealed a distinguishing pattern.

A TCP Signature?

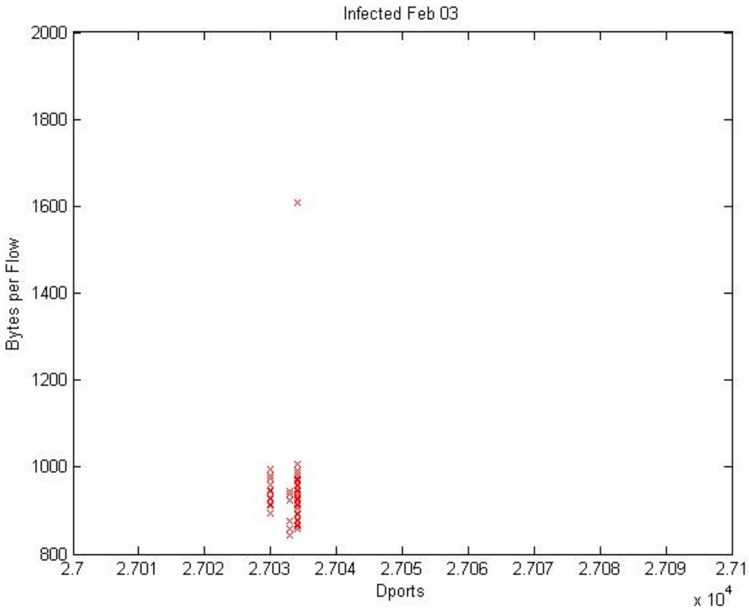
Infected Feb 03

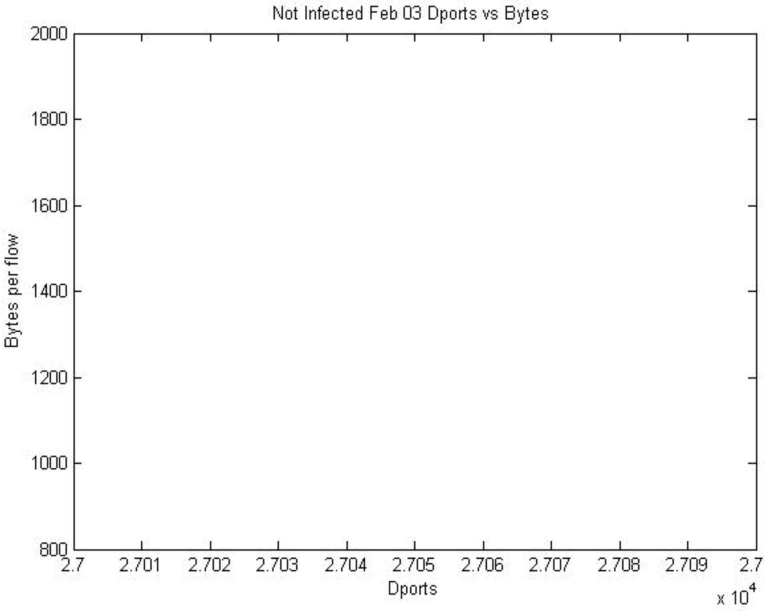
Not Infected Feb 03 Dports vs Bytes

TCP traffic contained in the Destination Port range 27,030, 27,033 and 27,034 with Bytes Per Flow sizes ranging from the low 800’s to slightly more than 1000 with a noticeable outlier at approximately 1600. This traffic is absent in the artificial normal data set.

A TCP Signature?

Probably Not

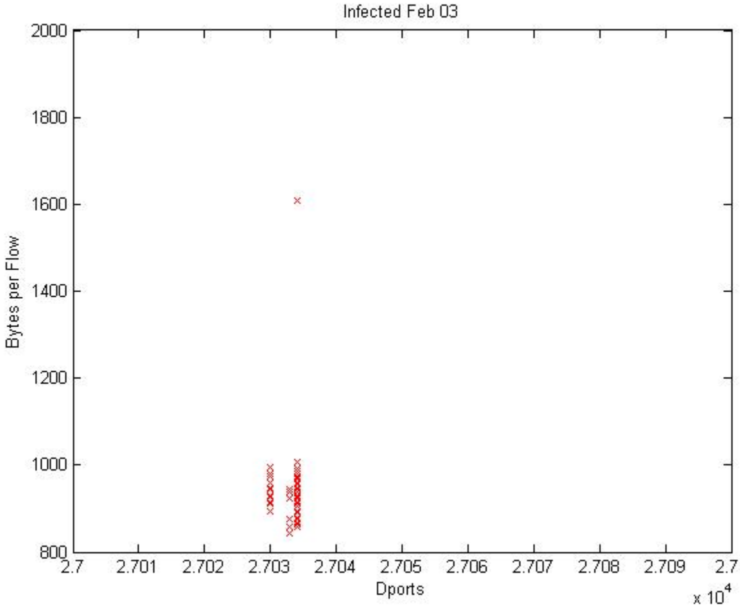

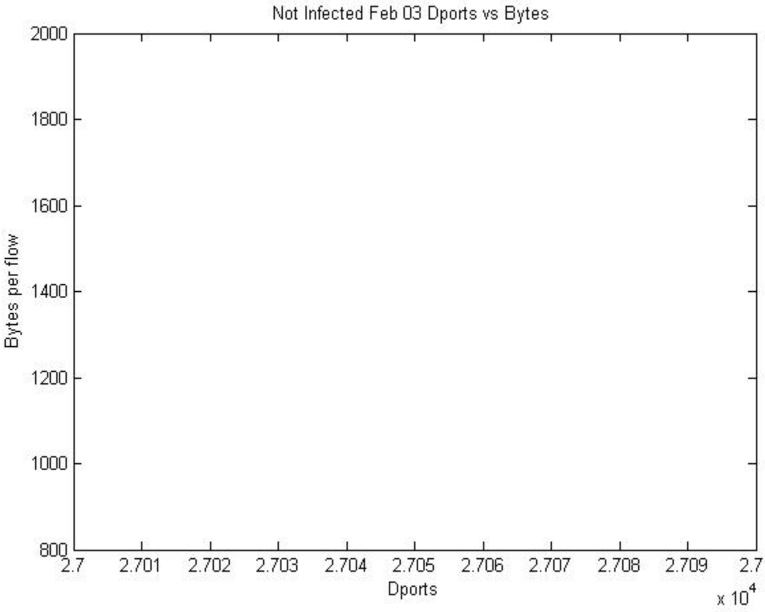

Be Careful of Assumptions

- This Host was not supposed to have any active users.
- At least half of the SourceIP’s creating the TCP Signature were immediately known to the Owner.

Be Careful of Assumptions

- This Host was not supposed to have any active users.
- At least half of the SourceIP’s creating the TCP Signature were immediately known to the Owner.
- However – An on-line discussion mentions

  Server to server communication to a European address range that exists in the data and communication on port 27010

Port Bag For Key Counts > 10,000

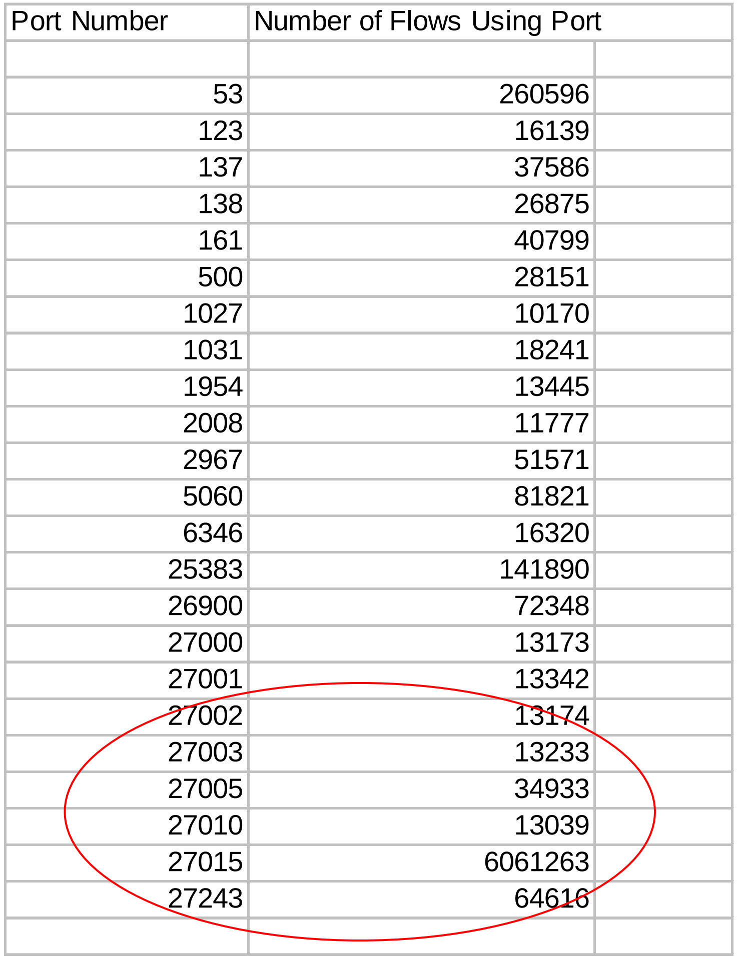

Port Number	Number of Flows Using Port	
53	260596	
123	16139	
137	37586	
138	26875	
161	40799	
500	28151	
1027	10170	
1031	18241	
1954	13445	
2008	11777	
2967	51571	
5060	81821	
6346	16320	
25383	141890	
26900	72348	
27000	13173	
27001	13342	
27002	13174	
27003	13233	
27005	34933	
27010	13039	
27015	6061263	
27243	64616	

The Search for a Scar

- Is there a unique traffic signature for a network that previously contained a game server host?
- Is there a residual SCAR in the traffic - Severed Connection Anomalous Records

The Search for a Scar

- Unfortunately, Game Server was disconnected from the network.
- A search of Null Traffic was conducted which revealed two interesting anomalies

SCAR Traffic?

dPort	pro	bytes	flags
27015	17	53	A
27015	17	106	A
27015	17	53	A
27015	17	106	A
27015	17	106	A
27015	17	106	A
27015	17	106	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	106	A
27015	17	106	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A

SCAR Traffic?

dPort	pro	bytes	flags
27015	17	53	A
27015	17	106	A
27015	17	53	A
27015	17	106	A
27015	17	106	A
27015	17	106	A
27015	17	106	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	106	A
27015	17	106	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A
27015	17	53	A

This same traffic exists in Null while Game server is active???

Null UDP Traffic on Apr 1

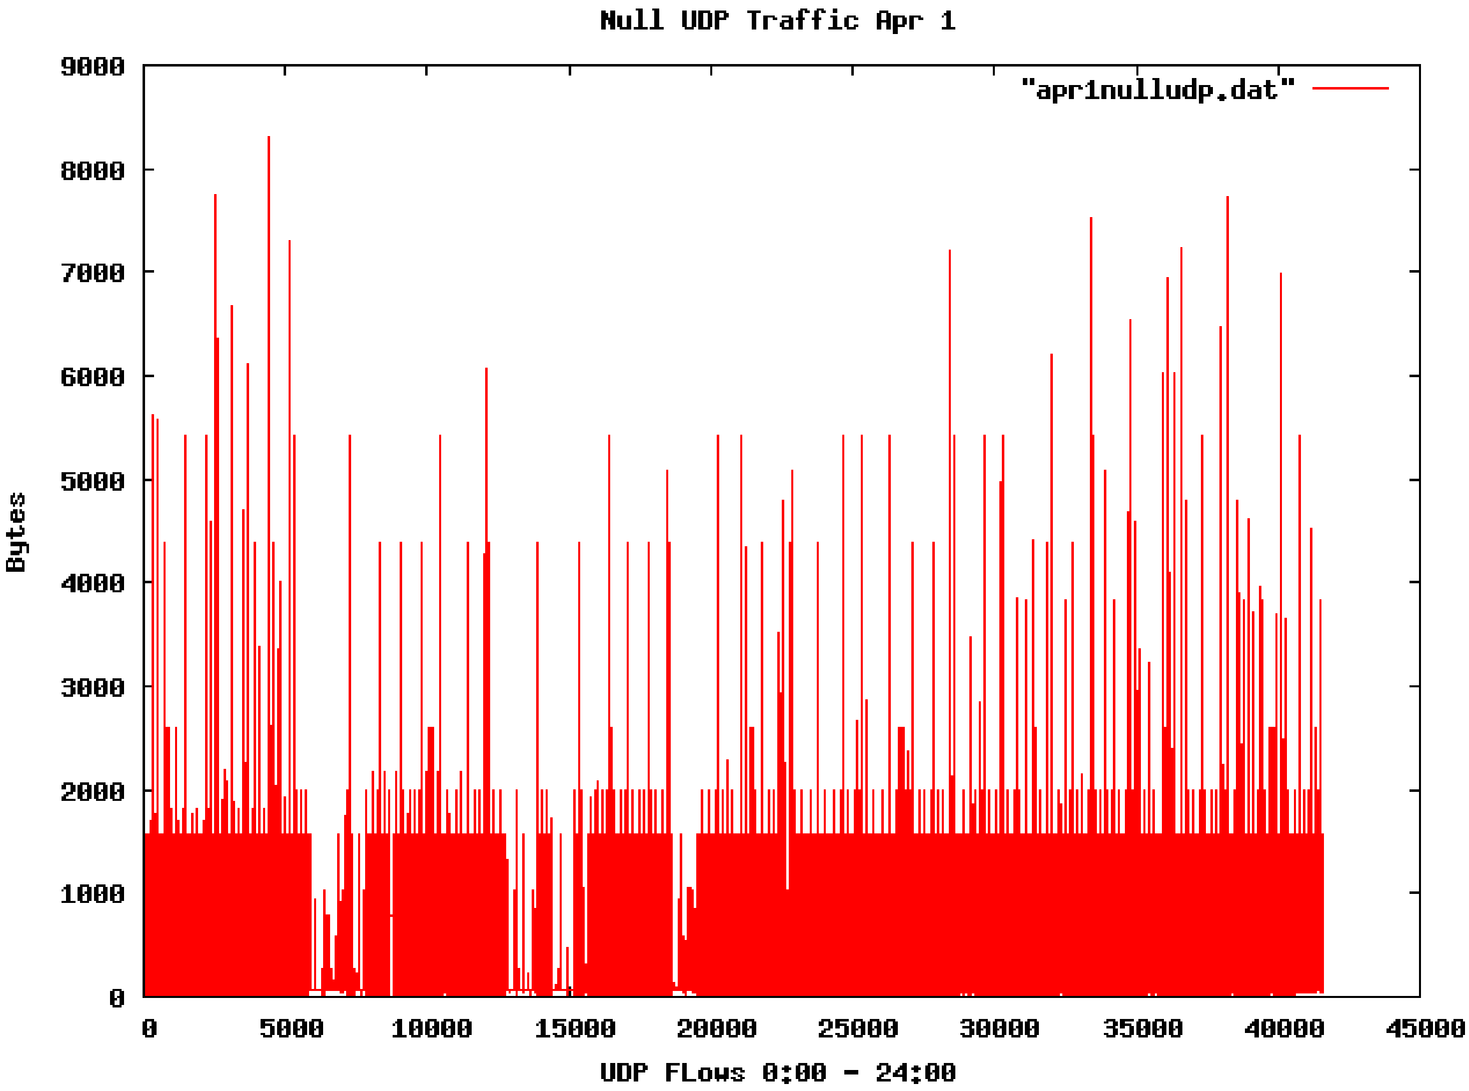

Null Game Traffic Only on Apr 1

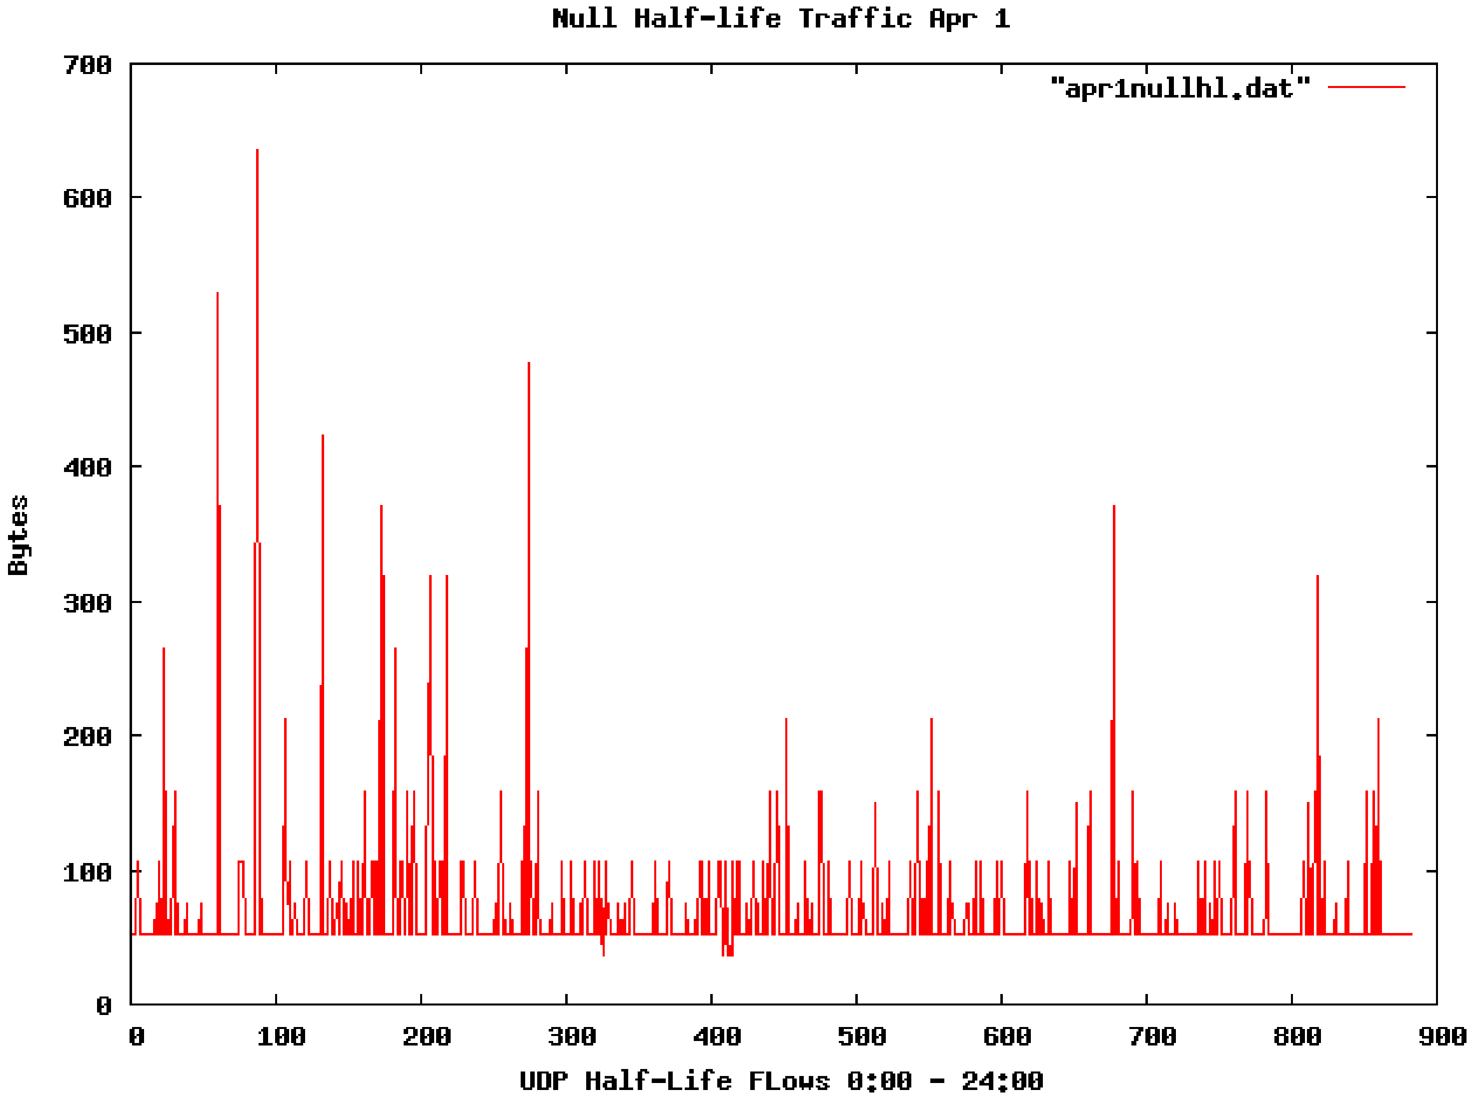

New Traffic Pattern
Non-Null Non-Port 80

April 1, 2006

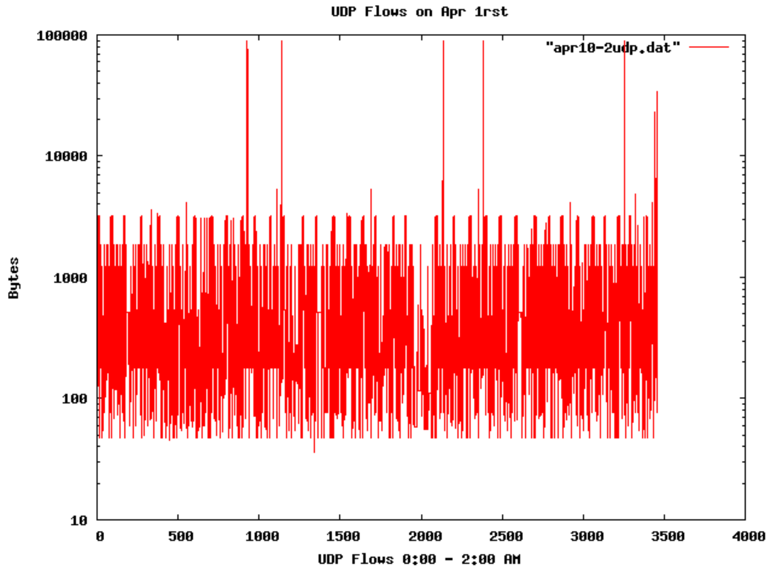

Feb 8, 2006

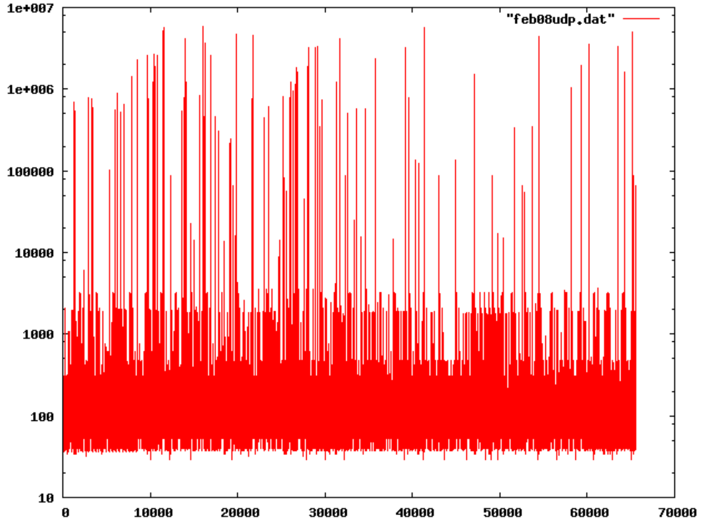

Observations on Gaming Traffic

- Much of the existing traffic profiling is aimed at providing a better game experience.
- Consumes considerable Resources.
- Represents a Level 7 WAN Network for Communication.
- Provides a channel to hide Malicious Traffic.

Future Work

- Anonomize the data so that it might be shared.
- Study the form and distribution of players, servers and meta-servers.
- A search for Management and other signatures continues.
- It was found that a virulent worm entered the network through this server. More on this in session 2.

Special Thanks

- To Dr. John McHugh for introducing me to a whole new world.
- To TARA for providing me with the support to conduct my research.
- To FloCon for inviting me to talk about both.