Abnormal traffic detection and alert

Yiming Gong
XO Communications

Flocon 2008

The problem and request

- XO network
  - OC-192 IP backbone with OC-12 uplinks in our markets and data centers, AS 2828
- Backbone level abnormal traffic detection
  - netflow

The problem and request

- Commercial product not good enough
  - You get what GUI gives you
  - Very likely to miss low volume traffic attack
    - (storm worm, scans)
  - By default, alert based on thresholds
  - Lacking data mining ability
  - Cost
- Free flow-based tool
  - Powerful but you need tell them what to do

So what we want

- Detect network abnormal traffic
  - both low and high volume
- Non-threshold based
- Automatically
- Fully controlled and customized
- Data mining
- Better be free

Perfect world

- In a perfect world, traffic shape should be very smooth

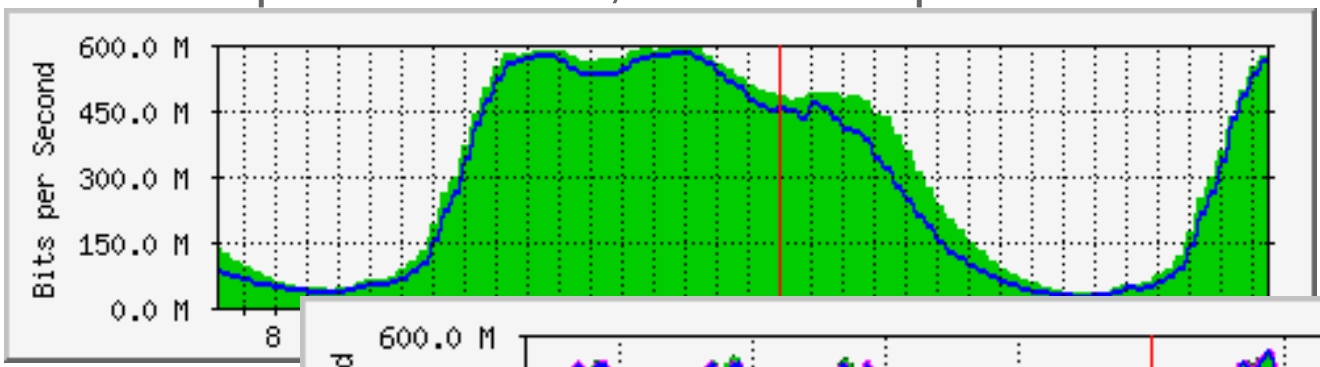

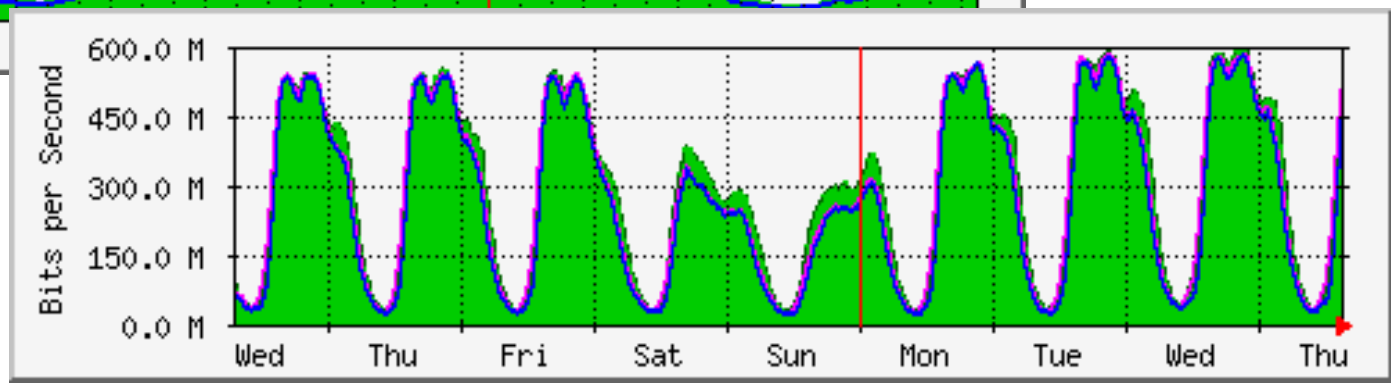

- Spike means……?

Detection at traffic level is not good

- Granularity is too coarse
  - real attack hides behind the huge traffic
- Not easy to tell what is going on
  - SYN attack? ICMP ping flood?

Our thoughts

- Netflow based
- Break down raw netflow records to
  - TCP SYN, UDP total, ICMP type|code, protocol on each IFIndex of each router
    - Session
    - Traffic
- For each element, establish a dynamic profile
- When there is spike, something is going on

Dynamic profile

- 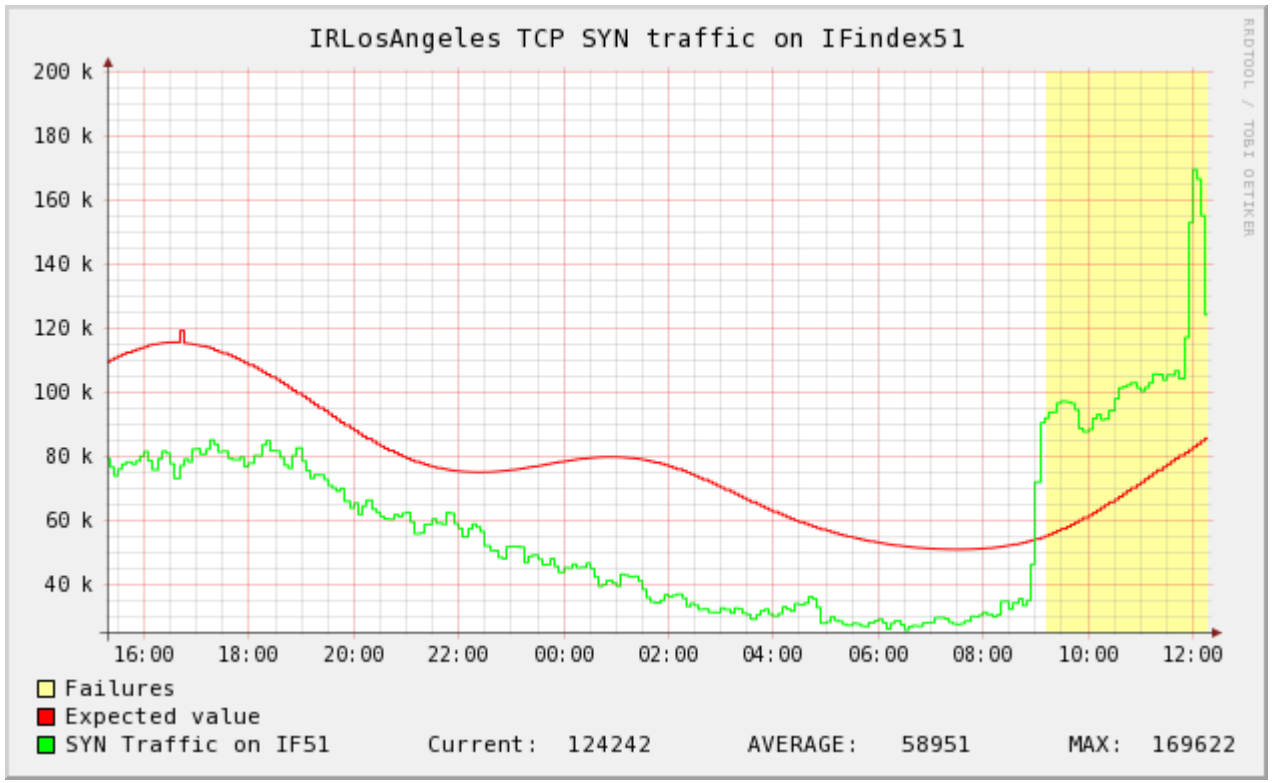

Dynamic profile





- Establishing a profile
  - Using NFDUMP receive, store and process netflow data
  - rrdtool with aberrant behavior module
  - rrdtool (http://oss.oetiker.ch/rrdtool/)
  - aberrant behavior module
    - Learns from past values and uses them to predict the future
    - Tolerance band

Dynamic profile





- Nfdump

```
yiming> more IR-syn-Amsterdam
13 1864
9 144
21 85
```

- Rrdtool

```
rrdtool create IR-syn-Amsterdam.rrd -s 300
DS:13:GAUGE:1200:0:U        \
DS:9:GAUGE:1200:0:U        \
DS:21:GAUGE:1200:0:U        \
RRA:HWPREDICT:2016:0.001:0.0035:288
```

Failure





- Only an entry
  - IR-syn-Amsterdam:[1196800800]RRA[FAILURES][1]DS[13] = 1.0000000000e+00
  - Need script do the trace back work
  - Every 10 minutes, scans the rrd output for failures
  - Short-life spike
    - window-length and failure-threshold
    - rrdtool tune x.rrd -–window-length 5 –failure-threshold 3

Failure

- Tracking down the failure
  - Nfdump + netsnmp + mysql + whois…
  - Narrowing down from flow and getting the suspicious host(s)

```
yiming> more IR-syn-Amsterdam
13 1864
9 144
21 85
```

  - Flow of “TCP + SYN bit only + IFindex 13 + router Amsterdam”
  - Finding the most ACTIVE host(s)
    - What is the definition of active?
      - Session number
      - Traffic volume

Finding active host

- Differences between these two pics?

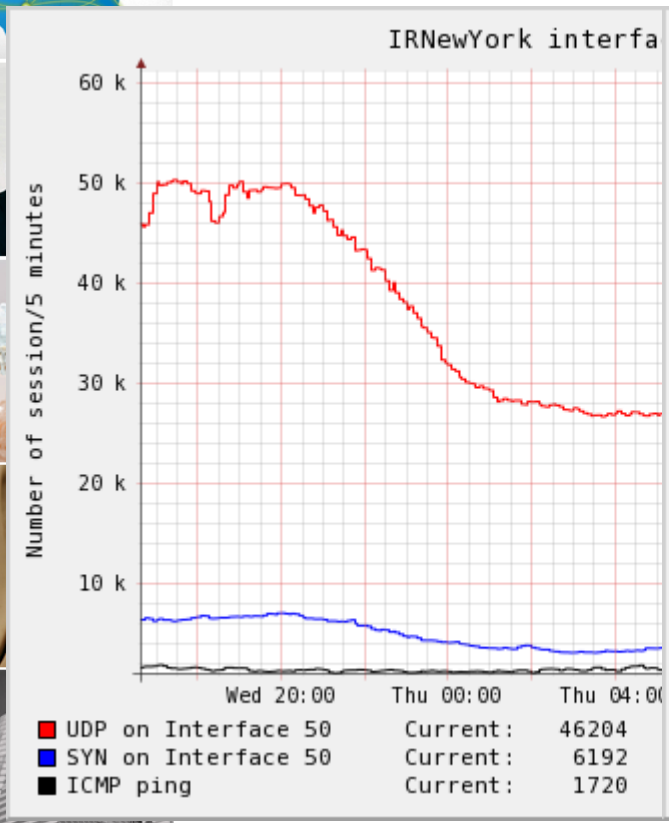

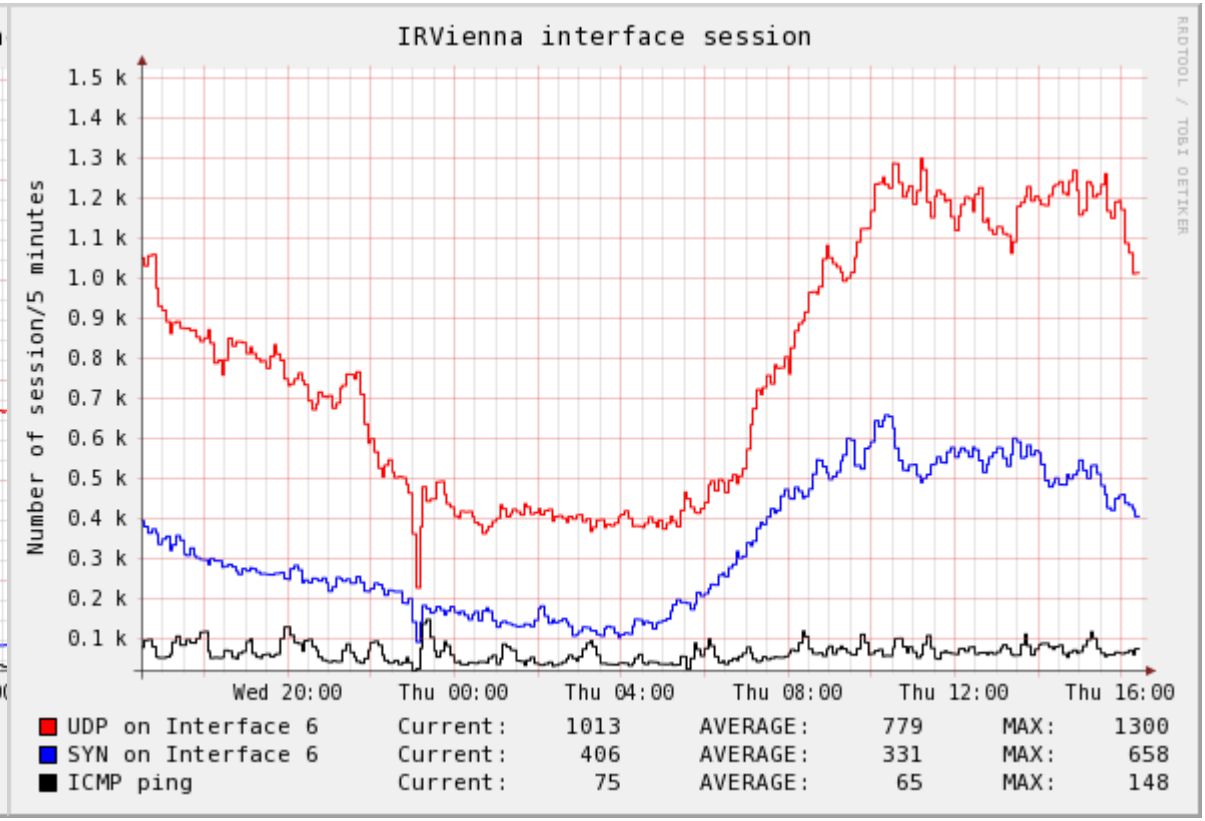

Finding active host





- Different criterion

```
session-icmp*)
        total-number="500";
        flowfilter="proto icmp and port 2048 and if $if";
        trigger-number="280";
;;
 session-syn*)
        total-number="2000";
        flowfilter="proto tcp and flags 2 and if $if";
        trigger-number="600“;
```

- Things we ignored
  - TCP SYN is supposed to be 1, but is 10 now
  - Low volume UDP spikes

Netflow records

- Pull out necessary data
- Generate alert
  - Picture, email

Alert

- Scan alert

```
>IR LosAngeles has 5462 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING: [redacted]
50 = STRING:

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto          IP Addr         Flows   Packets     Bytes       pps       bps
2007-12-05 08:51:02.520   289.718 any      218.233.1[redacted]    2114      2114     84560         7      2334
2007-12-05 08:51:20.493   130.413 any        218.234[redacted]     605       605     24200         4      1484

>Top IP info

* AS            IP              AS name                                    FQDN
[redacted]  | 218.233.[redacted]  | [redacted] Telecom Inc.  |
[redacted]  | 218.234.[redacted]  | [redacted] Telecom Inc.  |
```

Alert

- Day

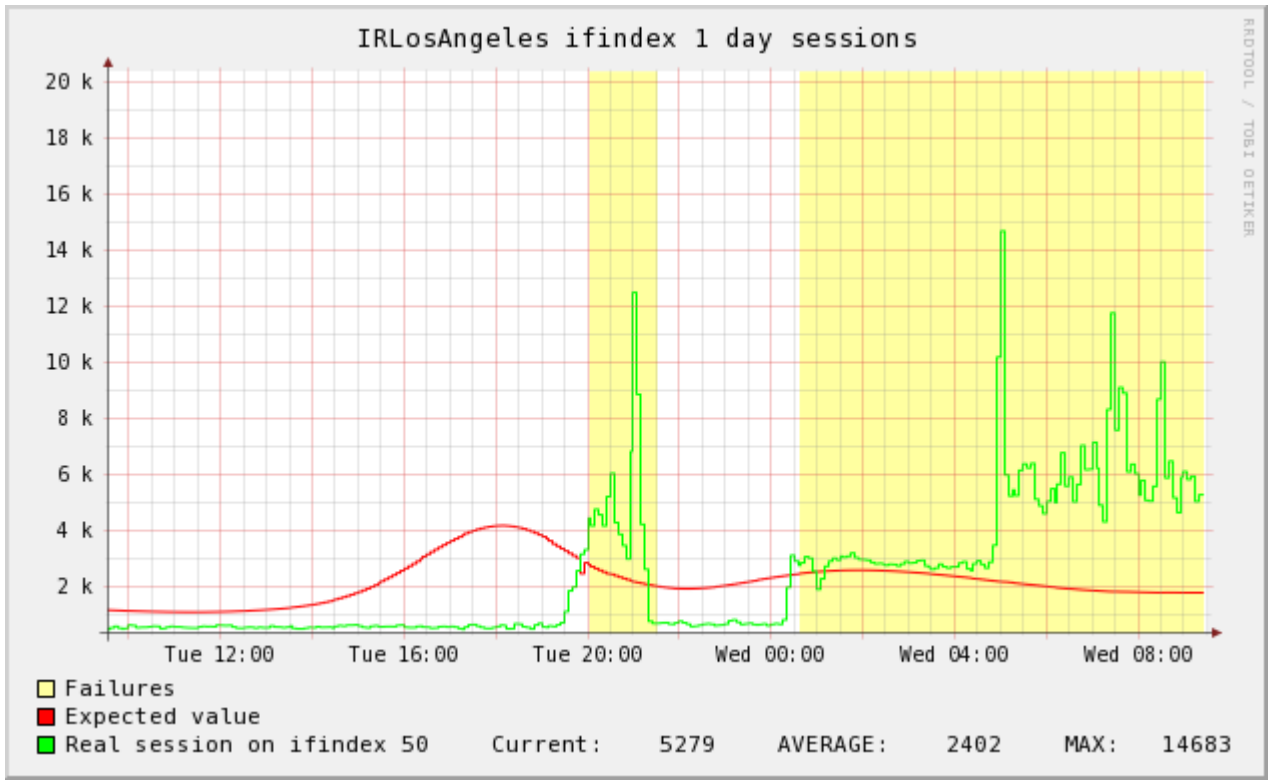

Alert

- Week

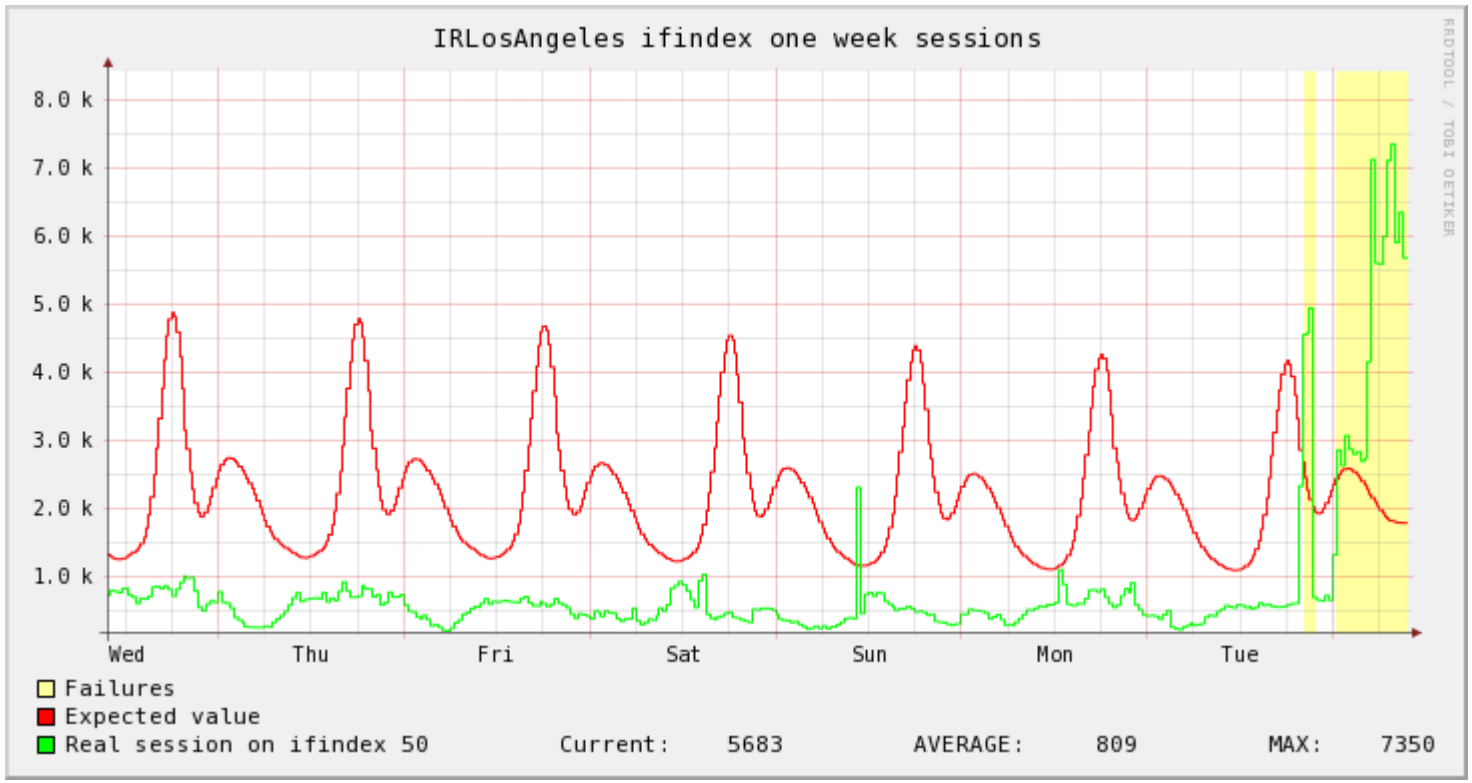

Alert

- Scan alert

```
>Top IP detail

/ ip 218.233.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]65.106.6.[redacted]ms
 6  [redacted]gin.ntt.n[redacted]19.12)  7.089 ms
 7  [redacted]ntt.net ([redacted]  7.317 ms
 8  [redacted].gin.ntt.[redacted]5.20)  73.034 ms
 9  [redacted]ntt.net ([redacted]  73.922 ms

**Protocol summary for 218.233.198.25

Proto  Flows  Packets    Bytes      pps      bps    bpp
 6     2116    2116    84640        7     2337     40
17        1       1      257        0        0    257

**Sampled netflow records

TCP     218.233.[redacted]:6000     65.99[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000     65.99[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000     65.99[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000     65.99[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000     65.99[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000   65.99.3[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000    65.99.[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
TCP     218.233.[redacted]:6000   65.99.3[redacted]:7212   ....S.    40   0  [redacted]  [redacted]   50   73
```

Alert

- Scan alert

```
/ ip 218.234.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]106.6.[redacted]ms
 6  [redacted]n.ntt.r[redacted]119.12)  7.135 ms
 7  [redacted].net ([redacted])  7.268 ms
 8  [redacted]n.ntt.[redacted].5.20)  79.209 ms
 9  [redacted].net ([redacted])  74.073 ms

**Protocol summary for 218.234.[redacted]

Proto  Flows  Packets   Bytes      pps      bps    bpp
 6      605     605     24200       4       1484   40

**Sampled netflow records

TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000      71.60[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000     71.60.[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000    71.60.1[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
TCP     218.234.[redacted]:6000    71.60.1[redacted]:6588   ....S.    40   0   [redacted]   [redacted]   50   73
```

alert

- Storm worm

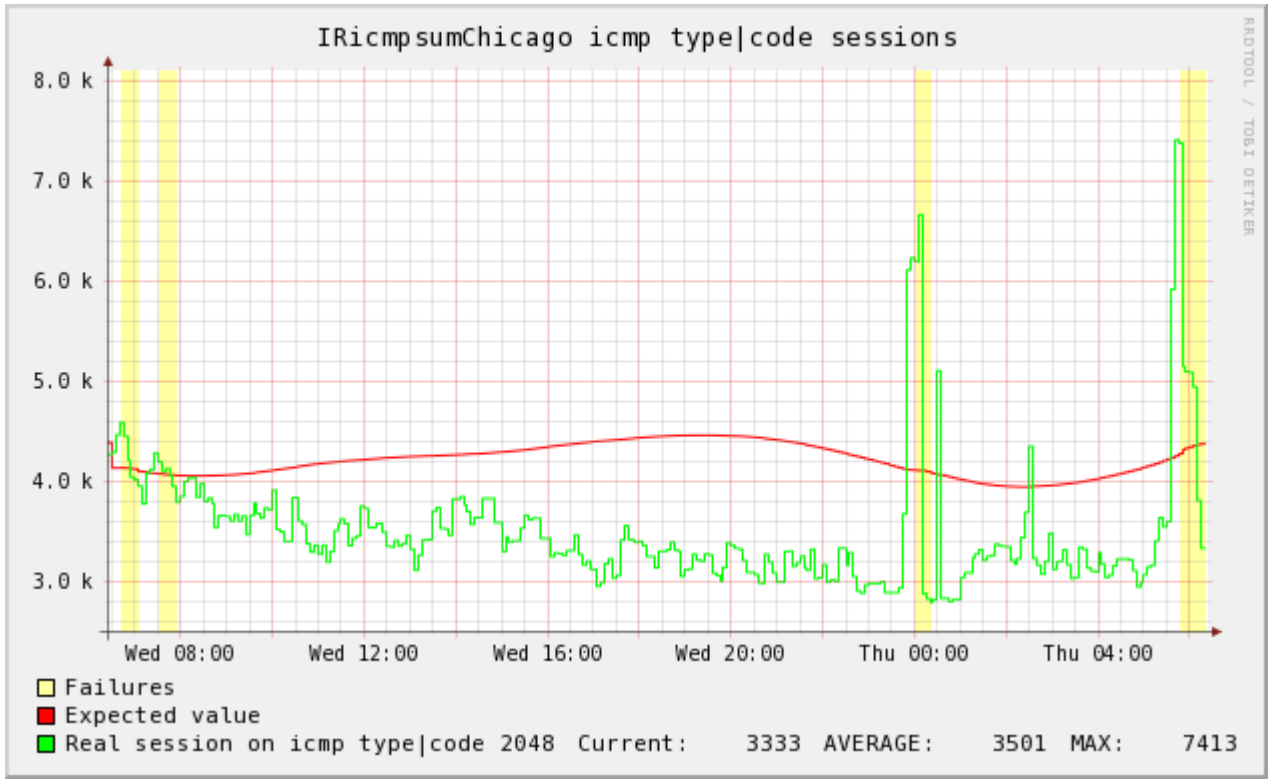

Alert - one week later

- DDos

```
IR LosAngeles has 177002 sessions on proto tcp and flags 2 and if 50 in 5 minutes

50 = STRING:[redacted]
50 = STRING:[redacted]

>Snapshot picture

http://[redacted]LosAngeles-50-abnormal.png

>One week|month picture

http://[redacted]LosAngeles-50-abnormal-week.png

http://[redacted]LosAngeles-50-abnormal-month.png

>Top IPs in 10 minutes

Date first seen          Duration Proto         IP Addr         Flows   Packets    Bytes      pps      bps
2007-12-12 09:00:23.705   320.317 any       89.144.[redacted]   173554    176183   10.0 M      550   261507
2007-12-12 09:00:28.361   297.573 any       211.211[redacted]7     1048      1056    50688        3     1362
2007-12-12 09:00:43.293   282.273 any       211.206[redacted]      692       708    33984        2      963
2007-12-12 09:00:43.269   291.093 any       211.44.[redacted]      658       667    42688        2     1173
2007-12-12 09:00:43.401   289.437 any       218.48.[redacted]      633       684    32832        2      907
2007-12-12 09:00:37.353   288.445 any       123.214[redacted]      627       640    30720        2      852
2007-12-12 09:00:23.705   311.869 any       58.127.[redacted]      603       618    39552        1     1014

>Top IP info

* AS          IP              AS name                                                FQDN
[redacted]6  | 89.144.[redacted]  | [redacted]Autonomus System number for [redacted]Net   |   ;; connectior
[redacted]   | 211.211[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
[redacted]   | 211.206[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
[redacted]   | 211.44.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
[redacted]   | 218.48.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
[redacted]   | 123.214[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
[redacted]   | 58.127.[redacted]  | [redacted]HANARO-AS Hanaro Telecom Inc.    |
```

Alert - one week later





- Traceroute returns nothing

```
>Top IP detail

/ ip 89.144.[redacted]

**Traceroute (from hop 5 to 9) <----------  no traceroute info here

**Protocol summary for 89.144.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6   173974   176610   10.0 M      551   261950    59

**Sampled netflow records

TCP     219.254[redacted]:2391      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    123.212.[redacted]:2735      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.233.[redacted]:3878      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.124.[redacted]:4375      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP    219.251.[redacted]:4049      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     58.123.[redacted]:3642      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP      123.21[redacted]:2340      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  10
TCP     211.200[redacted]:4256      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  11
TCP    221.143.[redacted]:4313      89.144.[redacted]:80    ....S.    64  0  [redacted]  50  10
TCP     218.234[redacted]:4353      89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```

Alert - one week later

```
/ ip 211.211.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]L.us.xo[redacted].85)  7.113 ms
 6  [redacted]et (65.[redacted]521 ms
 7  [redacted]net (20[redacted]6.505 ms
 8  [redacted]et (65.[redacted].604 ms
 9  [redacted].net (6[redacted]6.511 ms

**Protocol summary for 211.211.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
 6     1057     1065    51120        3     1374    48

**Sampled netflow records

TCP    211.211.[redacted]:29937    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:32301    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:30596    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:35573    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:26497    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:31263    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:27378    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:34829    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:28267    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
TCP    211.211.[redacted]:59695    89.144.[redacted]:80    ....S.    48  0  [redacted]  50  11
```

Alert - one week later

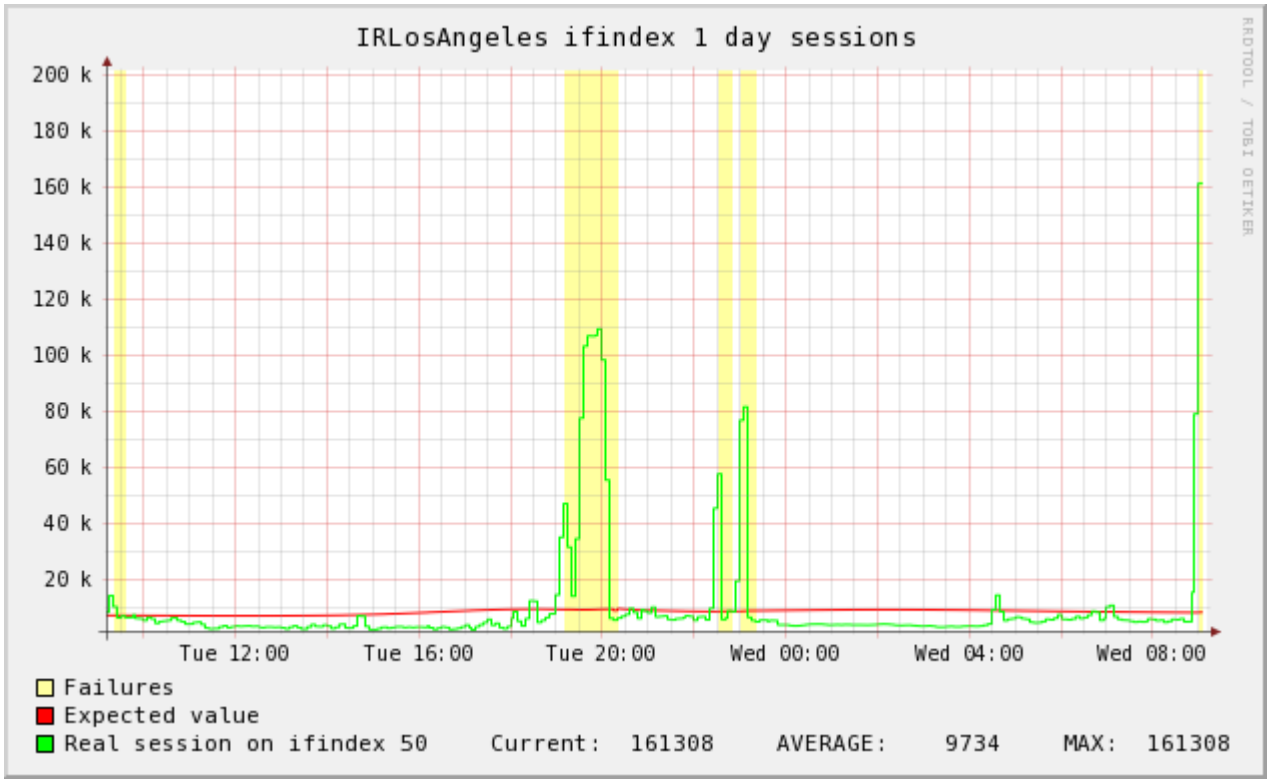

Alert - whitelist

- Special customers

```
>Top IP info

* AS            IP              AS name                                    FQDN
[redacted]  | 64.39.[redacted]  | [redacted] Inc.   |  scanner[redacted].com.
[redacted]  | 63.245[redacted]  | [redacted] Corporation  |  core2.[redacted].com.

>Top IP detail

/ ip 64.39.[redacted]

**Traceroute (from hop 5 to 9)

 5  [redacted]
 6  [redacted].com ([redacted])  6.814 ms
 7  [redacted].com ([redacted])  66.692 ms
 8  [redacted]com ([redacted])  66.836 ms
 9  [redacted]com ([redacted])  66.824 ms

**Protocol summary for 64.39.[redacted]

Proto  Flows  Packets    Bytes      pps      bps   bpp
  6      182      200     8584        0      231    42
 17        1        1       58        0        0    58

**Sampled netflow records

TCP      64.39.[redacted]:2681    63.245.[redacted]:25     ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:37672   63.245.[redacted]:35459  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:38206   63.245.[redacted]:47123  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:38318   63.245.[redacted]:2870   ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:39700   63.245.[redacted]:34739  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:40293   63.245.[redacted]:26733  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:40210   63.245.[redacted]:53606  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:40603   63.245.[redacted]:63822  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:41626   63.245.[redacted]:55450  ....S.    40  0  [redacted]  31  8
TCP      64.39.[redacted]:41565   63.245.[redacted]:2361   ....S.    40  0  [redacted]  31  8
```

Alert – whitelist and misc





- Whitelist <cont>
  - Email servers
  - We don't want to miss real attack even if an IP is on whitelist
- Alert email
  - Suppression period
  - Subject
    - 12-05 abnormal sessions at LosAngeles proto tcp and flags 2 and if 50

Data mining





- Database
  - 3 tables
    - IP,FQDN,AS
    - Summary
    - Raw netflow data
  - Data mining
    - Which peering neighbor sends out most attack traffic, who is the most attacked, which port is the most popular being scanned…etc.

Data mining





- Database
  - 3rd party outside data
    - Dshield TOP 10000
    - Dshield AS
    - CBL data
    - Mynetwatchman
    - Our own darknet project output
    - Other private outside data
  - If XO host involved, we will go through these table

problem

- Problem
  - Peering neighbor
  - Alert correlation
    - But you can do it in database.

What you need

- Nfdump, rrdtool, mysql, net-snmp, apache, some unix commands
- A box with linux installed

- For more info
  - yiming.gong@xo.com

- Thanks!