

Cloud-based Log Analysis and Visualization

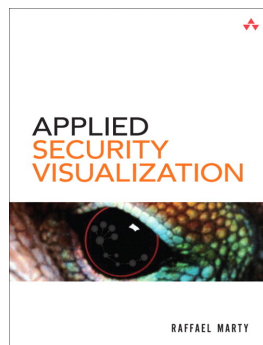
DeepSec 2010, Vienna, Austria



Raffael Marty – @zrlram

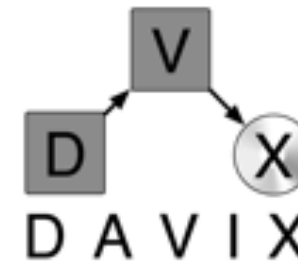
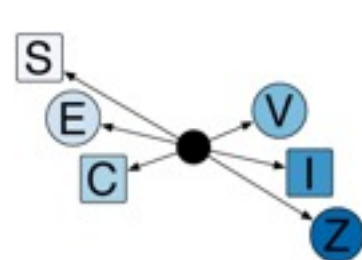
Raffael (**Raffy**) Marty

- Founder @ **loggly**
- Chief Security Strategist and Product Manager @ Splunk
- Manager Solutions @ ArcSight
- Intrusion Detection Research @ IBM Research
- IT Security Consultant @ PriceWaterhouse Coopers



Applied Security Visualization

Publisher: Addison Wesley (August, 2008)
ISBN: 0321510100



Agenda

- Introduction
- **Beaver Challenge**
- The Cloud
- Visualization
- Visualization **Tools**
- Visualization in the **Cloud**
- Visualization **Use-Cases**
- Visualization **Resources**

The Public Cloud

IaaS – Infrastructure

PaaS – Platform

SaaS – Software

LaaS – Logging



What is really new and has changed?

Visibility and Big Data



Visibility



- Monitoring

- Performance
- Availability
- Ephemeral Infrastructure

- Security

- New Threats
- New Vulnerabilities
- Different Risk Distribution

IaaS – Similar to before

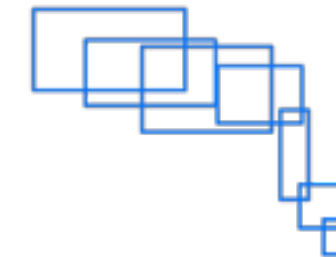
PaaS – Lack of Infrastructure

SaaS – Blind?

**Application Instrumentation
and Logging**

Big Data

- NoSQL
- Distributed data stores
- Distributed queues
- Map reduce
- ETL (Extract, Transform, Load)
- ...



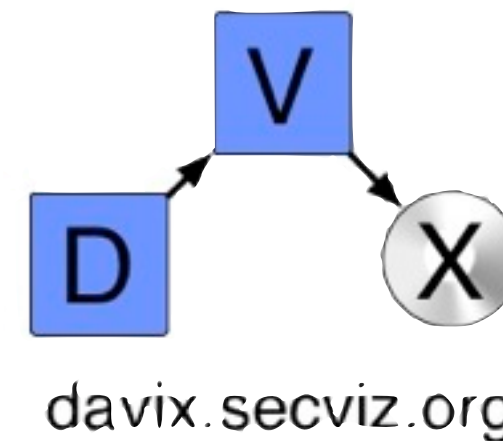
Logging as a Service

Information Visualization

- Better tools and capabilities



- Across disciplines
- More instrumentation
- Dichotomies



Open Your Eyes



Information Visualization?

A picture is worth a thousand log records.



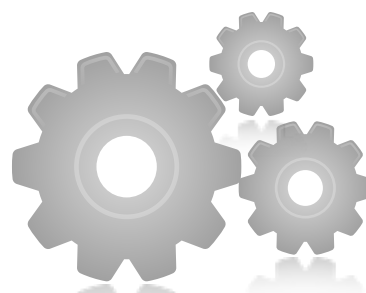
Explore and
Discover



Answer a
Question



Pose a New
Question



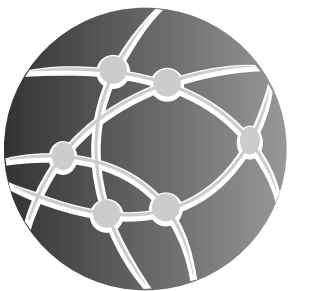
Increase
Efficiency



Communicate
Information



Support
Decisions



Inspire

Visualization Tools

Reporting vs. Visualization

- Reporting Libraries

- HighCharts
- Flot
- Google Chart API
- Open Flash Chart
- HTML5

- Visualization Libraries

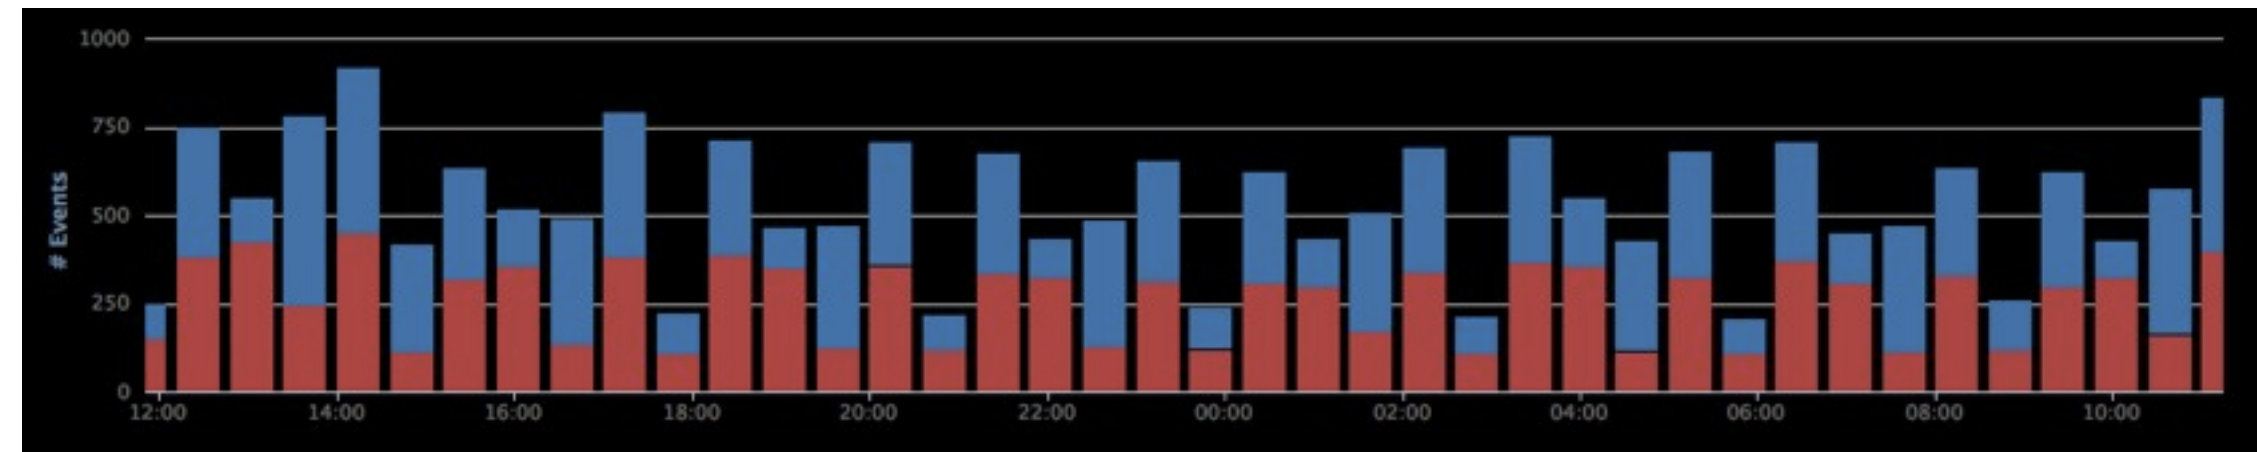
- TheJIT
- Graphael
- Protovis
- ProcessingJS
- Flare

JavaScript vs. Flash vs. XYZ

HighCharts



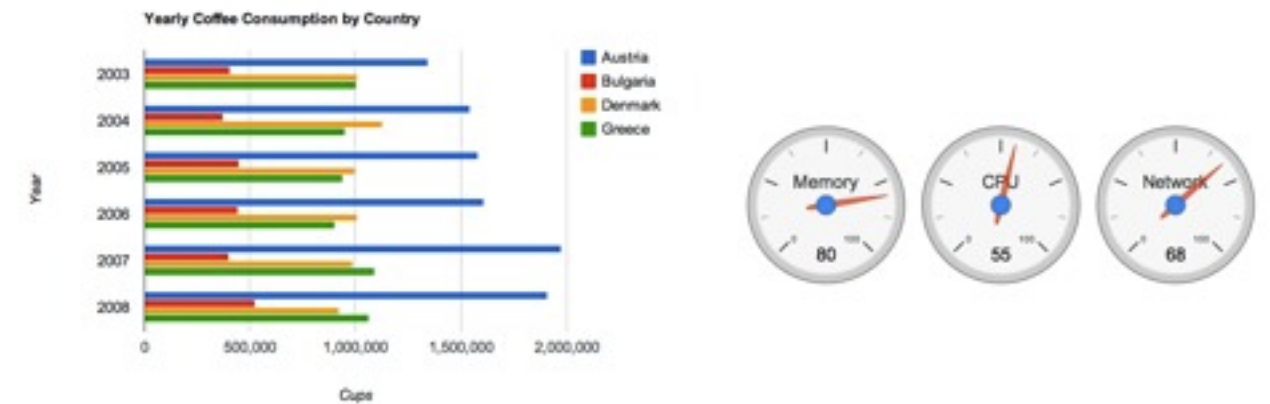
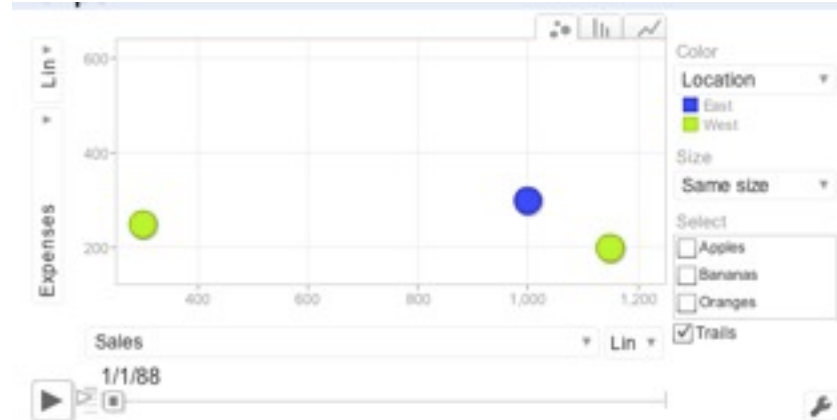
- Click-Through
- On load



- near real-time updates
- AJAX data input via JSON
- Zoom

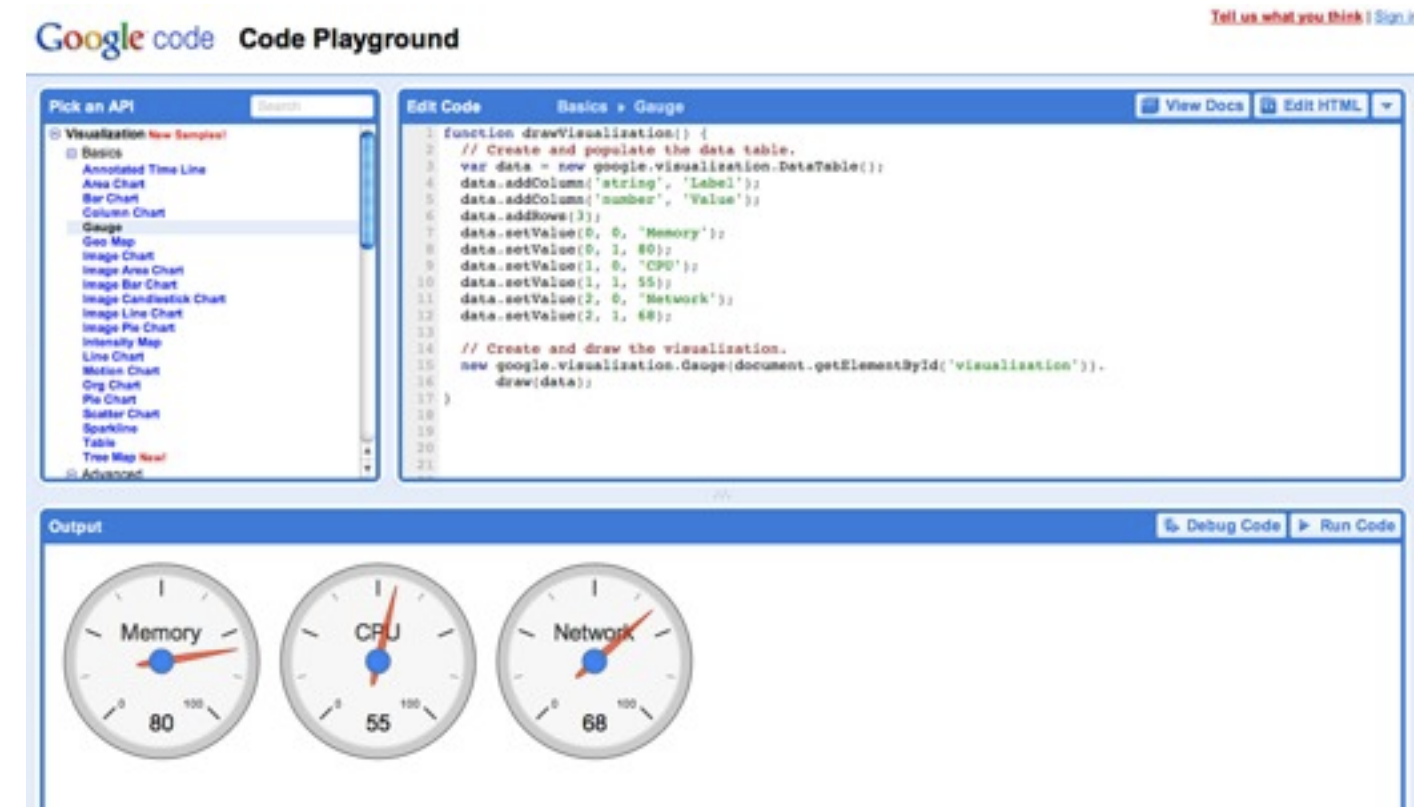
<http://www.highcharts.com/>

Google Visualization API



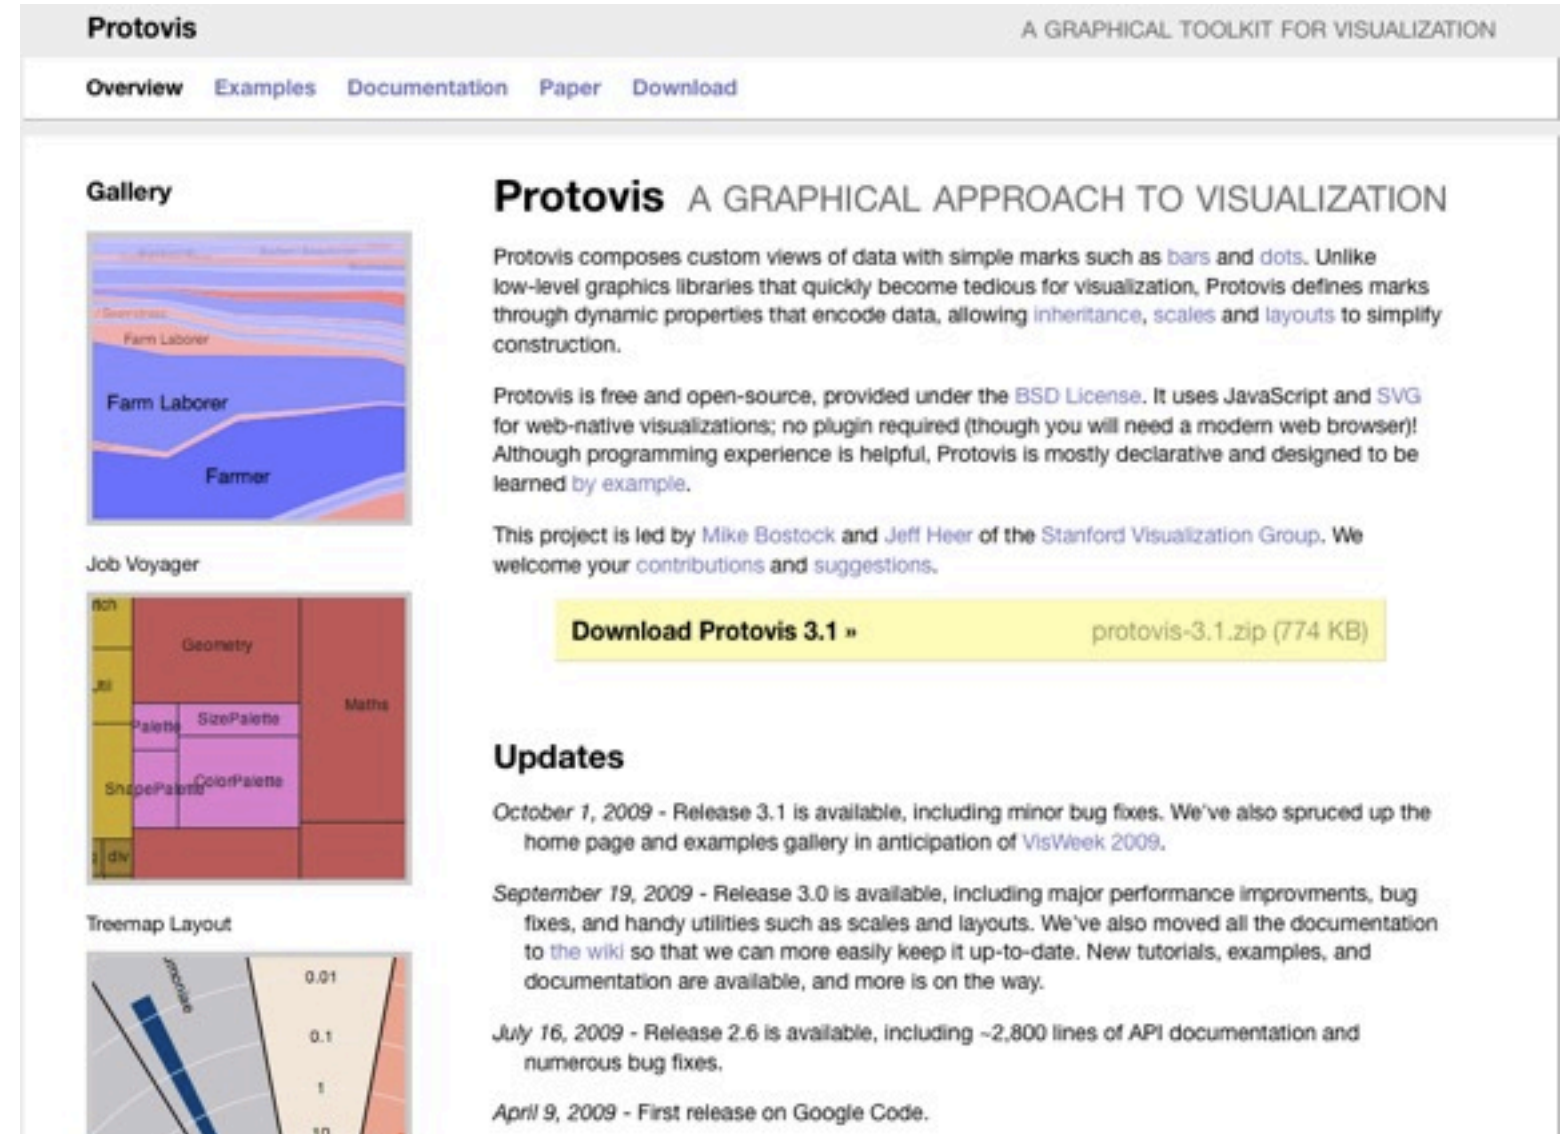
http://code.google.com/apis/visualization/interactive_charts.html

- JavaScript
- Based on DataTables()
- Many graphs
- Playground
 - <http://code.google.com/apis/ajax/playground>



ProtoVis

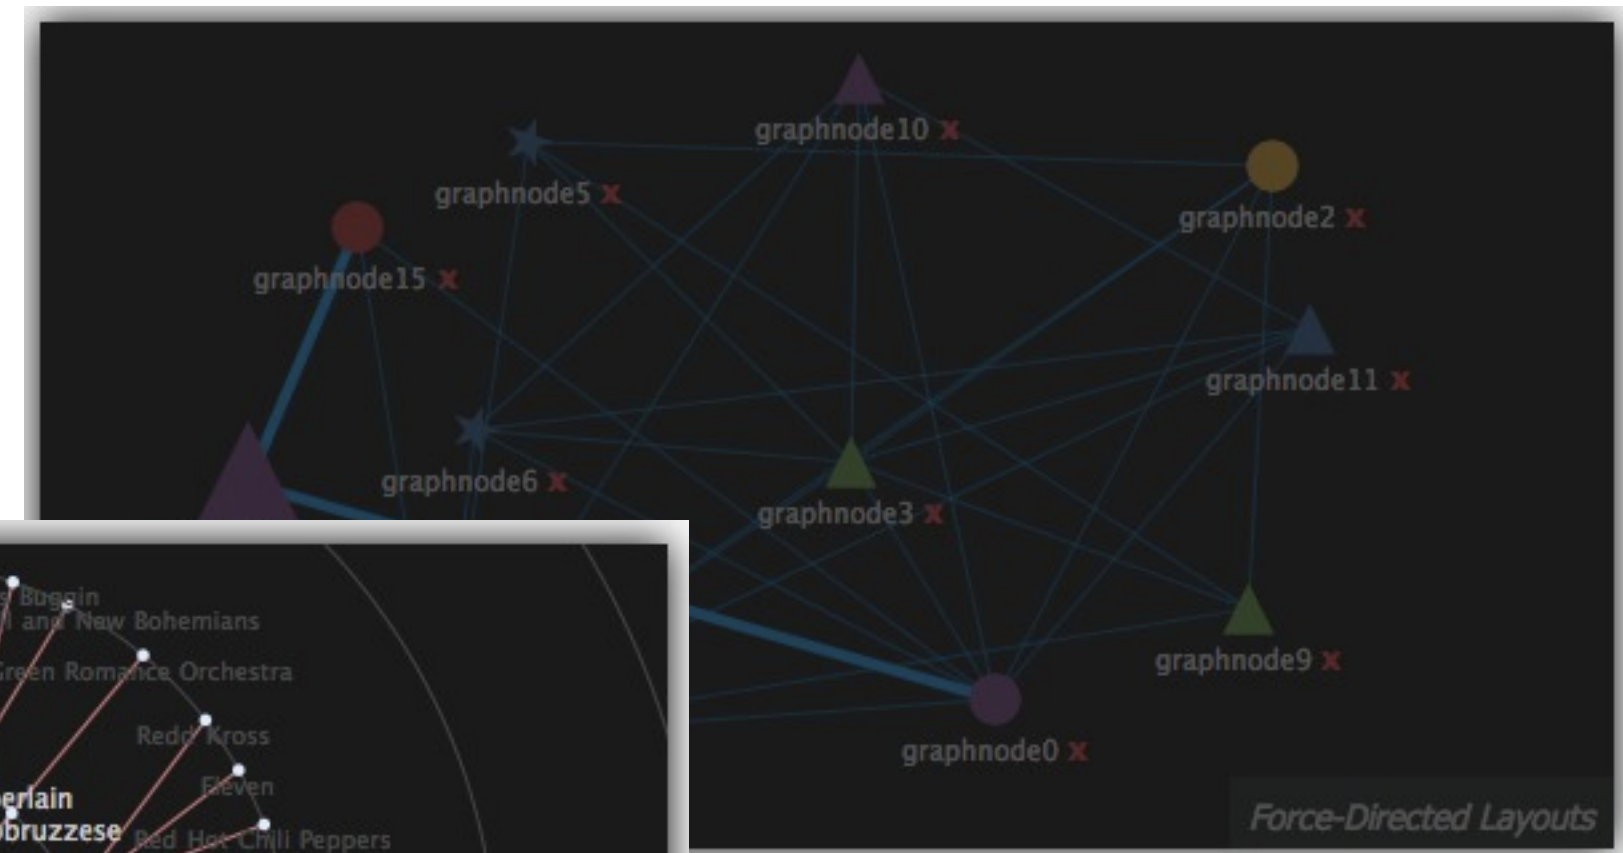
- JavaScript based visualization library
- Charting
- Treemaps
- BoxPlots
- Parallel Coordinates
- etc.



The screenshot shows the ProtoVis website homepage. At the top, the title "ProtoVis" is displayed in a large, bold, red font, with the subtitle "A GRAPHICAL TOOLKIT FOR VISUALIZATION" in a smaller, grey font to its right. Below the title is a navigation bar with links: "Overview", "Examples", "Documentation", "Paper", and "Download". The main content area is divided into two columns. The left column features a "Gallery" section with three visualizations: a Sankey diagram showing flows between "Farm Laborer" and "Farmer", a treemap titled "Job Voyager" showing a hierarchy of categories like "Geometry" and "Maths", and a "Treemap Layout" visualization. The right column contains a "ProtoVis" section with the subtitle "A GRAPHICAL APPROACH TO VISUALIZATION". It describes the library's purpose, its use of JavaScript and SVG, and its open-source nature under the BSD License. A yellow button labeled "Download ProtoVis 3.1" is prominently displayed, with the file name "protovis-3.1.zip (774 KB)" next to it. Below this, an "Updates" section lists several releases with their dates and key features, including the release of version 3.1 in October 2009 and the first release on Google Code in April 2009.

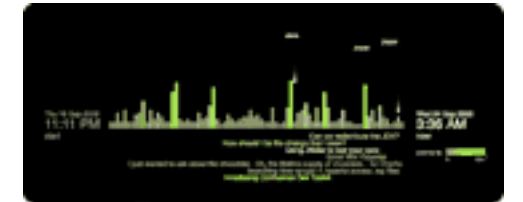
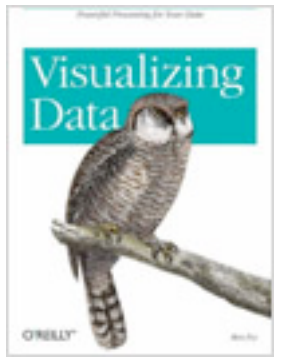
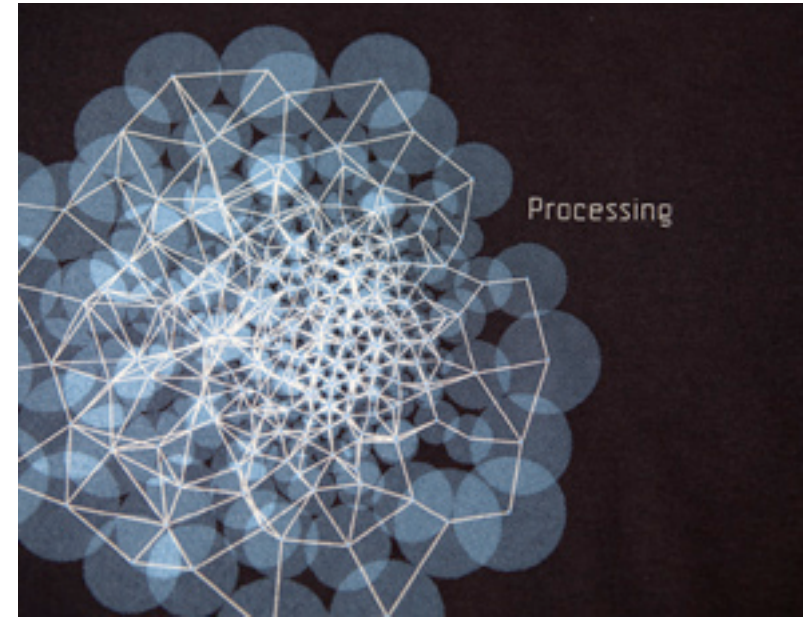
<http://vis.stanford.edu/protovis/>

- JavaScript InfoVis Toolkit
- Interactive
- Link Graphs



Processing

- Visualization library
- Java based
- Interactive (event handling)
- Number of libraries to
 - draw in OpenGL
 - read XML files
 - write PDF files
- **Processing JS**
 - JavaScript
 - HTML 5 Canvas
 - Web IDE



<http://processingjs.org/>

<http://processing.org/>



Data Visualization in the Cloud

LaaS – Logging as a Service



- **Log collection**
 - all data in one place
- **Log storage and management**
 - index, storage, archive
- **Extremely fast log search across all your data**
 - data source agnostic (no parsers)
 - innovative Web shell
- **API log access**
 - OAuth authentication
 - always on

Benefits

- No installation
- Easy configuration
- No maintenance
- Great scalability
- 7x24 availability
- Pay as you go



AfterGlow Cloud

loggly AfterGlow – Loggly Mashup

The interface displays a network graph with two main source nodes (red ovals) and several target nodes (yellow ovals). The first source node, labeled '81.253.46.252', connects to '/api/oauth/authorize/' and '/accounts/login/'. The second source node, labeled '63.200.141.142', connects to '/api/oauth/access_token/', '/api/search/', and '/api/oauth/request_token/'. A third, isolated node is also present. Below the graph, the 'Search Params' section includes a subdomain 'logdog', a search input 'inputname:logglyweb', a parser/field selector, a regex pattern '^(\d+\.\d+\.\d+\.\d+)\s+\S+\s+\S+\s+\S+\s+\"GET (.*)\"(?:\s+)?\$', and a rows limit of 20. The 'AfterGlow Params' section contains various configuration options: 'Split nodes?' (checked), 'Split Mode' (One unique event node), 'Print node count' (unchecked), 'Color Config', 'Filter Config', 'Threshold', 'Source fanout', 'Event fanout', 'Cluster Config' (Select Clustering for Source and Target), 'Node labels?' (unchecked), 'Two nodes only (not three)?' (checked), 'Default Edge Length' (1.2), and 'Which Viz Tool?' (neato). A 'Glow' button is located at the bottom left.

Search Params

Subdomain logdog

Search inputname:logglyweb

Parser/Field Select Parser

Regex $^(\d+\.\d+\.\d+\.\d+)\s+\S+\s+\S+\s+\S+\s+\"GET (.*)\"(?:\s+)?$$

Rows 20

AfterGlow Params

Split nodes? ☒

Split Mode One unique event node

Print node count ☐

Color Config

Filter Config

Threshold:

Source fanout:

Event fanout:

Cluster Config Select Clustering for Source Select Clustering for Target

Node labels? ☐

Two nodes only (not three)? ☒

Default Edge Length 1.2

Which Viz Tool? neato

Glow

Loggly



JSON



CSV



DOT



Graph

Visualization Use-Cases

Old Skewl

Yesterday

Today – Cloud

```

20:25:00.606664 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], seq 1:1440, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204146], length 1448
20:25:00.606722 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 1440, win 2184, options [nop,nop,TS val 4176204146 ecr 1236457188], length 0
20:25:00.679517 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [P.], seq 1440:1458, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204146], length 0
20:25:00.679607 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 1458, win 2184, options [nop,nop,TS val 4176204216 ecr 1236457188], length 0
20:25:00.696440 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], seq 1458:2006, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204146], length 1448
20:25:00.696578 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 2006, win 2008, options [nop,nop,TS val 4176204236 ecr 1236457188], length 0
20:25:00.705574 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], seq 1458:2006, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204146], length 1448
20:25:00.705878 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 2006, win 2008, options [nop,nop,TS val 4176204245 ecr 1236457188,nop,nop,sack 1 (1458:2006)], length 0
20:25:00.731543 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], seq 2006:4354, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204216], length 1448
20:25:00.731580 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 4354, win 3632, options [nop,nop,TS val 4176204271 ecr 1236457188], length 0
20:25:00.744852 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [P.], seq 4354:5554, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204271], length 1200
20:25:00.744926 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 5554, win 4356, options [nop,nop,TS val 4176204288 ecr 1236457188], length 0
20:25:00.744954 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [P.], seq 5554:5838, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204288], length 284
20:25:00.765432 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 5838, win 4356, options [nop,nop,TS val 4176204300 ecr 1236457188], length 0
20:25:00.765719 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], seq 5838:7206, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204288], length 1448
20:25:00.767350 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 7206, win 5008, options [nop,nop,TS val 4176204307 ecr 1236457188], length 0
20:25:00.777787 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [P.], seq 7206:7644, ack 601, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204307], length 358
20:25:00.777993 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 7644, win 5008, options [nop,nop,TS val 4176204317 ecr 1236457188], length 0
20:25:00.778258 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [P.], seq 601, ack 7644, win 5008, options [nop,nop,TS val 4176204317 ecr 1236457188], length 0
20:25:00.783978 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [C.], ack 602, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204317], length 0
20:25:00.787891 IP 205.206.231.13.80 > 172.16.1.222.51460: Flags [F.], seq 7644, ack 602, win 17376, options [nop,nop,TS val 1236457188 ecr 4176204317], length 0
20:25:00.787924 IP 172.16.1.222.51460 > 205.206.231.13.80: Flags [C.], ack 7645, win 5008, options [nop,nop,TS val 4176204327 ecr 1236457188], length 0
20:25:00.604582 IP 172.16.1.1.53 > 172.16.1.19.3713: 37440 11/9/1 ONAME liveupdate.symantec.dep.net., ONAME symantec-gaoredirector.oxada.net., ONAME a568.d.oxamail.net., A 66.28.46.145, A 66.28.46.136, A 66.28.46.151, A 66.28.46.142, A 66.28.46.144, A 66.28.46.143, A 66.28.46.134, A 66.28.46.150 (503)
20:25:00.926495 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [S.], seq 2833798679, ack 3063552958, win 17520, options [ssn 1440,nop,nop,sackOK], length 0
20:25:10.129046 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [C.], ack 235, win 17520, length 0
20:25:10.132592 STP 802.1d, Config, Flags [none], bridge-id 8001.00-8f:34:11:c3:c0.8004, length 35
20:25:10.352200 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [P.], seq 1:193, ack 235, win 17520, length 192
20:25:10.859334 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [P.], seq 193:390, ack 526, win 17520, length 197
20:25:11.578664 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [P.], seq 390:587, ack 779, win 17520, length 197
20:25:11.978837 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [C.], ack 5874, win 17520, length 0
20:25:12.056497 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [P.], seq 587:784, ack 1074, win 17520, length 197
20:25:12.434405 IP 66.28.46.145.80 > 172.16.1.19.3713: Flags [P.], seq 784:981, ack 1350, win 17520, length 197
20:25:14.841779 OTpV1, length 26
20:25:15.433067 IP 67.97.165.51.4500 > 172.16.1.157.4500: UDP-encap: ESP(cpi-b0e0f6f5d,seq-b0c47), length 84
20:25:19.142575 STP 802.1d, Config, Flags [none], bridge-id 8001.00-8f:34:11:c3:c0.8004, length 35
~/Data/logs$

```

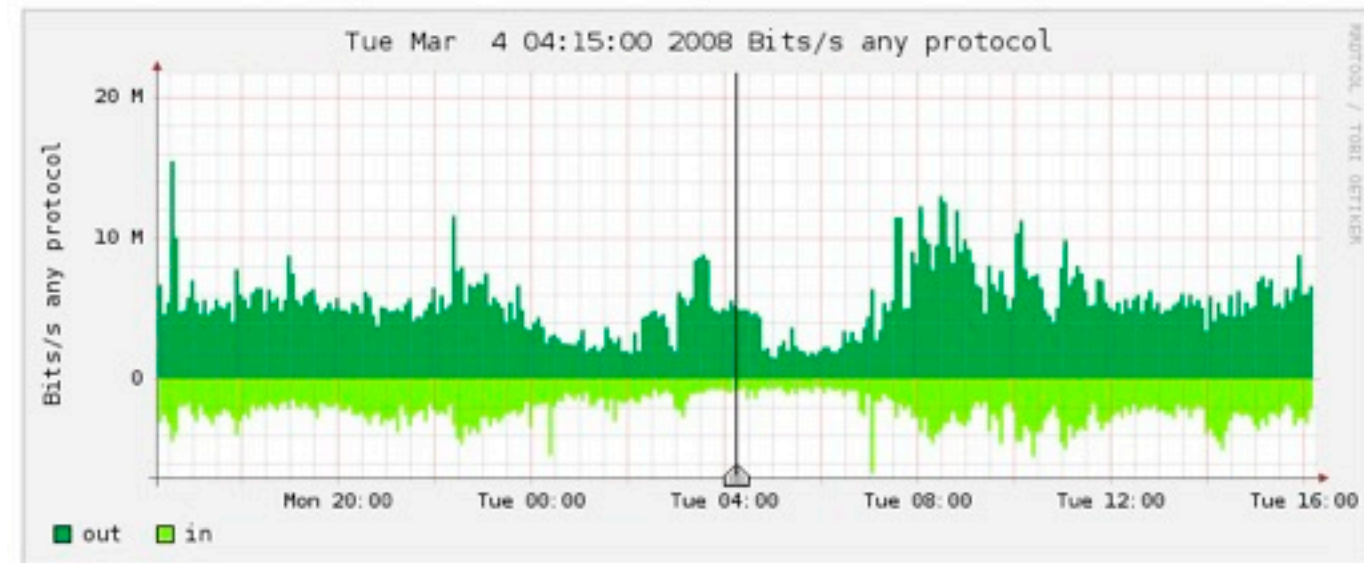
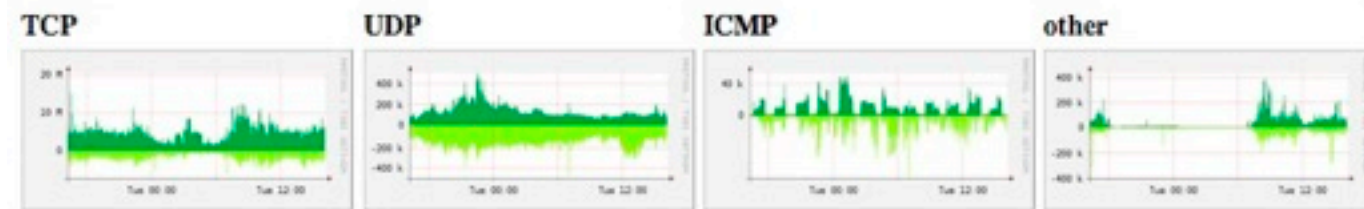
[illegible]

Traditional Style

Yesterday

Today – Cloud

Profile: inout



Profileinfo:

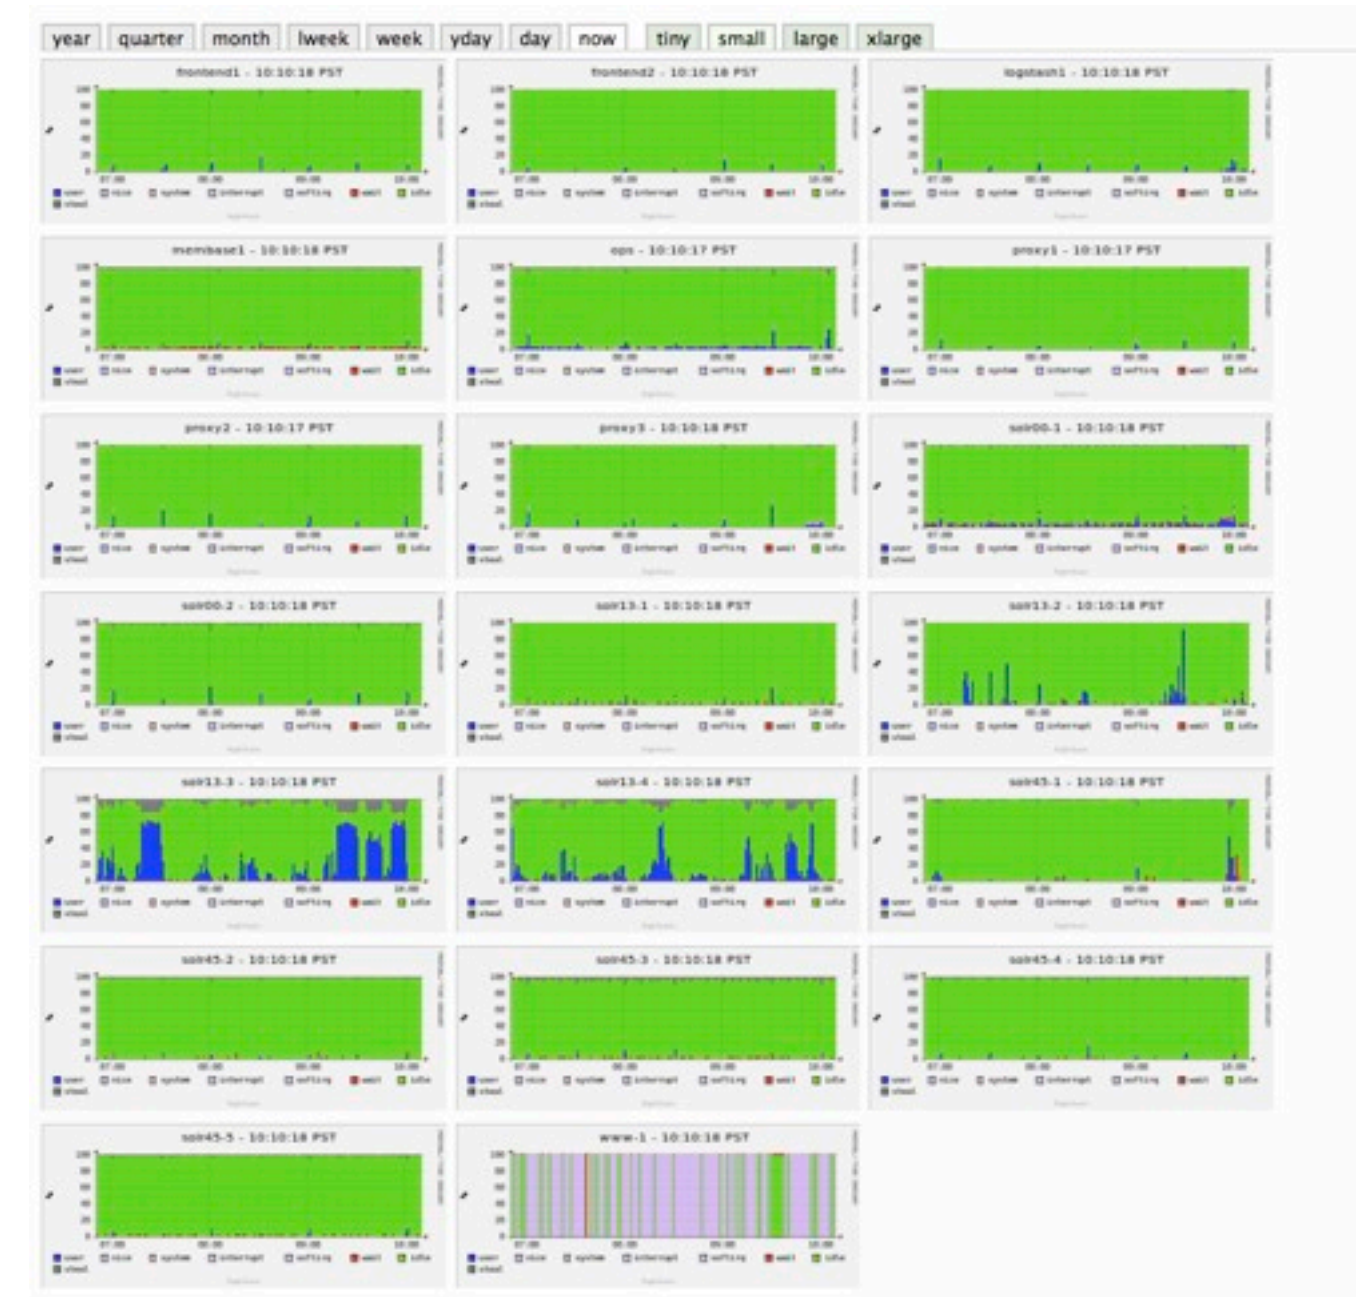
Type: continuous / shadow
Max: unlimited
Exp: never
Start: Jun 29 2007 - 02:20 MET
End: Mar 04 2008 - 16:15 MET

t_start 2008-03-04-04-15
t_end 2008-03-04-04-15

Packets



Flows



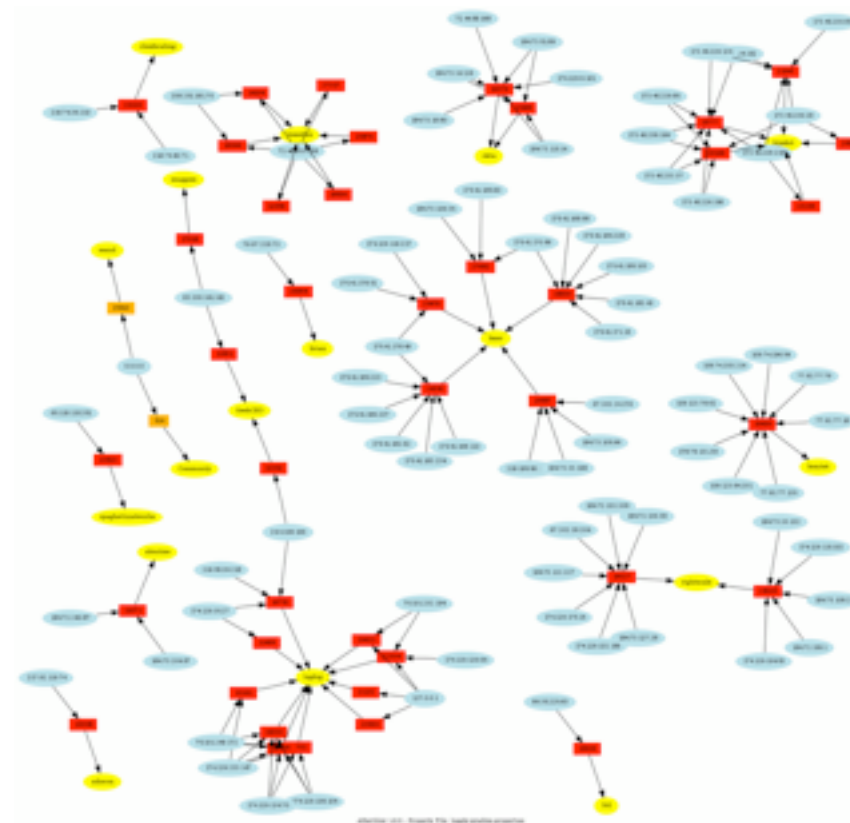
Logging as a Service

The Analysis Approach

Overview first

Zoom

Details on demand



```
zrlram@prod> search failure

2010 Nov 25 08:50:24.367 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:50:21.844 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:49:43.303 localhost loggly: severity=ERROR,requ
reason=invalid subdomain access,subdomain=prod
2010 Nov 25 08:49:26.390 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:49:24.193 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:49:21.774 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:48:24.350 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:48:21.797 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:48:05.407 localhost loggly: severity=INFO,AXES:
2010 Nov 25 08:47:24.882 localhost loggly: severity=INFO,AXES:

From:, Until:, Search:, Sort Order:, Results:
Pages: 1 2 3 ... 19 Jump to page: 
```

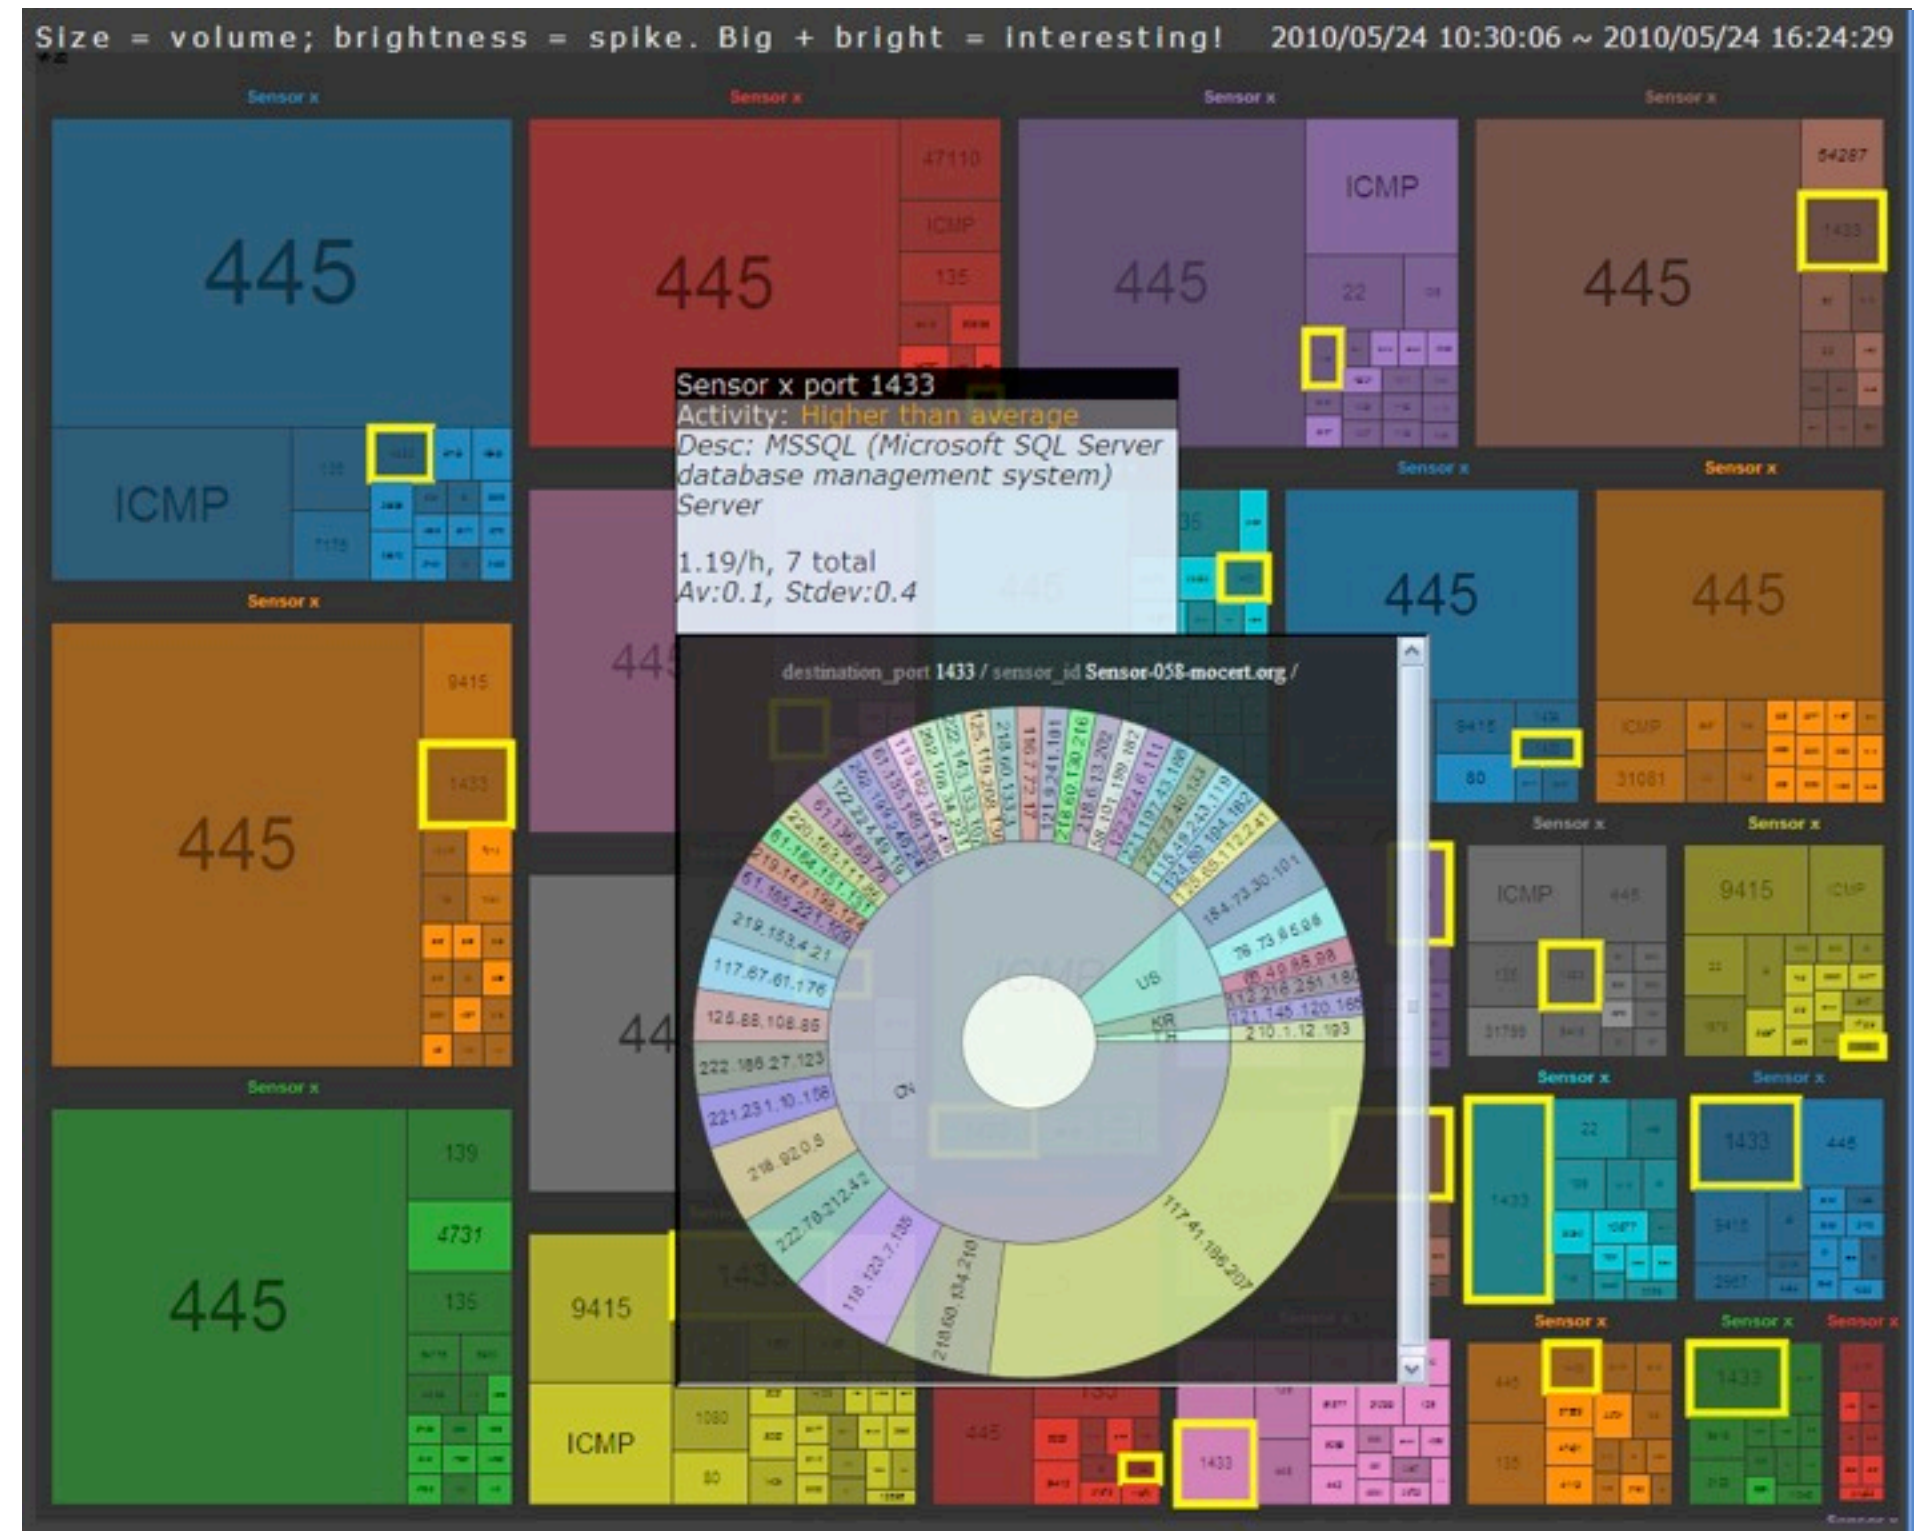
Principle by Ben Shneiderman



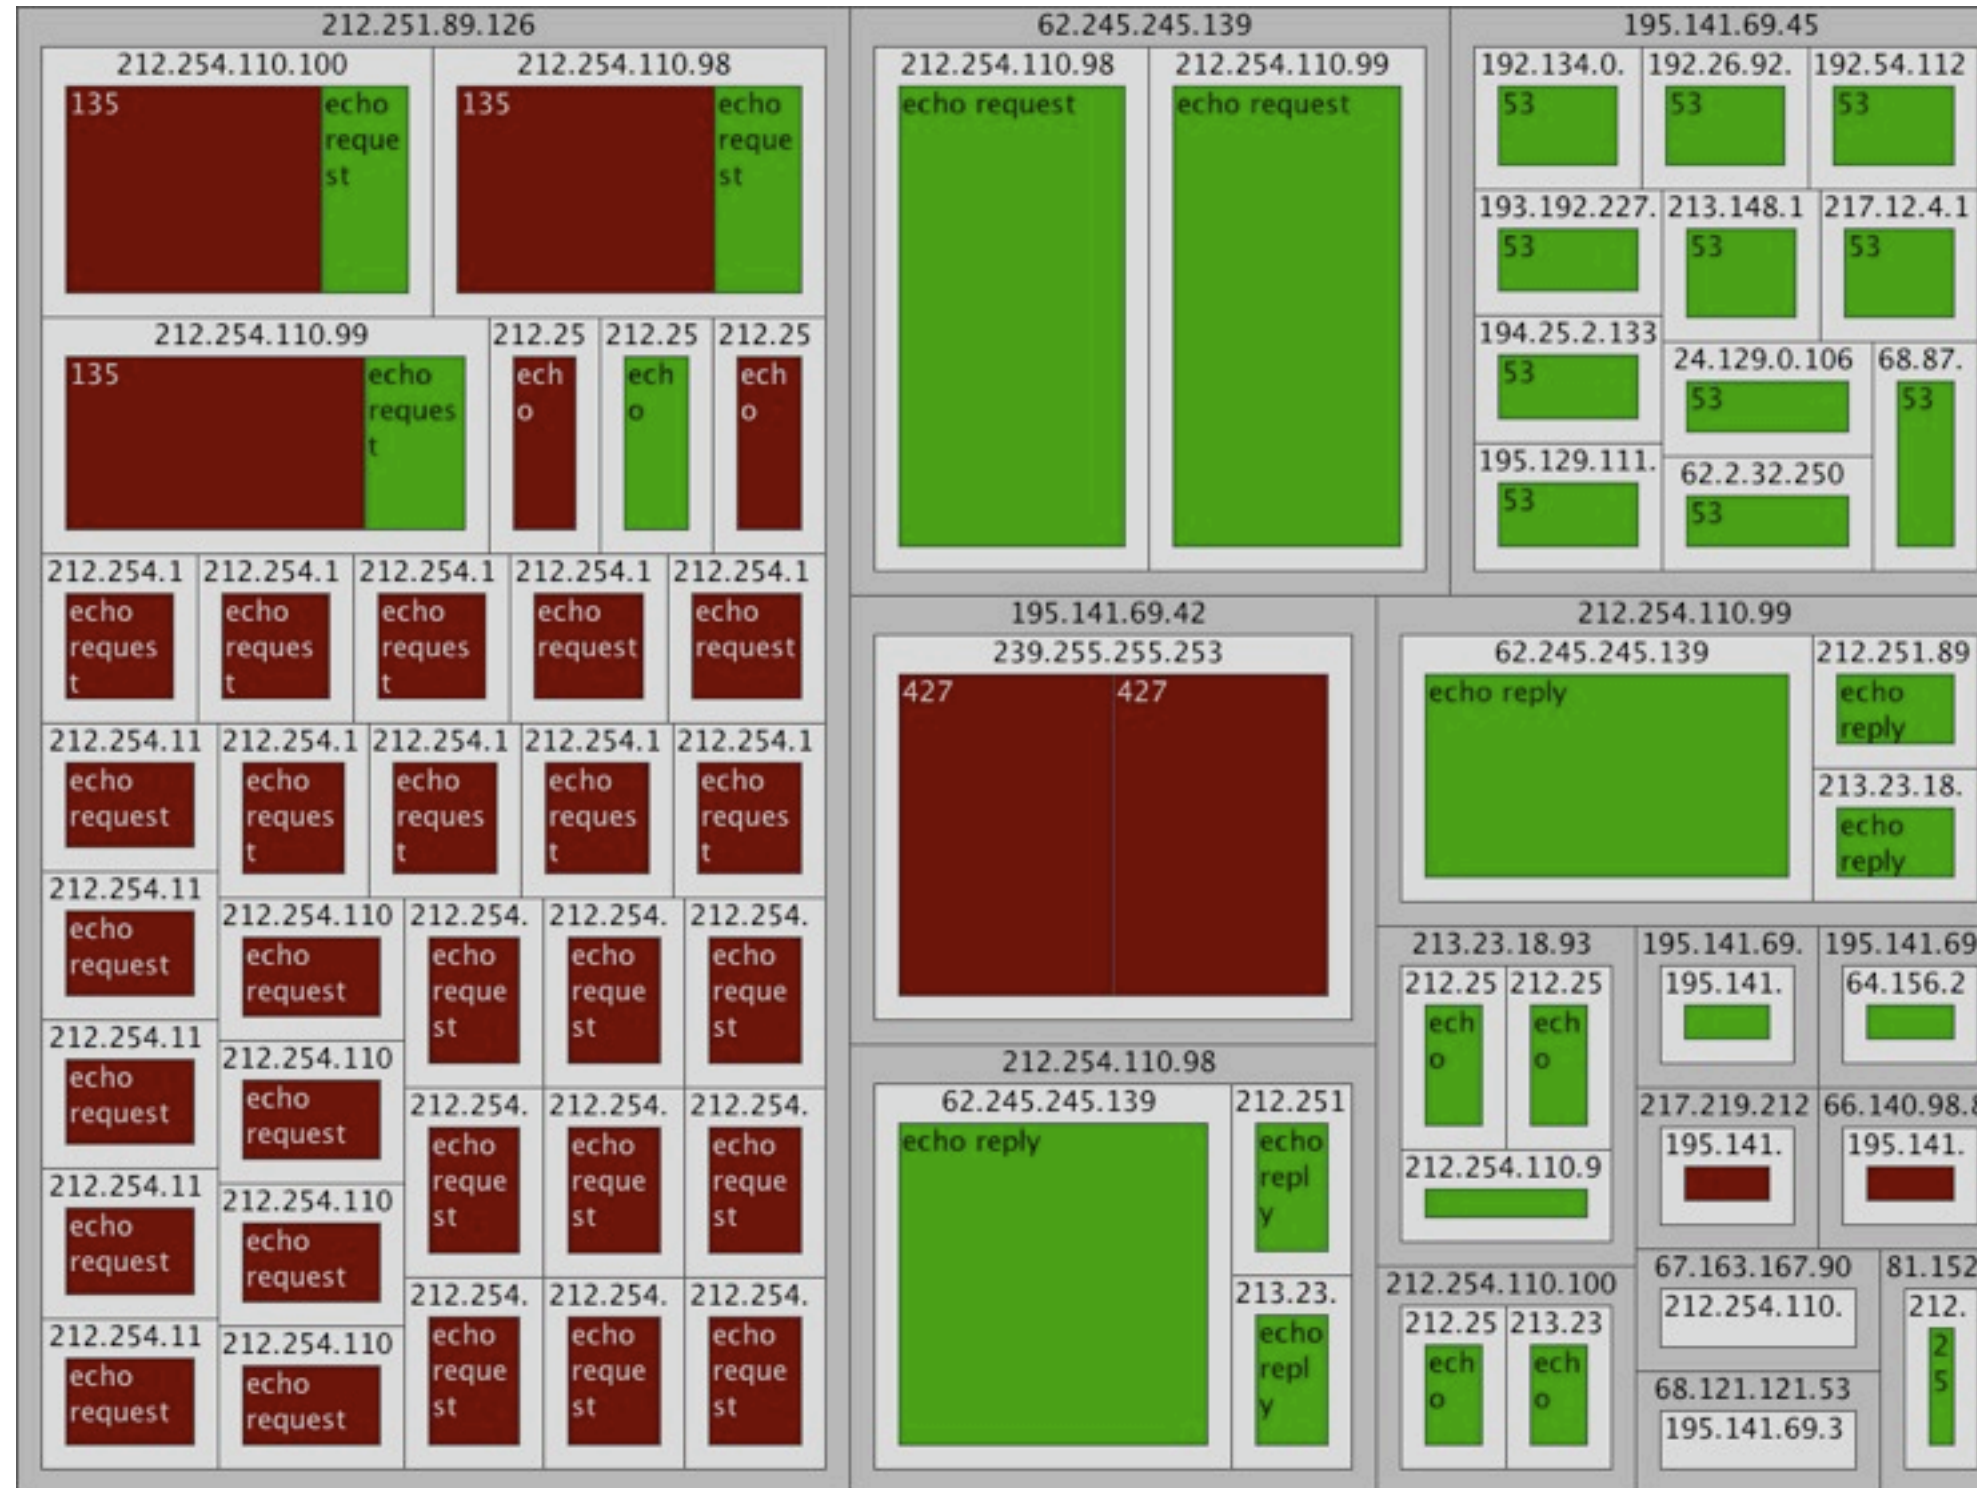
Logging as a Service

NetFlow Visualization

- Treemap
- Protovis.js
- Size: Amount
- Brightness: Variance
- Color: Sensor
- Shows: Scans – bright spots
- Thanks to Chris Horsley

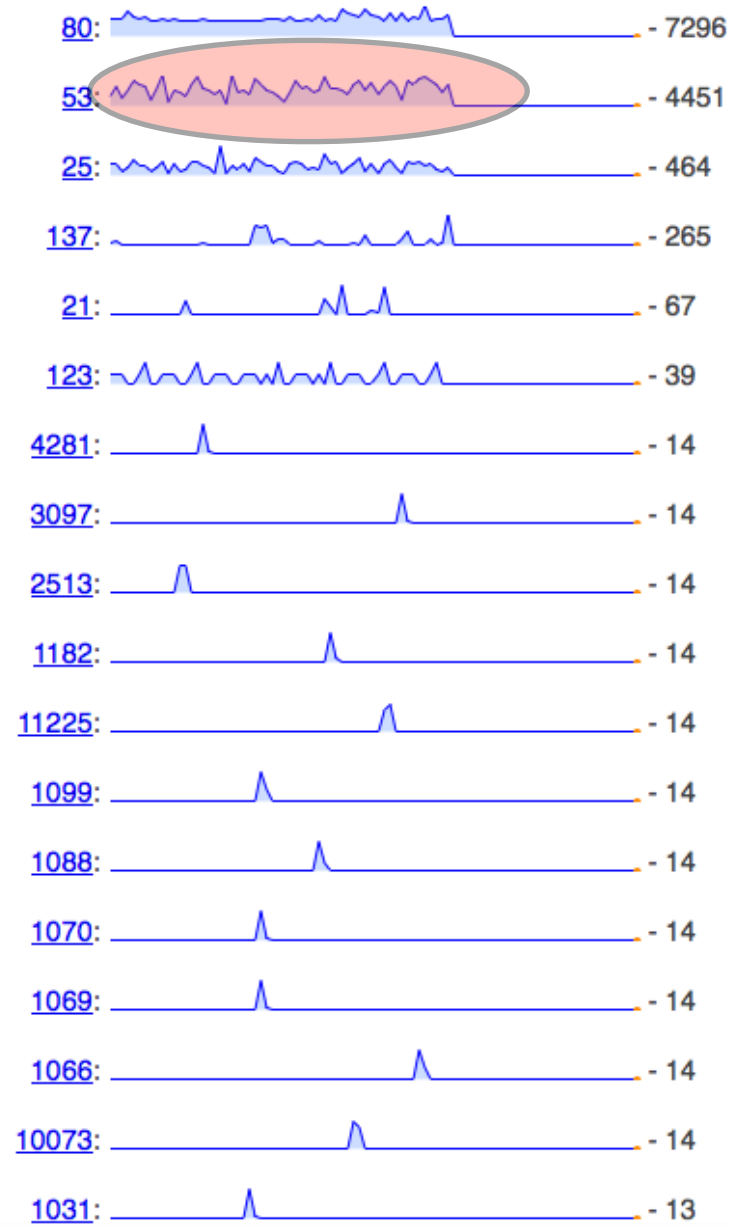


Firewall Treemap

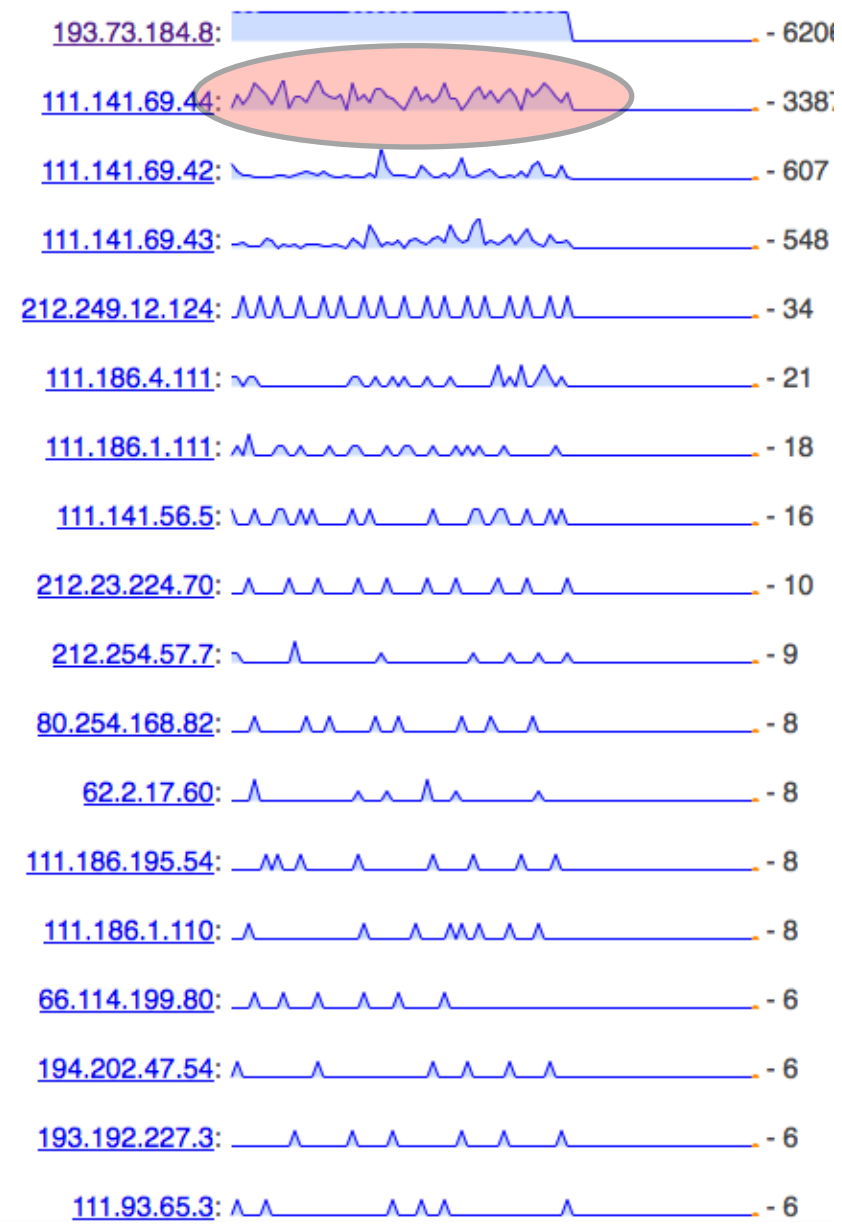


Firewall Log

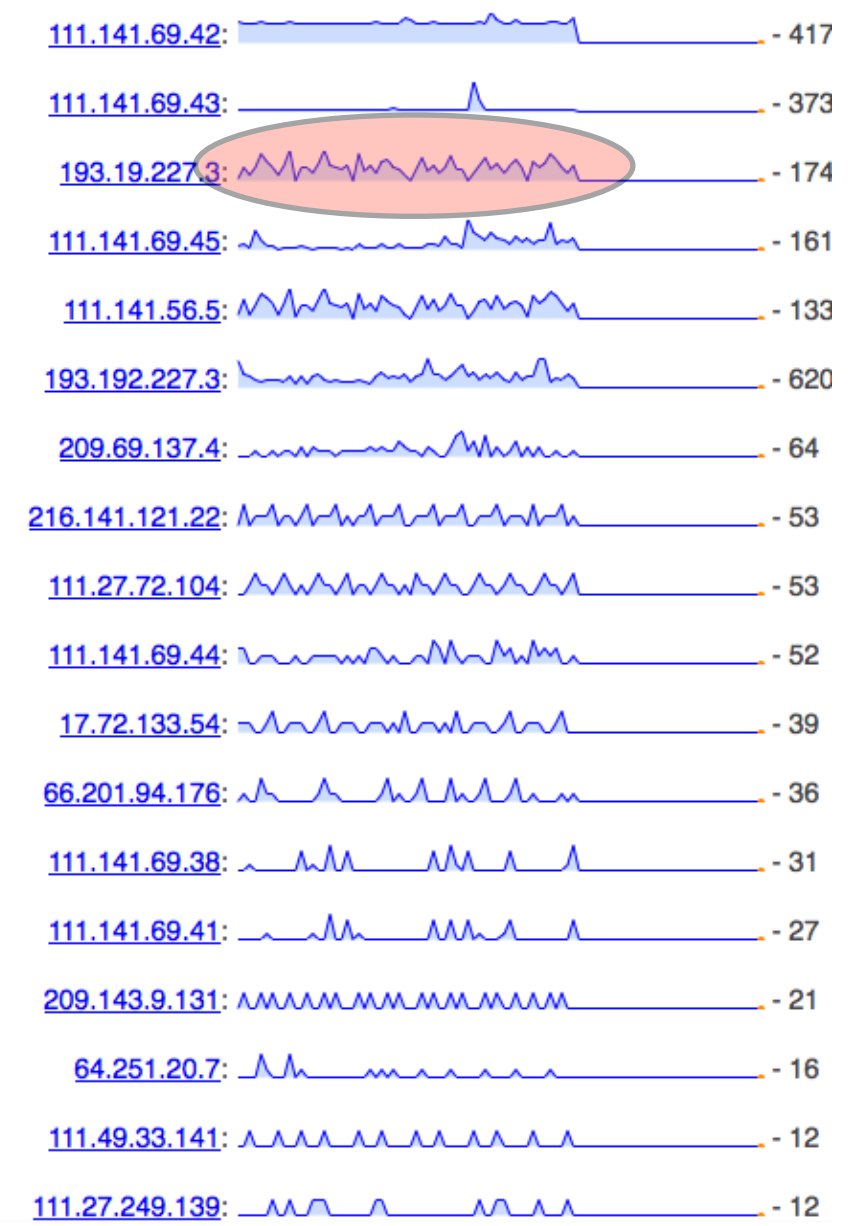
Port



Source IP

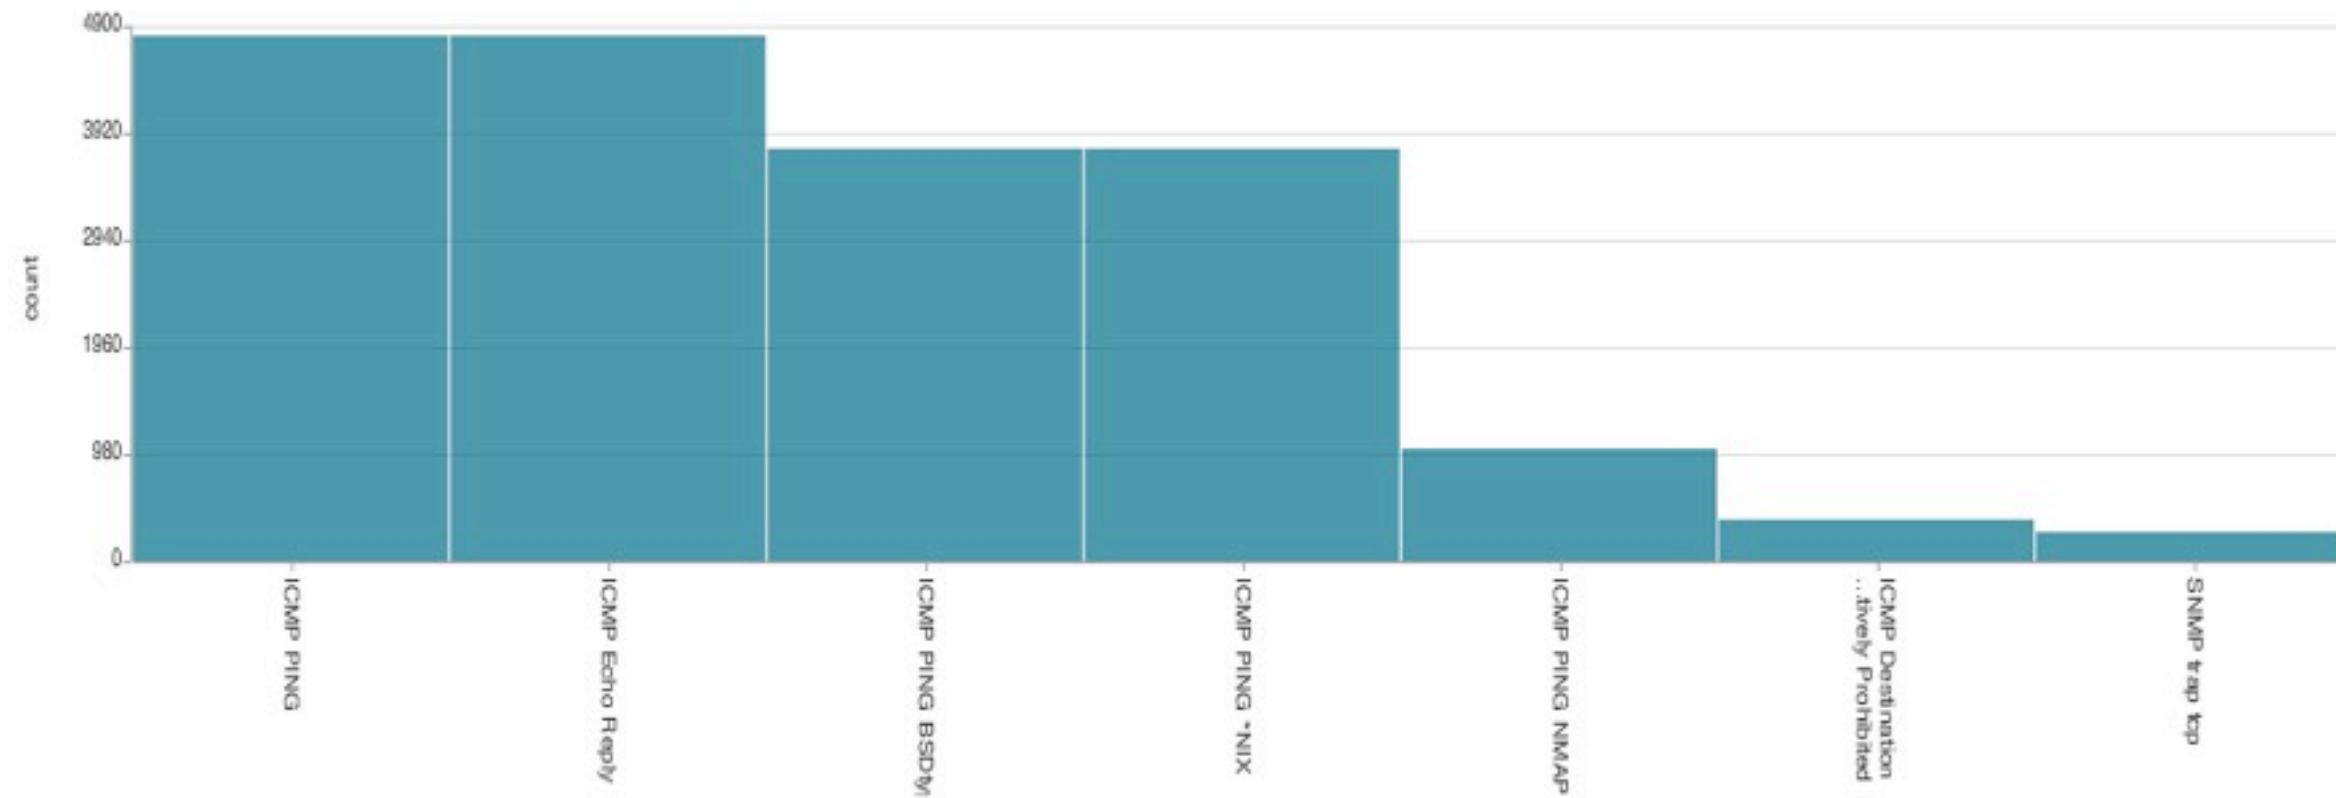


Destination IP



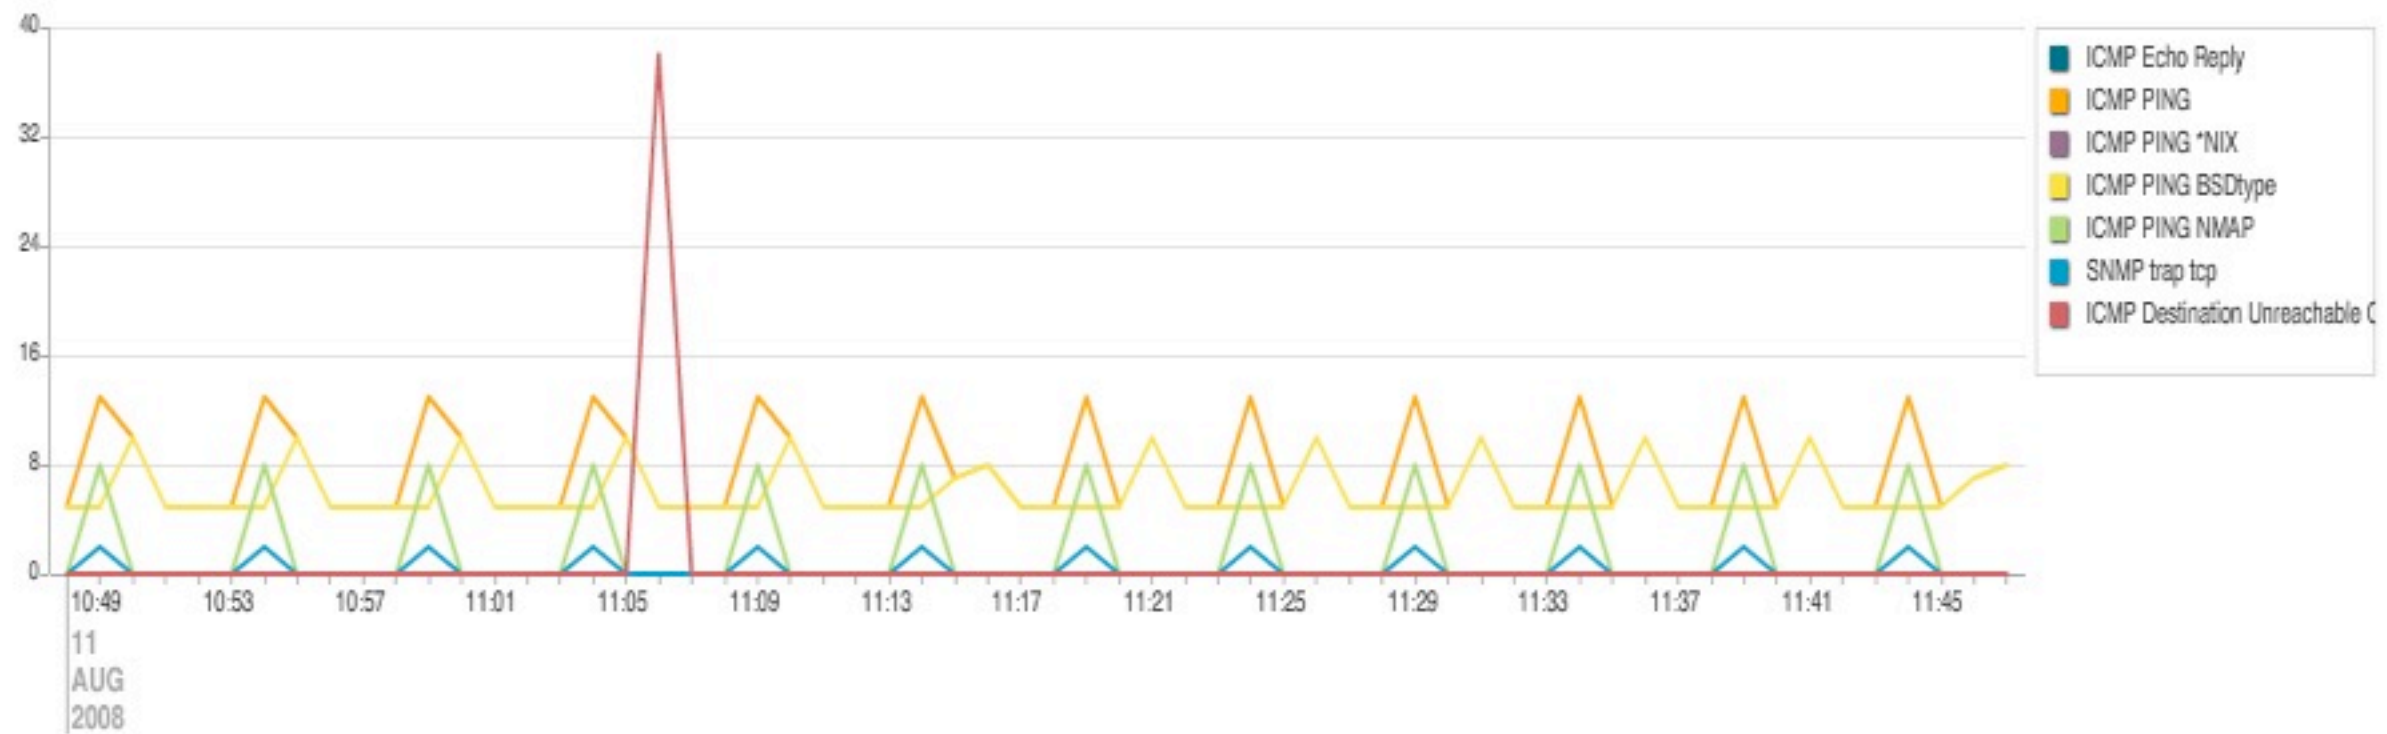
IDS Signature Tuning

Top signatures

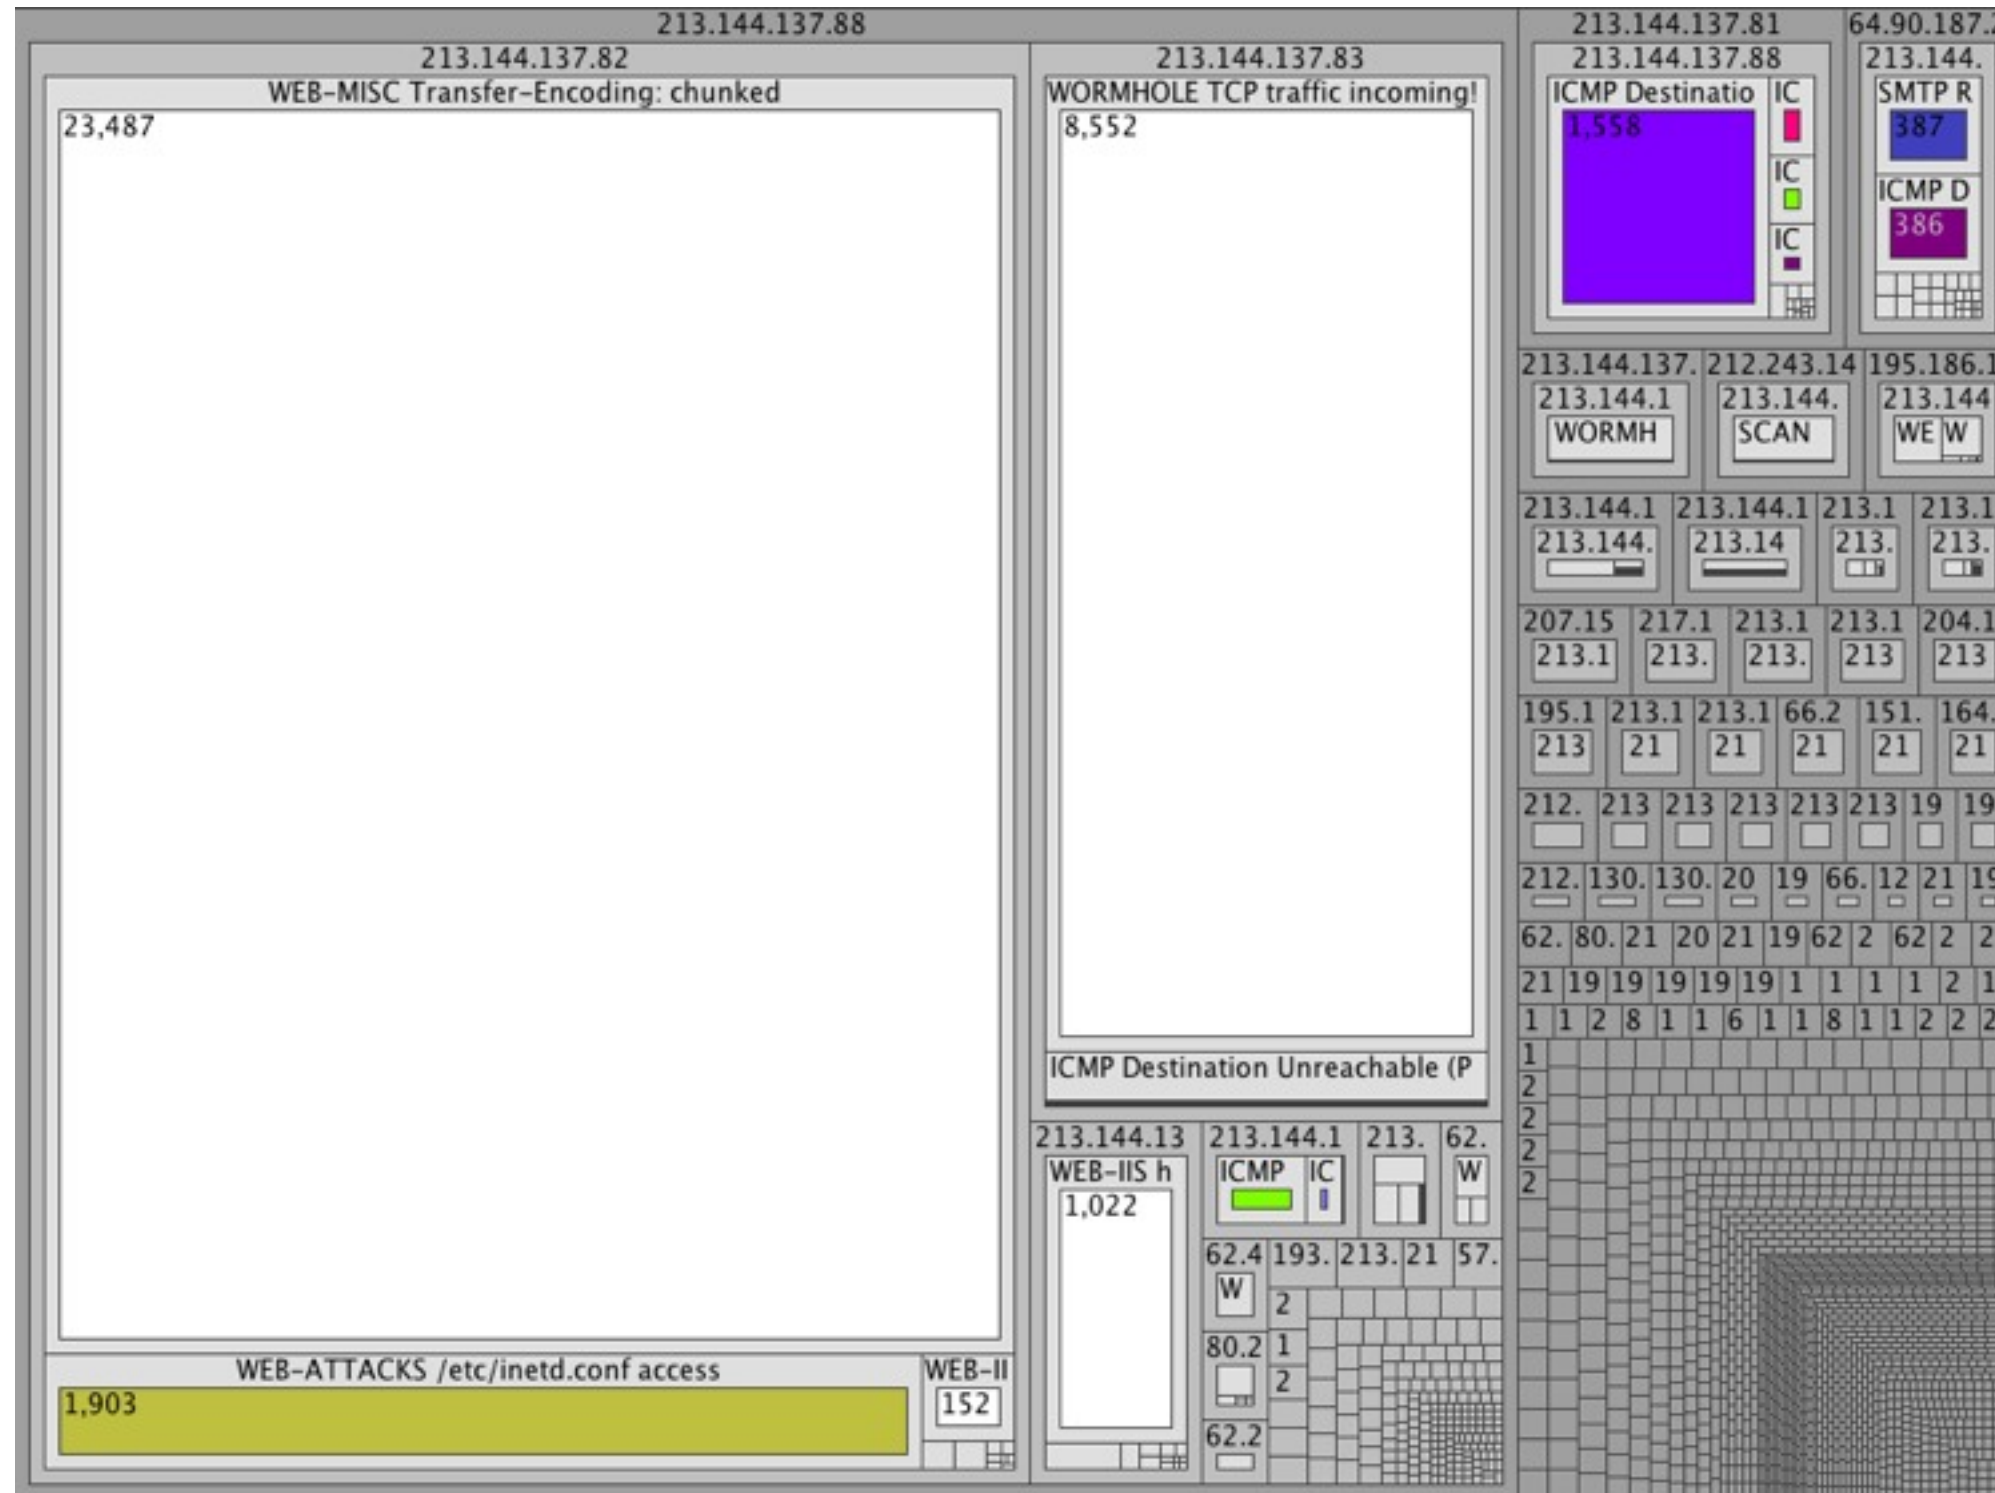


Signatures Over Time

count of _raw vs. time by name for results in the past 60 minutes



IDS Sig Tuning – Treemap



Hierarchy:

Source

Destination

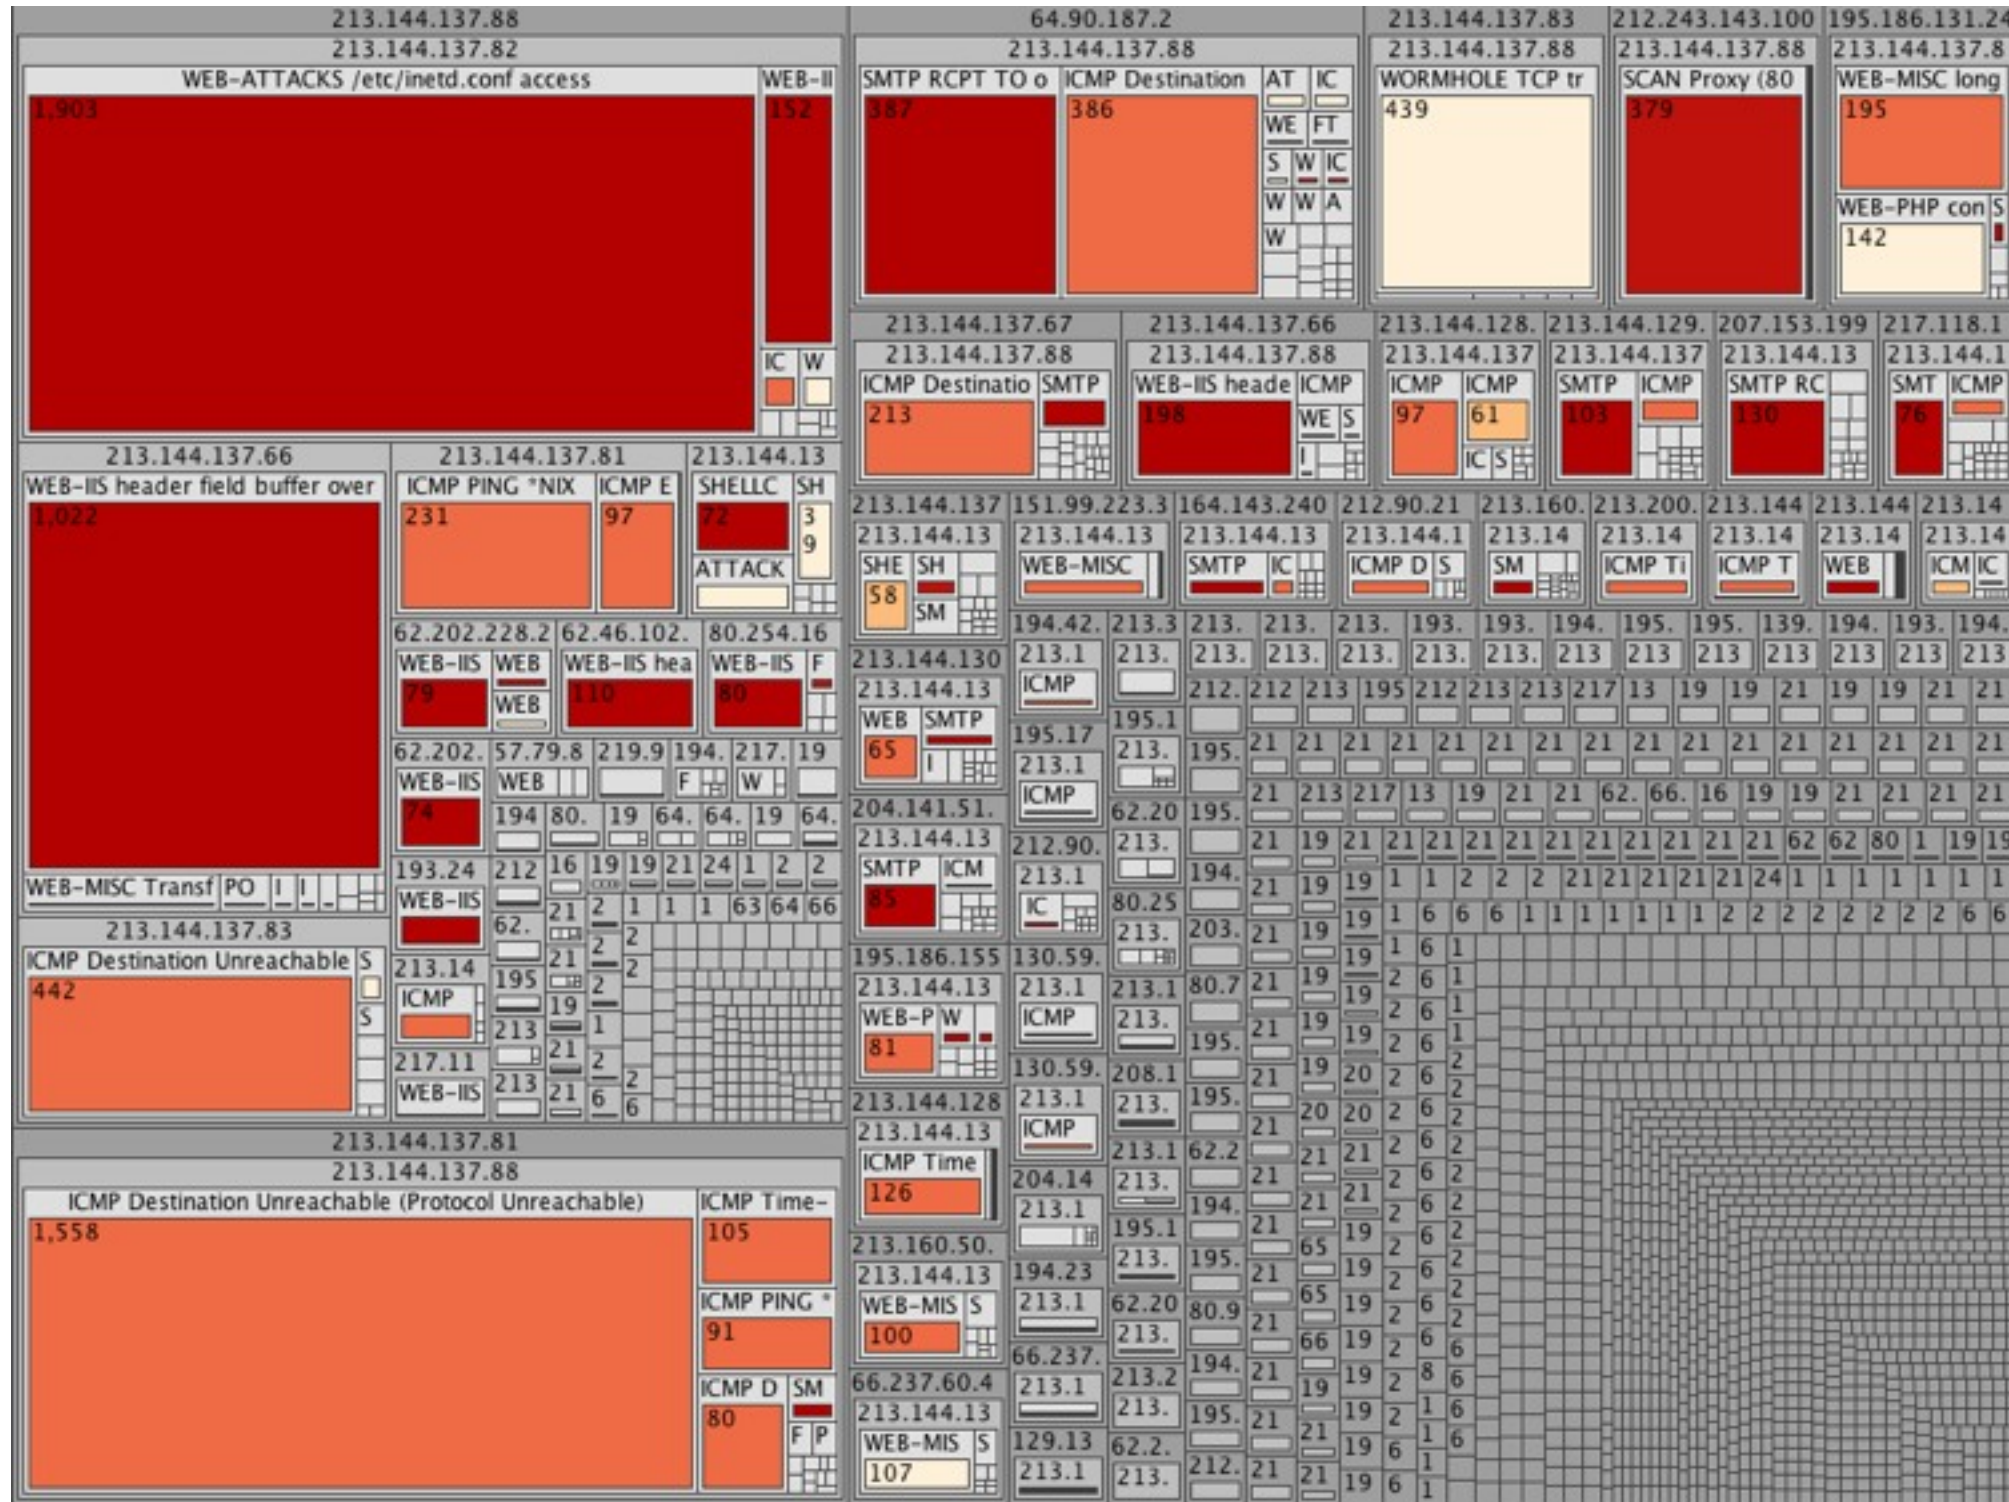
Signature

Number of Events

Color: Service

Size: Number of alerts

IDS Sig Tuning - Treemap



Hierarchy:

Source

Destination

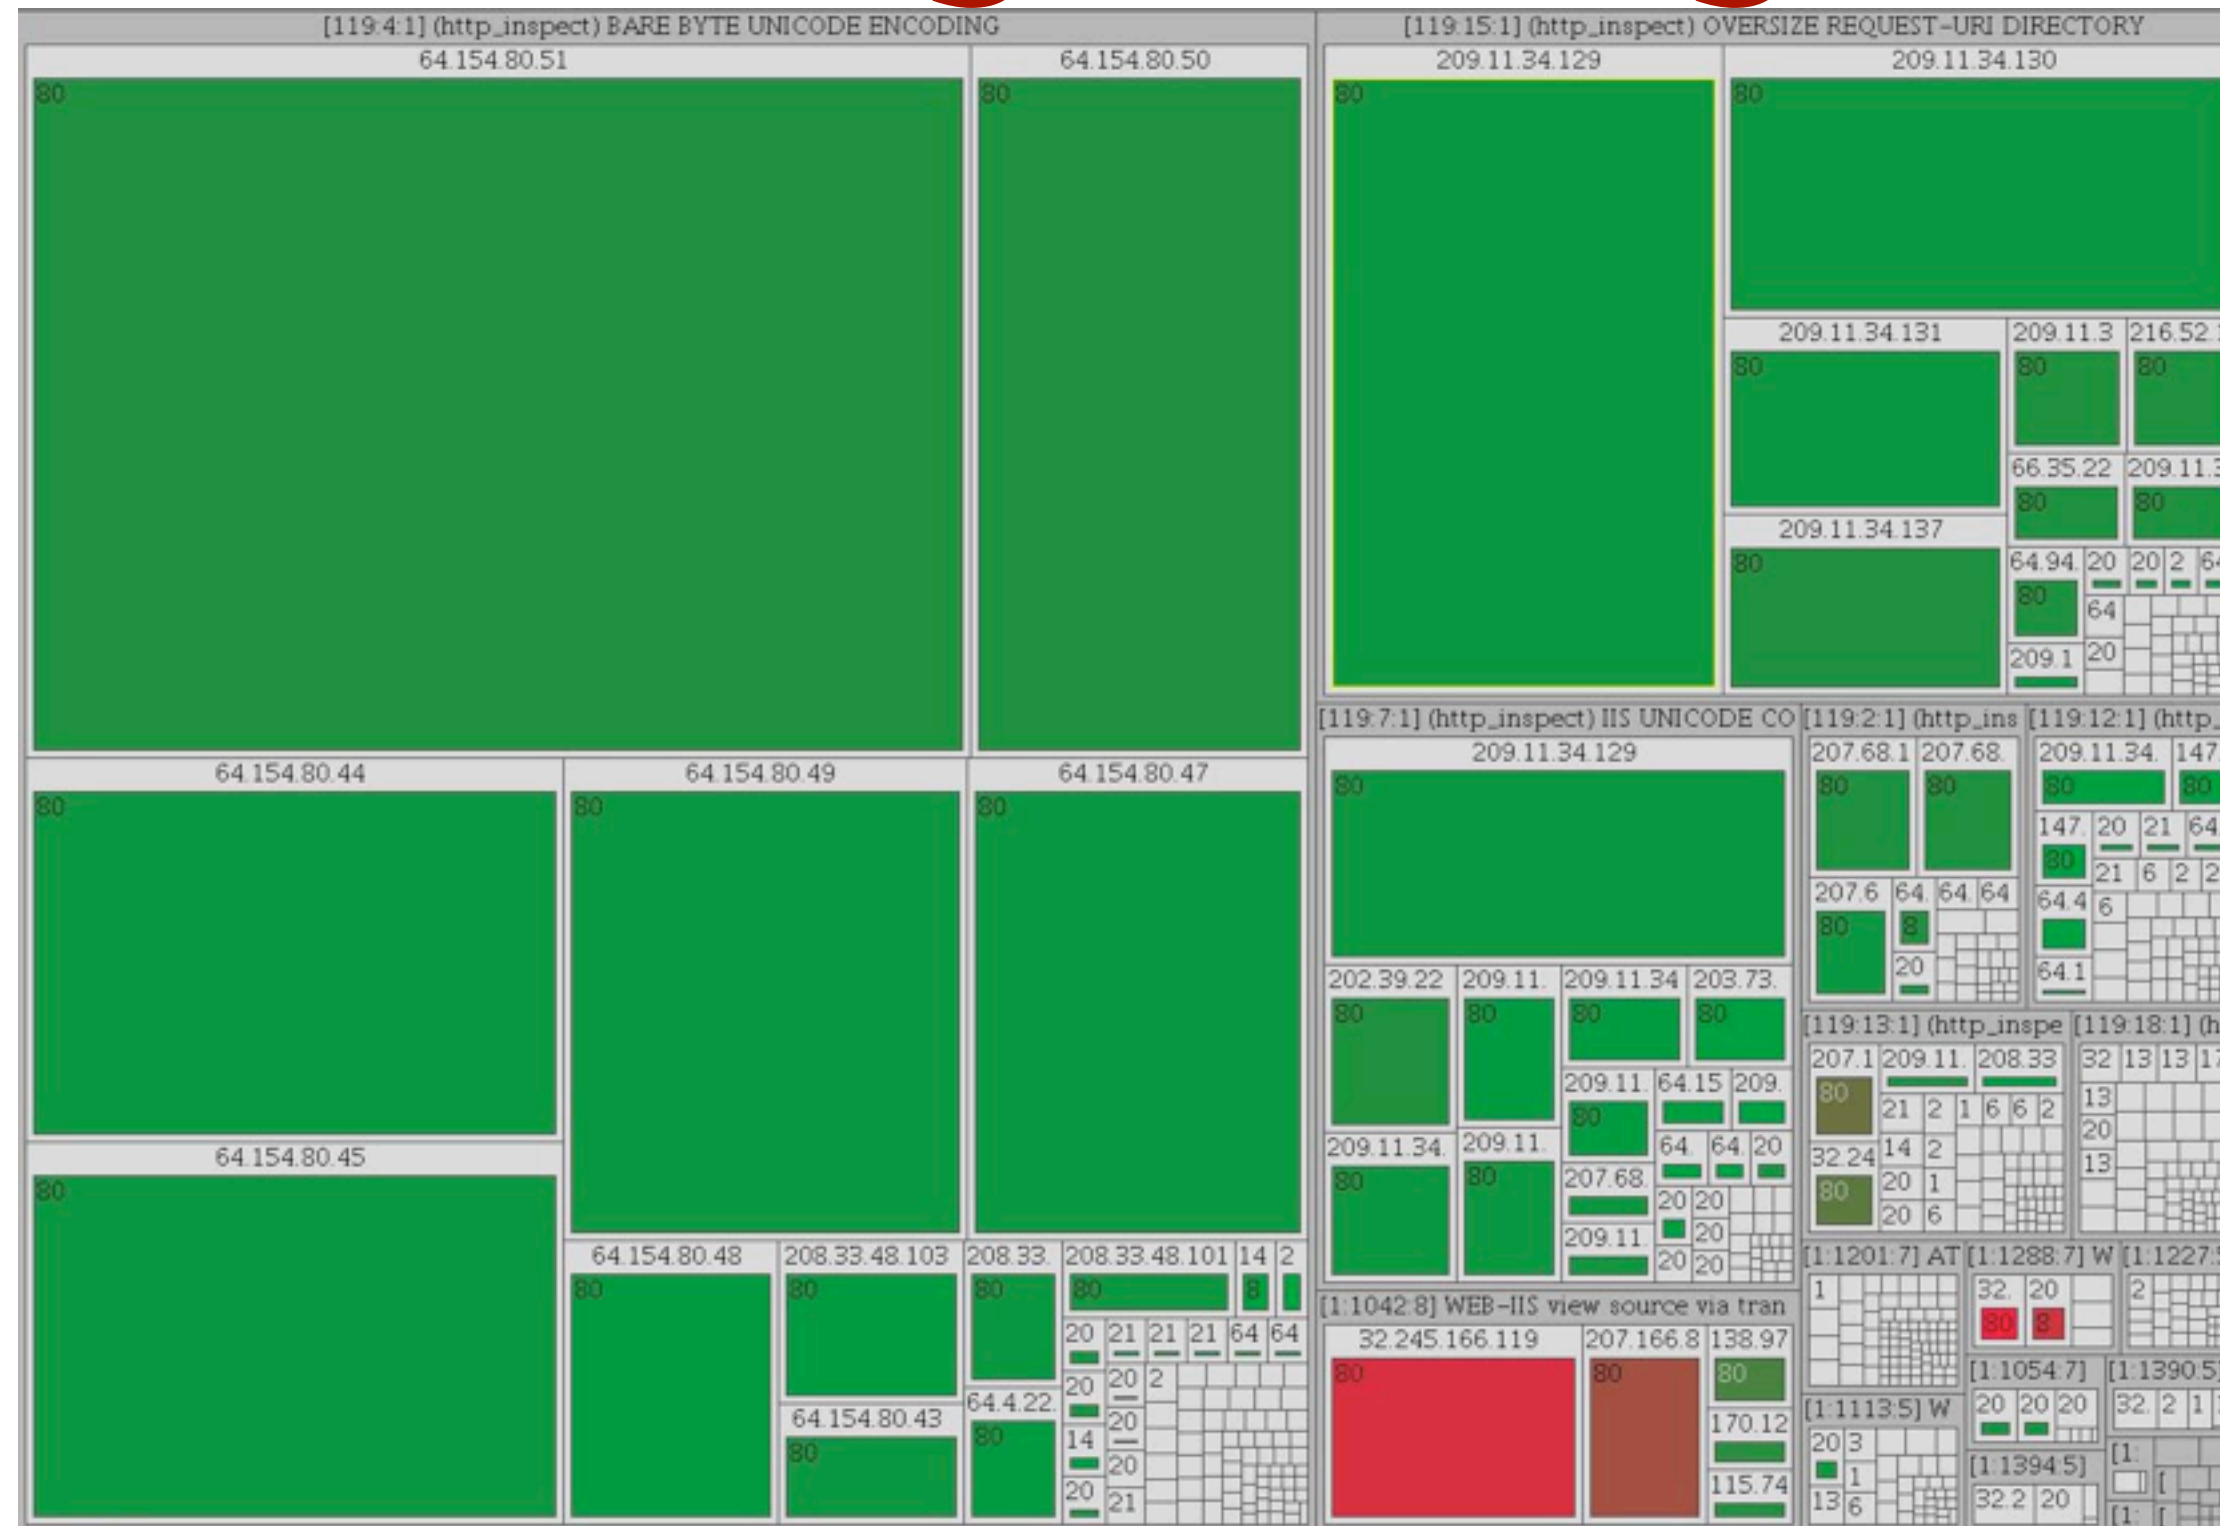
Signature

Number of Events

Color: Priority

Size: Number of alerts

IDS Sig Tuning – Treemap



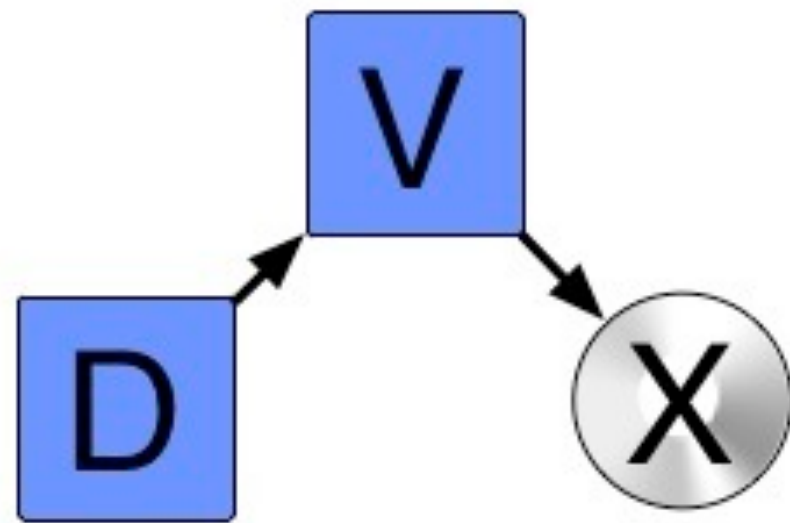
Hierarchy:
Signature
Source
Service (Port)
Color: Priority
Size: Number of alerts

Visualization Resources

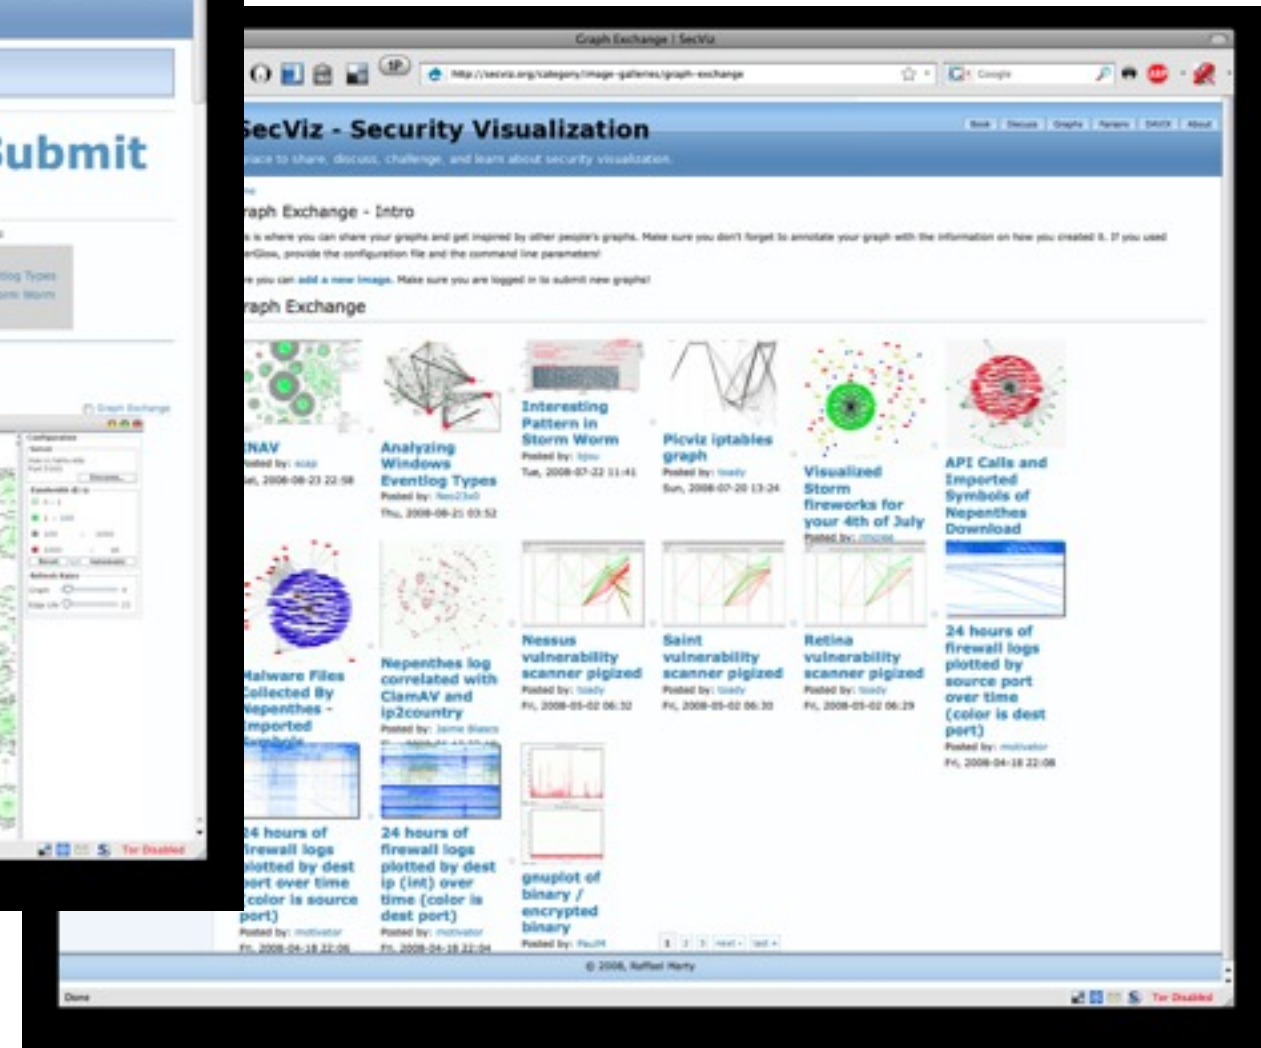
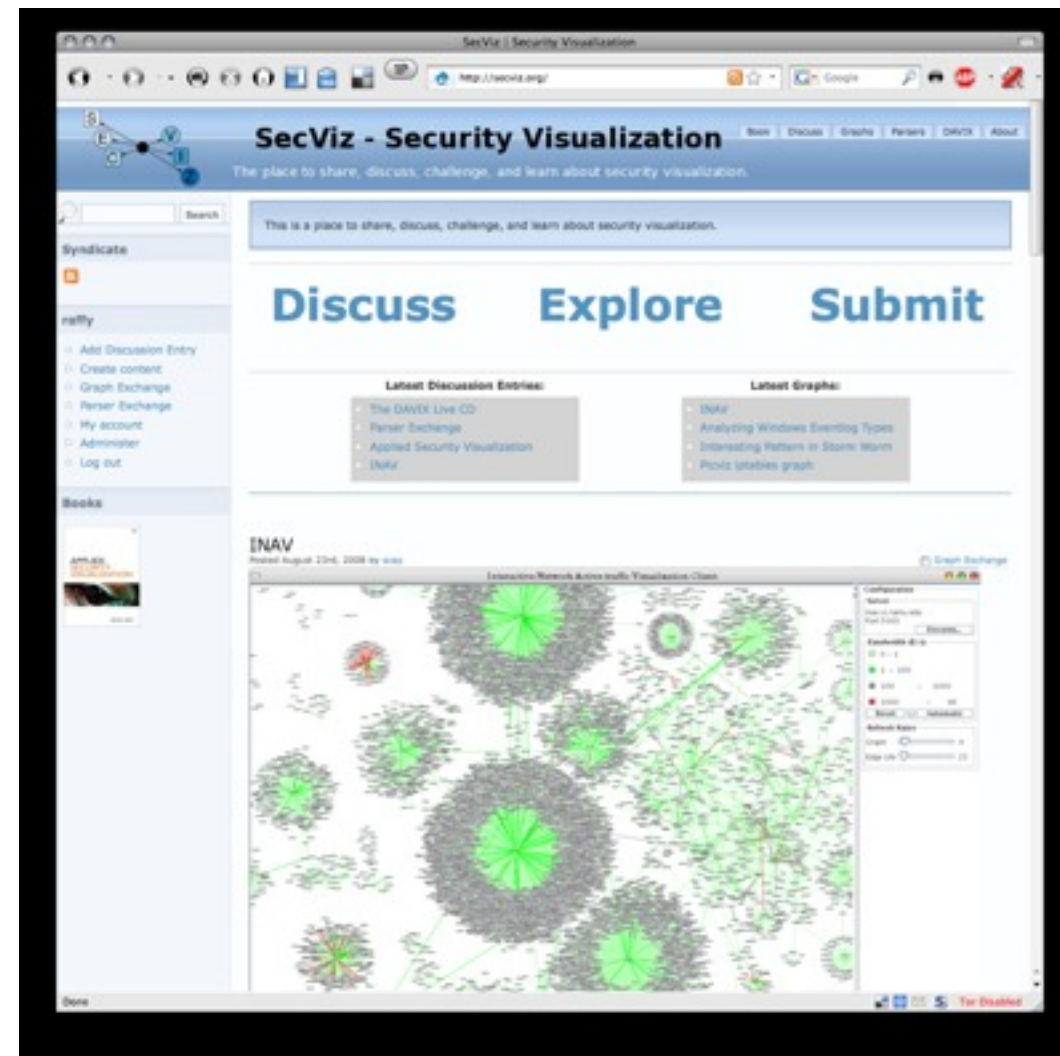
<http://secviz.org>

Share, discuss, challenge, and learn about security visualization.

- List: secviz.org/maillinglist
- Twitter: **@secviz**



davix.secviz.org

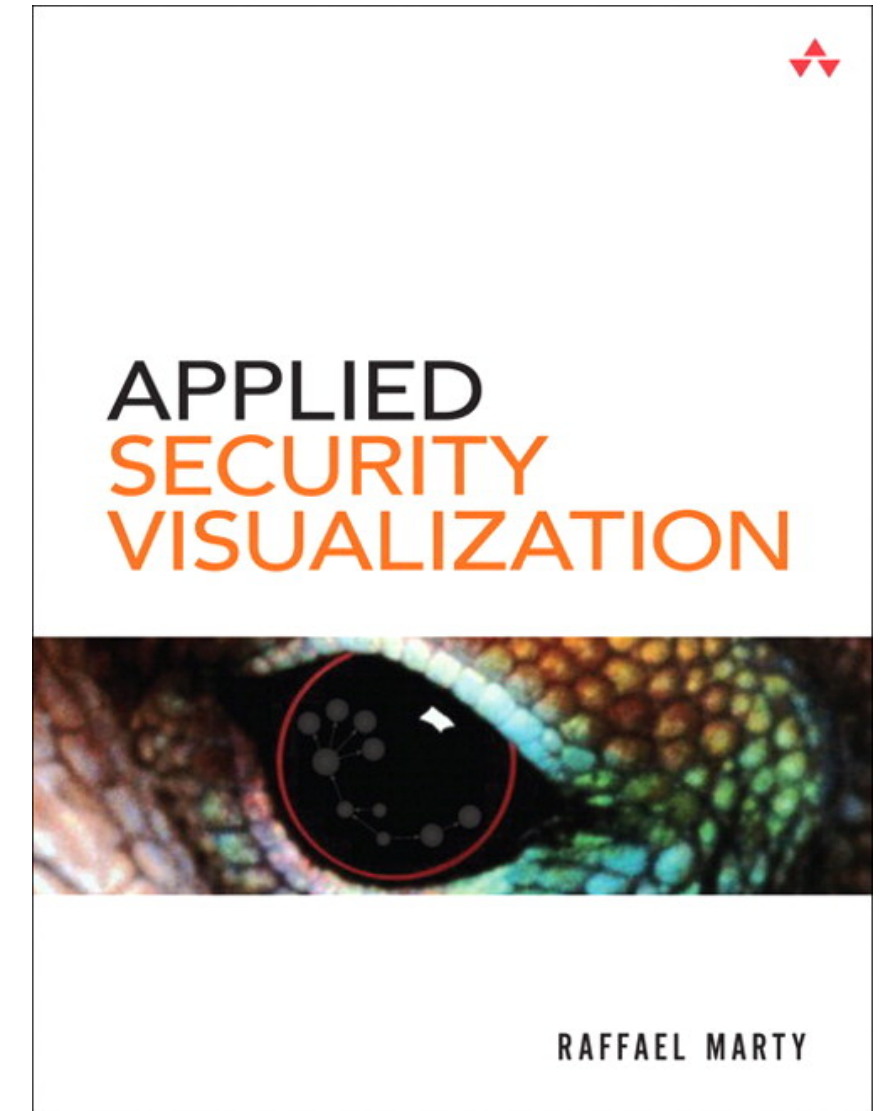


Applied Security Visualization

- Bridging the gap between security and visualization
- Hands-on, end to end examples
- Data processing and analysis

Chapters

- Visualization
- Data Sources
- From Data to Graphs
- Perimeter Threat
- Compliance
- Insider Threat
- Visualization Tools



Addison Wesley (August, 2008)
ISBN: 0321510100

loggly

about.me/raffy

We are **hiring!**

