

FloVis: Flow Visualization System

Teryl Taylor, Diana
Paterson, Joel Glanfield,
Carrie Gates, Stephen
Brooks and John McHugh

March 4, 2009

Carrie.Gates@ca.com



Need

- > Multiple organizations within a single visualization
- > Flow level data only – highly aggregated
- > Over 1.5 *billion* flows per day
- > 16 million IP addresses
- > Too few analysts with no time for training

sIP	dIP	sPort	dPort	pro	packets	bytes	flags	sTime	dur
168.192.2.25	10.10.15.223	1860	2100	6	2	96	S	2006/07/03T19:23:15.000	6.000
168.192.2.25	10.10.17.150	2164	2100	6	2	96	S	2006/07/03T19:23:25.000	6.000
168.192.2.25	10.10.15.225	2466	2100	6	1	48	S	2006/07/03T19:23:35.000	0.000
168.192.2.25	10.10.17.155	3681	2100	6	3	144	S	2006/07/03T19:24:12.000	9.000
168.192.2.25	10.10.14.48	3980	2100	6	2	96	S	2006/07/03T19:24:25.000	6.000
168.192.2.25	10.10.16.193	3982	2100	6	2	96	S	2006/07/03T19:24:25.000	6.000
168.192.2.25	10.10.14.49	4282	2100	6	2	96	S	2006/07/03T19:24:35.000	6.000
168.192.2.25	10.10.15.13	4858	2100	6	2	96	S	2006/07/03T19:24:45.000	6.000
168.192.2.25	10.10.17.159	1212	2100	6	2	96	S	2006/07/03T19:24:56.000	6.000
168.192.2.25	10.10.16.196	1211	2100	6	2	96	S	2006/07/03T19:24:56.000	6.000
168.192.2.25	10.10.15.15	1513	2100	6	2	96	S	2006/07/03T19:25:06.000	6.000
168.192.2.25	10.10.16.198	1818	2100	6	2	96	S	2006/07/03T19:25:16.000	6.000
168.192.2.25	10.10.14.54	2117	2100	6	2	96	S	2006/07/03T19:25:26.000	6.000
168.192.2.25	10.10.15.17	2118	2100	6	2	96	S	2006/07/03T19:25:26.000	6.000

Requirements

> System:

- Use SiLK data (unidirectional flow data only)
- Scalability – 1.5 *billion* flows per day!

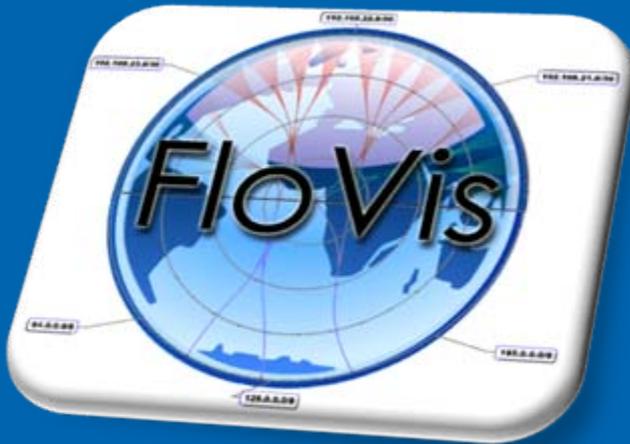
> Visualizations:

- Ability to detect patterns that are *not* detected currently through other means (e.g., NOT scans! 😊)
- Find the needle in the haystack!

Data Used for Examples

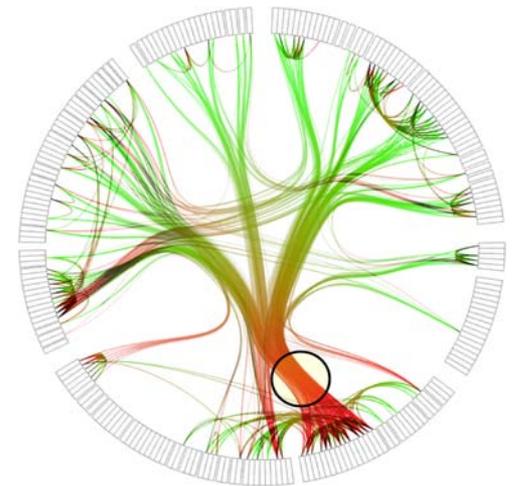
- > 1024 IP addresses with ~ 100 active hosts
- > Mix of research and workstation machines
- > Addresses have been anonymized
- > Contains 1 month of network activity from 2006

FlowBundle

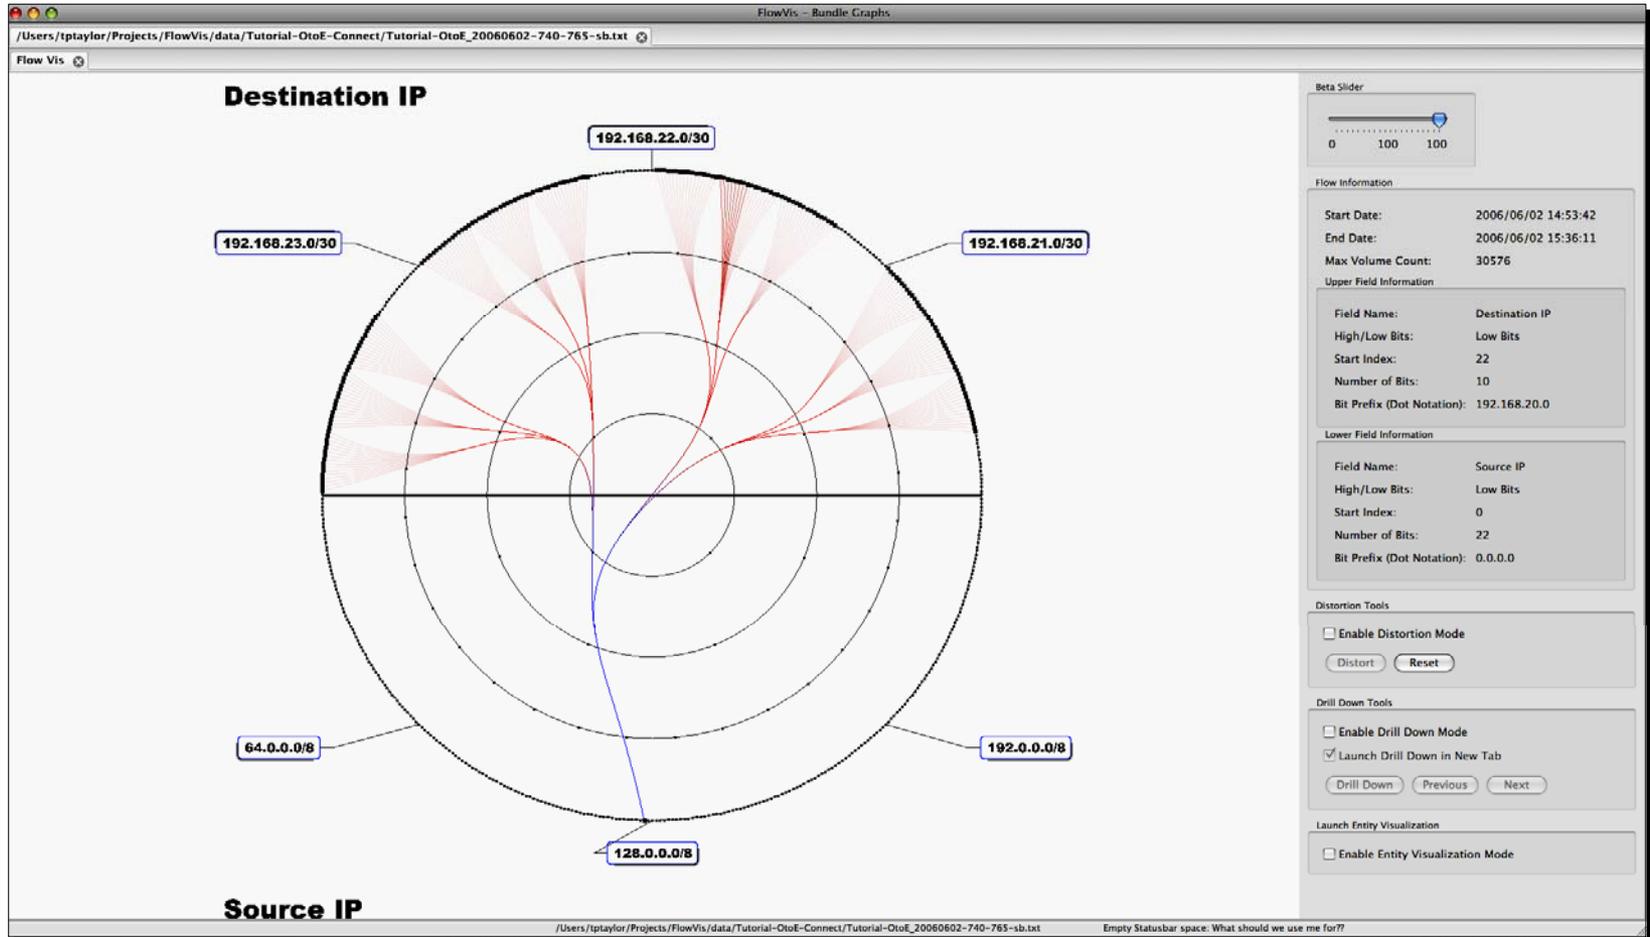


Purpose

- > Visualize the flow of traffic between entities on a network (host-to-host, subnet-to-subnet)
- > Deal with some key issues facing current connection-based visualizations: occlusion, drill down, labeling, etc
- > Incorporate other interactive features and visualizations
- > Introduced by D. Holten in IEEE Transactions on Visualization and Computer Graphics



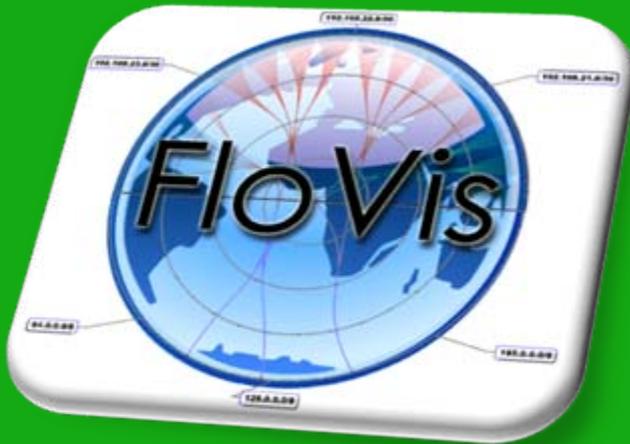
FlowBundle



Data Considerations

- > Takes a SiLK bag indexed by portions of any two scalar fields from NetFlow
- > Total length of scalars must add to 32 bits
 - e.g. Top 16 bits of source address/Lower 16 bits of destination address
- > Working towards creating full 64 bit indexes for full connections
- > Bag counts the number of flows/bytes/packets for the index over a specified time period (hours or days)

Activity Viewer

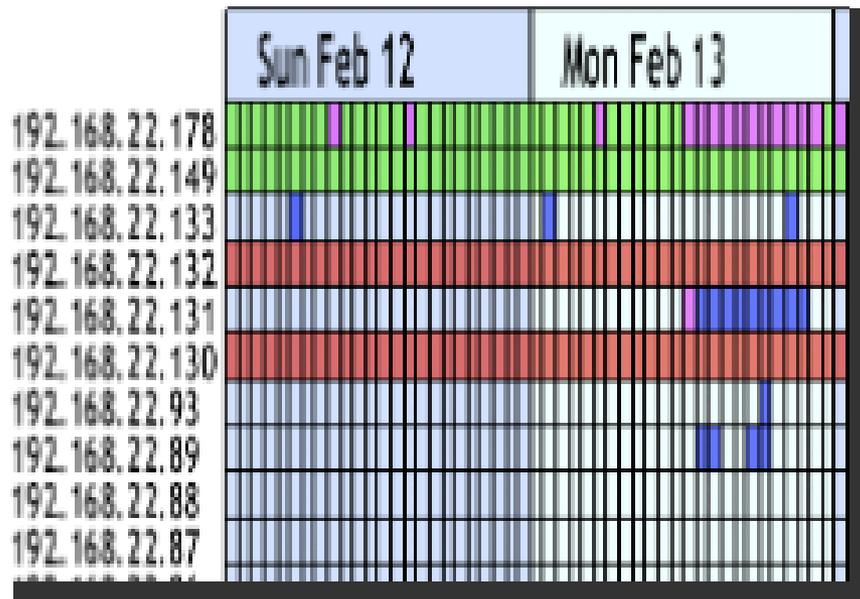


Activity Viewer

- > A visualization for displaying host activity as a function of time.
- > Activity can be host related, time related, simple, complex etc.
- > Individual hosts are plotted against time in a simple two dimensional grid.
- > Categorized hosts as servers / clients / both

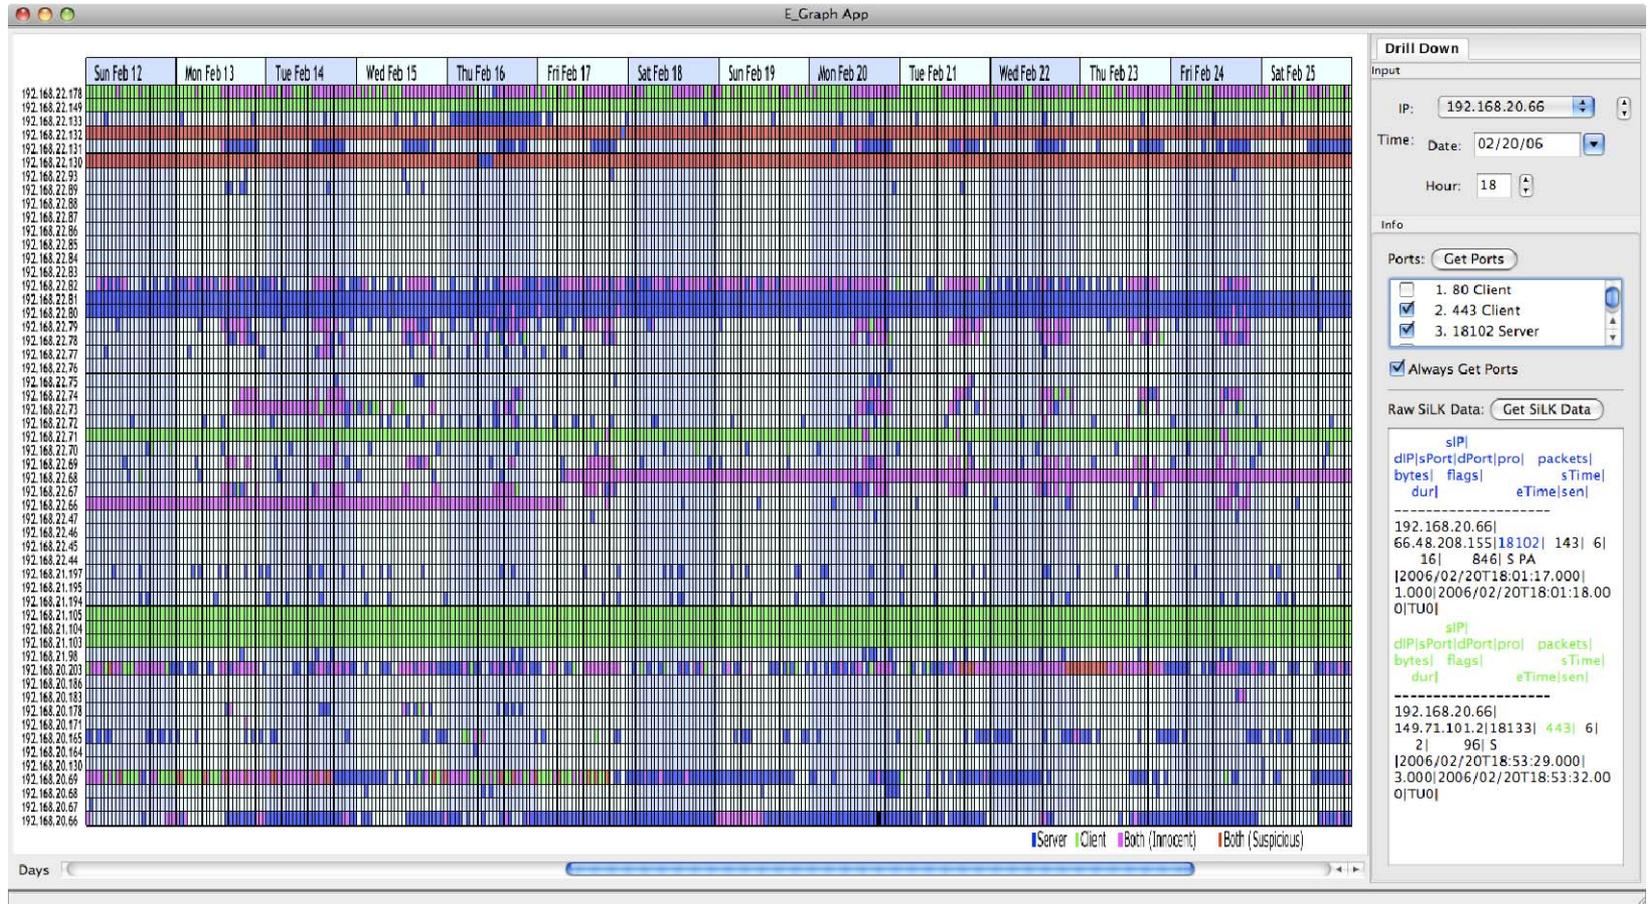
- > Based on work by Phil Groce and Jeff Janies (FloCon 2008)

Visualization Structure

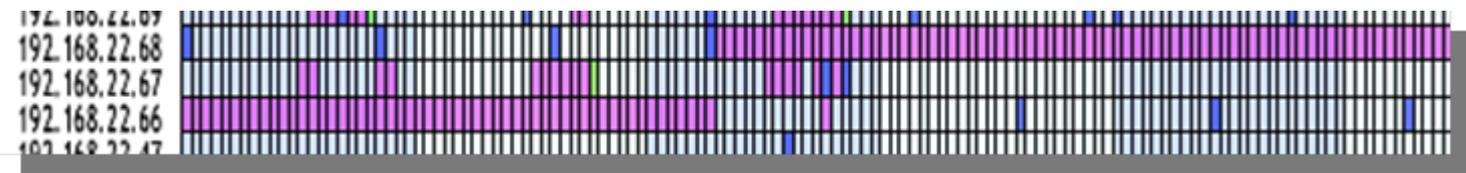


- > Use color to display a small number of activities
- > Grid format used to display activity and non-activity.
- > 14 days of data visible at one time at an hourly resolution.
- > Days of the week used in the time labels
- > Background highlighting used to group hours in a day.

Implementation: Activity Viewer

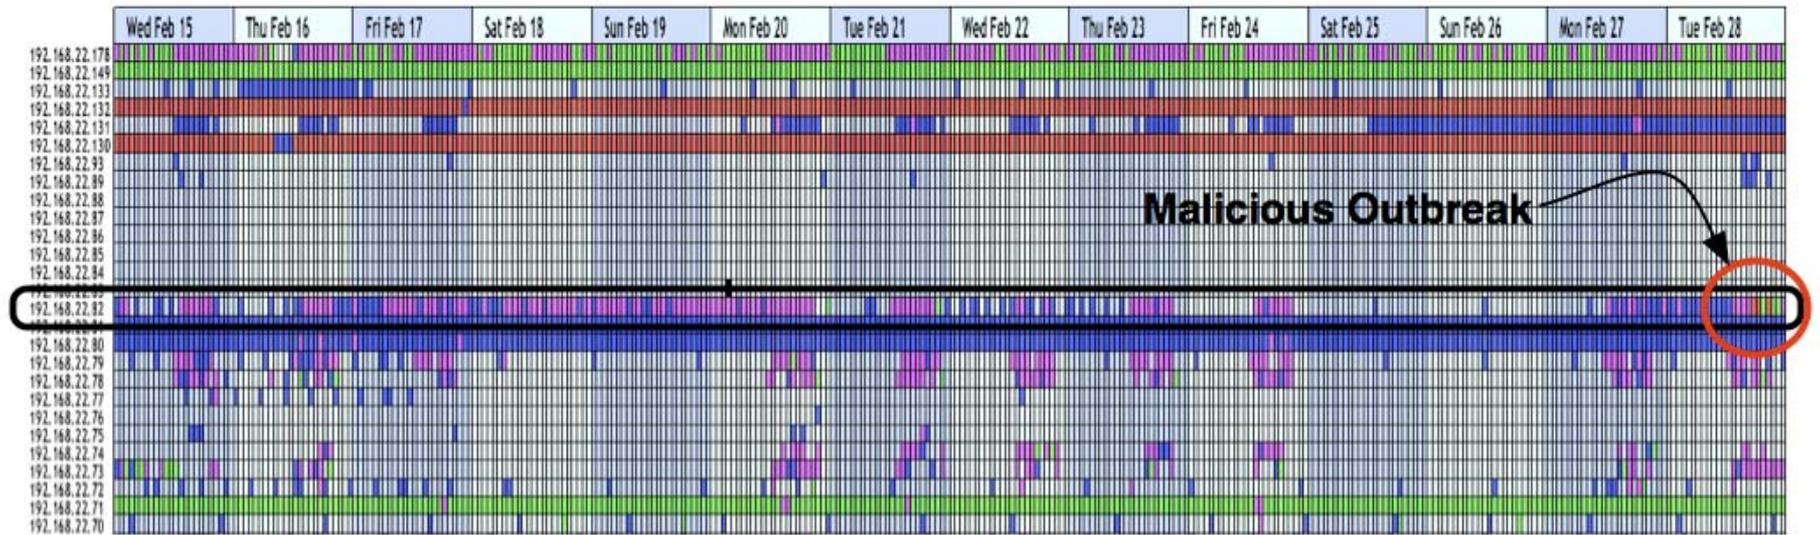


Case Study 1: Lease Switching



- > Host 192.168.22.68 and 192.168.22.66
- > Significant changes in their time based activity patterns of both hosts.
- > Further investigation with the SiLK tools supports the conjecture.

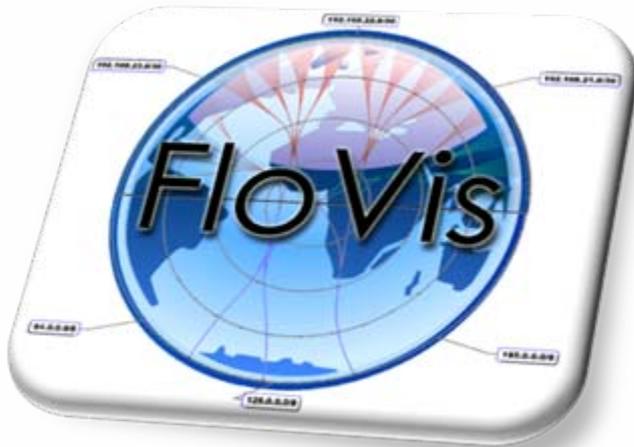
Case Study 2: Malicious Activity



- > February 28, 2007 at 17:00 host 192.168.22.82 started using the same port as a client and a server.



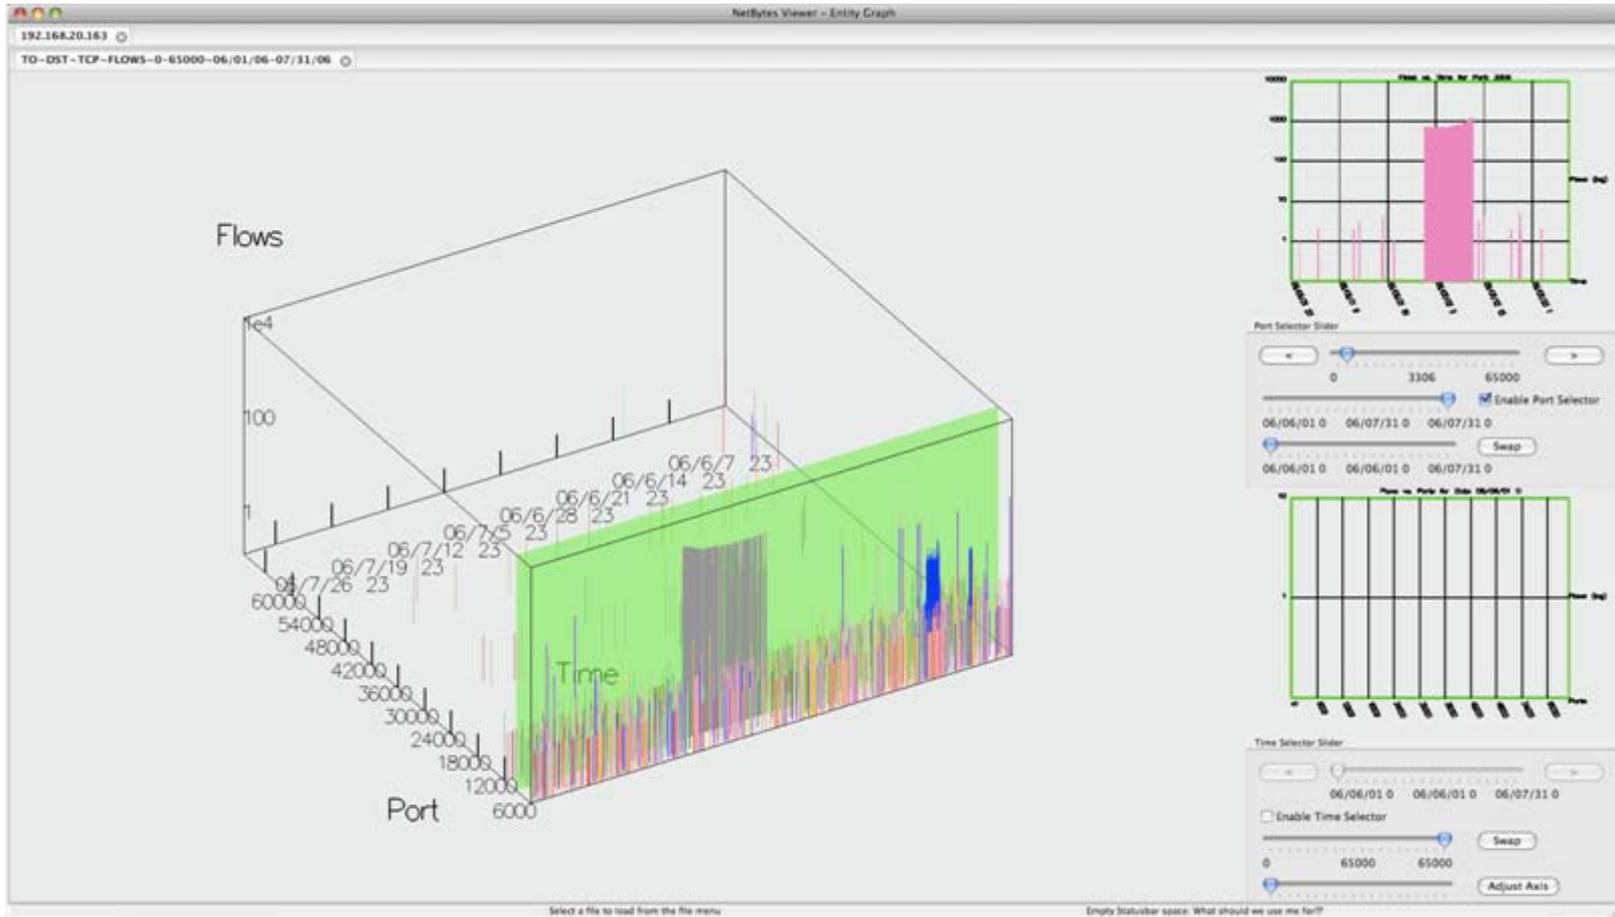
NetBytes Viewer



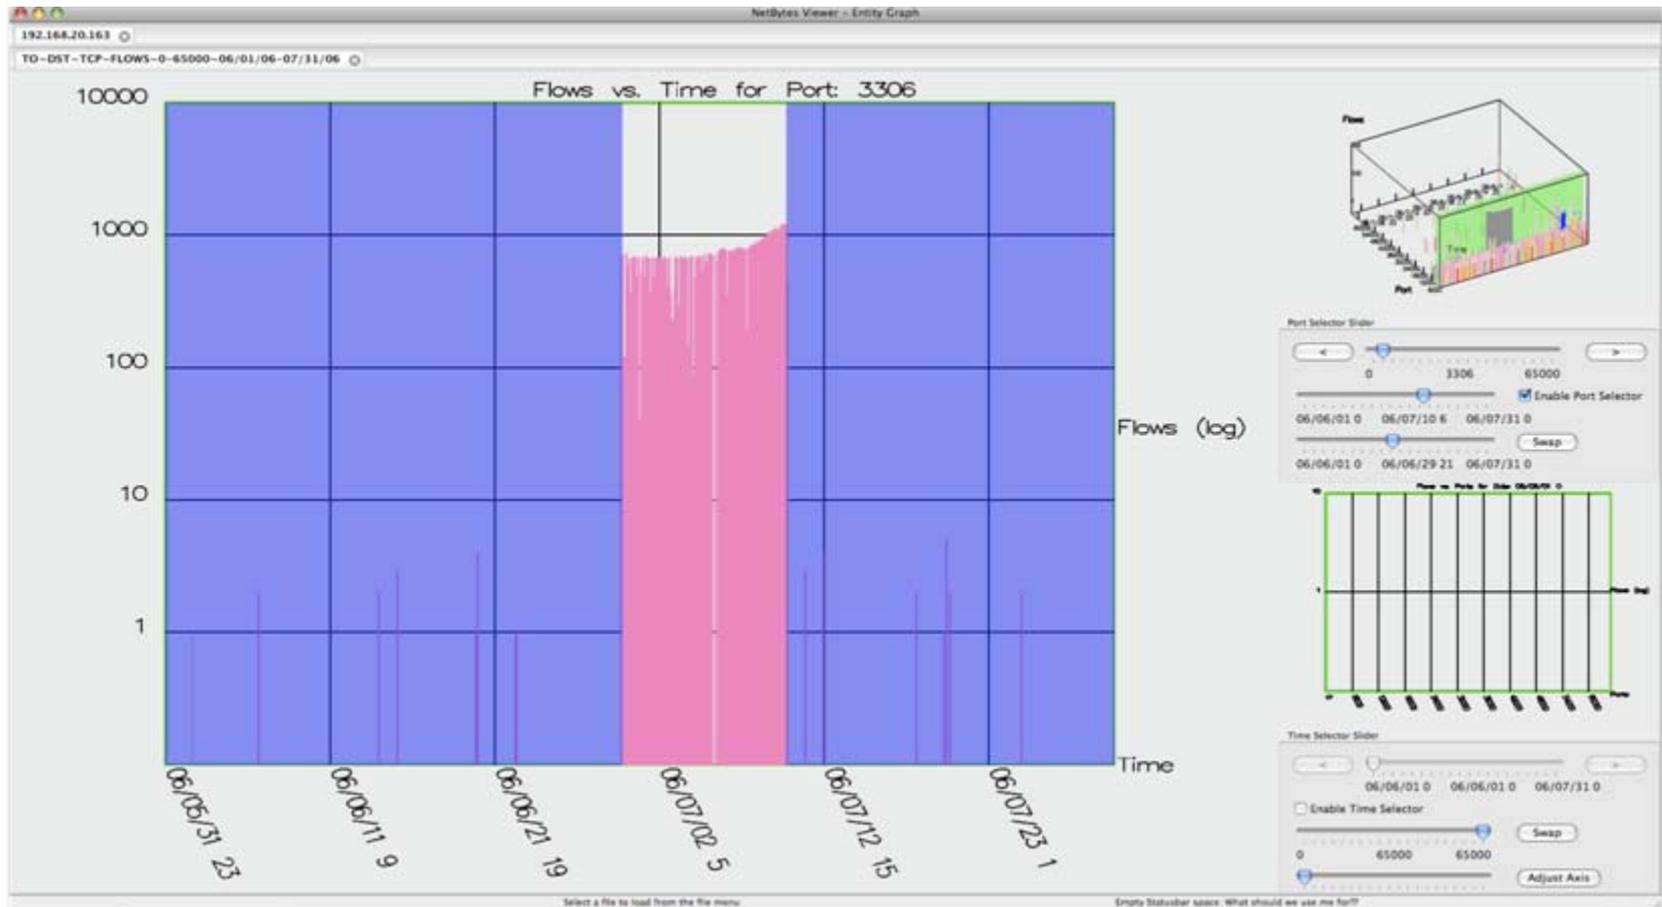
Purpose

- > Interactive visualization.
- > Part of the FloVis framework along with the Activity Viewer and FlowBundle.
- > Visualizes Netflow traffic using an entity-based approach.
- > Focuses on volumes (bytes, flows, packets).
- > Provides a historical context to traffic volume patterns using a 3D graph.

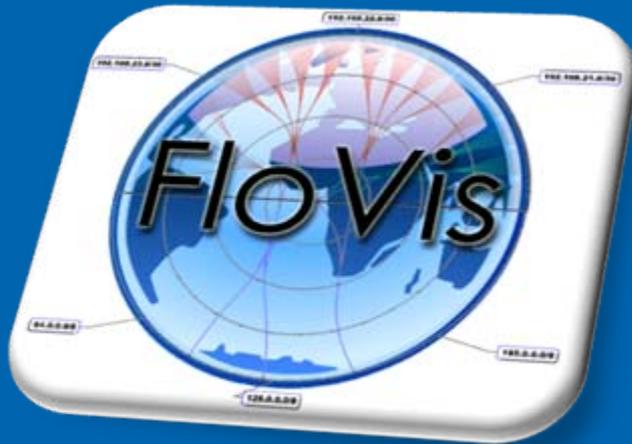
NetBytes Viewer



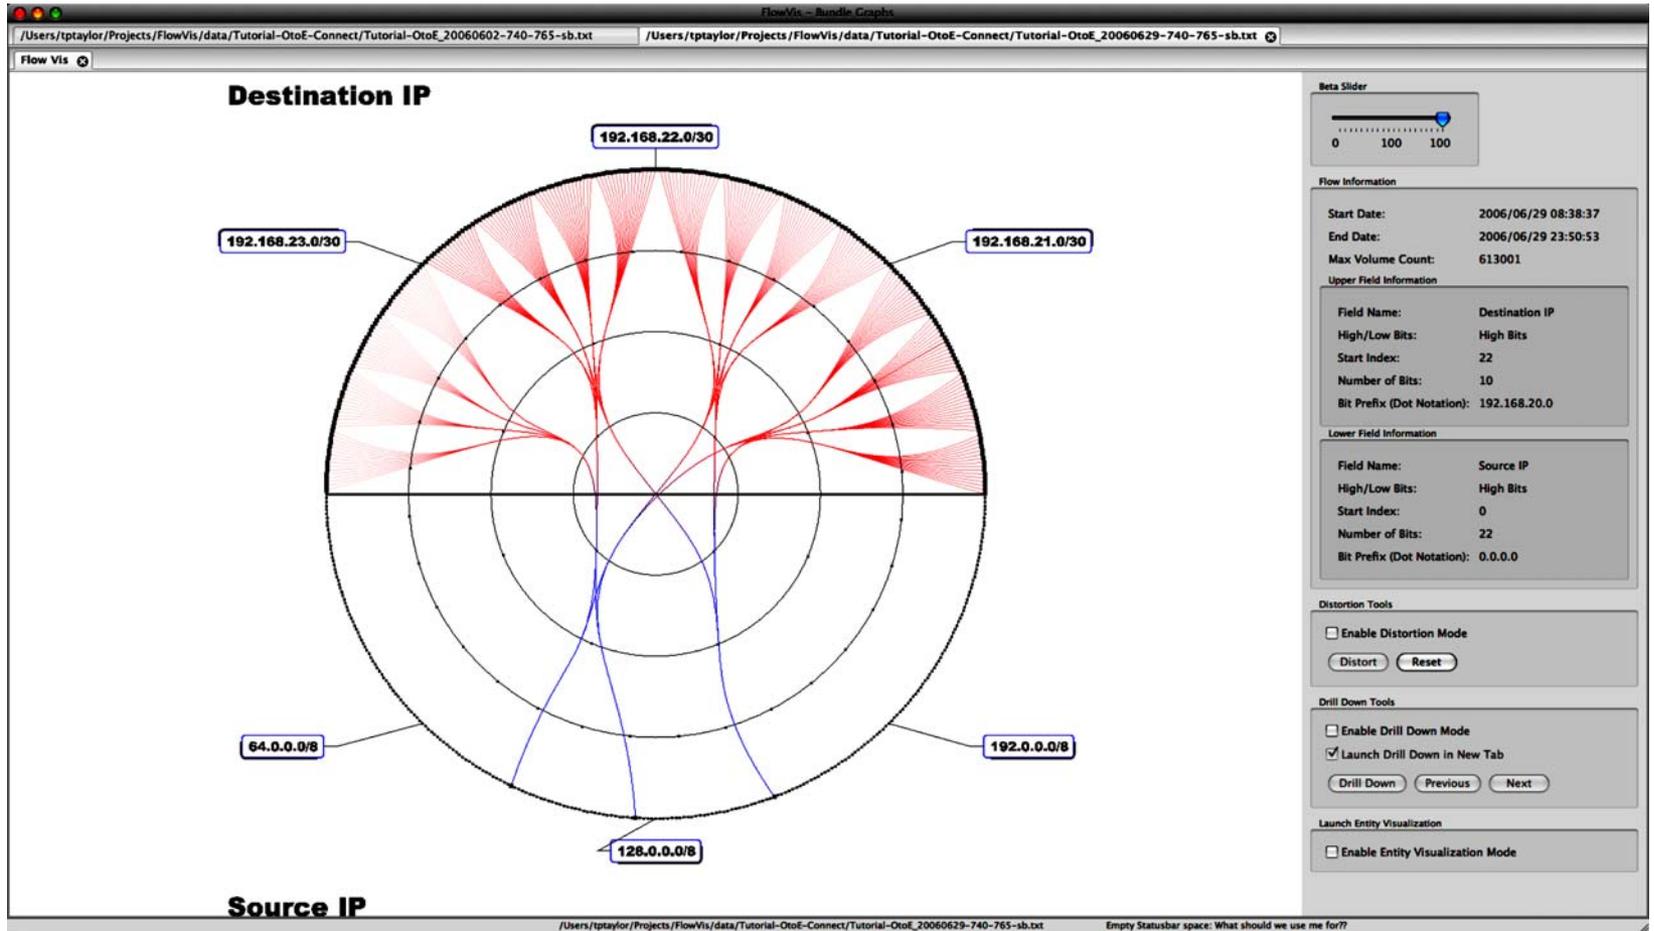
View Swap + Data Selection



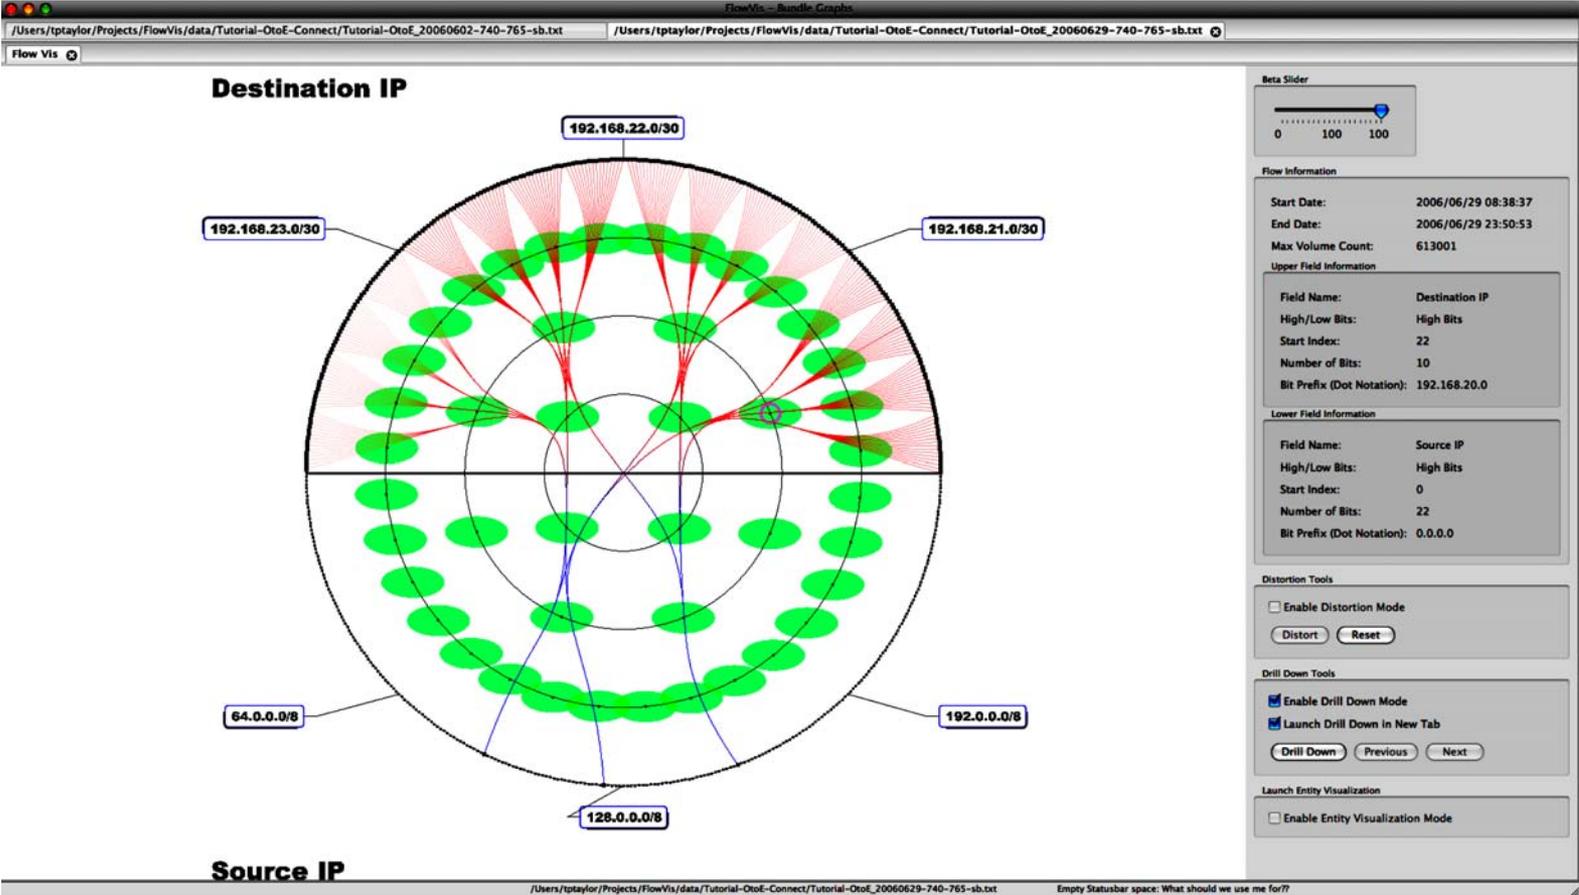
Use Case



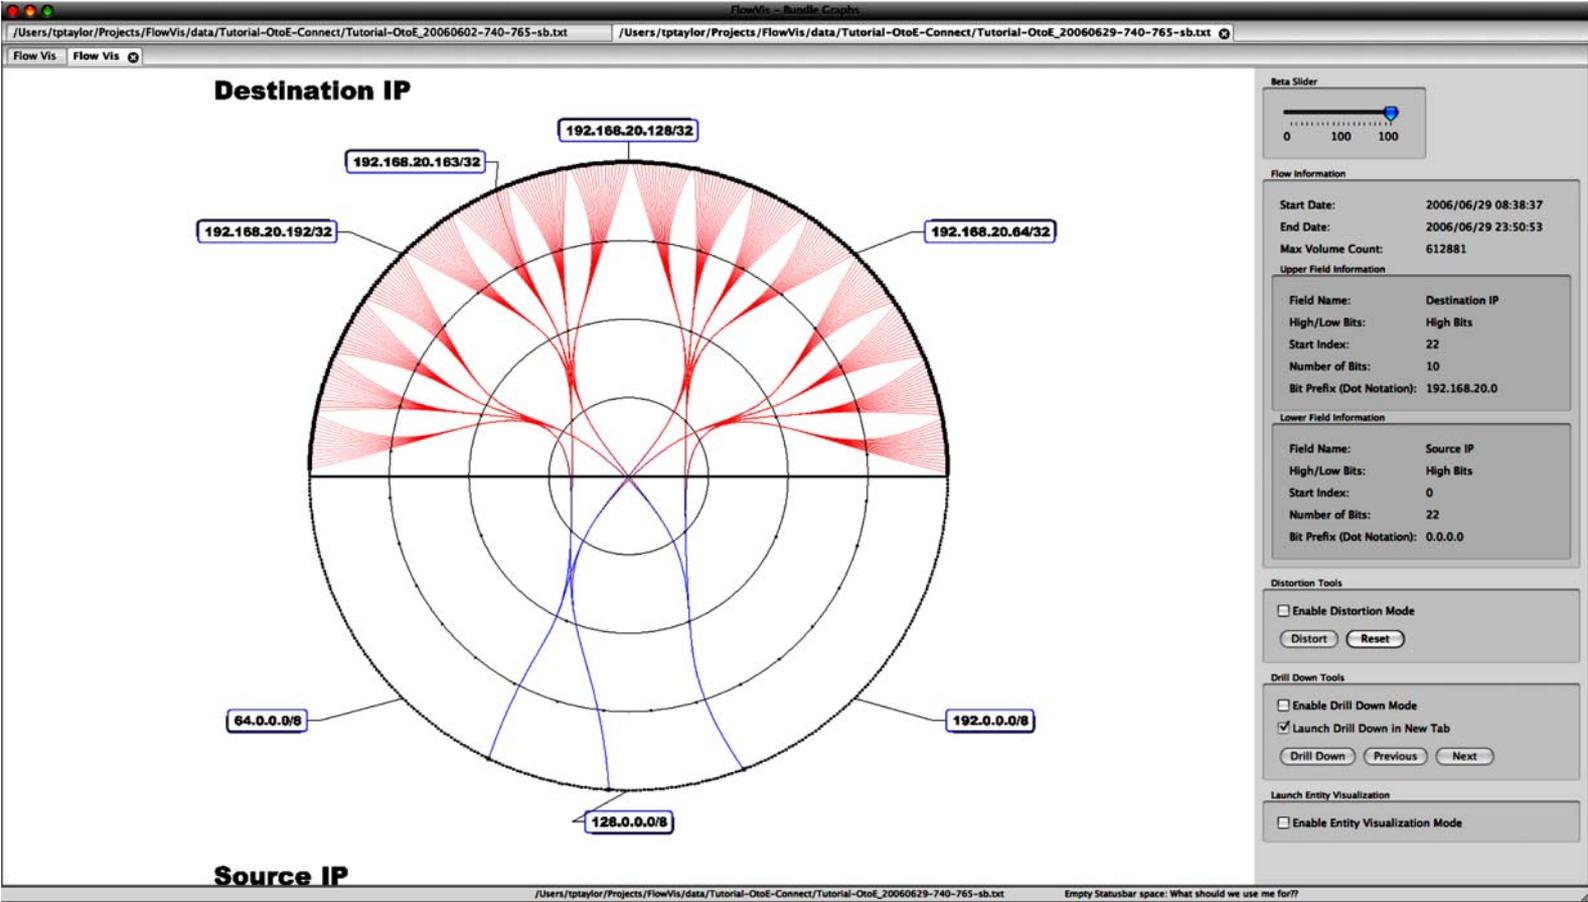
Scanning



Drilling Down



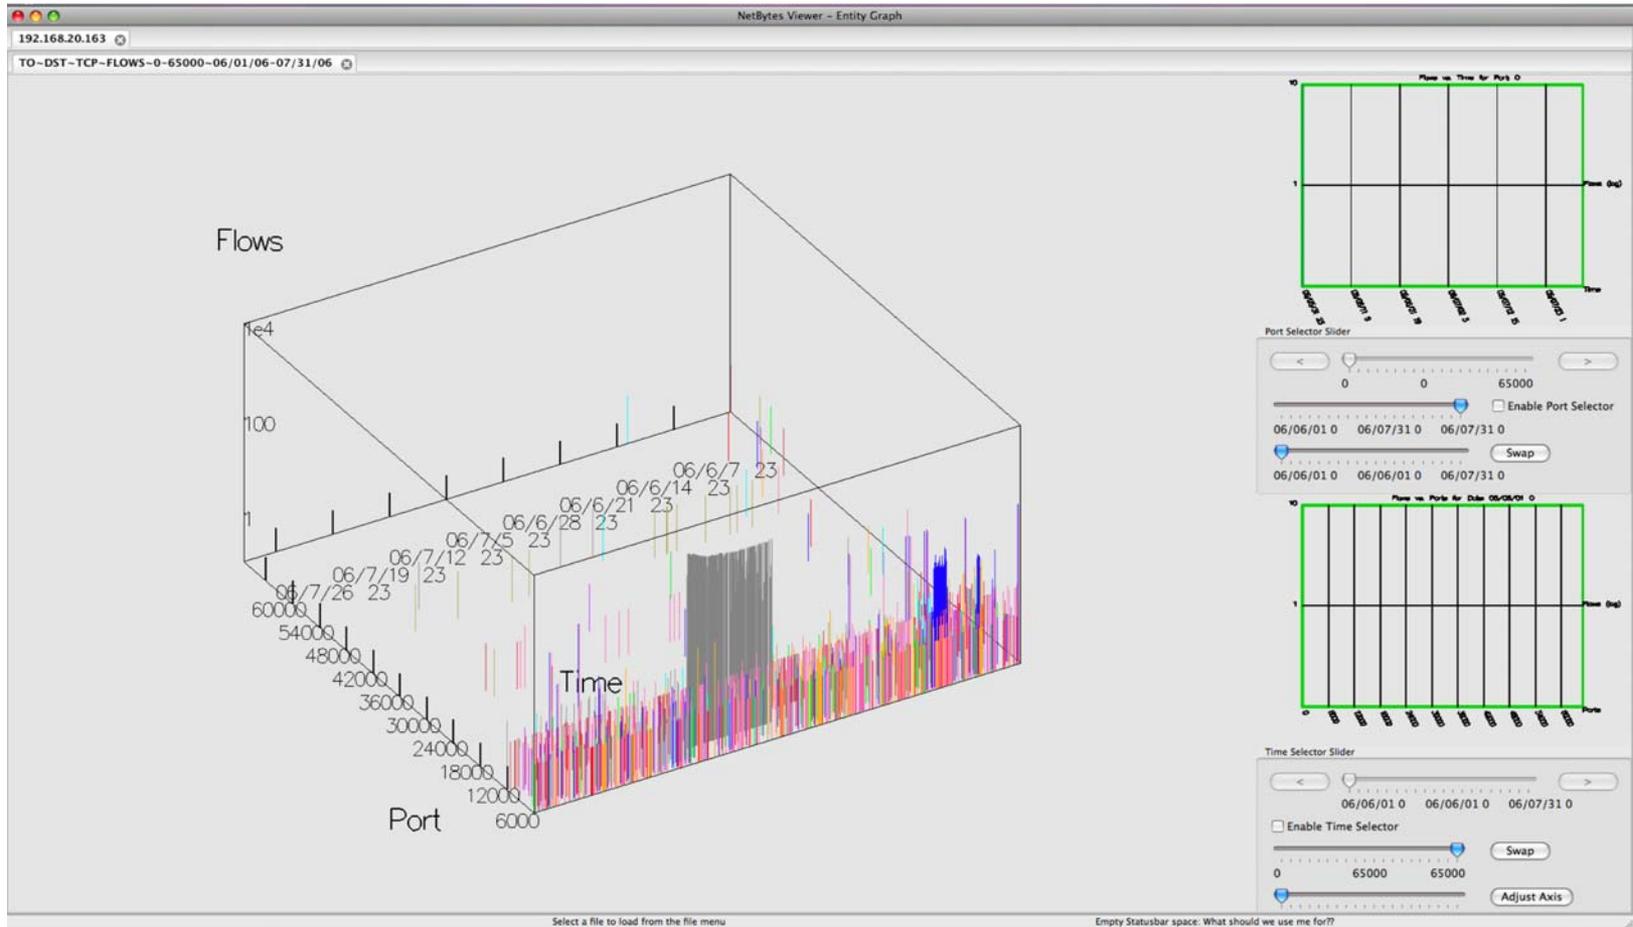
Result



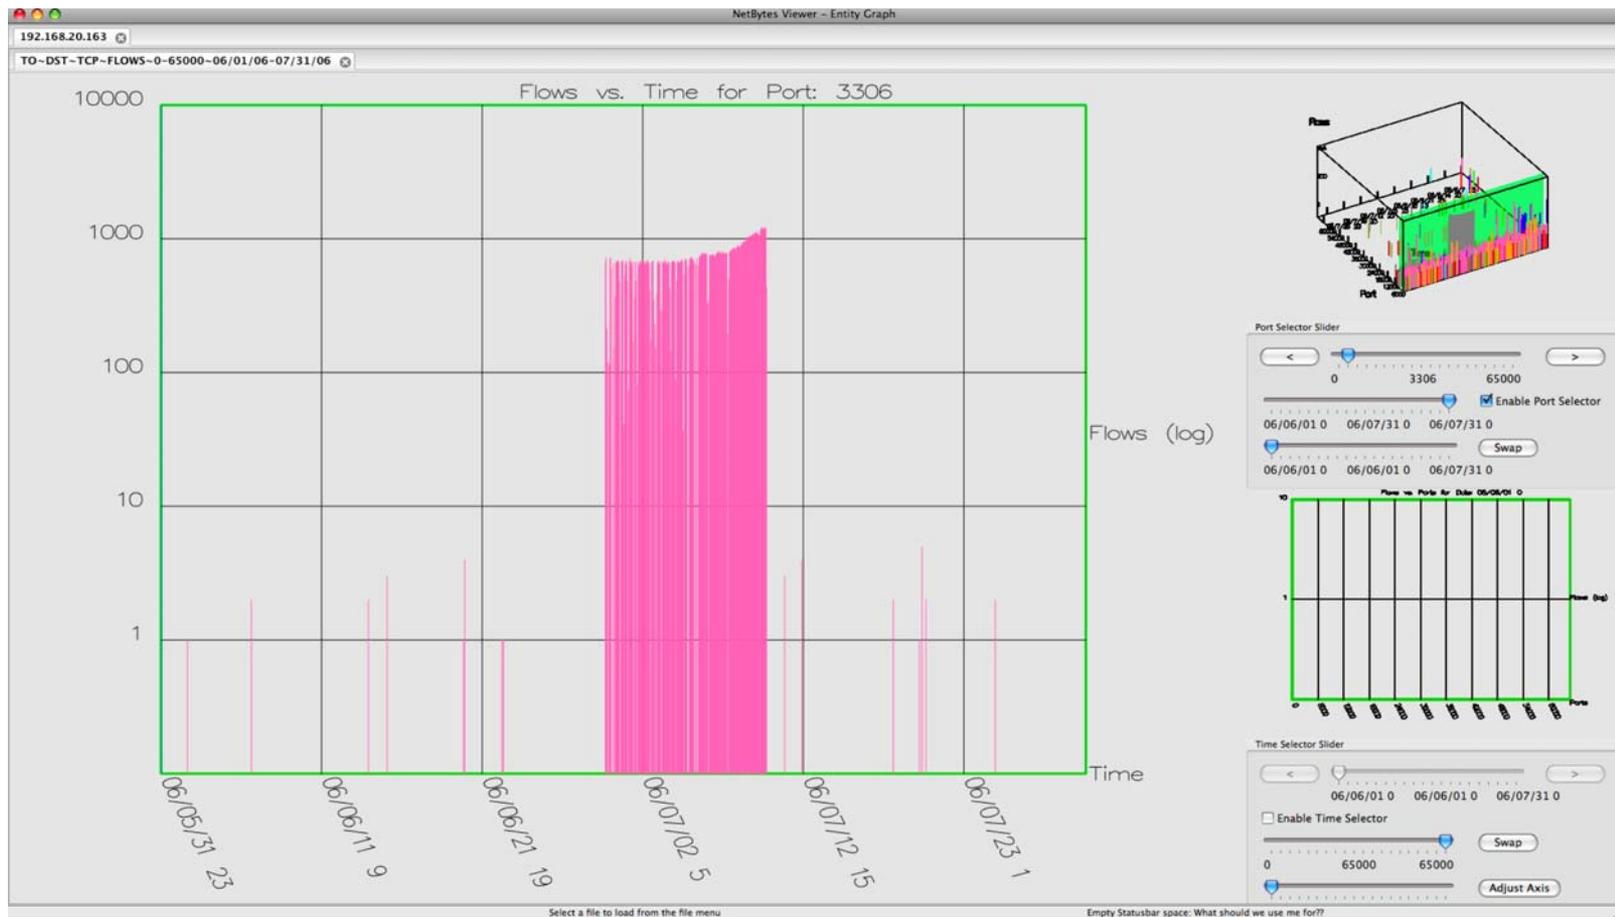
Launch NetBytes Viewer

The screenshot displays the NetBytes Viewer application window. On the left, a 'Configuration Dialog' is open, showing settings for 'Host IP Address' (192.168.20.16), 'Min Port' (0), 'Max Port' (65000), and 'Start Time' (02/02/2006). The 'Volume' section has 'Flows' selected, and the 'Protocol' section has 'TCP' selected. The 'Host' section has 'To' selected, and the 'Port Direction' section has 'Source' selected. The main window shows a circular visualization with 'Destination IP' at the top and 'Source IP' at the bottom. The visualization is divided into segments by concentric circles and radial lines. A red shaded area is visible in the lower half of the circle. Labels for source IP ranges include 64.0.0.0/8, 128.0.0.0/8, and 192.0.0.0/8. Labels for destination IP ranges include 8.20.192/32, 192.168.20.163/32, 192.168.20.128/32, and 192.168.20.64/32. On the right side, there is a 'Beta Slider' and a 'Flow Information' panel. The 'Flow Information' panel shows: Start Date: 2006/06/29 08:38:37, End Date: 2006/06/29 23:50:53, Max Volume Count: 612881. The 'Upper Field Information' panel shows: Field Name: Destination IP, High/Low Bits: High Bits, Start Index: 22, Number of Bits: 10, Bit Prefix (Dot Notation): 192.168.20.0. The 'Lower Field Information' panel shows: Field Name: Source IP, High/Low Bits: High Bits, Start Index: 0, Number of Bits: 22, Bit Prefix (Dot Notation): 0.0.0.0. Below these are 'Distortion Tools' (Enable Distortion Mode, Distort, Reset) and 'Drill Down Tools' (Enable Drill Down Mode, Launch Drill Down in New Tab, Drill Down, Previous, Next). At the bottom, there is a 'Launch Entity Visualization' panel with 'Enable Entity Visualization Mode' checked. The status bar at the bottom contains the path /Users/tpaylor/Projects/FlowVis/data/Tutorial-OtoE-Connect/Tutorial-OtoE_20060629-740-765-sb.txt and the text 'Empty Statusbar space: What should we use me for??'.

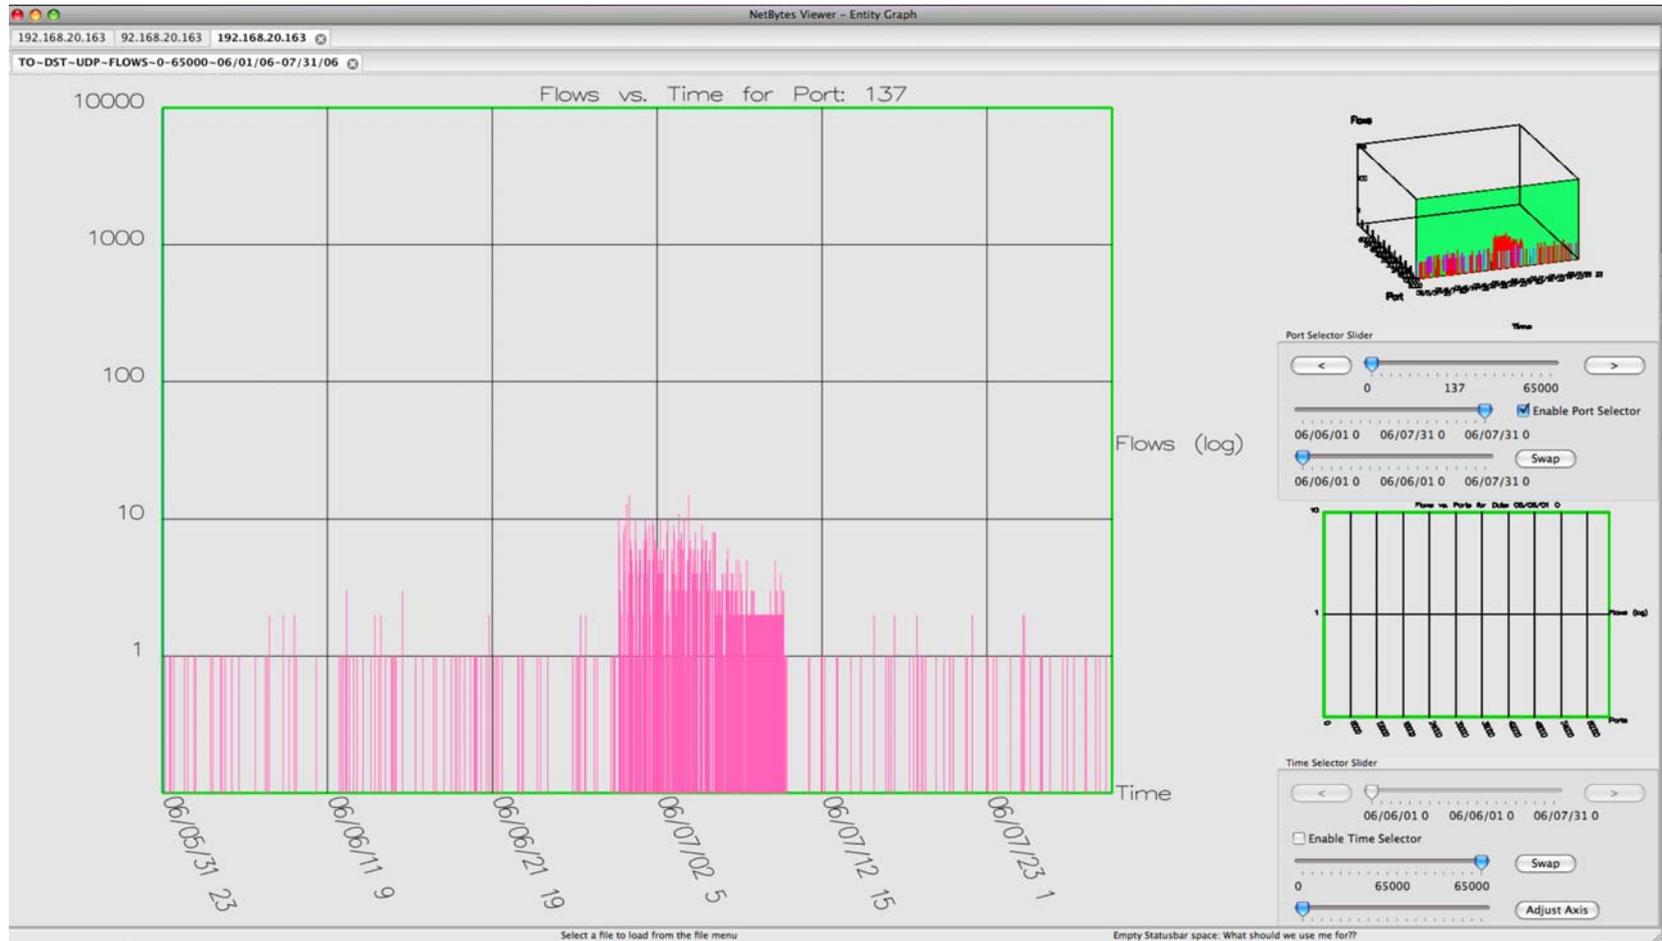
Suspicious Traffic



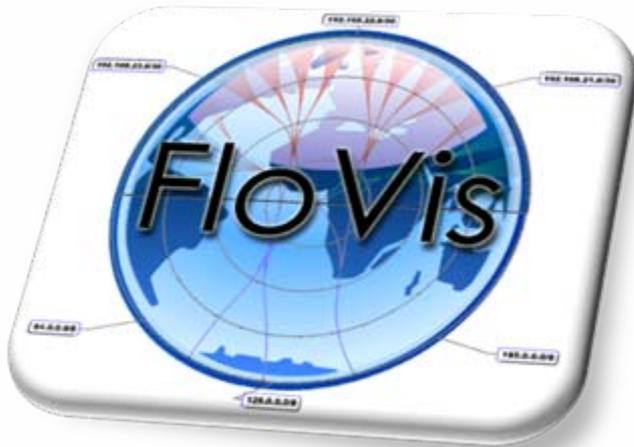
Traffic on Port 3306



Traffic on Port 137



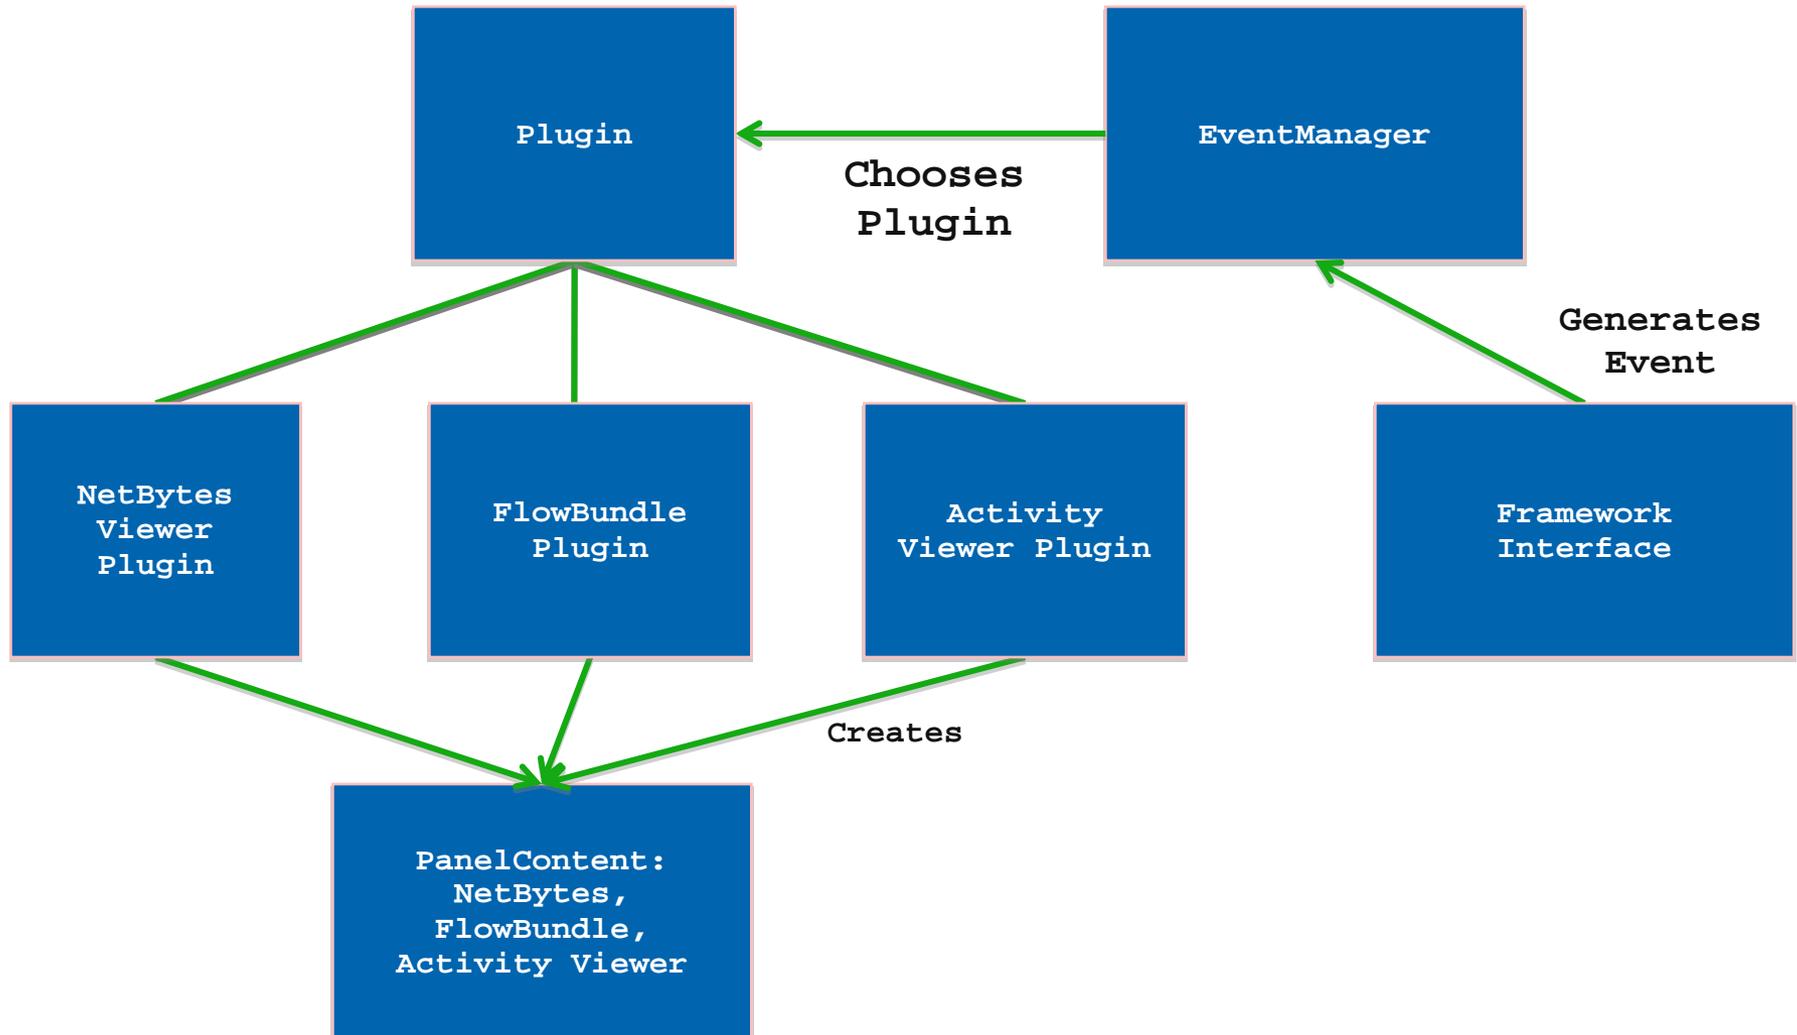
Putting It All Together



Approach

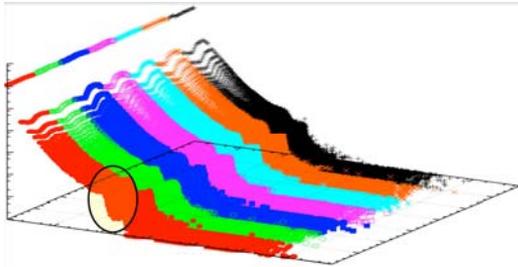
- > Create a visualization suite:
 - to help identify security events
 - to obtain a general understanding of the network
- > Develop a generic framework to support multiple visualizations and the ability to easily add more
- > Investigate and apply new visualization techniques to deal with typical problems of current techniques
- > Close integration (but not dependent) with the SiLK tool suite
 - Visualization approaches can generalize to other network data

FloVis Framework



Benefits of Visualization

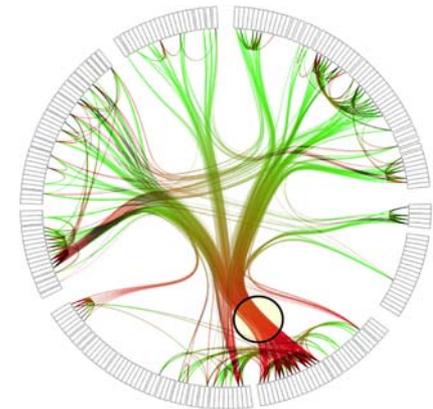
Gain Insight



- > Better understand own network
- > Recognize when something has changed
- > Know when 1 org is different from others

Easily Detect New Patterns / Attacks

- > Take advantage of people's strengths
- > Do not waste people's time



Carrie Gates
carrie.gates@ca.com

