PathScan: Detecting Attackers Within Your Networks

Joshua Neil, Curtis Storlie, Curtis Hash, Alex Brugh, Mike Fisk
Los Alamos National Laboratory
August 15, 2013

What keeps us up at night

- Los Alamos has been connected to the internet since 1984
  - First attack in 1986
- We worry about Advanced Persistent Threat Actors
- Penetrating perimeter defenses
- Compromising core network assets
- Requiring the shutdown of the entire network
  - Perhaps for weeks!
- Financial industry networks should worry about the same thing
  - You are rich targets just like Los Alamos

Hackers are in your networks!

Hackers get into networks

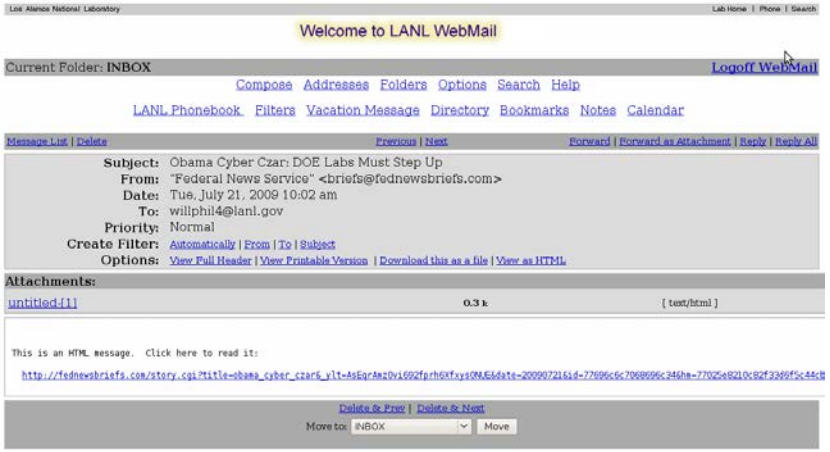

- One way is through phishing
  - 10% of users click (140 clicks/1400 recipients)
  - 15% of management (210/1400)
- There is a need to rapidly detect hackers once they are inside a network
  - Rapid detection is the difference between minor incidents and major compromises

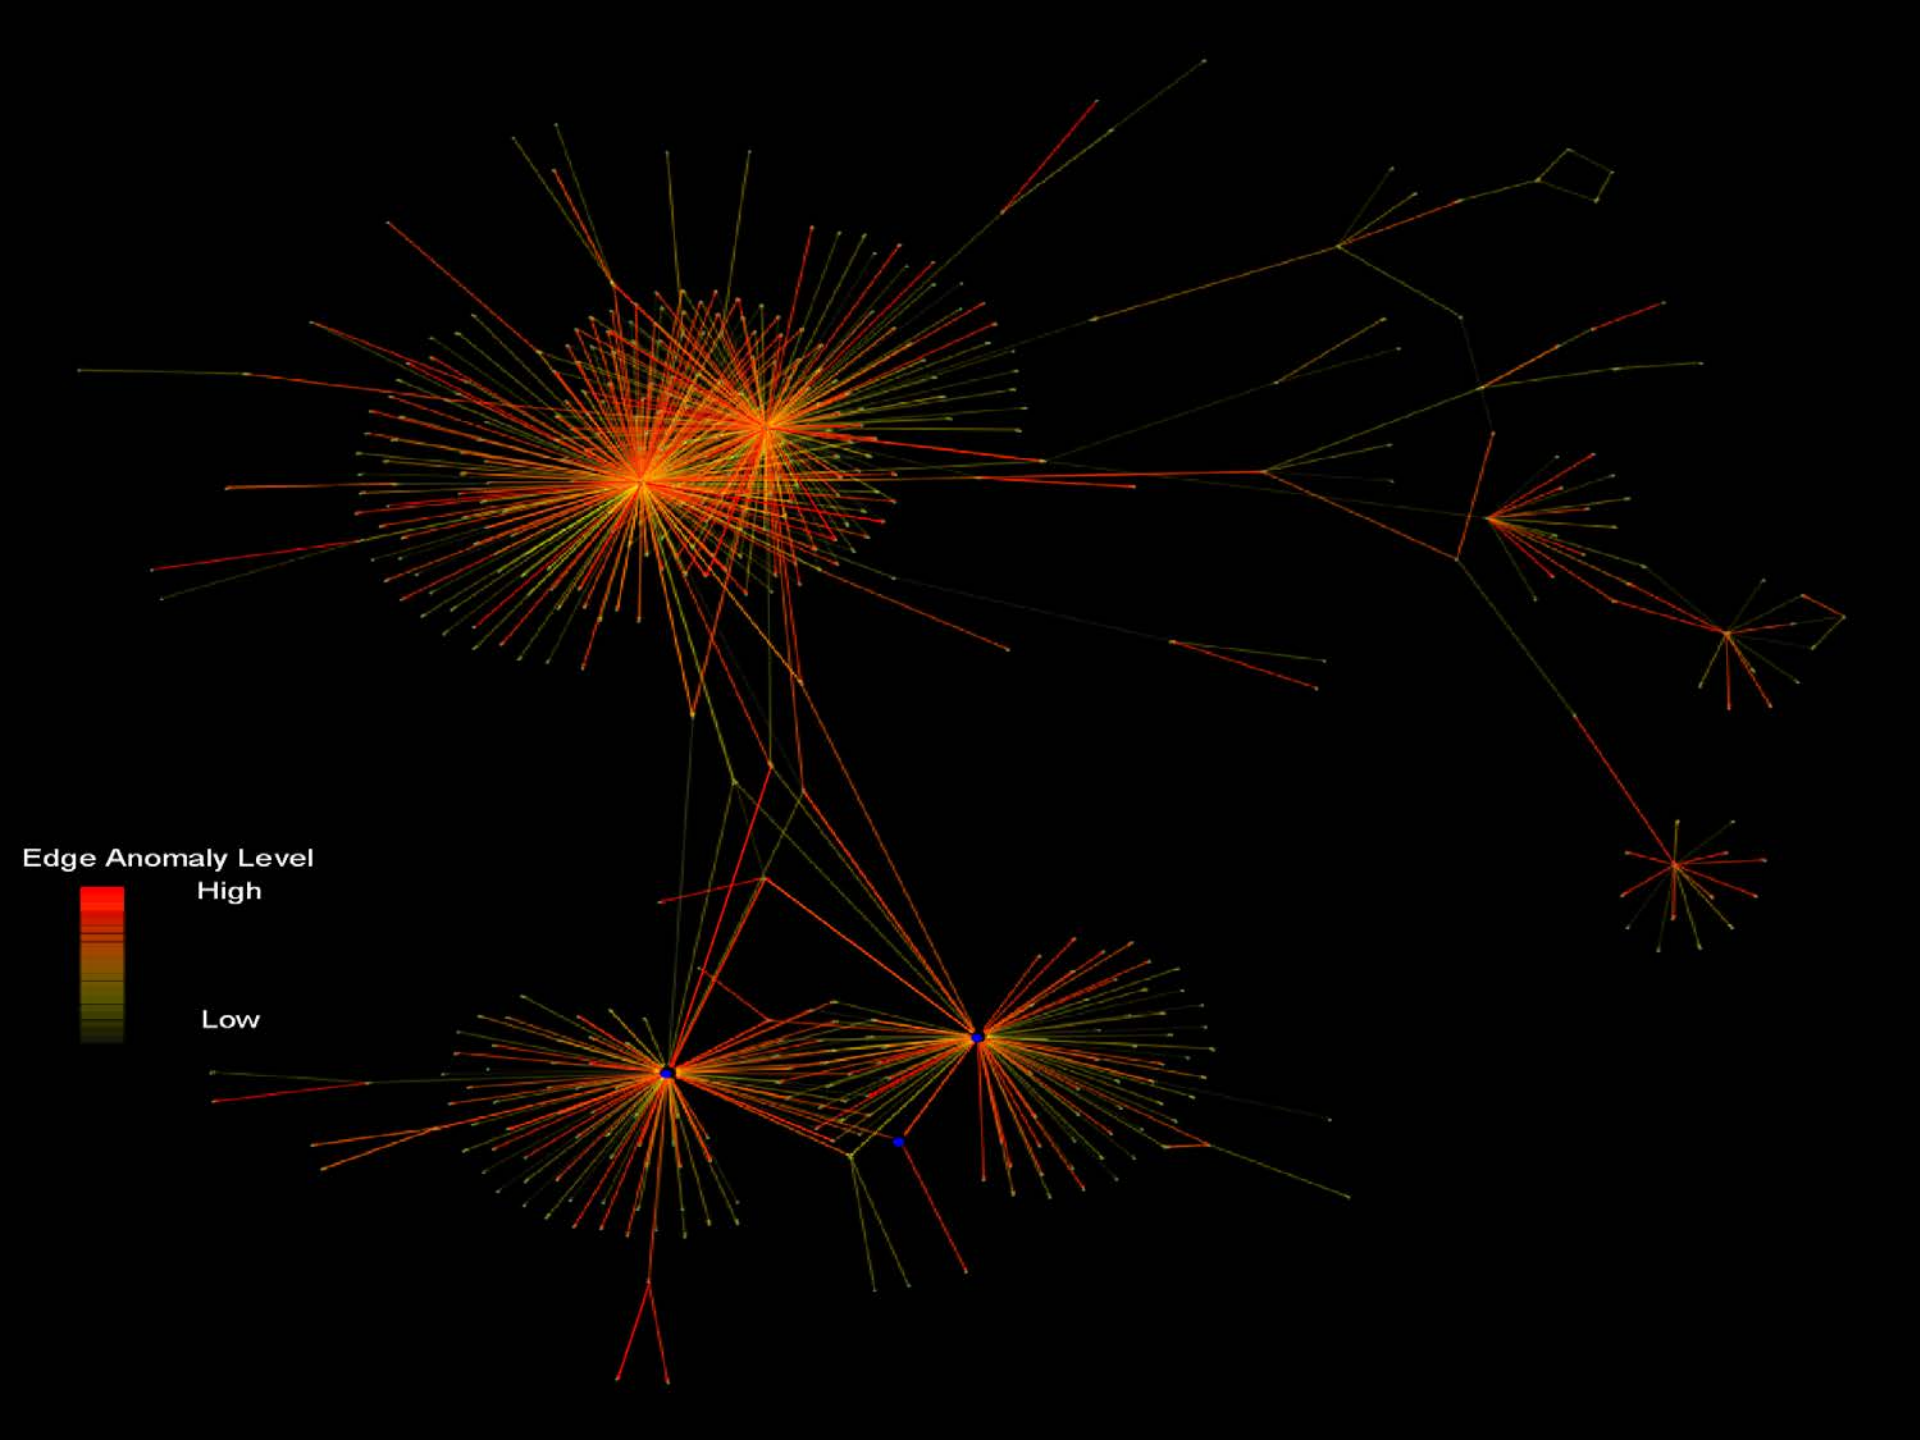
Edge Anomaly Level
High
Low

Approach

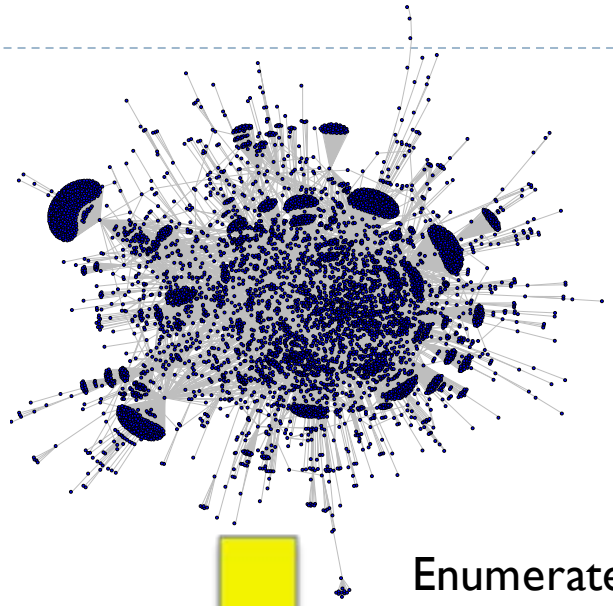

Enumerate All the Paths!

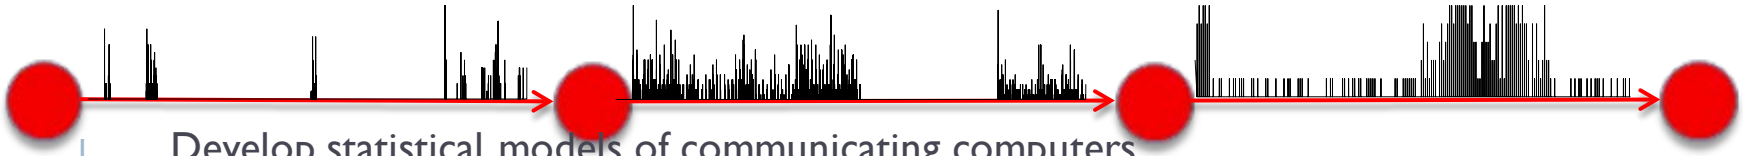

1. Develop statistical models of communicating computers
2. Split the network up into millions of small shapes
3. For each shape ask:

Does the data in this shape look like it comes from our historic model or from some alternative indicating hacking?

Approach (cont'd)

- Tool has two modes
- Online: Currently running online on LANL's network
- Forensic: Able to "flesh out" attacks initially discovered by cyber analysts
  - Has found additional compromised machines in real attacks
- Resources required:
  - With a single 2U server (48 cores) we can keep up with LANLs network
    - 65K computers, > 1M comms
    - Analyzing the network every 5 minutes
  - Needs a network information feed such as DNS or NetFlow data
- Output: Ranked list of IP addresses along with heat maps to allow the analyst to quickly focus on areas of the network

Key Concept

- Simplest approach: For each computer, is it acting weird?
- Better: for each pair of communicating computers are they acting weird together
- Key Concept: For local areas of the network, are all of the pairs of computers in the area acting weird when considered together
- This approach increases statistical power:
  - Reduces false alarms and increases true detection rates

PathScan has no competition

- Existing products look for signatures of previous attacks
  - We use statistical models to identify both previously seen attacks and 0-day behavior
- Existing products only monitor data at the perimeter
  - We use internal data to enable the detection of insider attacks
- Academic approaches mainly propose solutions that don't scale
  - We have scaled to very large (DOD scale) networks

Close
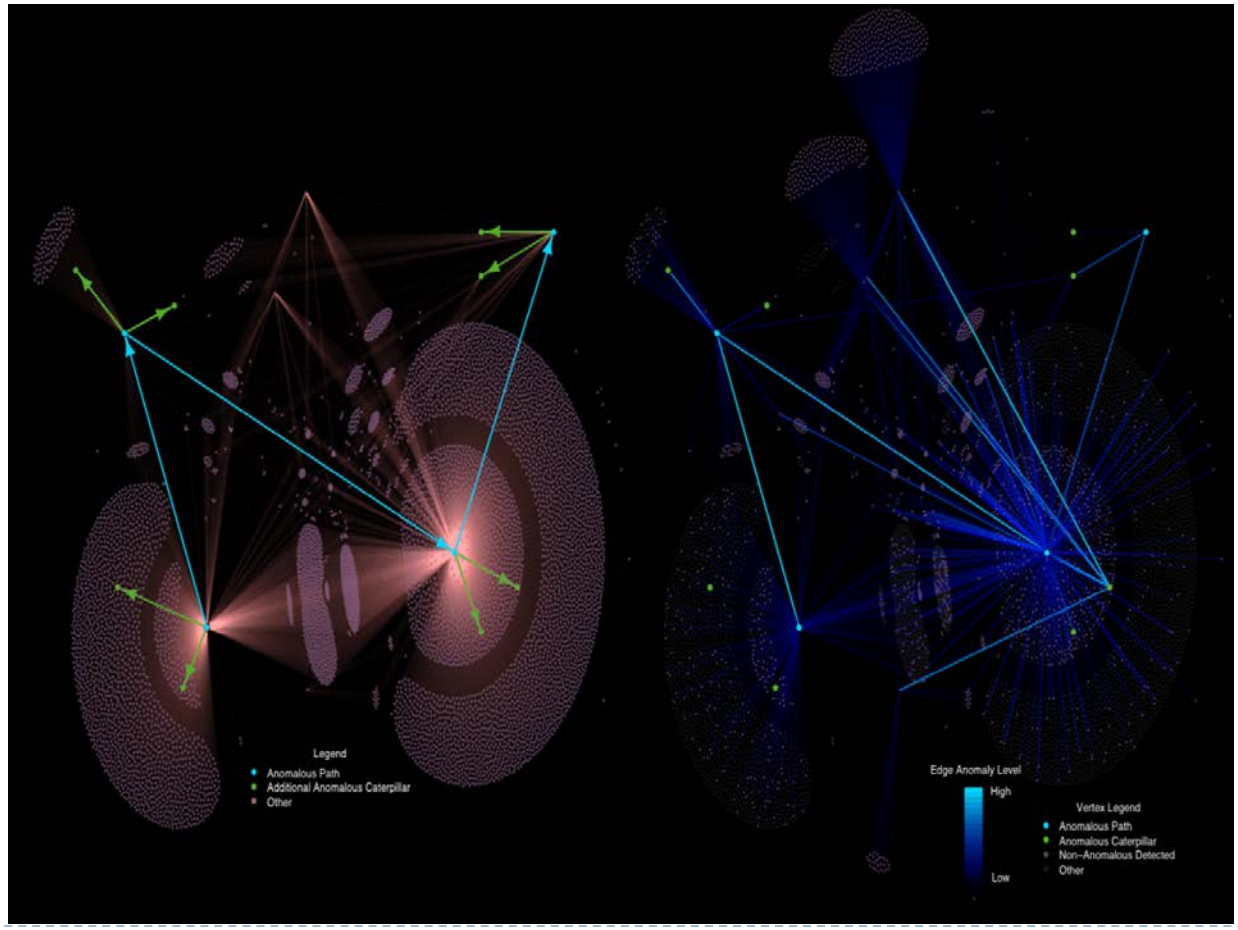
We need partners to help commercialize this work
and pilot networks to continue to validate and further develop the methods