

# Iseb: Trace Driven Modeling of Internet- Scale BGP Attacks and Countermeasures



Systems and Internet  
Infrastructure Security



**CSE**

DETER/EMIST Workshop  
September 28th, 2005 - Newport Beach, CA  
Patrick McDaniel

- Community in the midst of comprehensive effort to evaluate and address the security limitations of BGP, the Internet interdomain routing protocol.
- Where are we?
  - Enormous body of analysis of the behavior of BGP
  - Many, many security protocol proposals
  - Time-tested operational (protective) procedures
  - Community funding, interest, and involvement

**Q:** Why don't we have secure BGP?

**A:** Massively complex protocol interactions involving thousands of independent (*often untrustworthy*) entities

**A:** No complete testbed to evaluate proposals/practices

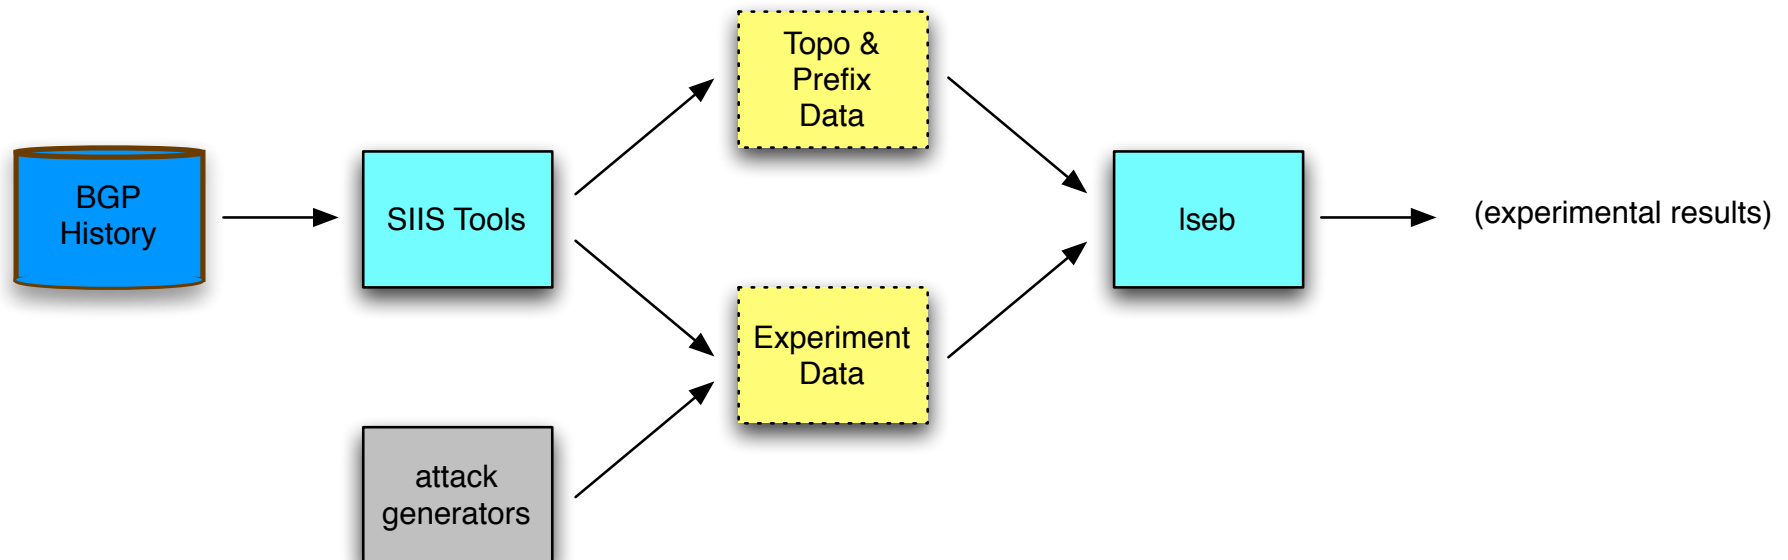
# Evaluation Testbed: an opportunity?

- We have a *global, diverse, detailed* record of BGP behavior over extended period
  - PREDICT, PCH, Routeviews, RIPE, ...
- Q: How can we exploit this unique repository?
  - Analysis
  - Simulation



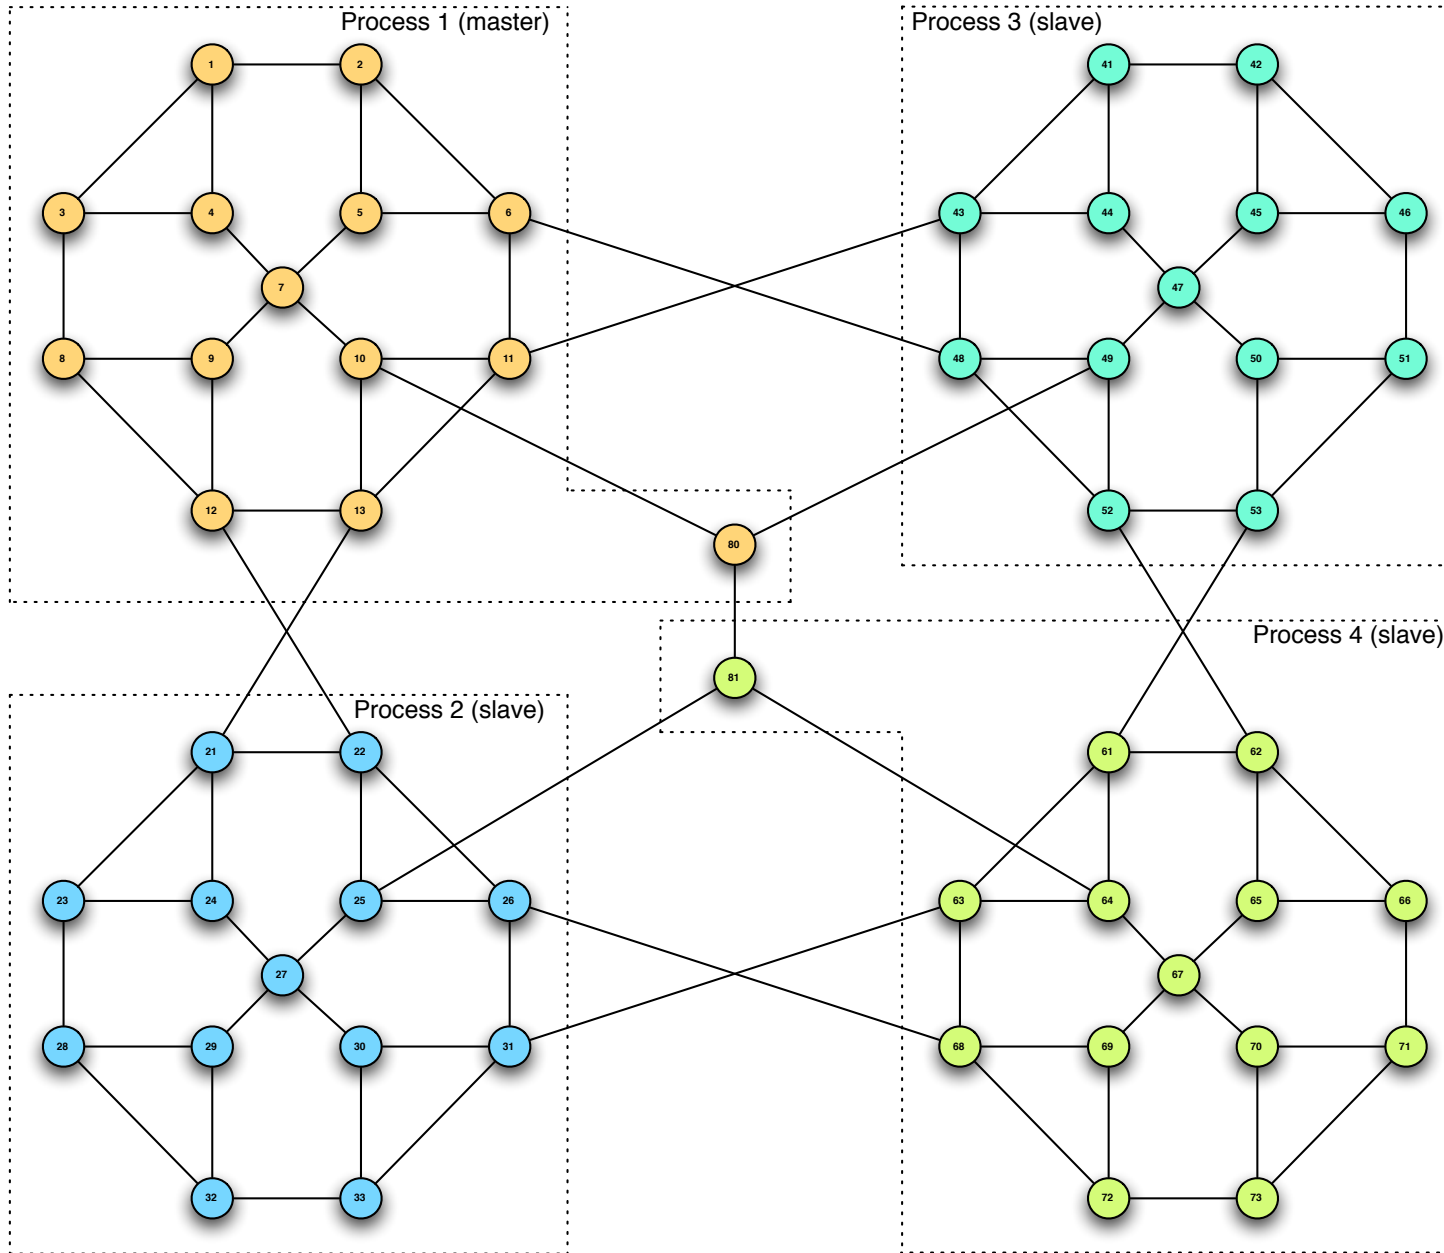
# Iseb: a BGP “way-back” machine?

- *Large-Scale eBGP Simulation* (Iseb)
- Goal: an experimental apparatus to *insert any attack* in *any place* at *any time* in the recorded history of BGP.
- Massive-scale distributed simulation



- *Domain*: understanding global effects of BGP attacks and countermeasures
  - protocol simplifications (e.g., iBGP, local prefs)
- Hypothesis: real traffic traces will preserve realism of simulation BGP behavior
  - Extract meaningful BGP structure and events from history
  - *Advantage*: verifiable via traces
- Simulation of *entire* Internet eBGP protocol
  - e.g., 20,000+ simulated autonomous systems
  - Each AS executes eBGP and maintains BGP table (**166Mb**)
  - Model **every** BGP message on Internet
- Scalability lies in the clustering of simulated ASes

# Simulation Topology



# Simulation Setup

- Single file configuration (process self organization)
  - identify the ASes, links, and topology
  - As gleaned from BGP data (e.g., SIIS-topo-extract)

```
% Demo network Setup
...
4 AS 72
4 AS 73
4 AS 81
% Links
4 PEER 61 62
4 PEER 61 63
4 PEER 61 64
4 PEER 62 65
...
% Prefixes
4 PREFIX 61 61.0.0.0/8
4 PREFIX 73 73.0.0.0/8
...
```

# Scripting the experiment ...

- Single file script delivered to all via master
  - identify the network events, experimental signaling
  - As gleaned from BGP data (future work)

```
% Demo network Experiment
START
SLEEP 5000
DUMP
DROP 1 1.0.0.0/8
...
ADD 62 62.0.0.0/8
...
FAIL 80 81
...
RECOVER 80 81
ADD 1 1.0.0.0/8
...
STOP
```

- A process is started at each host associated with the simulation
  1. The hosts *self-organize* by contacting the master process (this is the control channel)
  2. The master process advertises the mapping of process IDs to IP address/port numbers
  3. The master informs the processes of the setup of the experiment (AS assignments, topology, prefixes)
  4. Each host contacts the peer processes as needed by experimental topology
- The experimental script is executed in order
  - All data is logged in master log (over control channel)
  - Models “real time” BGP via messaging
  - Limited synchronization (*scalability*)

# Experimental Visualization

The screenshot displays a network simulation window titled "bgp-wor - EMIST\_ESVT". The main area shows a network graph with nodes numbered 1 through 73, connected by green lines. A "BGP EVENTS" window is open on the right, showing a list of events. The events are filtered by AS (53) and Event Type (Recover). The events list includes:

- 4033-32 RECV UPDATE 41.0.0.0/8 path=[33 31 63 61 53 52 48 43]
- 4035-28 RECV UPDATE 41.0.0.0/8 path=[32 33 31 63 61 53 52 48]
- 4037-33 RECV UPDATE 41.0.0.0/8 path=[30 31 63 61 53 52 48 43]
- 4039-33 RECV UPDATE 41.0.0.0/8 path=[32 29 27 25 26 68 69 67]
- 4041-29 RECV UPDATE 41.0.0.0/8 path=[32 33 31 63 61 53 52 48]
- 4043-29 RECV UPDATE 41.0.0.0/8 path=[28 32 33 31 63 61 53 52]
- 4044-23 RECV UPDATE 41.0.0.0/8 path=[28 32 33 31 63 61 53 52]
- 4047-33 RECV UPDATE 41.0.0.0/8 path=[32 29 27 30 31 63 61 53]
- 4050-21 RECV UPDATE 61.0.0.0/8 path=[13 10 80 81 64 61 ]
- 4052-22 RECV UPDATE 61.0.0.0/8 path=[12 13 10 80 81 64 61 ]
- 4054-21 RECV UPDATE 73.0.0.0/8 path=[13 10 80 81 64 67 69 72]
- 4056-22 RECV UPDATE 73.0.0.0/8 path=[12 13 10 80 81 64 67 69]
- 4058-21 RECV UPDATE 73.0.0.0/8 path=[13 10 80 81 64 67 70 73]
- 4060-22 RECV UPDATE 73.0.0.0/8 path=[12 13 10 80 81 64 67 70]
- 4062-21 RECV UPDATE 53.0.0.0/8 path=[13 10 80 81 64 61 53 ]
- 4064-22 RECV UPDATE 53.0.0.0/8 path=[12 13 10 80 81 64 61 53]
- 4066-21 RECV UPDATE 41.0.0.0/8 path=[13 10 80 81 64 61 53 52]
- 4068-22 RECV UPDATE 41.0.0.0/8 path=[12 13 10 80 81 64 61 53]
- 4070-21 RECV UPDATE 33.0.0.0/8 path=[13 10 80 81 64 67 69 68]
- 4072-22 RECV UPDATE 33.0.0.0/8 path=[12 9 7 10 80 81 64 67 69]
- 4074-22 RECV UPDATE 33.0.0.0/8 path=[12 13 10 80 81 64 67 69]
- 4076-22 RECV UPDATE 33.0.0.0/8 path=[12 9 7 10 80 81 64 63 31]
- 4078-21 RECV UPDATE 33.0.0.0/8 path=[13 10 80 81 64 63 31 33]
- 4080-22 RECV UPDATE 33.0.0.0/8 path=[12 13 10 80 81 64 63 31]
- 4082-21 RECV UPDATE 41.0.0.0/8 path=[13 11 43 41 ]
- 4084-24 RECV UPDATE 41.0.0.0/8 path=[21 13 11 43 41 ]
- 4086-23 RECV UPDATE 41.0.0.0/8 path=[21 13 11 43 41 ]
- 4088-23 RECV UPDATE 41.0.0.0/8 path=[24 21 13 11 43 41 ]

# Experimental Visualization (cont.)

The screenshot displays a network simulation window titled "bgp-wor - EMIST\_ESVT". The main area shows a network graph with nodes numbered 1 through 53, connected by green lines. Two dialog boxes are open over the graph:

- AS Information**: Shows AS 53. The AS Prefix List contains "53.0.0.0/8". The AS Routing Table lists various prefixes and their paths, including:
  - Prefix 62.0.0.0/8 \*\*\* PATHS \*\*\*: [52 62 ], length 2; [61 62 ], length 2
  - Prefix 53.0.0.0/8 \*\*\* PATHS \*\*\*: [ ], length 1
  - Prefix 33.0.0.0/8 \*\*\* PATHS \*\*\*: [61 63 31 33 ], length 4
  - Prefix 1.0.0.0/8 \*\*\* PATHS \*\*\*: [50 47 49 48 6 11 10 7 4 1 ], length 10; [52 49 48 6 11 10 7 4 1 ], length 9; [51 50 47 49 48 6 11 10 7 4 1 ], length 11; [61 62 52 49 48 6 11 10 7 4 1 ], length 11
  - Prefix 13.0.0.0/8 \*\*\* PATHS \*\*\*: [61 63 68 26 22 21 13 ], length 7; [52 48 6 11 13 ], length 5; [50 47 44 43 11 13 ], length 6; [61 46 42 41 43 11 13 ], length 7
- BGP EVENTS**: Shows a list of events for AS 53, including RECV UPDATE and RECV WITHDRAW messages with their respective paths. The log ends with "25065-AS 53 Dumping Finished".

At the bottom of the window, the status bar shows "Ready", "Steps: 0", and "Net Package: 0".

# Experimental Visualization (cont.)

The screenshot displays a network simulation window titled "bgp-wor - EMIST\_ESVT". The main area shows a network graph with nodes numbered 1 through 73, connected by green lines. A red line highlights a specific path between nodes 80 and 81. An "BGP EVENTS" window is open on the right, showing a list of events. The events include "Dumping Finished" for various ASes and "RECV WITHDRAW" messages for specific IP addresses. The status bar at the bottom indicates "Ready", "Steps: 0", and "Net Package: 0".

**BGP EVENTS**

Close Continue Events:  AS  Event Type  All

0 Fail Recover Apply

23011-AS 53 Dumping Finished  
23050-AS 52 Dumping Finished  
23090-AS 50 Dumping Finished  
23130-AS 49 Dumping Finished  
23177-AS 47 Dumping Finished  
23222-AS 48 Dumping Finished  
23259-AS 46 Dumping Finished  
23300-AS 44 Dumping Finished  
23343-AS 45 Dumping Finished  
23381-AS 51 Dumping Finished  
23421-AS 32 Dumping Finished  
23459-AS 31 Dumping Finished  
23498-AS 30 Dumping Finished  
23537-AS 24 Dumping Finished  
23582-AS 22 Dumping Finished  
23621-AS 27 Dumping Finished  
23665-AS 25 Dumping Finished  
23709-AS 21 Dumping Finished  
23752-AS 29 Dumping Finished  
23790-AS 23 Dumping Finished  
23828-AS 33 Dumping Finished  
23867-AS 28 Dumping Finished  
23910-AS 26 Dumping Finished  
**23918-Processing : FAIL 80 81**  
23920-10 RECV WITHDRAW 62.0.0.0/8 path=[80 81 64 61 62 ]  
23922-10 RECV WITHDRAW 53.0.0.0/8 path=[80 81 64 61 53 ]  
23924-10 RECV WITHDRAW 33.0.0.0/8 path=[80 81 64 63 31 33 ]  
23926-10 RECV WITHDRAW 61.0.0.0/8 path=[80 81 64 61 ]

Ready Steps: 0 Net Package: 0

- Working simulator
  - Large Java-based implementation
  - Multiprocessor
  - Self organization
  - Complete recreation of BGP messaging
  - Forensic visualization (thanks Peng/Lunquan)
  - Preliminary version run on DETER (thanks Ihab)
- Working experiments
  - 54 AS topology executing on single processor and 4 processor setup
  - Executed experiments in real time (scaling?)

- *Realism*
  - Extension of path selection algorithm (policy extraction)
  - Delay modeling (trace analysis, PREDICT)
  - PSU: Kevin Butler starting this year
- *Attack Modeling*
  - Extending/using experimental lexicon to model real attacks (Sandy Murphy -- SPARTA)
  - Tools integration
- *Analysis*
  - Building measurement apparatus
  - Modeling countermeasures
  - DETER: Experiment, experiment, experiment, ...

# Selected Related Publications/WIPs

All Publications available at <http://www.patrickmcdaniel.org/> or by request

- K. Butler, T. Farley, P. McDaniel, and J. Rexford, A Survey of BGP Security Issues and Solutions. Technical Report TD-5UGJ33, AT&T Labs - Research, Florham Park, NJ, September 2005. (*work in progress*)
- S. Qiu, P. McDaniel, F. Monroe, and A. Rubin, Characterizing Address Use Structure and Stability of Origin Advertisement in Interdomain Routing. Technical Report NAS-TR-0018-2005, Network and Security Research Center, Department of Computer Science and Engineering, Pennsylvania State University, University Park, PA, USA, July 2005. (*in submission*)
- W. Aiello, K. Butler, and P. McDaniel, Path Authentication in Interdomain Routing. Technical Report TR NAS-TR-0002-2004, Network and Security Center, Department of Computer Science and Engineering, Penn State University, September 2005. (*in submission*)
- W. Aiello, J. Ioannidis, and P. McDaniel, Origin Authentication in Interdomain Routing. *Proceedings of 10th ACM Conference on Computer and Communications Security (CCS)*, ACM, pages 165-178, October 2003. Washington, DC.
- G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing. *Proceedings of Network and Distributed Systems Security 2003 (NDSS)*, Internet Society, pages 75-85, February 2003. San Diego, CA.