NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Trustworthy Networking Program

Autonomous System Isolation Under BGP Session Attacks with RFD Exploitation

K. Sriram, D. Montgomery, O. Borchert, O. Kim, and R. Kuhn
National Institute of Standards and Technology
Gaithersburg, MD 20878
Contacts: ksriram@nist.gov, dougm@nist.gov

November 9, 2005

This research was supported by the Department of Homeland Security under the Secure Protocols for the Routing Infrastructure (SPRI) program and the NIST Information Technology Laboratory Trustworthy Networking Program.

BGP Vulnerabilities and Risks

- Much speculation..
  - Potential vulnerabilities and consequences.
  - Most threatening might be “bugs” – can cripple a router with a single packet.

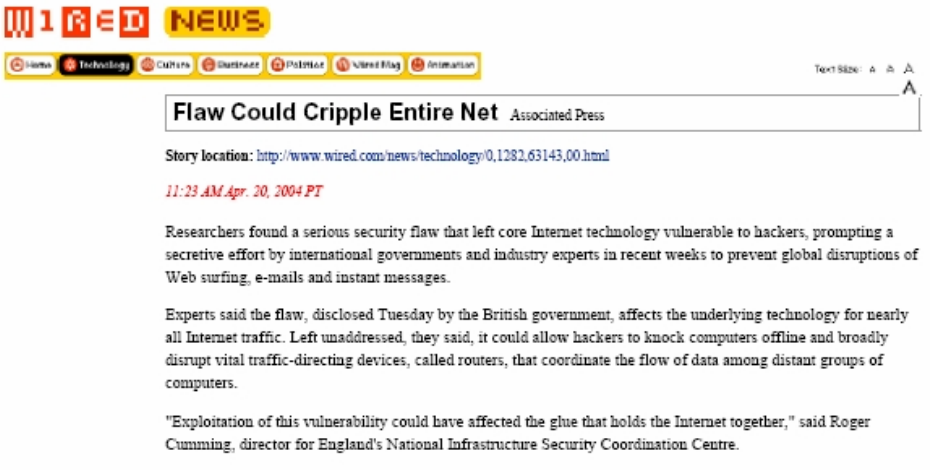

WIRED NEWS

Flaw Could Cripple Entire Net Associated Press

Story location: http://www.wired.com/news/technology/0,1282,63143,00.html

11:23 AM Apr. 20, 2004 PT

Researchers found a serious security flaw that left core Internet technology vulnerable to hackers, prompting a secretive effort by international governments and industry experts in recent weeks to prevent global disruptions of Web surfing, e-mails and instant messages.

Experts said the flaw, disclosed Tuesday by the British government, affects the underlying technology for nearly all Internet traffic. Left unaddressed, they said, it could allow hackers to knock computers offline and broadly disrupt vital traffic-directing devices, called routers, that coordinate the flow of data among distant groups of computers.

"Exploitation of this vulnerability could have affected the glue that holds the Internet together," said Roger Cumming, director for England's National Infrastructure Security Coordination Centre.

- Little public analysis or data ….
  - Empirical analysis of vulnerabilities and their potential consequences.
  - Trace data of actual attacks on the routing infrastructure.

Efforts to Understand the Risks and Possible Solutions

Long term solutions in a state of flux.

- S-BGP, SO-BGP, MD5/IPsec, GTSM, Route Verification, Filtering, Listen & Whisper, etc.
- Range of technologies that may, or may not, be viable.
- It depends on what **your** view of the **risks and benefits vs. costs**.

Lack of shared understanding of both the problem & solution space.

- Need to raise community awareness of potential threats, **risks**, mitigation techniques and their **cost**.
- Need to take "systems view" of improving routing's survivability.
- DHS – "need some way of characterizing benefit vs. cost of various solution techniques."

NIST Objectives:

- **Expedite Research** - Help researchers characterize the design space: risks, mitigation techniques and deployment costs.
- **Expedite Development** - Evaluate the effectiveness and impact of proposed technical solutions.
- **Expedite Adoption** - Help users / decision makers understand threats & mitigations.

NIST Efforts

Near Term Efforts:

- ***DHS - "Focus on the problem / design space."***
- **Large Scale Modeling of BGP Attacks**
  - Most modeling / analysis focused on post-mortem analysis of recent worms/viruses, but **"what if" scenarios** of yet unseen attacks may be more important.
  - **Risk analysis** of the potential impact of successful attacks on BGP.
  - Discover and evaluate **new vulnerabilities**.
  - Look for **emergent behaviors** – e.g., cascading failures, congestion collapse, degraded routing.
  - Framework for **characterization of proposed solutions** & deployment scenarios
- **Modeling and Analysis of Proposed Solutions**
  - Characterizing the effectiveness and cost of the various combinations of countermeasures.
  - Characterize the risk associated with the deployment of proposed solutions.
- **Issue Federal Guidance**
  - FISMA guidance on BGP Security.

BGP Attack Tree Enumeration

- Broad classification of attacks (IETF drafts):
  - Establish Unauthorized BGP Session with Peer
  - Originate Unauthorized Prefix/Attribute into Peer Route Table
  - Change Path Preference of a Prefix
  - Conduct Denial/Degradation of Service Attack Against BGP Process
  - Reset a BGP Peering Session
  - Send Spoofed BGP Message

BGP Peering Session Attacks

- There are many different attack possibilities on the BGP routing infrastructure (IETF ID: draft-ietf-rpsec-bgpattack-00)
- We focus on attacks that cause BGP peering sessions to be reset
- Common way to reset a BGP peering session is to reset or attack the underlying TCP connection
- Multiple TCP/ICMP vulnerabilities documented - may be exploited to launch TCP connection-reset attacks
  - "Slipping in the window" TCP reset attack (requires correctly guessing a TCP sequence number within a flow control window)
  - ICMP error messages spoofed to cause TCP reset (IETF ID, Dec. 2004)
    - ✓ Does not require guessing the TCP sequence number
    - ✓ Hard ICMP error messages (spoofed)
    - ✓ Soft ICMP error messages (spoofed)

MRAI: Minimum Route Advertisement Interval

- A BGP router sends route advertisements/withdrawals to a peer at intervals no smaller than MRAI
- Jittered MRAI: randomly chosen from a range of 22.5s to 30s (independently at each node)
- MRAI is a sender side discipline for neighbor overload avoidance

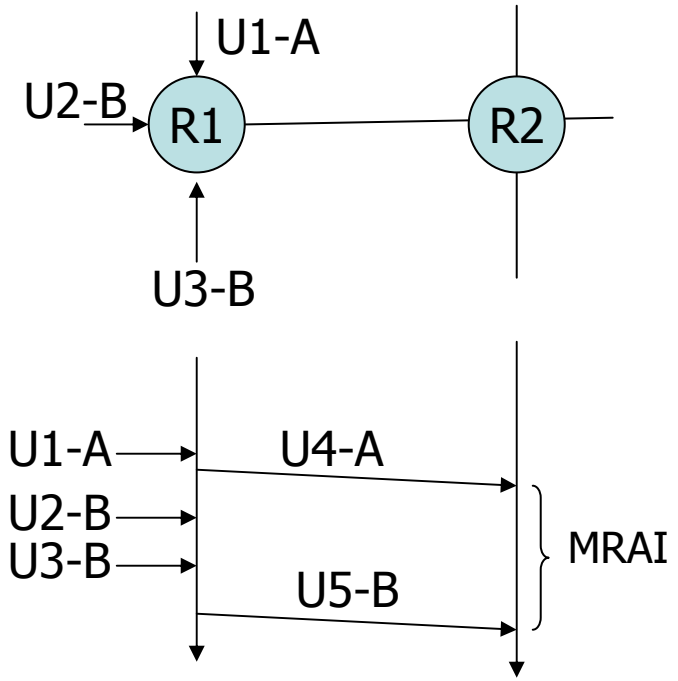

RFD: Route Flap Damping

- An upstream router assigns an incremental RFD penalty to a peer and destination (prefix) combination each time an update is received from that peer for that destination
- If the RFD penalty exceeds a preset cutoff threshold, then the route is suppressed
- RFD is a method for receiver side route monitoring and suppression in the event of frequent updates

RFD Parameter	Vendor A	Vendor B	
Withdrawal penalty	1000	1000	
Re-advertisement penalty	0	1000	
Attribute change penalty	500	500	
Cutoff threshold	2000	3000	
Half-time	900	900	sec
Reuse threshold	750	750	
Max supress time	3600	3600	sec
Max penalty	12000	12000	

- **The two sets of numbers correspond to two commercial implementations**
- **Use the numbers for sensitivity study in our numerical examples**

Exploitation of Route Flap Damping

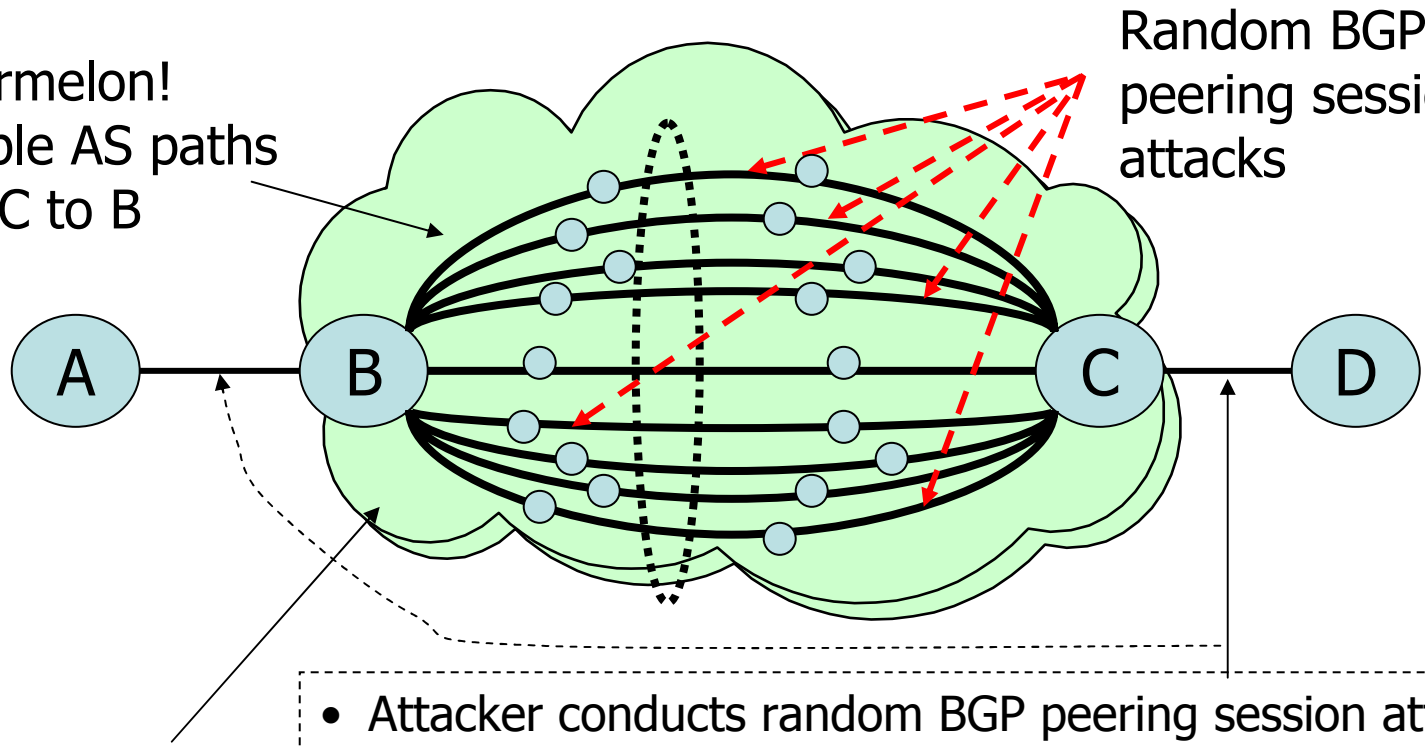

- Attacker conducts random BGP peering session attacks into the cloud with some probability of success
- RFD behavior on either of these links is exploited by the attacker

Illustration: How It Works (MRAI = 30 s)

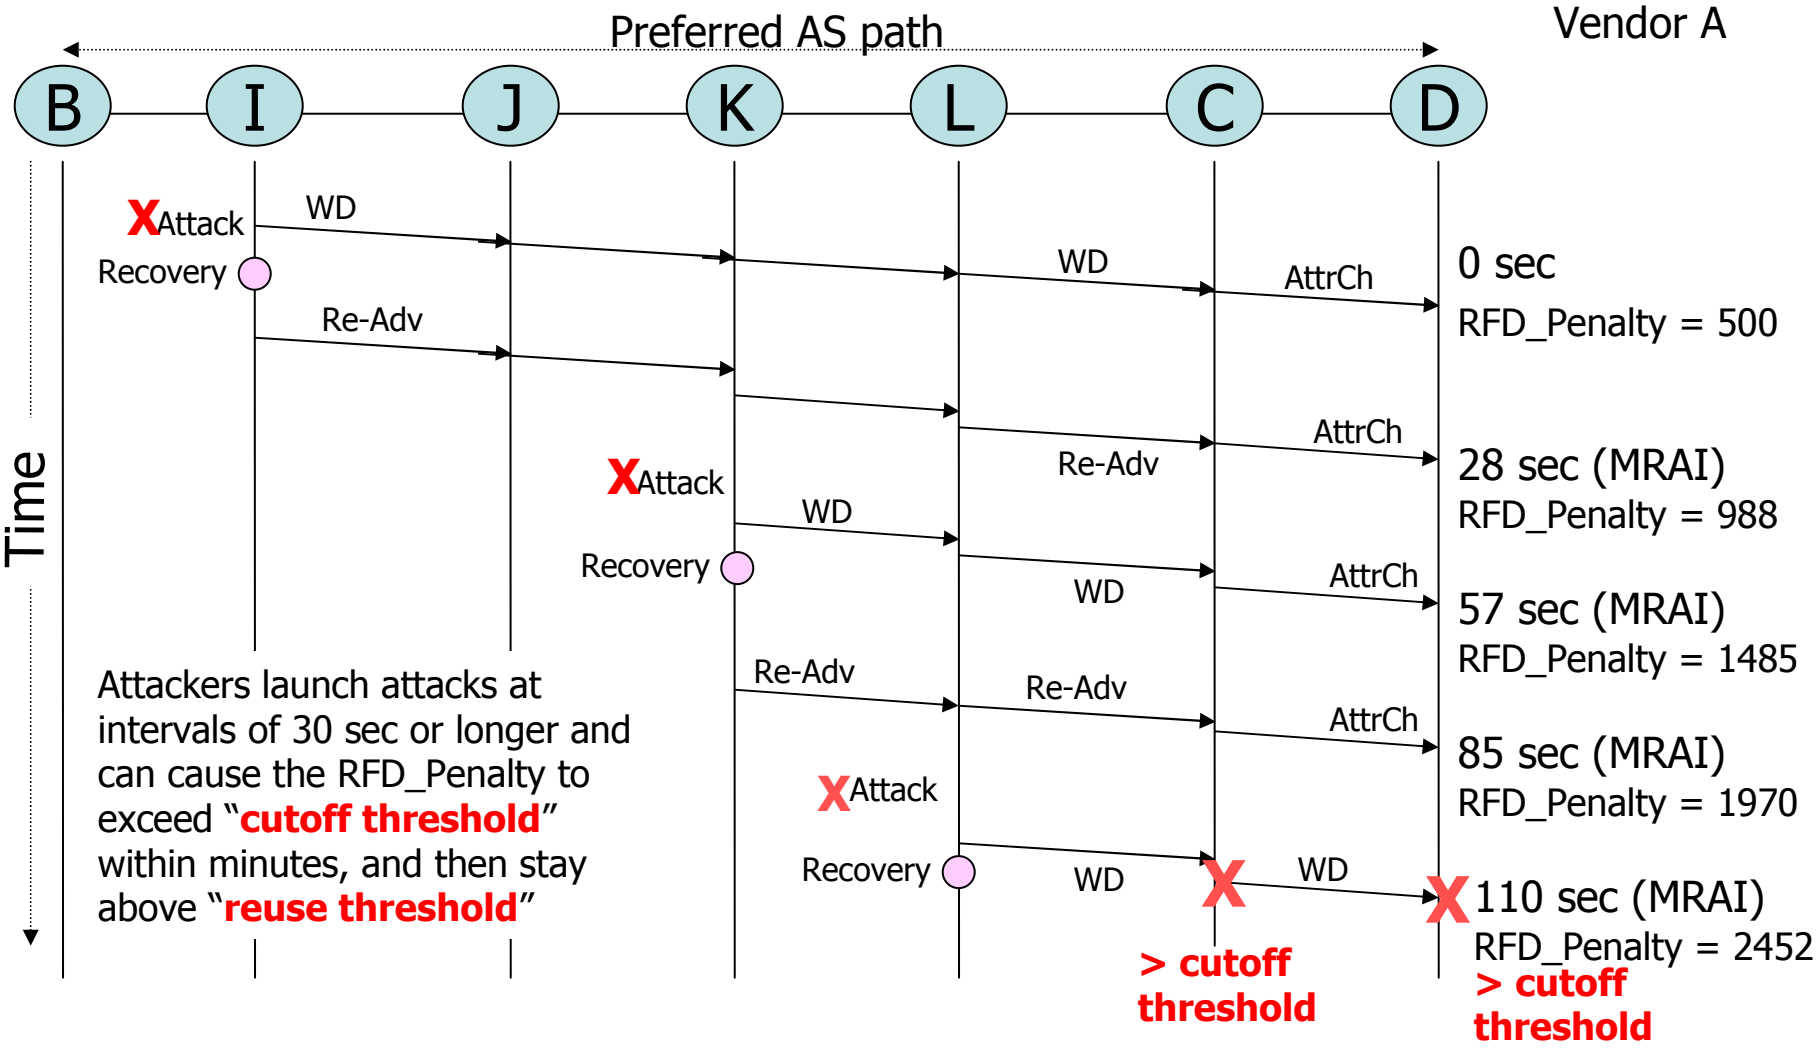

Illustration: How It Works (MRAI = 30 s)

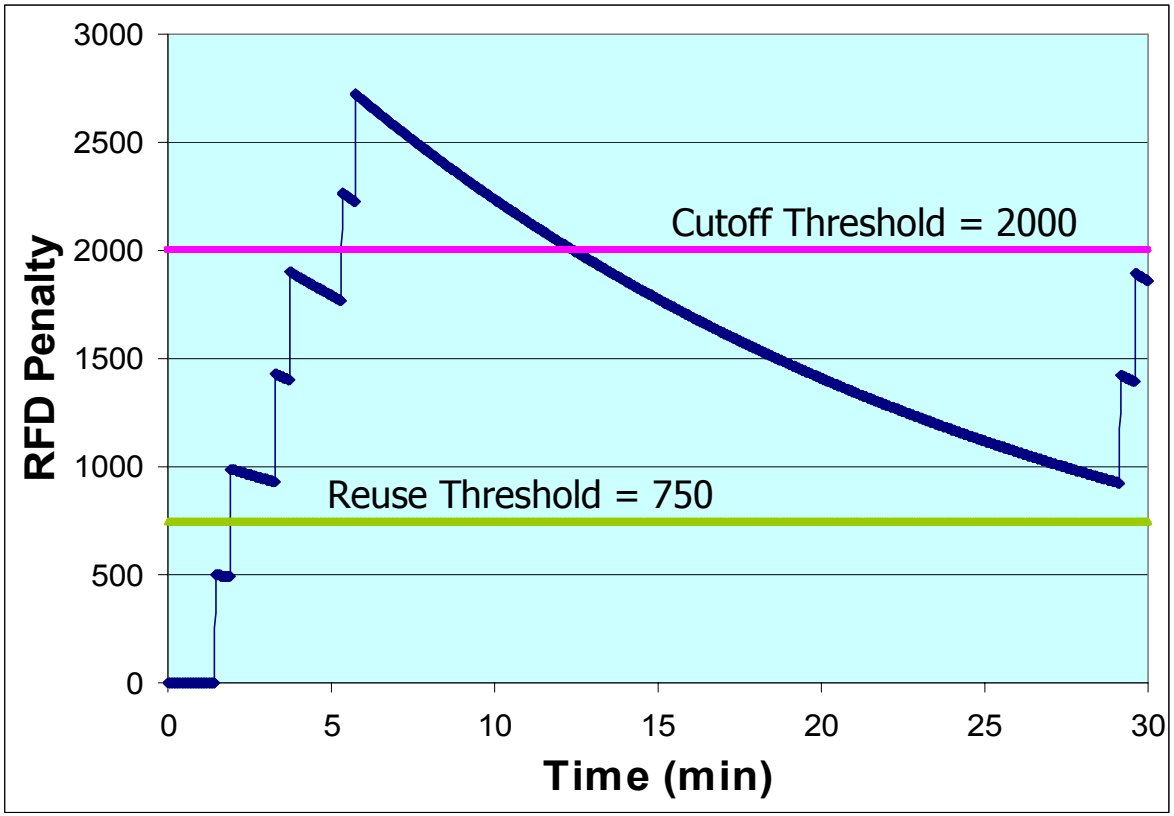

- The update interval is effected by MRAI
- Attackers need to successfully attack one of the BGP peering sessions on the preferred path for the penalty to go higher
- 30 sec MRAI allows enough time for the damaged BGP session to recover within the MRAI
- The waves of attacks would be spaced at intervals equaling approximately MRAI
- To achieve prolonged AS isolation, it is enough if only some of the attacks succeed
- Once RFD penalty is exceeded, the attack interval can be made larger (although attackers don't know when they have succeeded)

Analytical Model for AS Isolation Probability

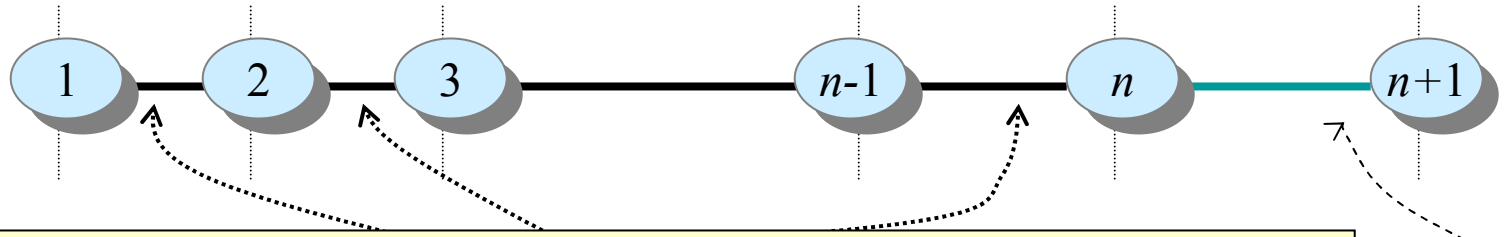

- *n*-1 BGP peering sessions
- Attacks are assumed to be spaced at roughly MRAI intervals
- Each router is subjected to an attack with probability *p* in each interval
- Each BGP peering session can be attacked with probability *q* if there is a router at either end that is subjected to attack

- Model predicts the probability that update rejections due to Route Flap Damping are imposed at router *n*+1 for peer *n* and destination 1
- Model also predicts the sustenance probability that the attackers can sustain the RFD in update rejection state and thus cause prolonged isolation between router *n*+1 and destination 1 (also all subsequent destinations reachable via router 1).

Attacks and RFD Penalty Accumulation Model

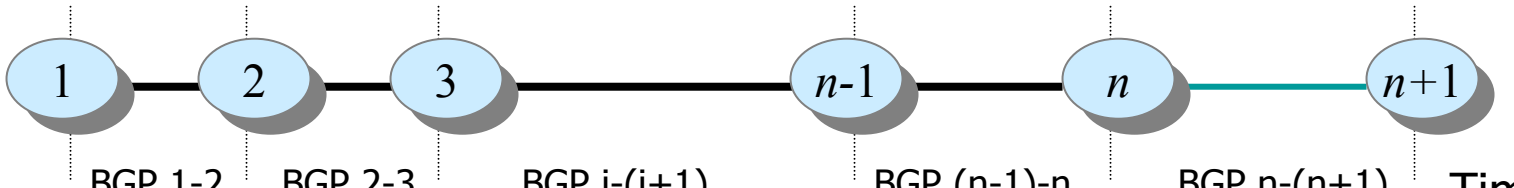

BGP 1-2	BGP 2-3	BGP i-(i+1)	BGP (n-1)-n	BGP n-(n+1)	Time
	X		Withdrawal Re-Adv	AttrCh AttrCh	MRAI i
					MRAI i+1
			X	AttrCh AttrCh	MRAI i+2
X		X	Withdrawal Re-Adv	AttrCh AttrCh	MRAI i+3
		X		⋮	MRAI i+4
				RFD cutoff state	
				⋮	

Time (n x MRAI)

X = Successful BGP peering session attack

Note: Router *n* has alternate routes to Router 1

Estimation of Attacks Needed to Push Penalty Above Cutoff

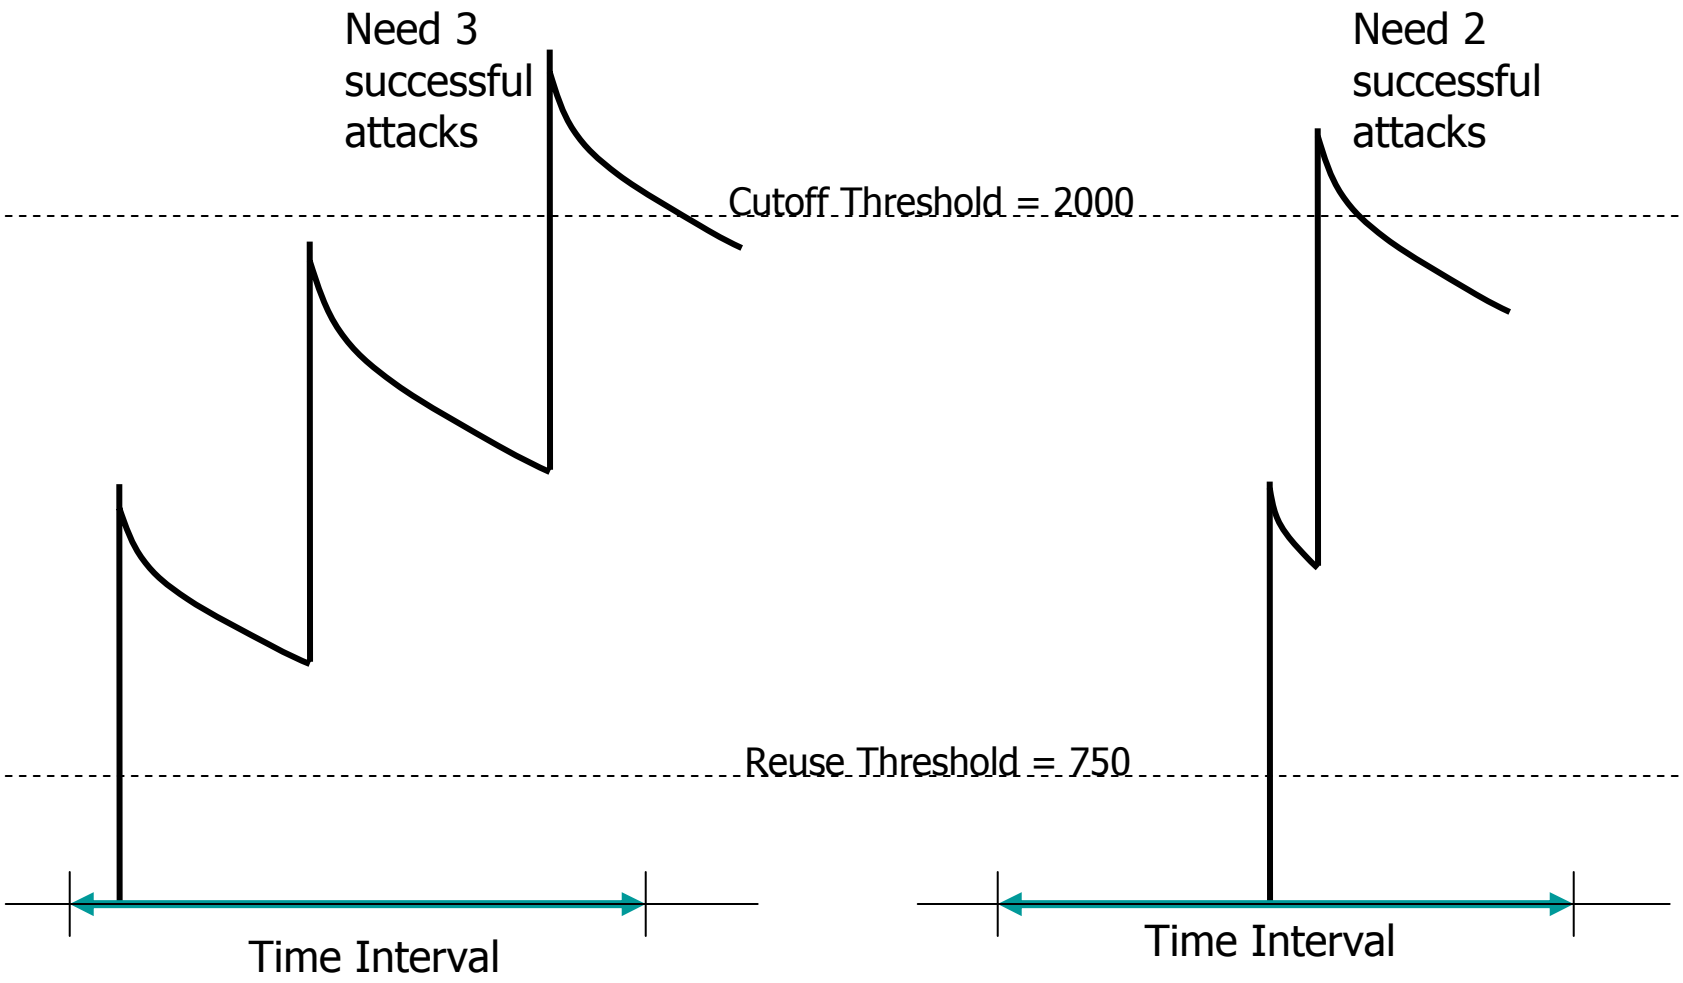

Attacks and RFD Penalty Accumulation Model

$C = cutoff\ threshold$,

$R = reuse\ threshold$,

$H = half\ time$ (decay parameter),

$T =$ MRAI time (≈ 30 sec),

$P =$ incremental penalty incurred per successful attack event,

$n =$ number of BGP nodes in the AS path subject to attacks,

$Q = \Pr\{$a BGP peering session attack is successful$\}$,

$\theta = \Pr\{$AS path of n ASes is successfully attacked at one or more BGP peering sessions$\}$,

$E =$ Elapsed time from the time of beginning of BGP session attacks (in multiples of MRAI)

$R_P(n+1; n, 1; iT) =$ RFD penalty at router $n+1$ for peer n and destination 1 at time iT

$\alpha(n,k) = \Pr\{R_P(n+1; n, 1; iT) > C$ for some $i \in (0,k) \mid E = kT\ )$

Attacks and RFD Penalty Accumulation Model

$$\theta = 1 - (1 - Q)^{n-1}$$

RFD cutoff threshold check (for j attacks in k MRAI intervals):

$$P\sum_{i=0}^{j-1} 2^{\left\{-\frac{ikT}{(j-1)H}\right\}} > C$$

Let $j_{\min}(k)$ be the smallest j that satisfies the above inequality.

Then,

$$\alpha(n,k) = \sum_{j_{\min}(k)}^{k} \beta_i(n,k)$$

where,

$$\beta_i(n,k) = \frac{k!}{i!(k-i)!}\theta^i(1-\theta)^{k-i}$$

==========

AS/Peer Isolation Sustenance Probability:

$$P_{sus} = 1 - (1-\theta)^{\left\lfloor H\left(\log_2 \frac{C}{R}\right)/t_M \right\rfloor}$$

Probability of AS-Prefix Isolation

Probability that AS-Prefix isolation occurs within t sec from start of attacks:

- Sensitivity to vendor settings of RFD parameter values is quite significant
- $n = 4$

(#ASes in AS path)

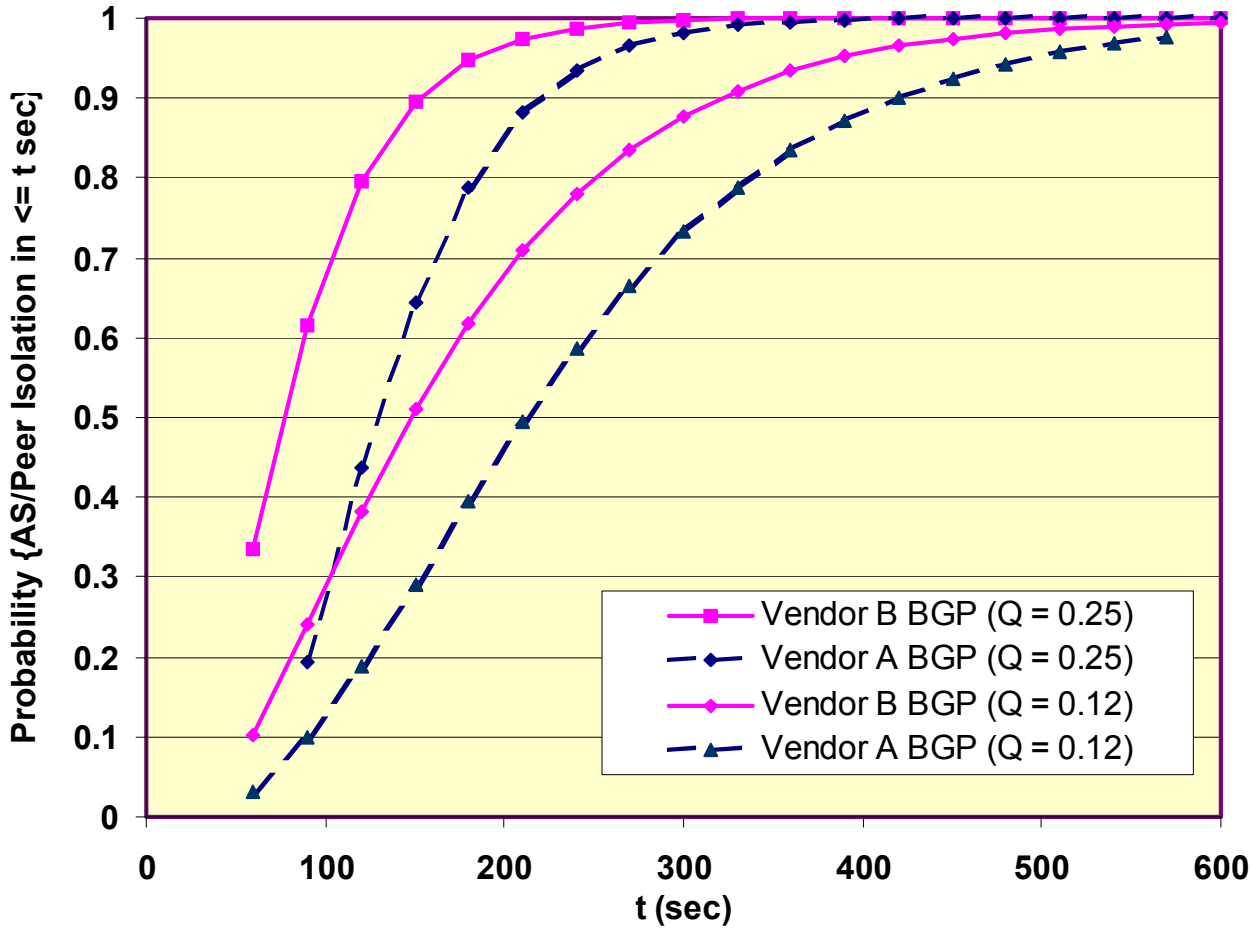

Probability of AS-Prefix Isolation

- Vulnerability is higher if AS path-lengths within the attack area are higher
- $Q = 0.25$

Probability that AS-Prefix isolation occurs within t sec from start of attacks:

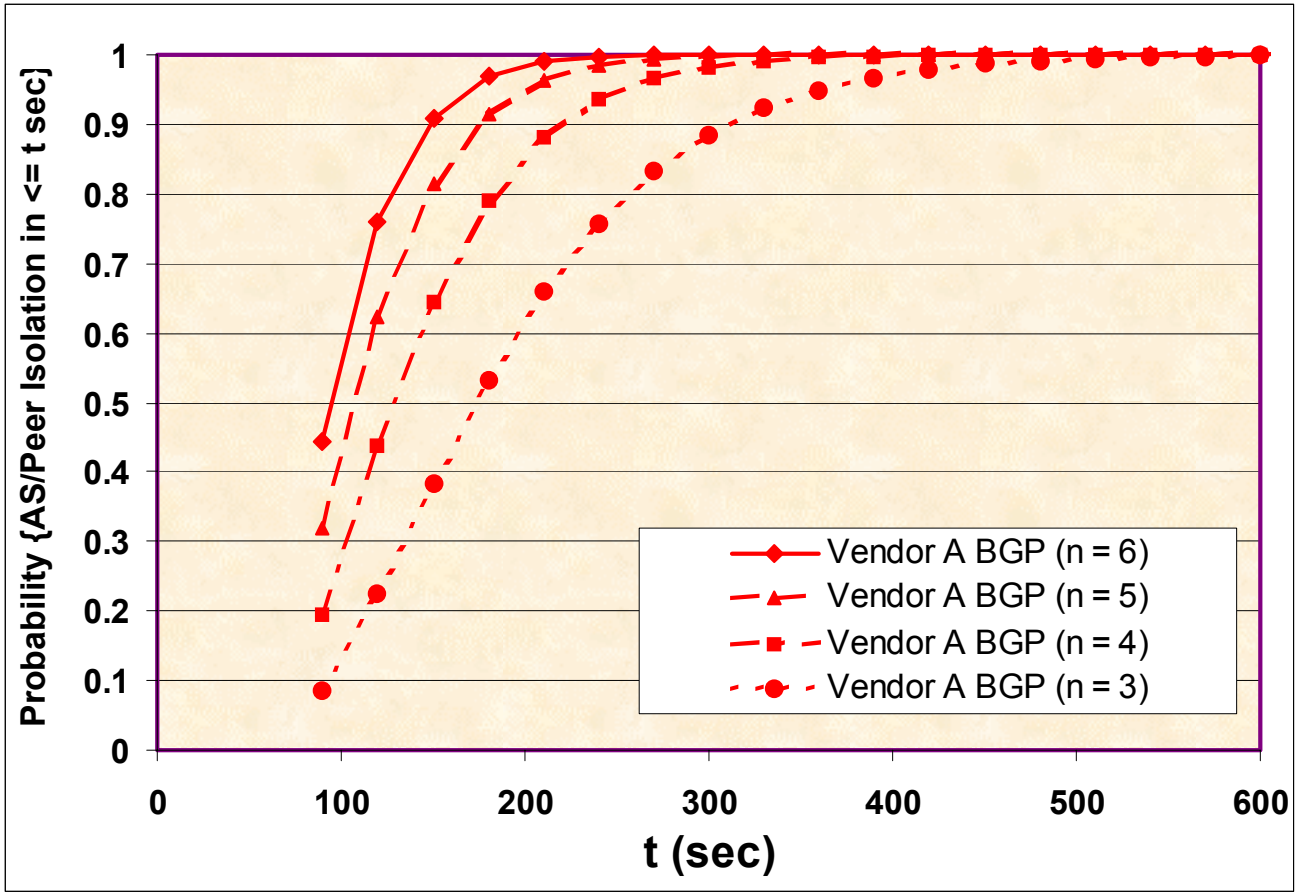

Probability of AS-Prefix Isolation

Probability that AS-Prefix isolation occurs within t sec from start of attacks:

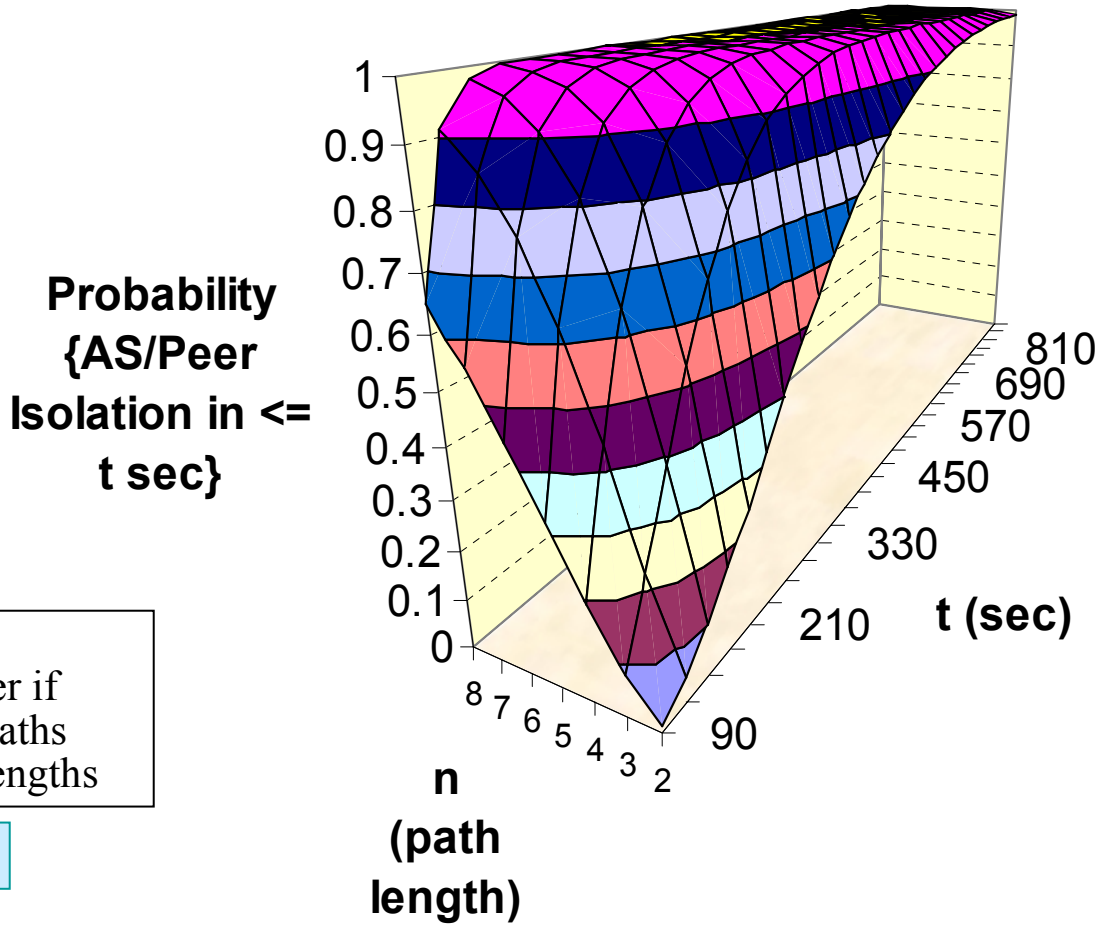

- Attack goal is reached sooner if targeted AS paths have longer lengths

$Q = 0.25$

Probability of Sustenance of AS-Prefix Isolation

Given that an AS-Prefix isolation occurred, what is the probability that it can be sustained for a prolonged period by the attackers:

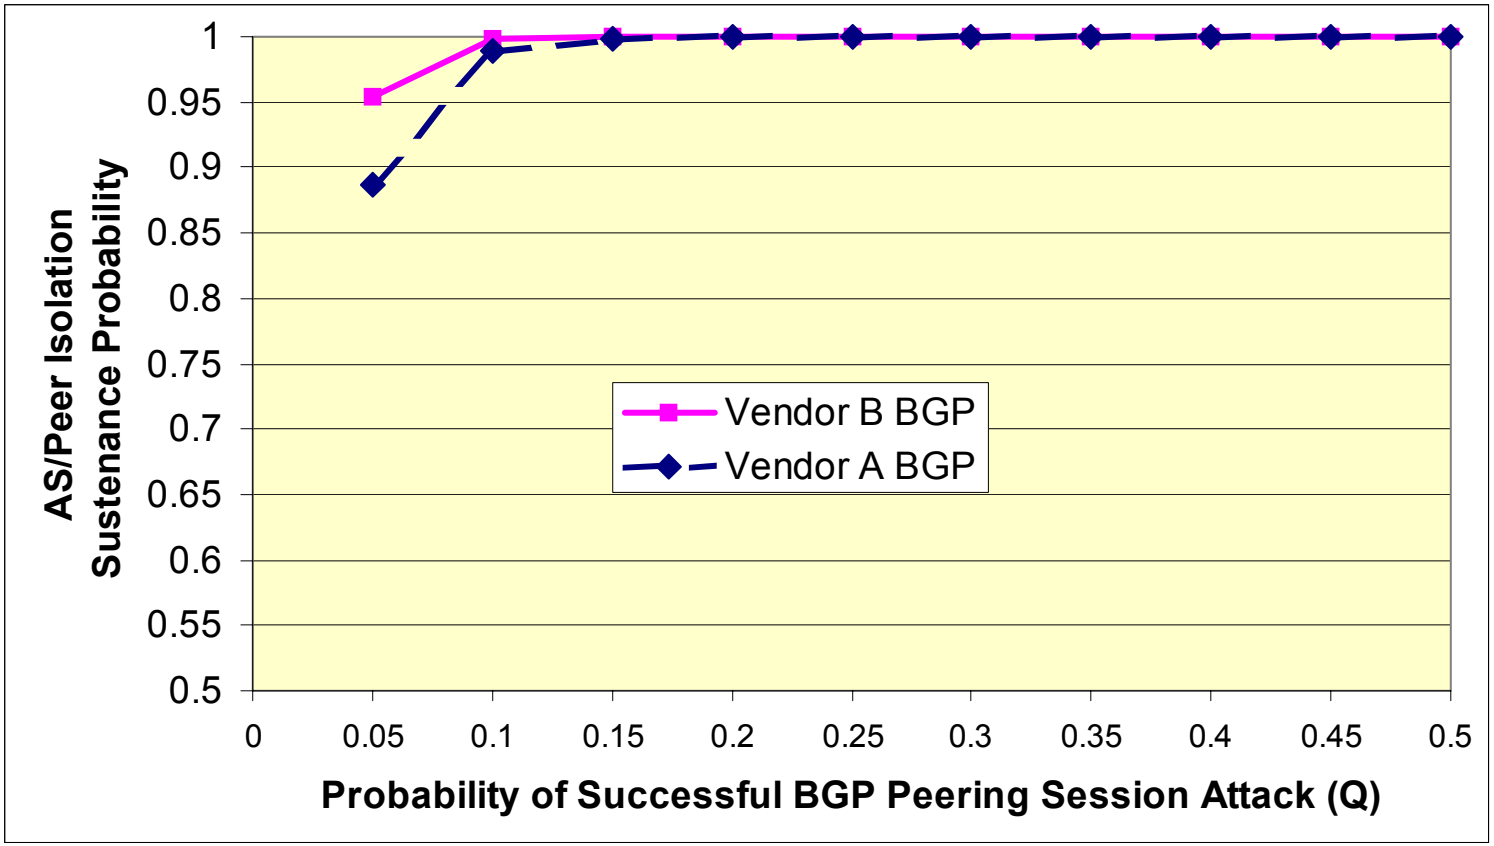

$n = 4$

BGP Graceful Restart: Brief Description

- Gives downed router time to restart without peers withdrawing its routes
- Option negotiated at OPEN
- Two flag bits in capability advertisement
  - Restart bit = router has restarted
  - Forwarding bit = preserved forwarding state
- During restart, peers do not send withdrawals for the restarting router; prevents route flapping
- Restart timer:
  - Restart-time determines how long peer routers will wait to delete stale routes before a BGP open message is received
- If restart-time expired: restart failed, routes deleted, withdrawals sent

BGP Graceful Restart: Mitigation of RFD Exploitation Attacks and Avoidance of AS Isolation

- Without BGP-GR, the RFD exploitation attack resulting in AS isolation is much more feasible
- BGP-GR helps mitigate this type of attack
- With BGP-GR, the attackers need a lot more effort (100 times or more) to even induce route withdrawals at a peer
- BGP-GR restart time = 120 s
- BGP session recovery time = 4 s

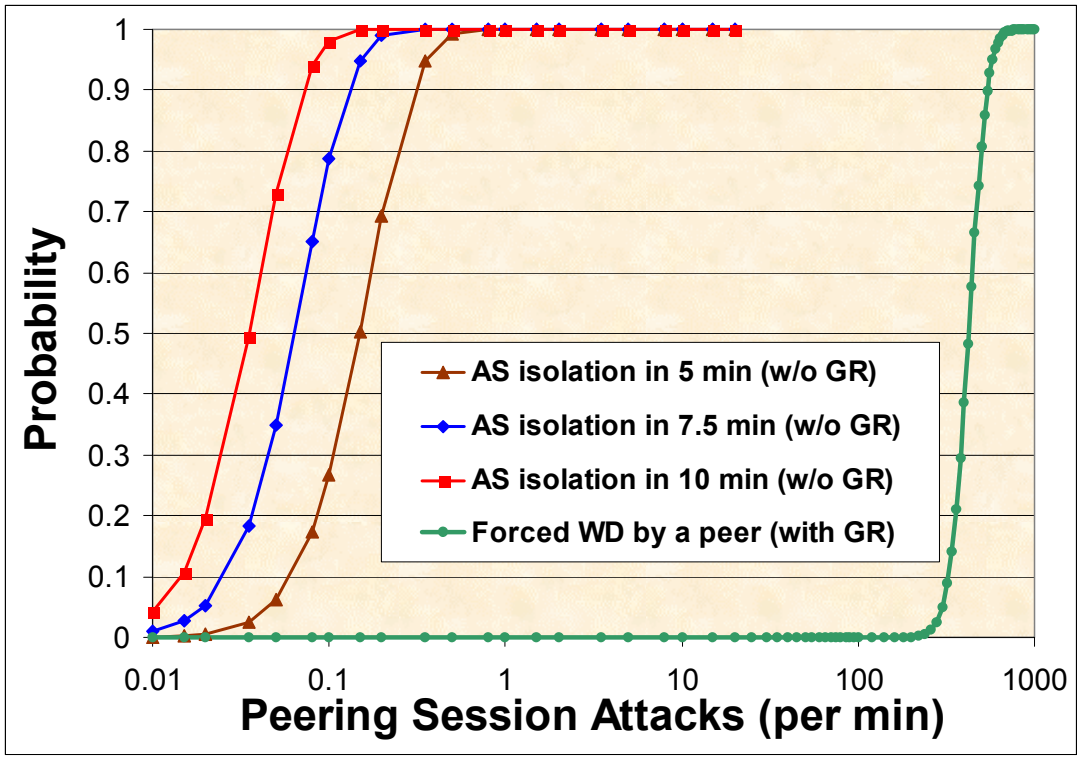

$n = 4$

$Q = 0.1$

- "Several providers (US) suggest that the cost of implementing this feature outweighs the benefit." – NISCC (UK govt) BGP Best Practices

RFD Attacks: Simulation Results

Grid Topology of Size 8x8

- 64 node grid
- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%

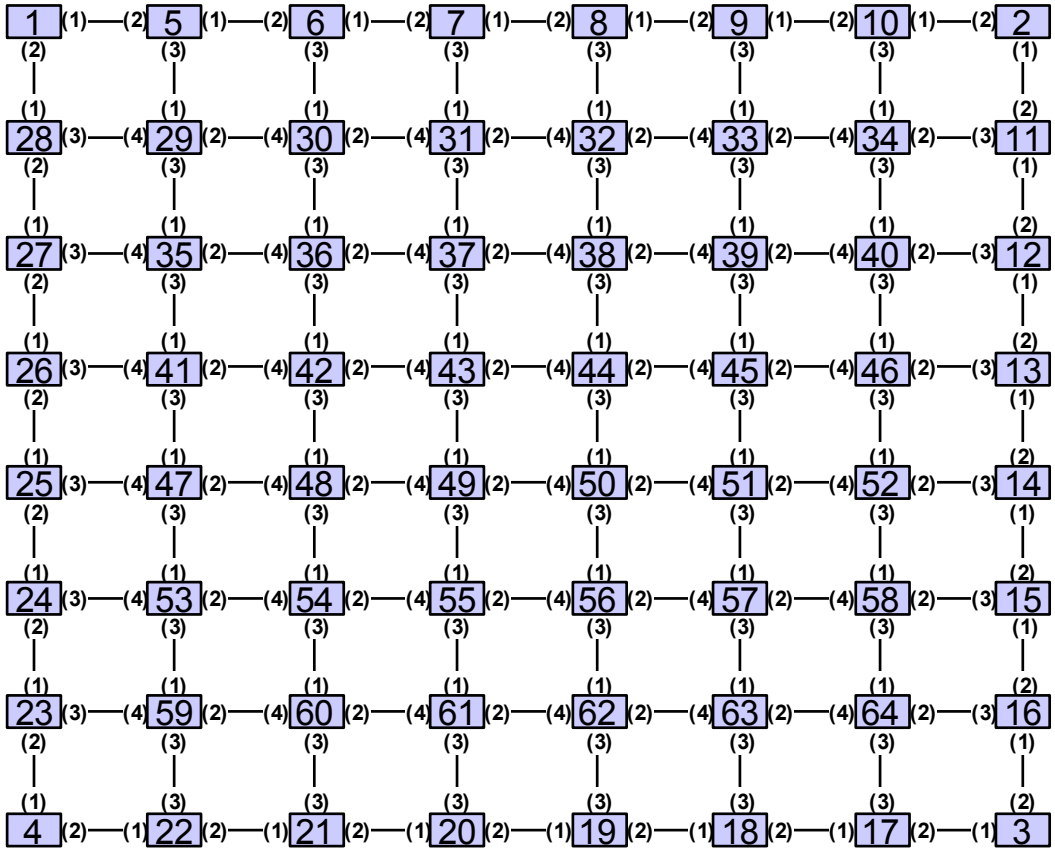

8X8Matrix.dml

Measured # BGP Session Resets vs. Node ID

- 64 node grid
- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%

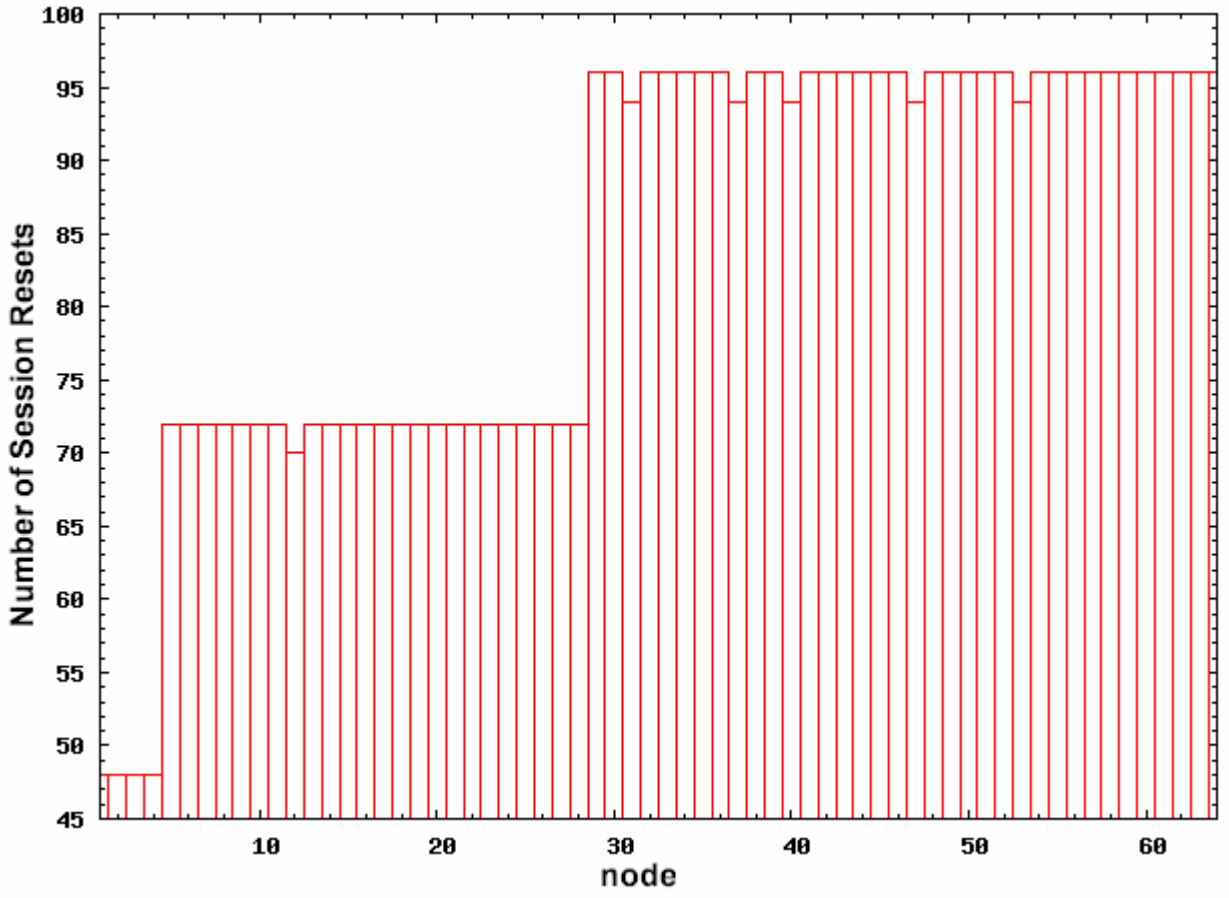

Comparison of Unreachability Time

(a) Without RFD

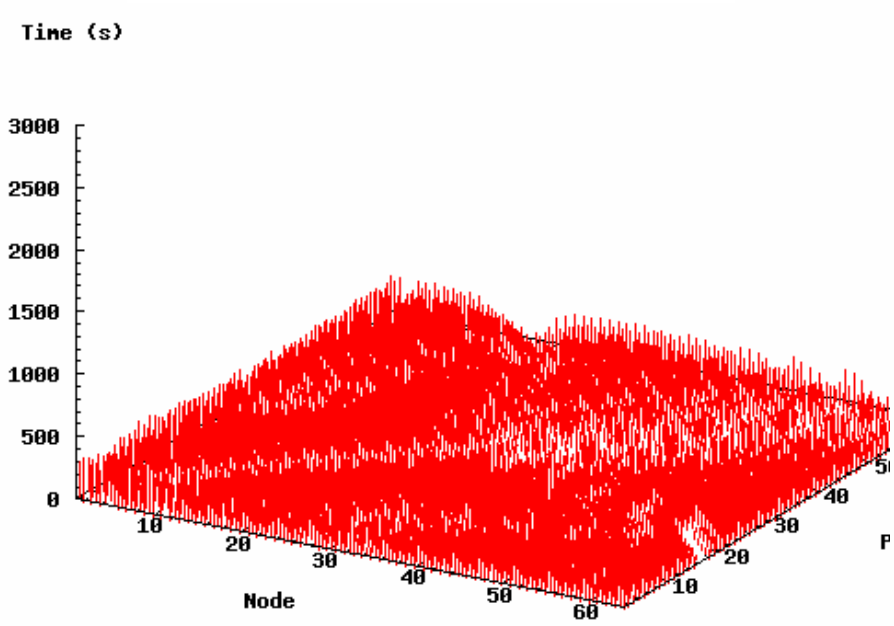

(b) With RFD

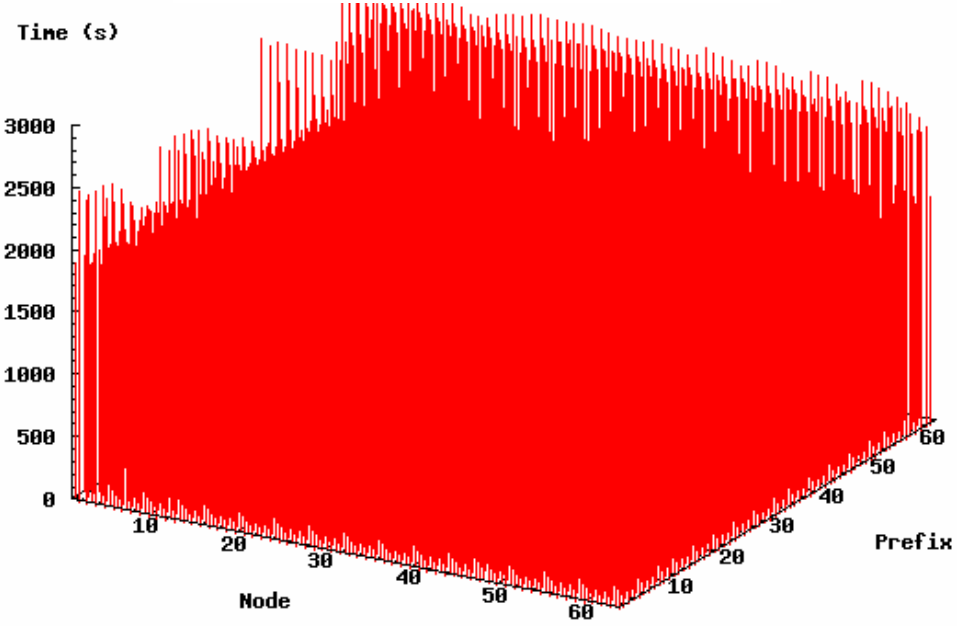

- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%
- 64 node grid

Comparison of Update Count

(a) Without RFD

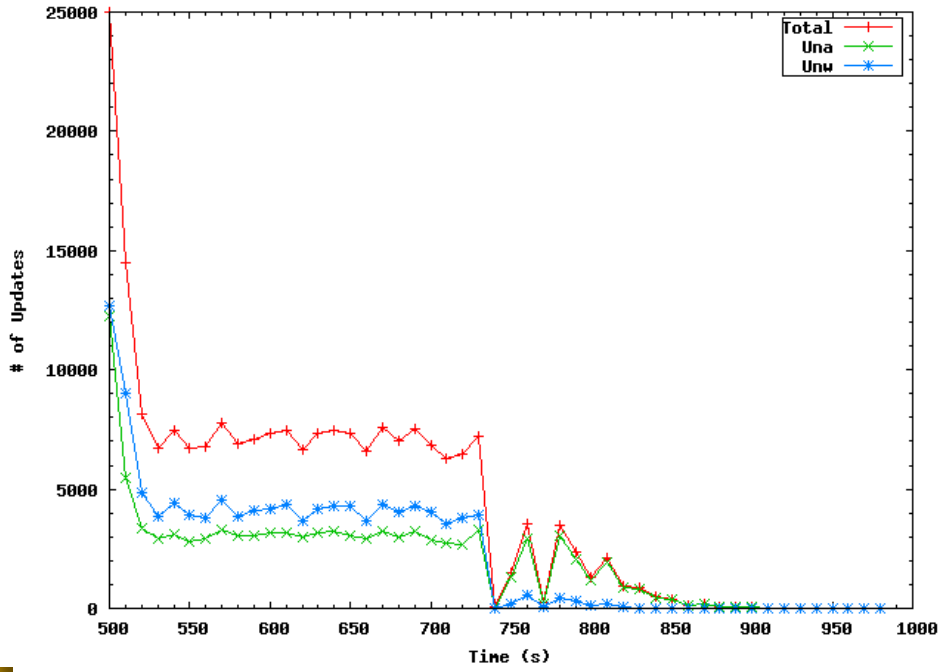

(b) With RFD

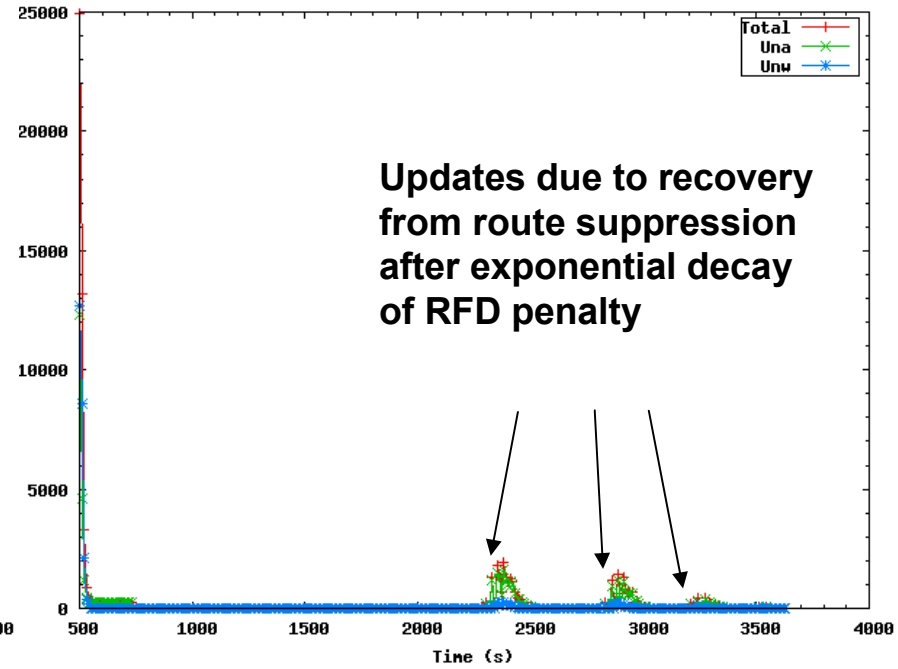

- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%
- 64 node grid

Count of (*i,j*) Pairs Unreachable

(a) Without RFD

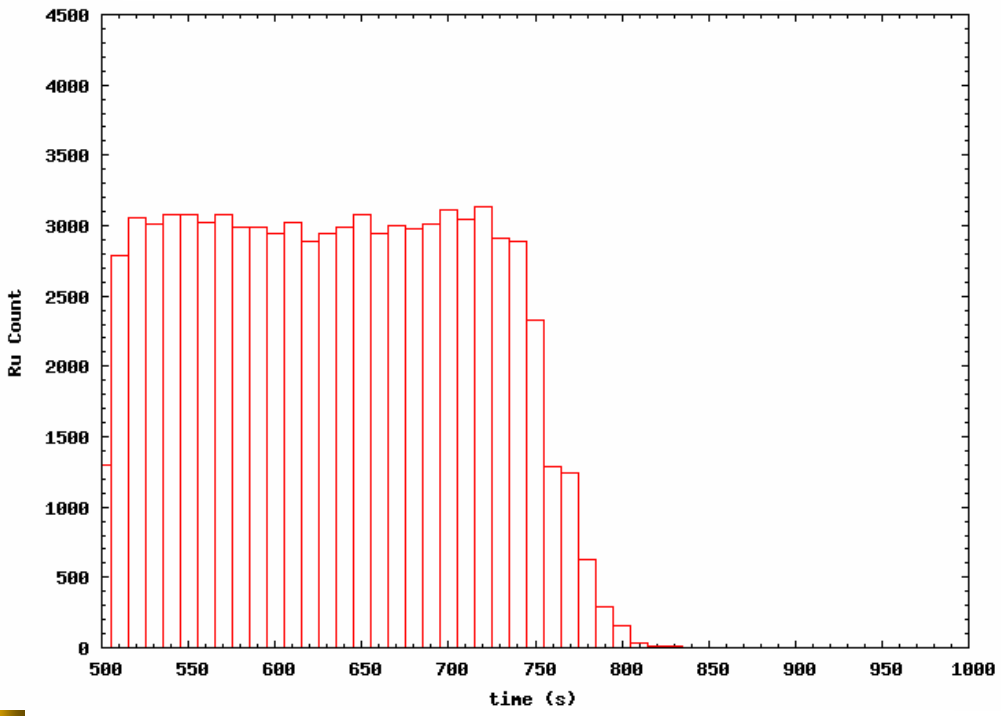

(b) With RFD

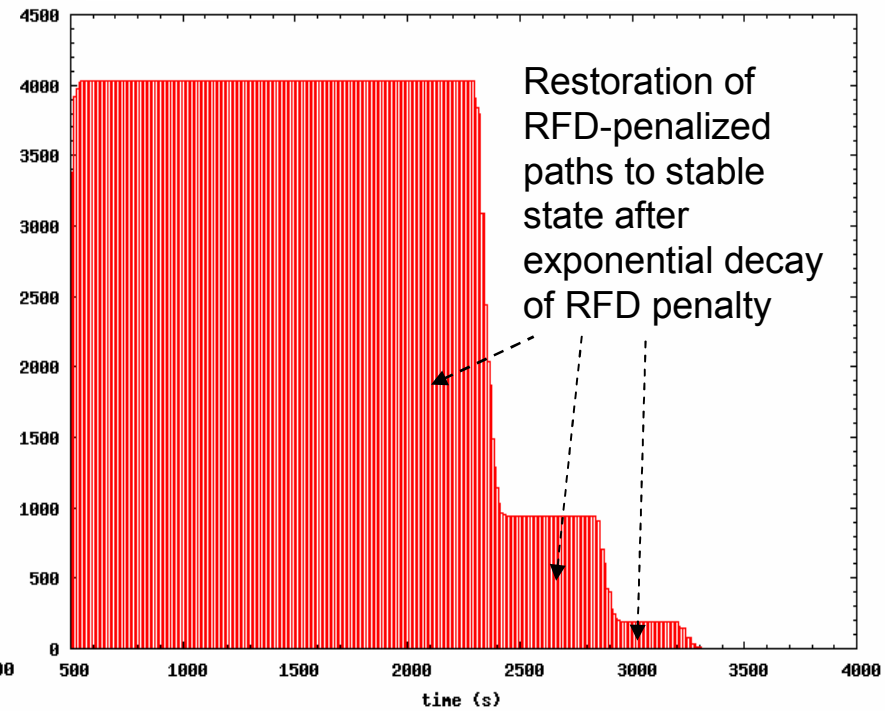

- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%
- 64 node grid

Restoration to Stable Route: Time & Count

(a) Without RFD

Cumulative time away from stable path

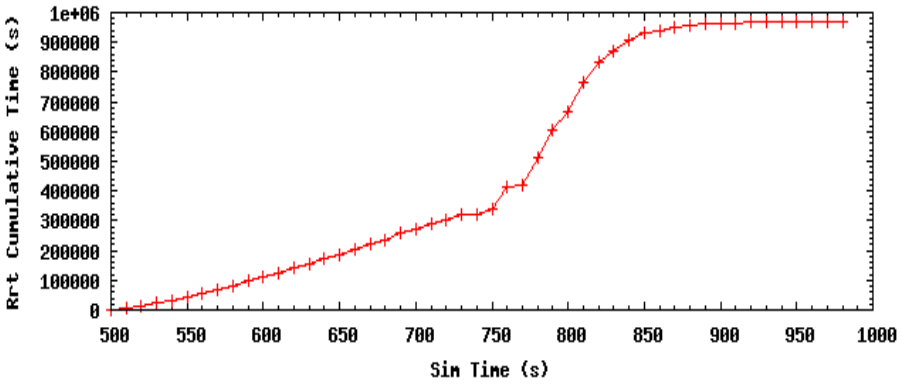

 # Retuned to stable path

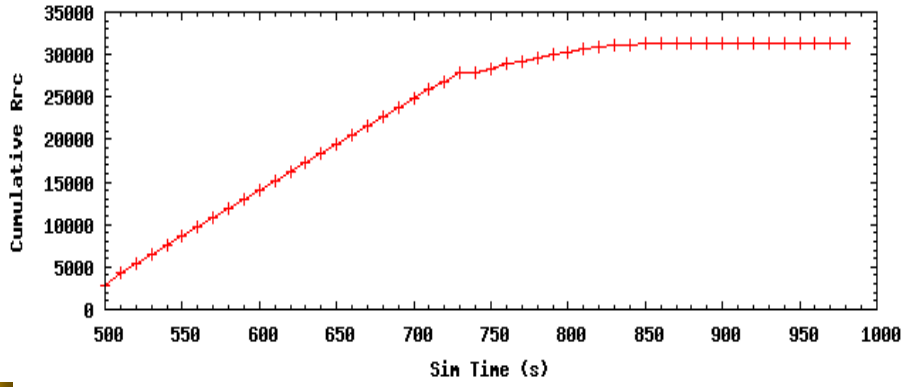

(b) With RFD

Cumulative time away from stable path

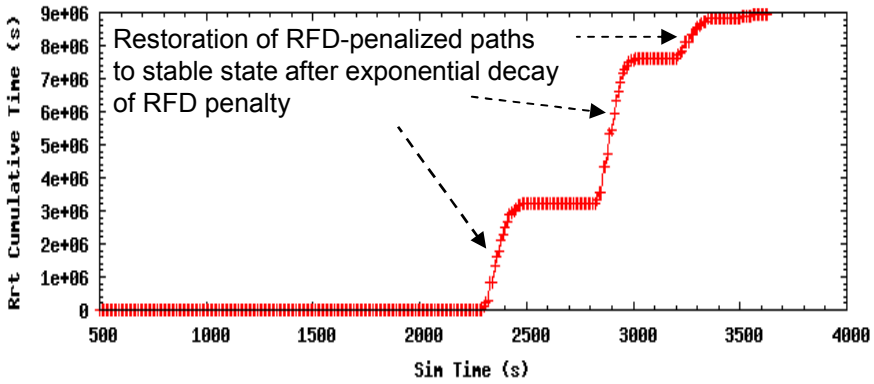

 # Retuned to stable path

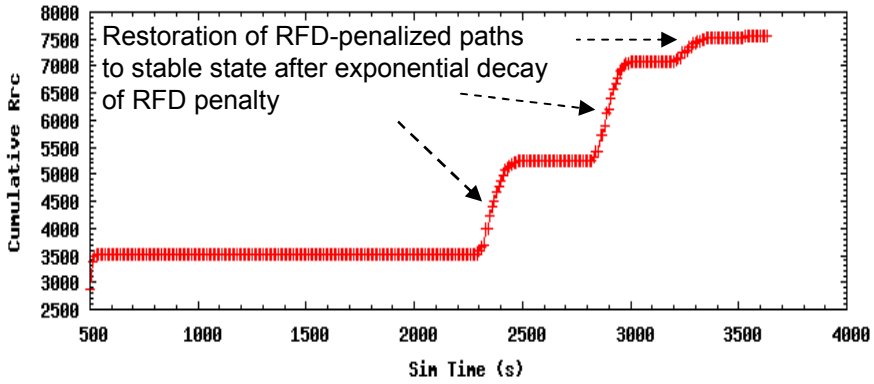

- Total attack duration = 240 sec
- # Attack intervals = 24 (each is 10 sec)
- Prob. of success for each attack = 100%
- 64 node grid

Measured # BGP Session Resets Plotted over Topology

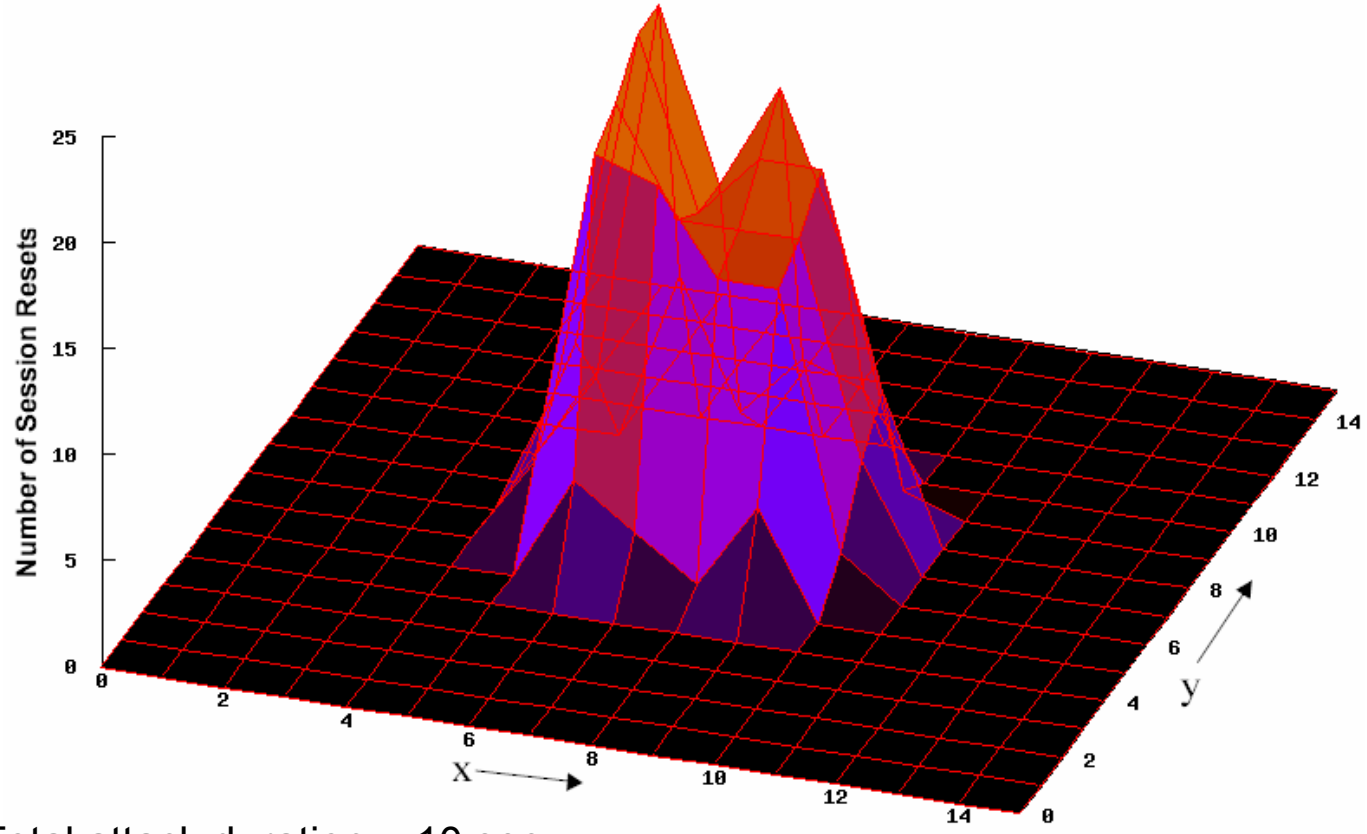

- Total attack duration = 10 sec
- 4x4 sub-grid under attack
- # Attack intervals = 16 (each is 5/8 sec)
- Prob. of success for each attack = 25%
- 256 node grid

Comparison of Unreachability Time

(a) Without RFD

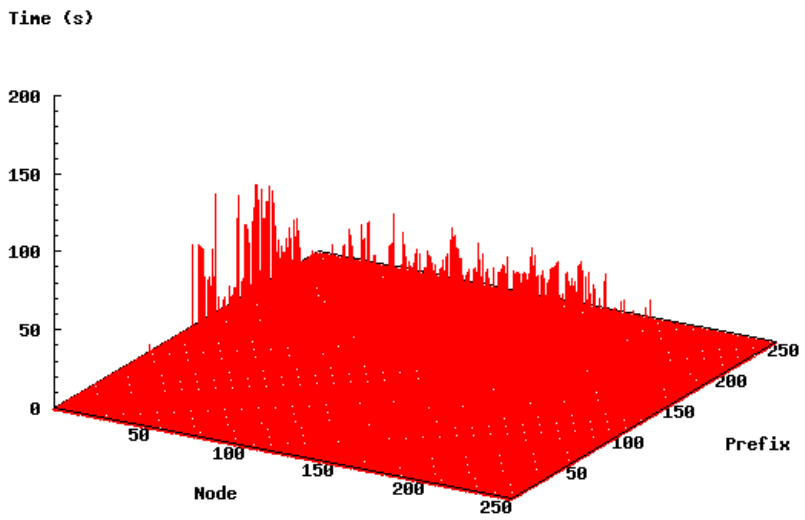

(b) With RFD

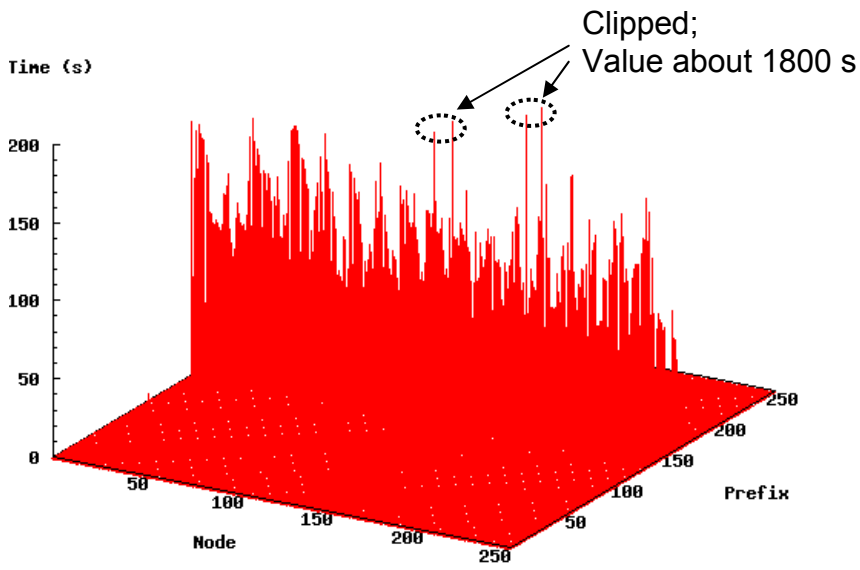

- Total attack duration = 10 sec
- 4x4 sub-grid under attack
- # Attack intervals = 16 (each is 5/8 sec)
- Prob. of success for each attack = 25%
- 256 node grid

Update Count

(a) Without RFD

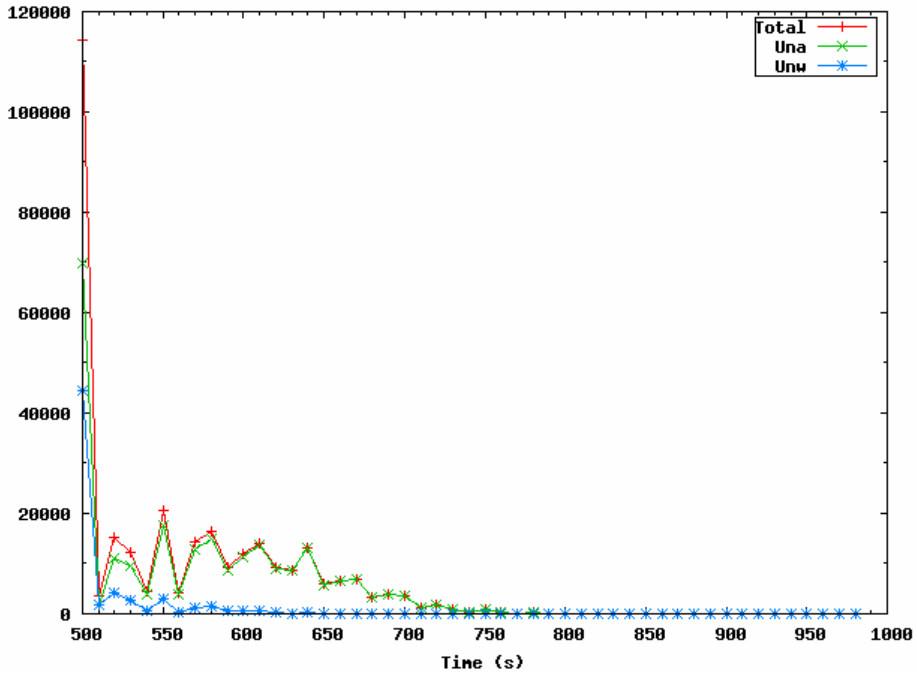

(b) With RFD

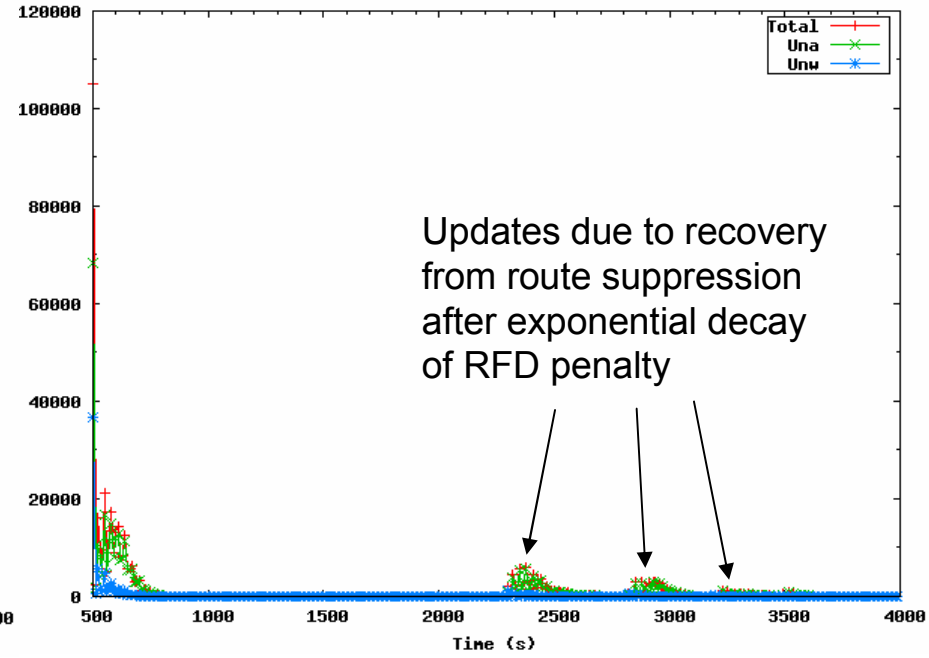

- Total attack duration = 10 sec
- 4x4 sub-grid under attack
- # Attack segments = 16 (each is 5/8 sec)
- Prob. of success for each attack = 25%
- 256 node grid

Restoration to Stable Route: Time & Count

(a) Without RFD

Cumulative time away from stable path

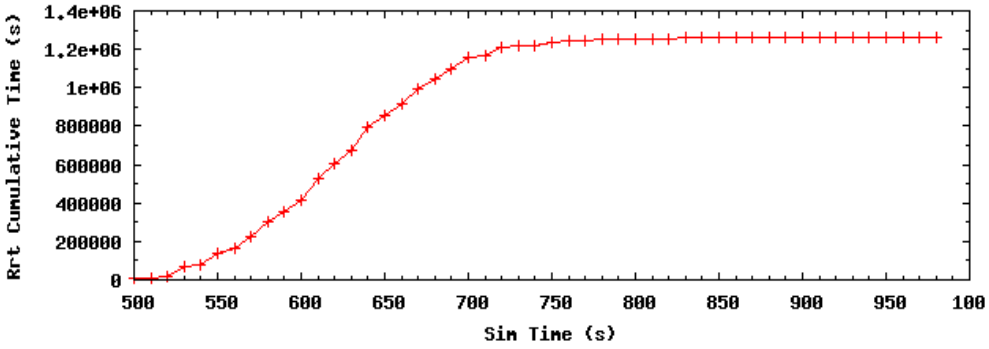

 # Retuned to stable path

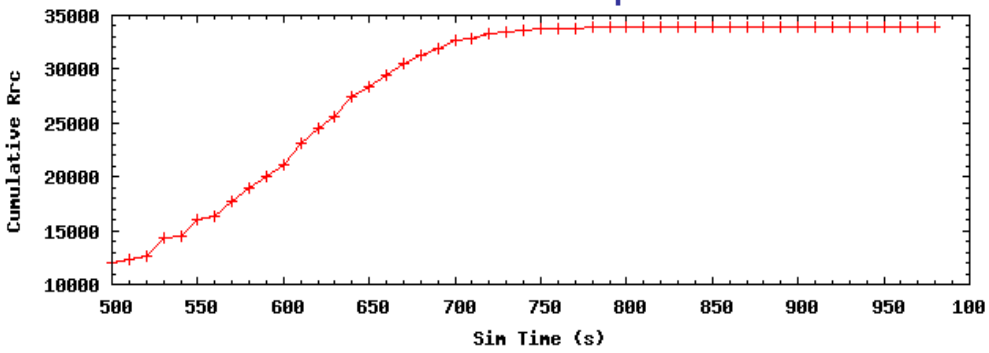

(b) With RFD

Cumulative time away from stable path

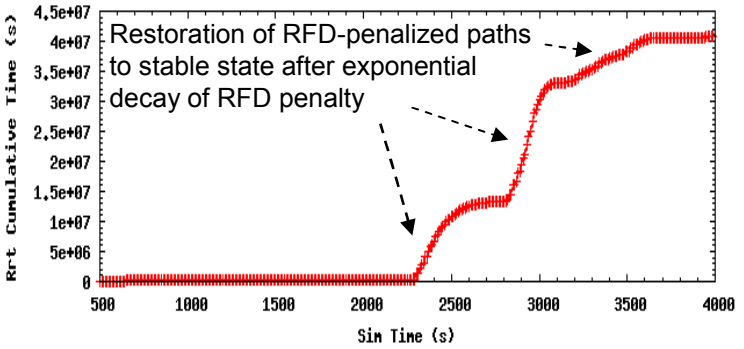

 # Retuned to stable path

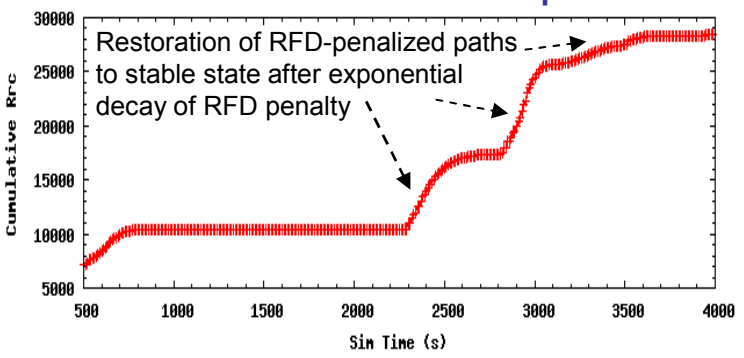

- Total attack duration = 10 sec
- 4x4 sub-grid under attack
- # Attack segments = 16 (each is 5/8 sec)
- Prob. of success for each attack = 25%
- 256 node grid

"Realistic" Topology

- Generated using BRITE and ported into the SSF BGP simulation tool
- 200 nodes; minimum connectivity = 2, maximum connectivity = 8
- Plan to create such networks with hierarchy (access, metro, core)
- Introduce policy based routing (e.g., core can not route through access BGP routers)

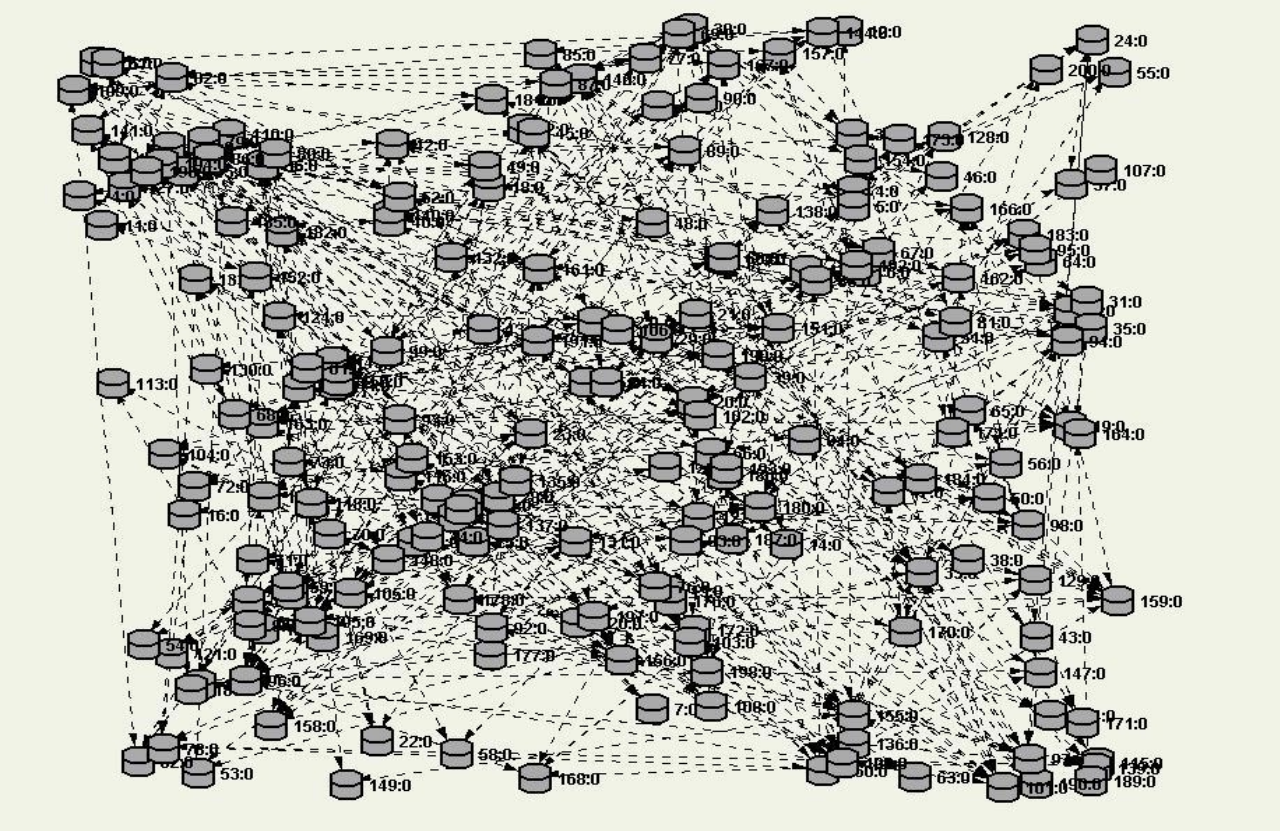

Conclusion on RFD Exploitation Attacks

- Attackers can exploit RFD behavior to cause extended AS isolation
- The attack rate need be no more than about one successful attack every few MRAI intervals
- With Graceful Restart (GR), the effort involved goes several orders of magnitude higher; so use of GR can add significant resiliency
- ISP's reluctant to enable GR?
  - "Several providers (US) suggest that the cost of implementing this feature outweighs the benefit." – NISCC (UK govt) BGP Best Practices
  - "Customers prefer to use an alternate route rather than GR because staleness of FIB issue with use of GR" – one source from an ISP says