Securing Grid Data Transfer Services with Active Network Portals

Onur Demir[1] **Michael R. Head**[2] Kanad Ghose[3]
Madhusudhan Govindaraju[4]

Department of Computer Science
Binghamton University (SUNY)
{onur[1], mike[2], ghose[3], mgovinda[4]}@cs.binghamton.edu

The 8th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing 2007

*This work supported by in part by the NSF under award numbers EIA 9911099 and CNS 0454298

BINGHAMTON UNIVERSITY
State University of New York

Outline

Outline

Requirements Driving Grid Data Transfer Services

- Higher bandwidth infrastructure
  - TCP and FTP as-is are unsuitable to connections with a high bandwidth delay product
  - Network bandwidth can outperform raw disk access
- Enormous data files
  - Output from high energy physics experiments
  - Large databases: protein sequencing databases, human genome project
- Need to authenticate and authorize in a globally scalable manner

GridFTP's Answers to Grid Data Requirements

GridFTP...

- Supports high bandwidth-delay-product infrastructure
  - Allows tunable TCP window sizes
  - Supports multiple parallel streams
- Supports striping to increase disk bandwidth
- Supports reliable resumption of canceled or dropped transfers
- Integrates with Globus's GSSAPI authentication

There are other solutions, but GridFTP is available everywhere

Outline

Threats to Grid Services

Working assumption: As Grid services become more broadly available, they will increasingly become targeted. . .

- (Distributed) Denial of Service
- “Flash crowding” (not strictly an attack)
  - Services should prepare for near instantaneous explosions in (legitimate) activity
- Abusive users (with greater or lesser degrees of competence/intent)
  - This is hard to solve without service-specific solutions
  - Also ultimately requires some heuristics to classify abusers

Hypothesis: Prioritizing requests for different file sizes can improve performance for classes of users while maintaining overall throughput, even under attack.

BINGHAMTON UNIVERSITY
State University of New York

Outline

Previous Work

- A number of solutions supporting QoS and differentiated service involve packet inspection in the server
  - We propose inspection in an intelligent router
- A number of solutions specifically for DDoS exist
  - We also help protect against load attacks
- Our own Grid 2005 work
  - This work involves testing new policies and new results
  - Specifically, differentiating service for different classes of requests

Outline

Architectural Outline

- Connections from clients to GridFTP servers are mediated by an Active NIC (programmable gateway)
- Gateway examines packets and performs destination NAT and balances connections across the GridFTP server pool
- When attack conditions are detected may implement different policies to drop packets from clients or distribute connections differently

Server Architecture

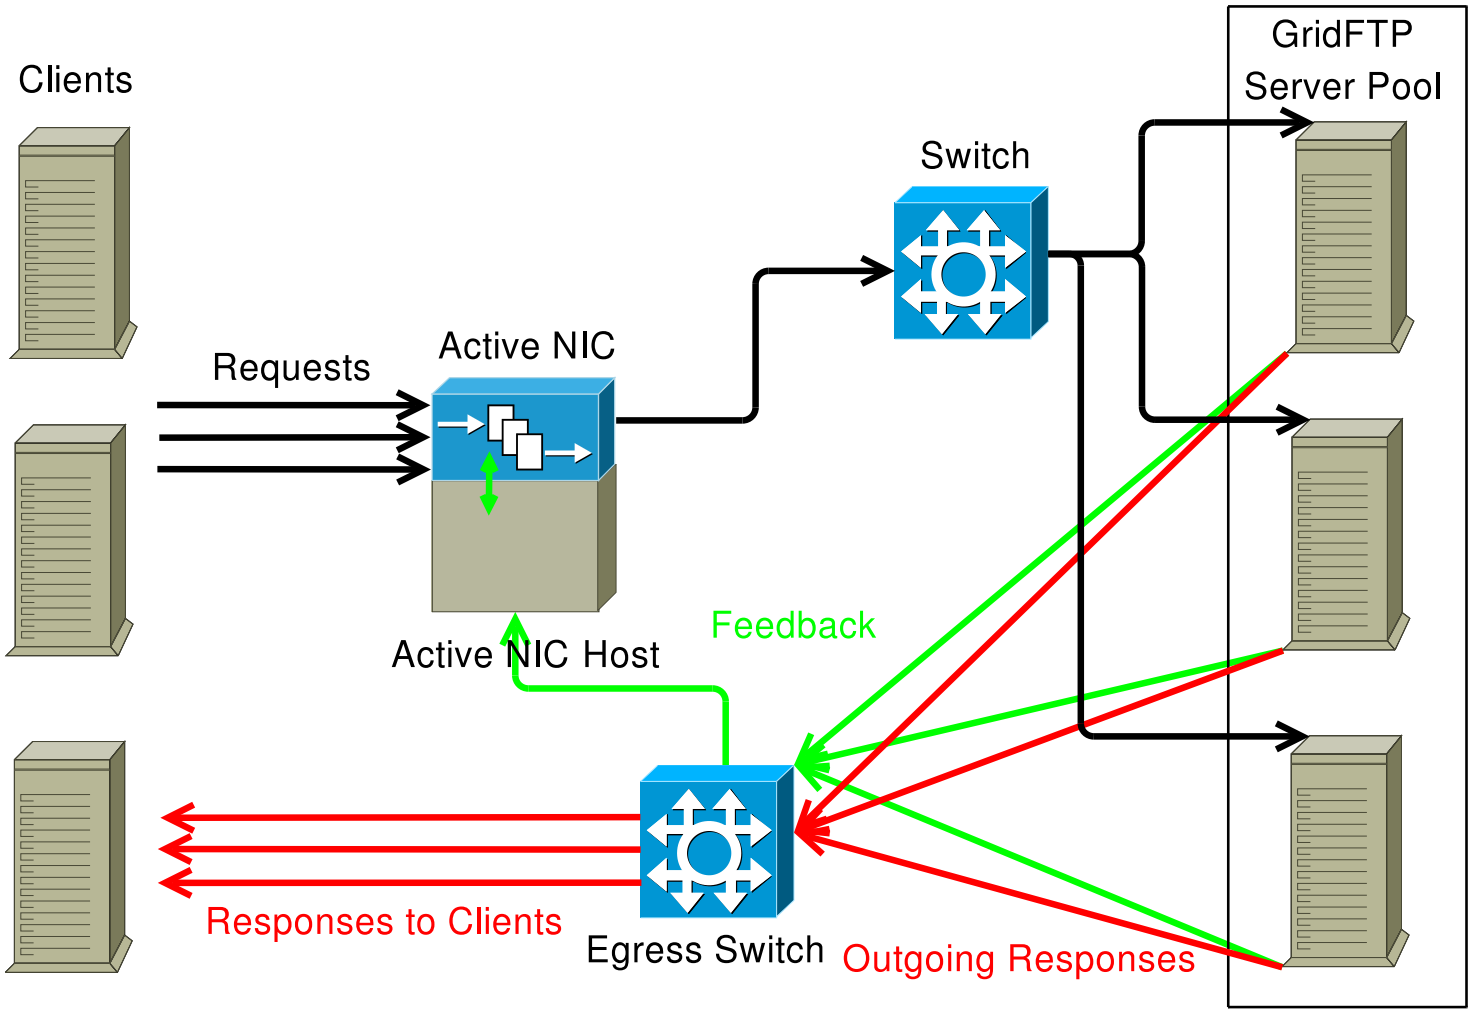

Processing Logic Inside the Gateway

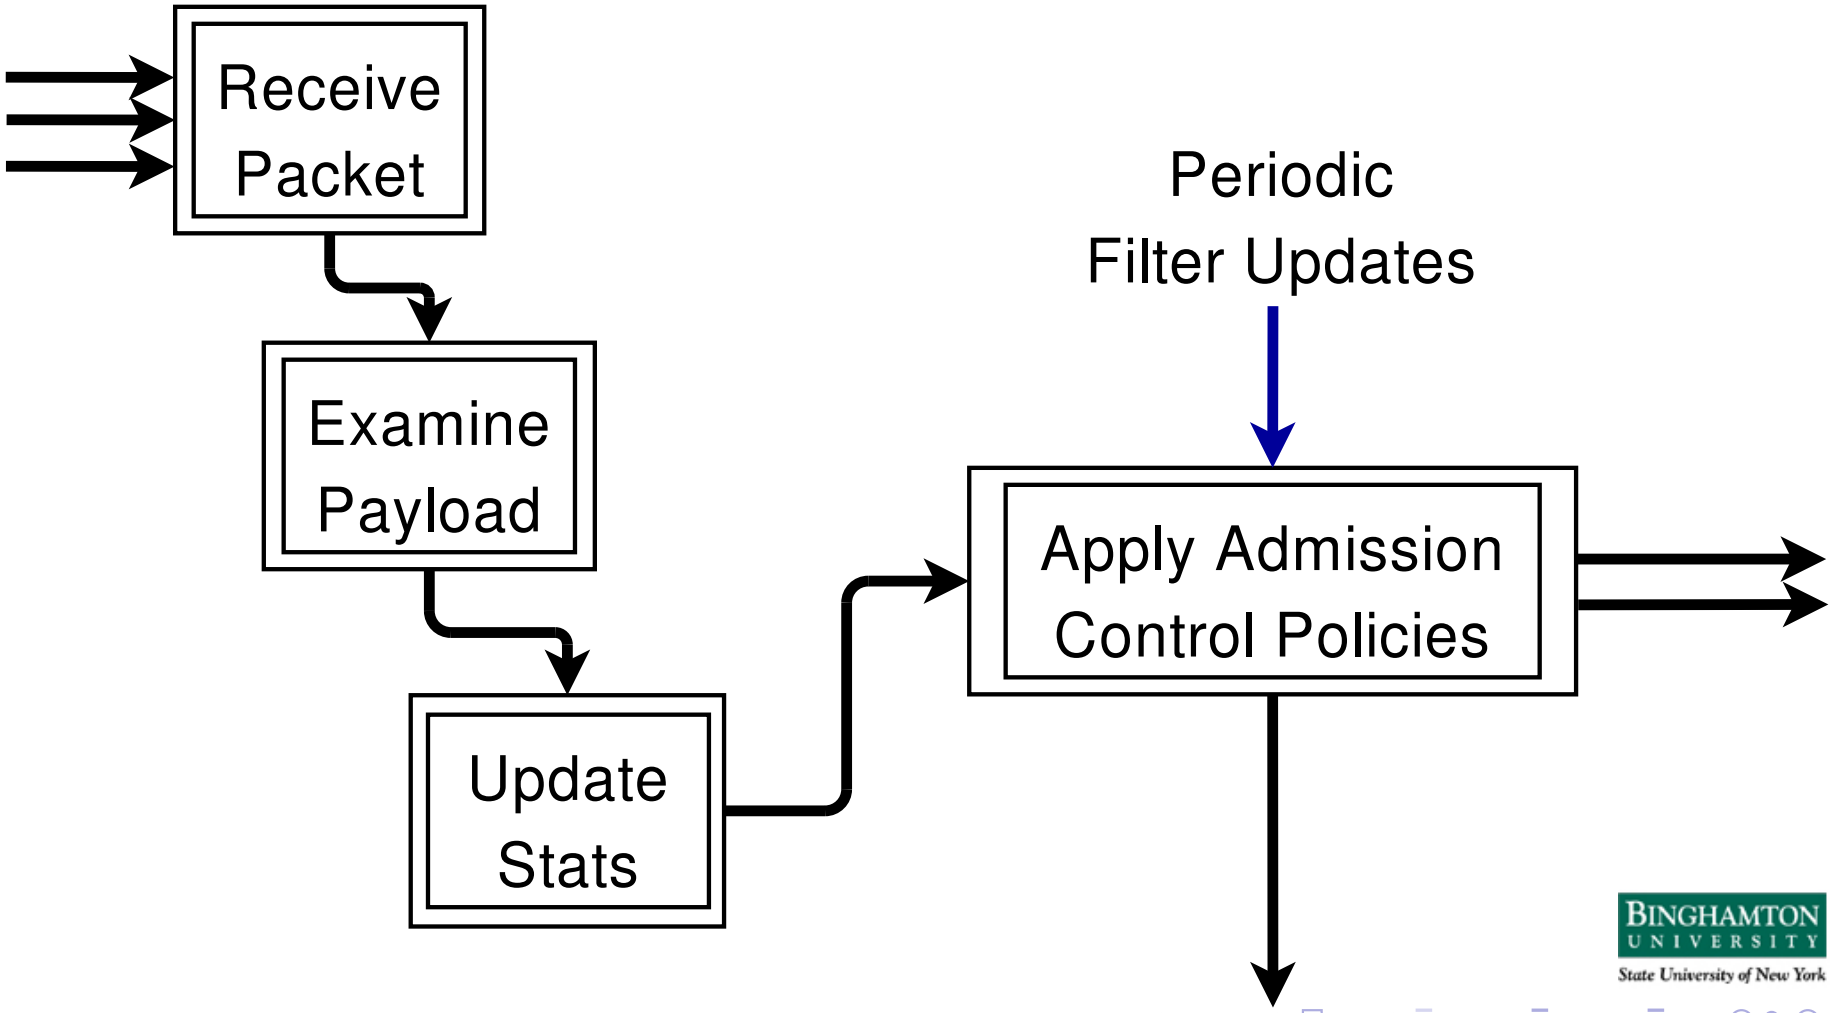

Policies Examined

We examined the results of privileging certain classes of requests during attack

- *Small Requests Favored*: Clients requesting 48KB size files have priority
- *Medium Requests Favored*: . . . 2MB . . .
- *Large Requests Favored*: . . . 64MB . . .

Note: $(64MB * 10^6)/(10^8 bps/8)$

$=5.120s$ to transfer $64MB$ over a 100Megabit link

Implications

- It is possible to favor a file class by modifying thresholds.
- Depending on the file type character of the server, completion rate of transfer can be improved.
  - Small files can be favored for servers that hold source code, small images
  - Large files can be favored for servers that hold multimedia files, large data files, . . .
- The system itself is dynamic and allows custom policies.
- The server feedback support provides implementation of new policies based on other application level criteria.

Outline

Experimental Design

Five cases . . .

1. Base (“NORMAL”) case
2. Attack, no policy
3. Attack, small favored
4. Attack, medium favored
5. Attack, large favored

Note: plotting the results of running a completely unprotected server pool is uninteresting

Test Scripts

- Each script instance repeatedly requests file several time (using `globus-url-copy`)
  - Uses “Extended Block Mode” and four parallel streams
- Scripts requesting a given file size all run on a specific client machine
- 20-50 script instances per file size class

Outline

Effect of Policies on Large File Class

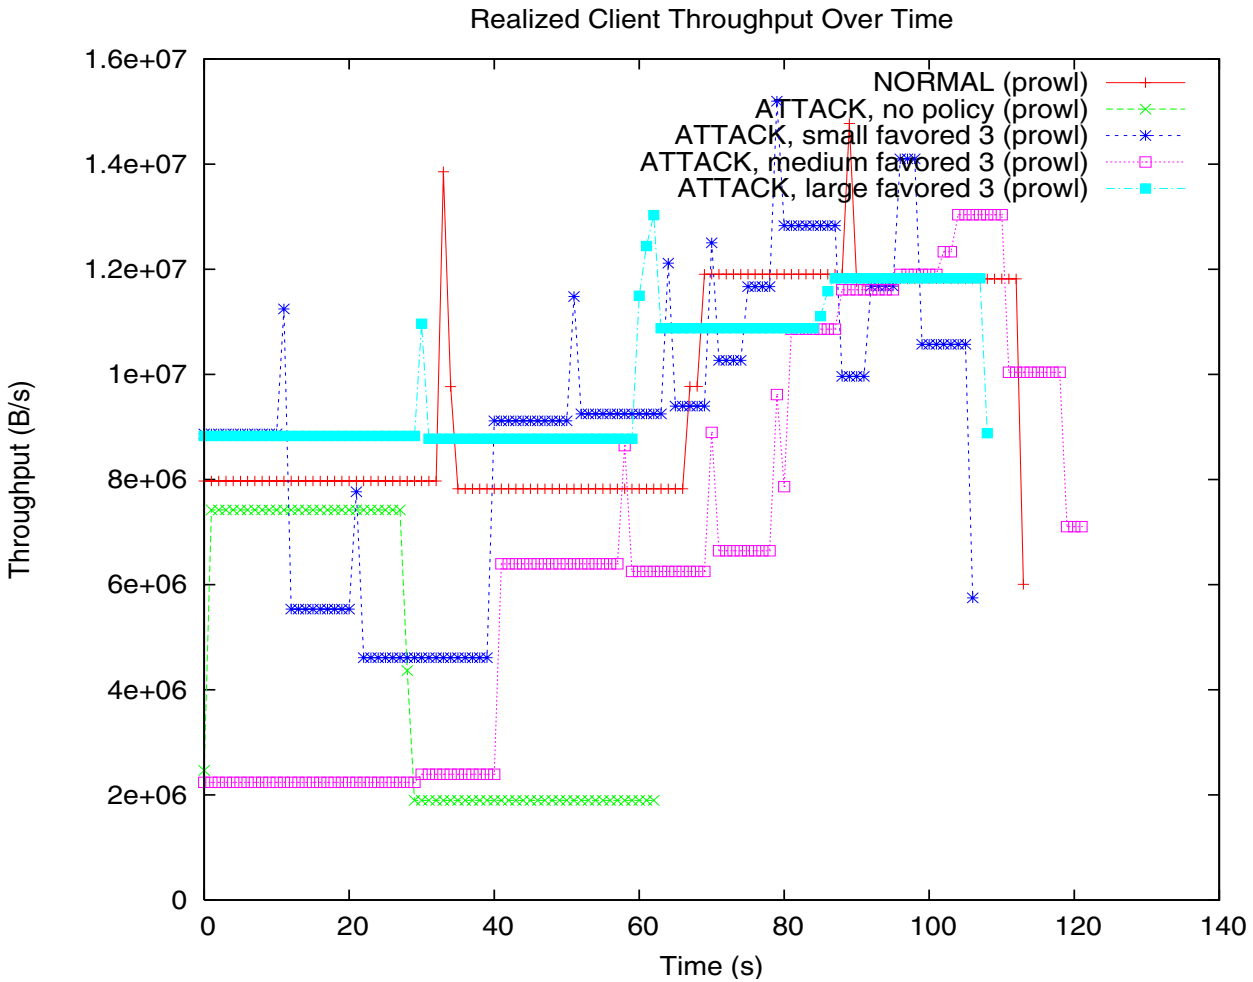

Effect of Policies on Small File Clients

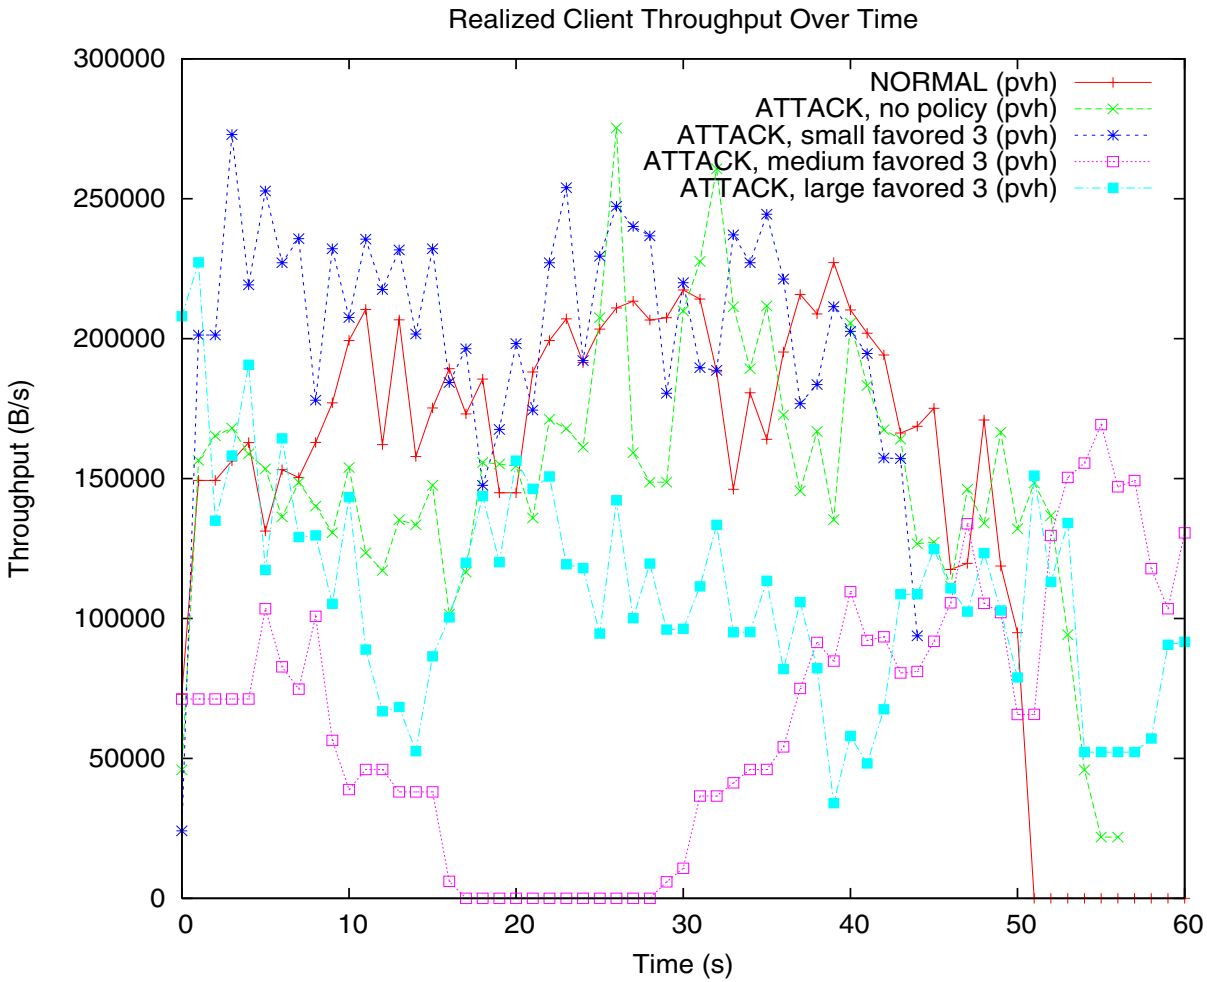

Overall Server Throughput

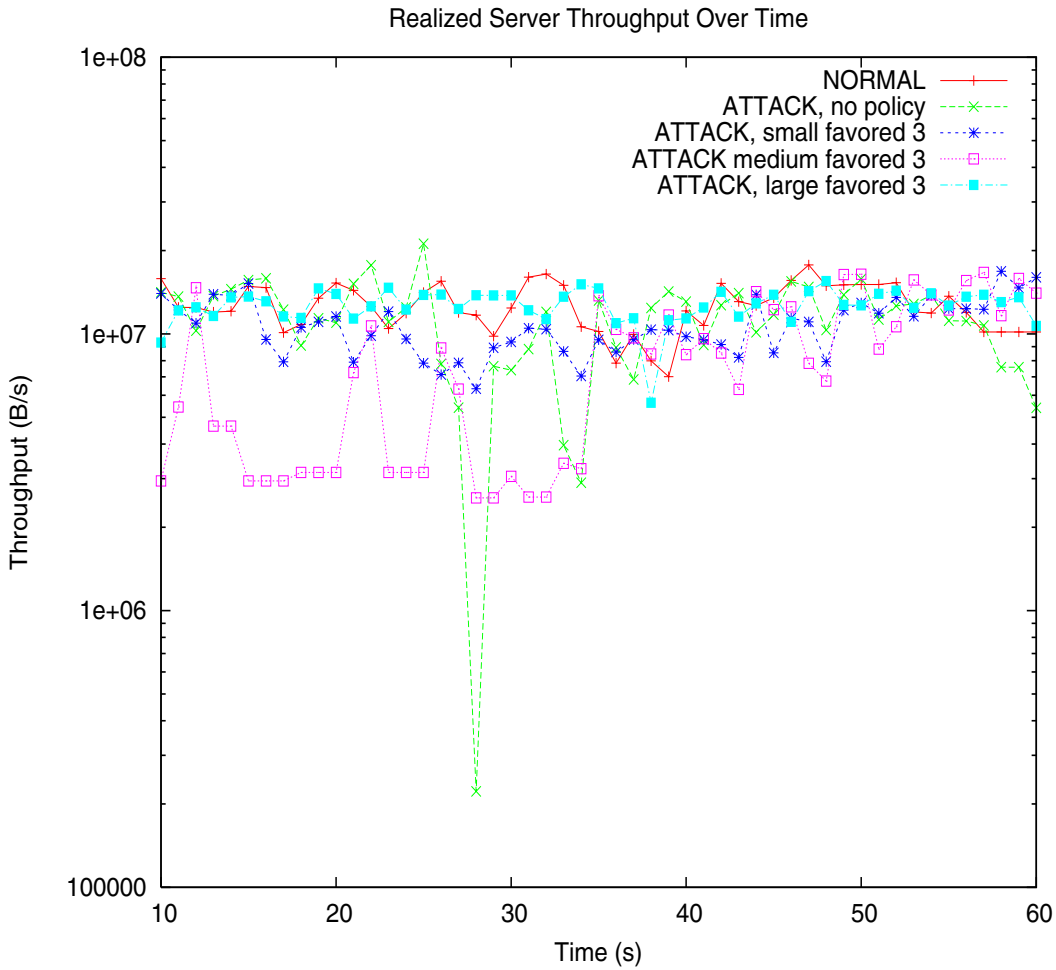

Limitations

- Currently uses static knowledge about the requests to determine the size of the file associated with the request
  - May need to decrypt the control stream in the Active NIC (expensive operation)
- Data about ongoing connections are from GridFTP logs
  - May need to write a GridFTP module or otherwise modify GridFTP to provide more detailed connection information
- Constants for the quota multiplier and window size were determined experimentally
- Only shapes incoming packets – outgoing packets (and GridFTP data connections) go through separate egress switch

Summary

- Adaptive traffic management improves server throughput
- Active NIC based gateway serves as an unobtrusive mechanism for classifying requests and shaping incoming traffic
- Implements fast response to attacks
- Careful choice of thresholds essential
  - Identify through experimentation