A Case Study in Internet Pathology: Flawed Routers Flood University's Network

LISA '03, San Diego, October 29, 2003

http://net.doit.wisc.edu/~plonka/lisa/lisa2003/

THE UNIVERSITY of WISCONSIN MADISON

Dave Plonka
<plonka@doit.wisc.edu>
Division of Information Technology (DoIT) &
Wisconsin Advanced Internet Lab (WAIL)

Agenda

- Preface: some Definitions
- The Initial Flood, Background, & Investigation
- The Review Process
- Current Status
- Consider Possible Operational Strategies
- The Plan
- Recommendations for the Community
- Afterthoughts
- Acknowledgements

Preface: Some Definitions

- *Internet* Engineering:
  - to contrive or plan out usually with more or less subtle skill and craft ***with the goal of obviating failure***
- *Internet* Pathology:
  - the study of the essential nature of ~~disease~~ ***Internet anomalies*** and especially of the structural and functional changes produced by them
- based on definitions from Merriam-Webster Dictionary & Henry Petroski

The "Initial" Flood

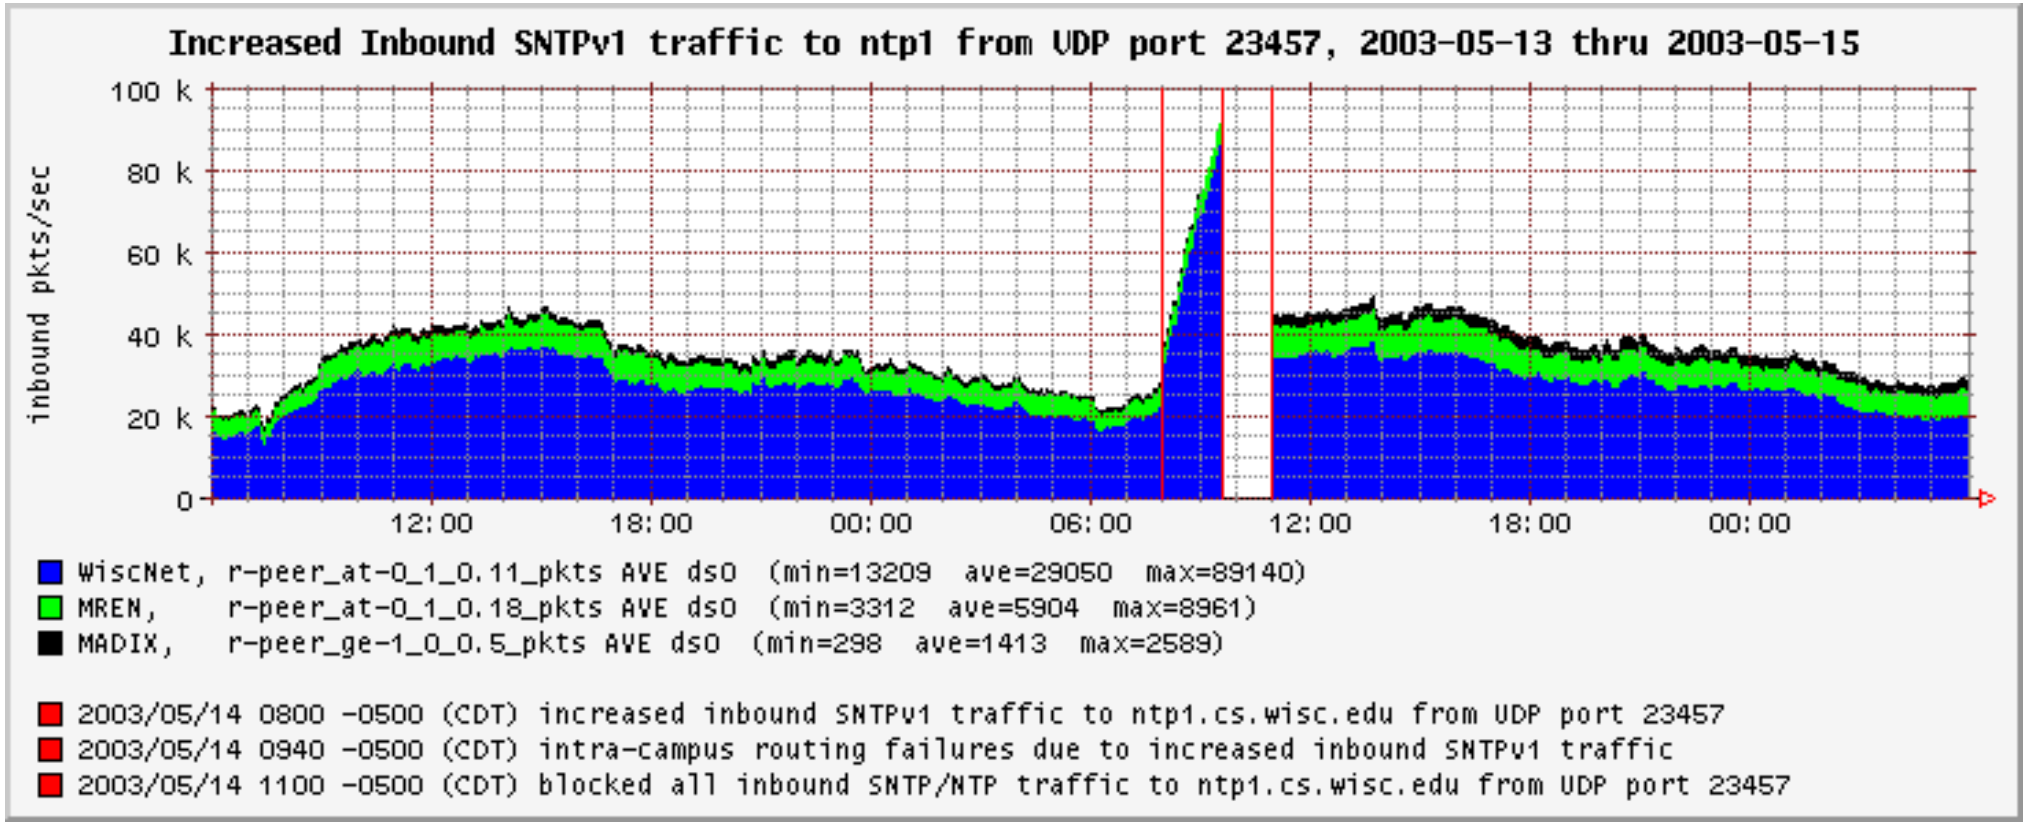

Simple Network Time Protocol (SNTP)

- ***Simple***, stateless Remote Procedure Call
- Expectations similar to the UDP TIME protocol
- For use when full performance of NTP is not necessary
- Defined by *Informational* RFC 2030, c. 1996
- Utilizes same packet format as NTP defined by RFC 1305, *DRAFT STANDARD*, c. 1992
- Designed so that SNTP clients can use NTP servers

A SNTP Request Packet

LI =0	VN = 1–4	Mode = 3	Stratum = 0	Poll = 0	Precision = 0
Root Delay = 0					
Root Dispersion = 0					
Reference Identifier = 0					
Reference Timestamp = 0					
Originate Timestamp = 0					
Receive Timestamp = 0					
Transmit Timestamp = n (some number, zero, or the time of the request sent by client)					

A Unicast SNTP Reply Packet

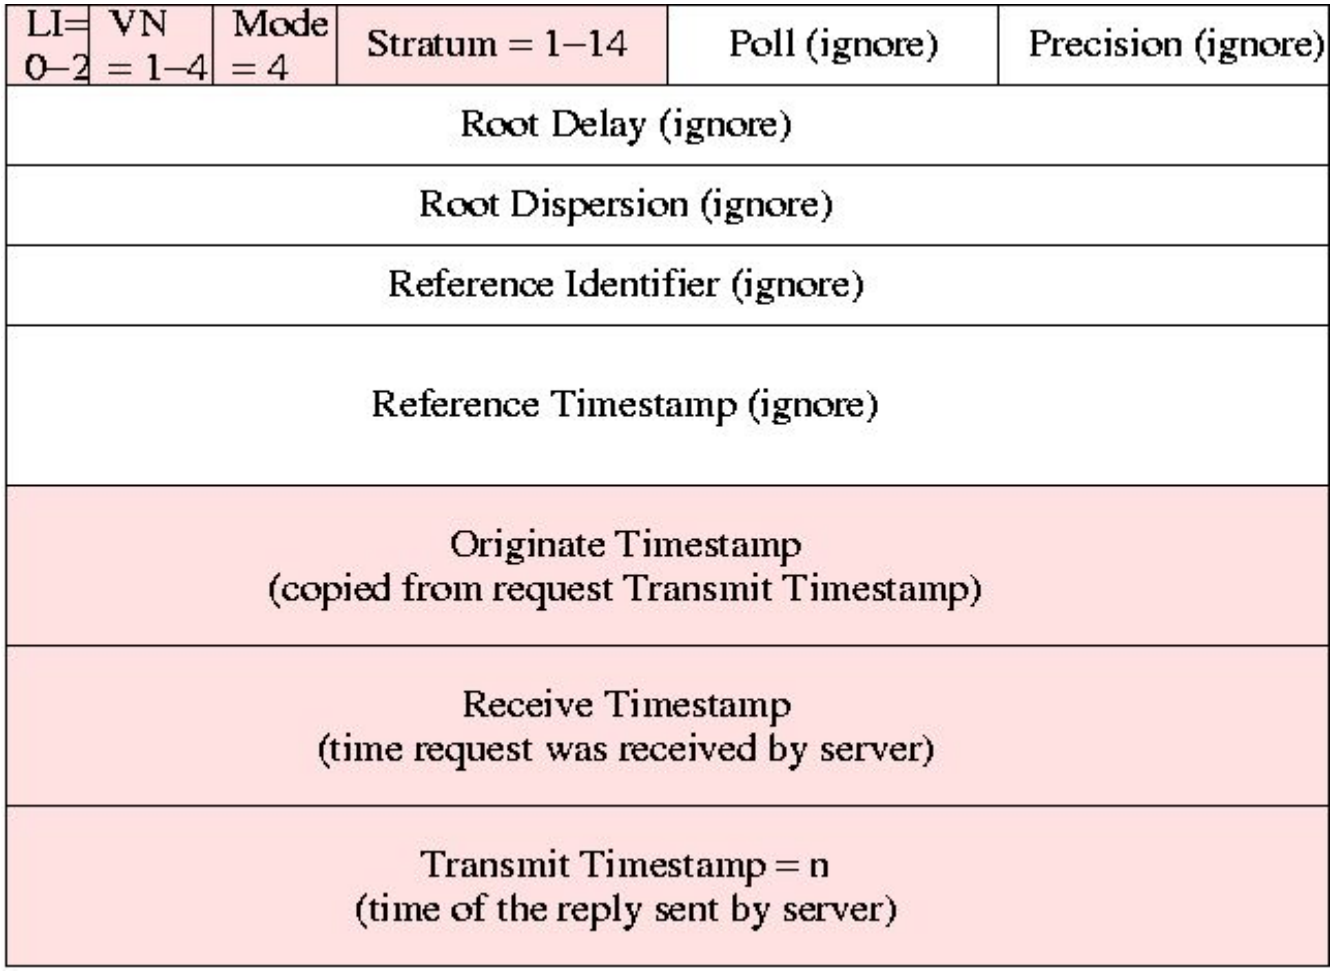

The Flood Continues: One Month Later

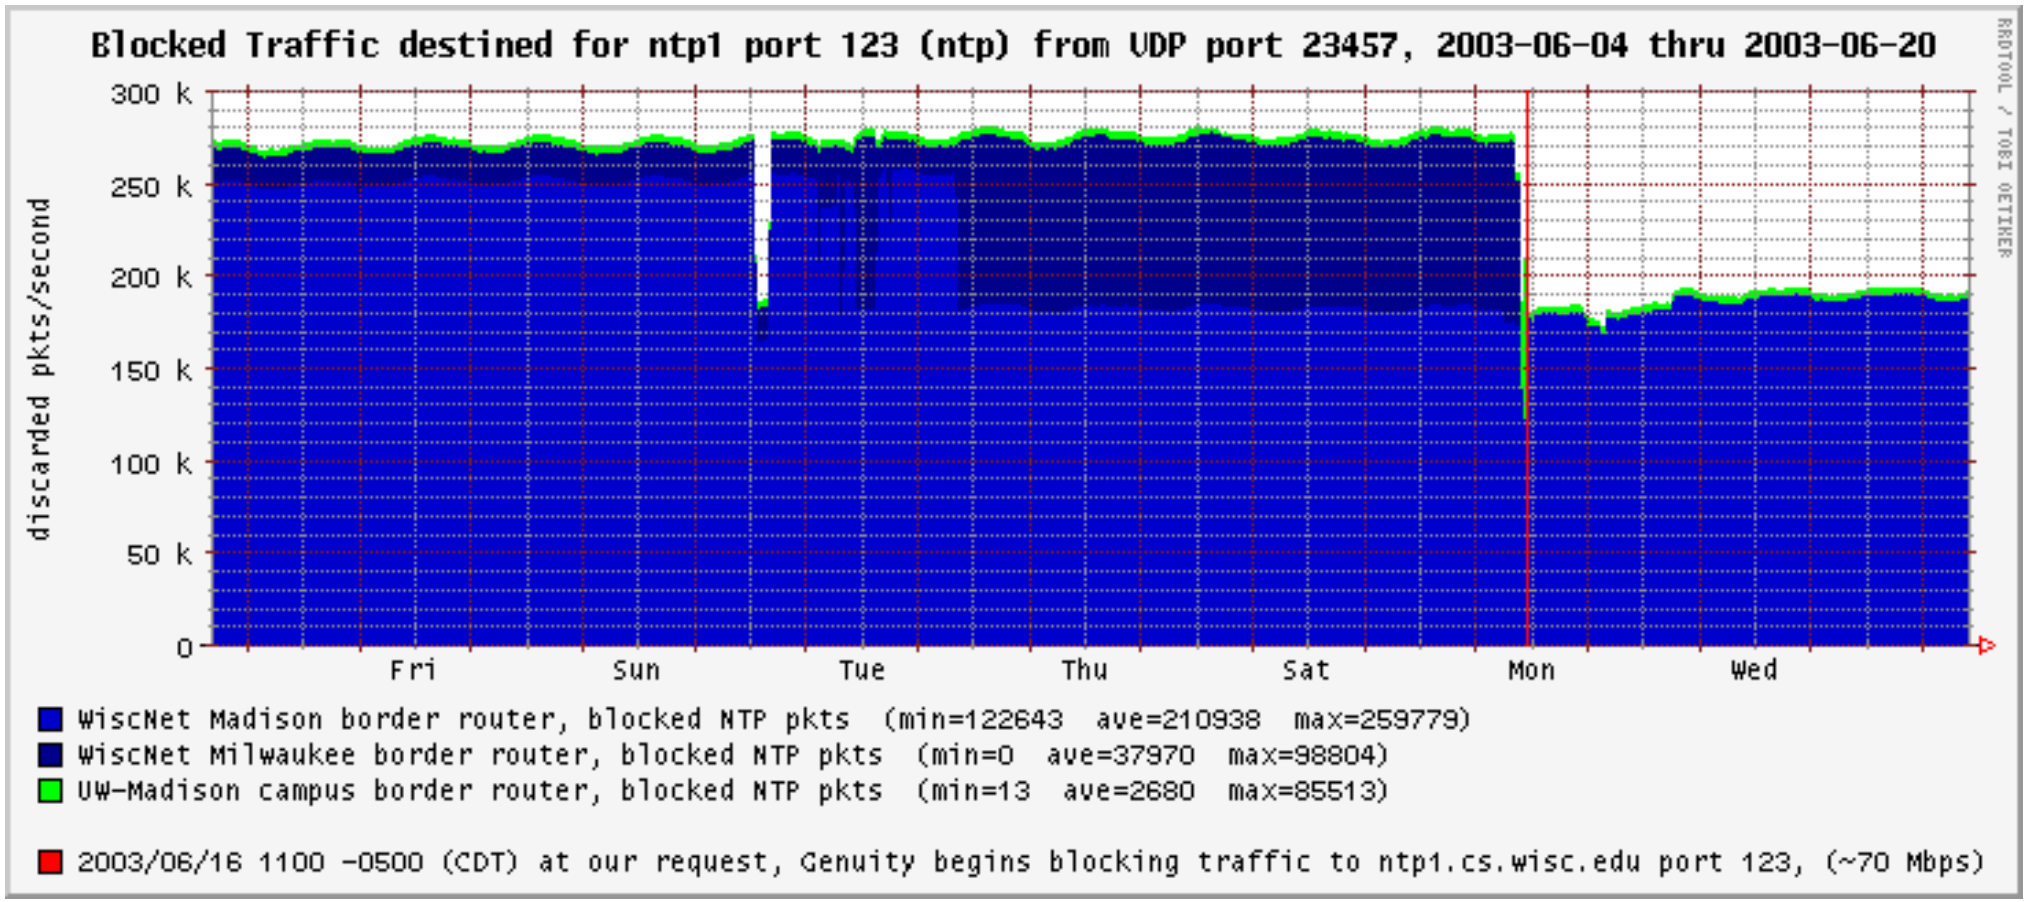

Investigation: Contacting Source Networks

- Selected two SNTP "top talkers" from a packet trace, such that both were from other Universities likely to have experienced staff familiar with responding to network abuse incidents
- On Saturday, June 14, sent email missive to abuse@{ufl.edu,harvard.edu}
- Included IP flow records and tethereal output (packet decomposition)

Investigation:
Contacting Source Networks

- Within hours I received initial responses...
- By Monday, both sources had been identified as Netgear brand products, one a model MR814 - a broadband router with wireless LAN side.
- Eureka! Many sources all using the same source port could be due to an embedded SNTP client in which the programmer hard-coded the source port number of 23457.

Investigation: Gather Background Info

- Web searches yielded almost nothing...
- Except: an ICSA Labs report on an unrelated Netgear product which was curious in that it did not have a battery-backed clock, and therefore was very reliant on a hard-coded NTP server.
- As an aside, that product - the FR114P - still met the certification criteria.

Investigation: Examining the Netgear Code

- Downloaded the code for Netgear's "Platinum" products including RP614 and MR814
- The "strings" command is your friend.
  - Verified that the code referred to 23457 as a port number
  - Found numerous embedded IP addresses including that of ntp1.cs.wisc.edu: 128.105.39.11

Examining the Netgear Code

```
$ strings RP614_4_12.bin |grep 23457
 on 23457 port.
$ strings MR814_4_11.bin |grep 23457
 on 23457 port.
$
```

Examining the Netgear Code

```
$ strings RP614_4_12.bin |perl ...
128.105.39.11 # ntp1.cs.wisc.edu
192.168.1.101
66.37.215.43
12.234.94.14
192.168.0.1
```

Examining the Netgear Code

```
$ strings MR814_4_11.bin |perl ...
0.0.0.0
12.234.94.142
66.37.215.43
92.168.0.102
128.105.39.11 # ntp1.cs.wisc.edu
192.168.0.1
192.168.1.101
```

Contacting Netgear

- On Monday, June 16, sent email missive to support@netgear.com, *Subject: NETGEAR products abusing University of Wisconsin time server*.
- Received an automated response as expected.
- I also sent same to various netgear.com email addresses culled from the web, asking the recipient to communicate it to engineering and have someone contact me by phone or email ASAP.

Contacting Netgear

- On Tuesday, June 17, called Netgear's toll free support and, after 15 minutes of waiting, left a voicemail asking for them to return my call.
- Emailed the contact for the domain netgear.com.
  - postmaster@netgear.com promptly replied that the address is undeliverable.
- Called Netgear's HQ, asked for VP of Engineering. After a short explanation, I was transferred to voicemail where I again explained and supplied contact information.

Contacting Netgear

- On Wednesday, June 18, emailed five members of Netgear's executive team by guessing their email addresses based on company convention found by searching groups.google.com.
  - I was notified that all were delivered since I had added a Return-Receipt-To: header.
  - I wrote, "*I absolutely need to hear from responsible parties at NETGEAR immediately, if NETGEAR wishes to begin a dialogue before this goes public. We're not expecting an immediate solution; in fact, I'm fairly certain there is no complete solution without UW-Madison's involvement.*"

Contacting Netgear

- On Thursday, June 19, I received a voicemail message from the director of support for Netgear. He confirmed that they have located "*some fault in [their] code*."
- As an aside, 23 days after the original email to support, I received a canned response:
  - "Thank you for your email. We apologize for the delay in responding."
  - "Your issue may have already been resolved. Please reply to this email if you still require assistance and we will respond as quickly as we can."

The Review Process

- A review team was formed with about fifteen members:
  - Netgear Employees
    - Engineers
    - Management, VP
  - University Employees
    - DoIT Network Services
    - Computer Sciences CSL
  - Independent experts from their respective fields:
    - Regional Internet Registries
    - Internet Measurement Research
    - Network Time Protocol

The Flawed Netgear SNTP Client

- Uses a *hard-coded* IP address for the NTP server 128.105.39.11
- Uses a fixed UDP source port number 23457
  - Incredibly advantageous - can identify most Netgear clients
  - However Network Address Port Translation (NAPT) messes this up
- Polls at *one second* intervals until it receives a response. Polls at various intervals thereafter: one minute, ten minutes, two hours, or 24 hours, depending on firmware version

Impact to Netgear Customers

- Currently, customers should not notice due to this flaw because UW-Madison is servicing the time requests on a best-effort basis.
- Customers that do not make use of the time-related features are unlikely to notice any problem:
  - logging
  - firewall policy scheduling
  - email notifications
- For most of the affected products, customers can upgrade to newer firmware versions which do not abuse UW-Madison's time server.

Affected Netgear Products

- RP614, RP614v2:
  - 4-Port Cable/DSL Router with 10/100 Mbps Switch
  - CNET Editors' Choice, July 2002, ZDNet Editors' Choice
  - High Production Volume, Currently Produced
  - Firmware fix released ~July 11 for RP614v2

Affected Netgear Products

- MR814:
  - 802.11b Cable/DSL Wireless Router
  - Innovations International CES, Design & Engineering Showcase Honors, 2003, CNET & ZDNet Editors' Choice, November 2002
  - High Production Volume

Affected Netgear Products

- DG814:
  - DSL Modem Internet Gateway
  - Macworld Editors' Choice, MacFormat Editor's Choice, October 2002, Internet Magazine Best Buy, Sep/Oct 2002, PC Answers Platinum award, September 2002

Affected Netgear Products

- HR314
  - 802.11a Cable/DSL High-Speed Wireless Router
  - no known awards ;^)

The Initial Fix: "Instant" Code

- Released in July, August, September (varies by model)
- Apparently in the works before notification of the problem by UW-Madison.
- Now requires a DNS server to be configured or learned via DHCP.
- Resolved "time-a" and "time-b.netgear.com", alternately at ten minute intervals until success.
- Sends an SNTP query to the resulting IP address at ten minute intervals until success or five retries
- Any configuration change seems to cause the clock to be zeroed and this process repeated.

The Initial Fix: "Instant" Code Bugs

- Does not appear to validated the NTP response payload at all - it will accept any UDP payload delivered to port 23457 as a valid response, for instance even if that packet is a misdirected request rather than response.
- While the SNTP client is awaiting a response from time-a or time-b, it will accept any UDP packet even if the source IP address is not that of the server that it queried.

Engineering Flaw #1

- SNTP implementation uses a *hard-coded* IP address for the NTP server 128.105.39.11 (known as ntp1.cs.wisc.edu)
- Violates intention of "INTERNET REGISTRY IP ALLOCATION GUIDELINES", RFC 2050, c. 1996
- *Internet* products rather than *Network* products?

Engineering Flaw #2

- SNTP implementation violates NTP Rules of Engagement / netiquette for public NTP servers:
- ntp1.cs.wisc.edu' listing says "IP adress may change, please use DNS"
- NTP Community asks that server administrator be notified, *even for public servers* - an email address for our NTP administrator is provided

Engineering Flaw #3

- SNTP implementation polls at *one second* intervals until it receives a response
- Clients should not increase their query rate to a service that seems unresponsive; it leads to collapse - a lesson learned some time ago.

Engineering Flaw #4

- Communication channel to engineering is inneffectual:
- Communication, especially regarding failures, is imperative for successful engineering
- There is a unacceptable mismatch between deployed hosts' impact and support organization's ability.

Current Status

- Netgear has cooperated with us on the initial steps of this process and we are forging an agreement that will enable us to implement a suitable solution.
- Netgear is intending to support the development and operation of what we deem to be the best approach for the long term.

Current Status: Occasional Floods

- UW-Madison continues to service Netgear SNTP requests in spite of receiving occasional large-scale floods of traffic from Netgear products.

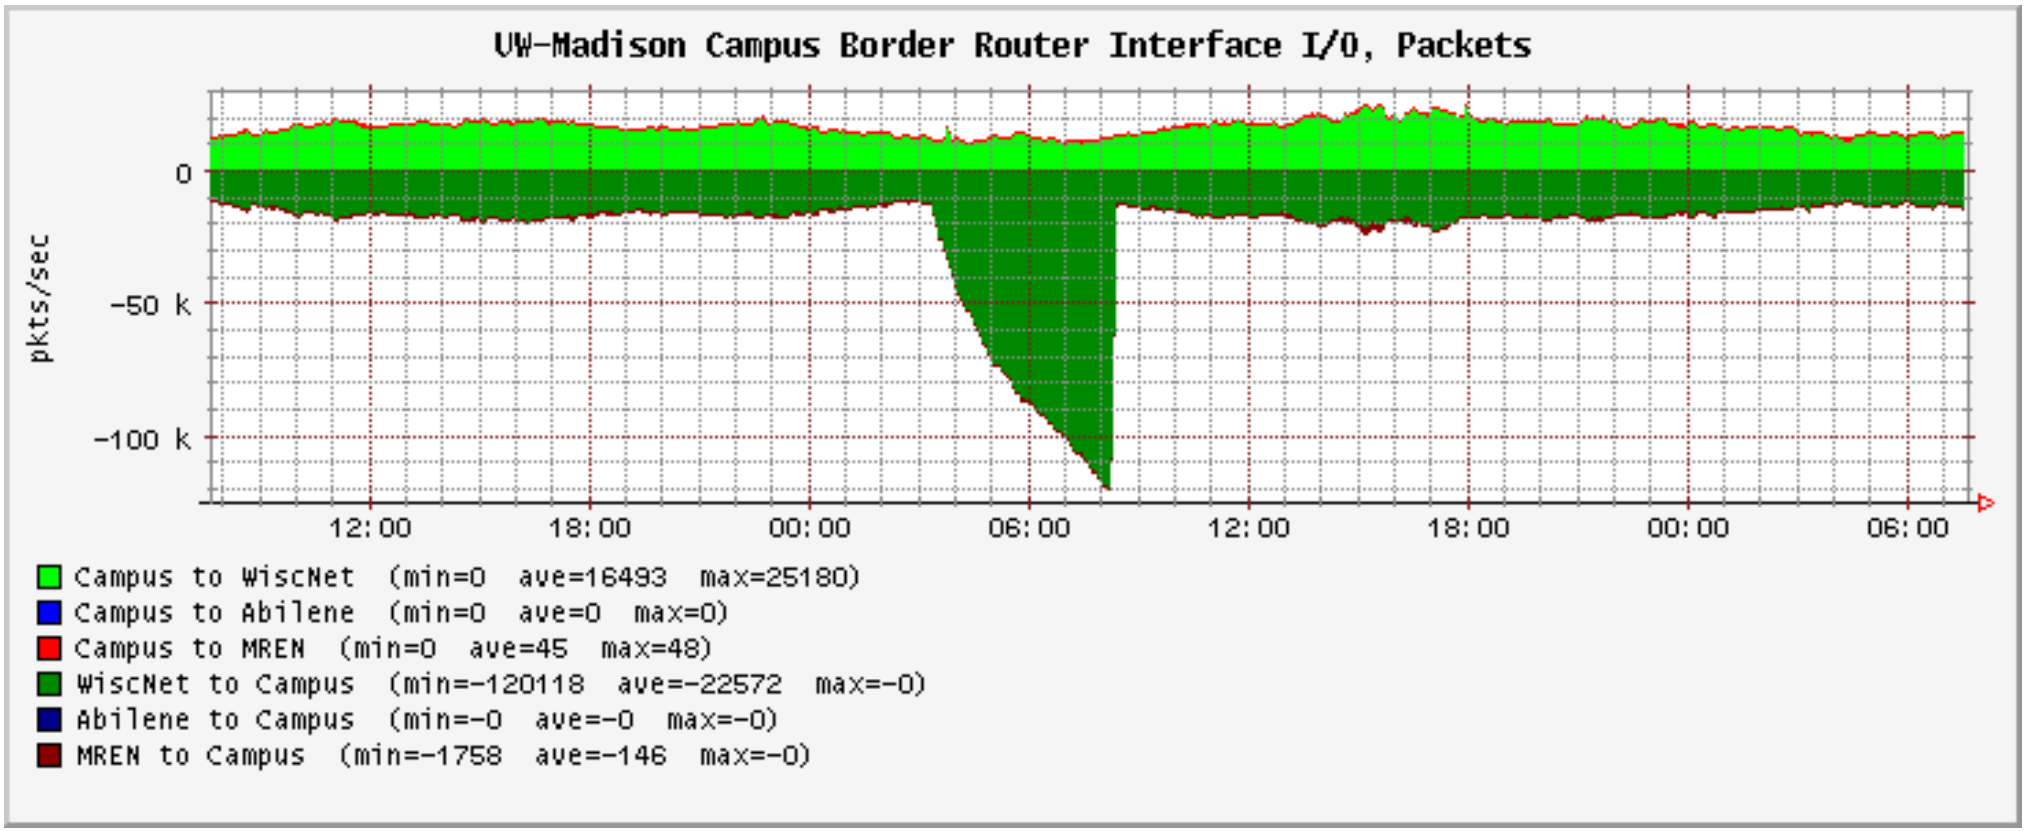

Current Status: Abusive Client Count

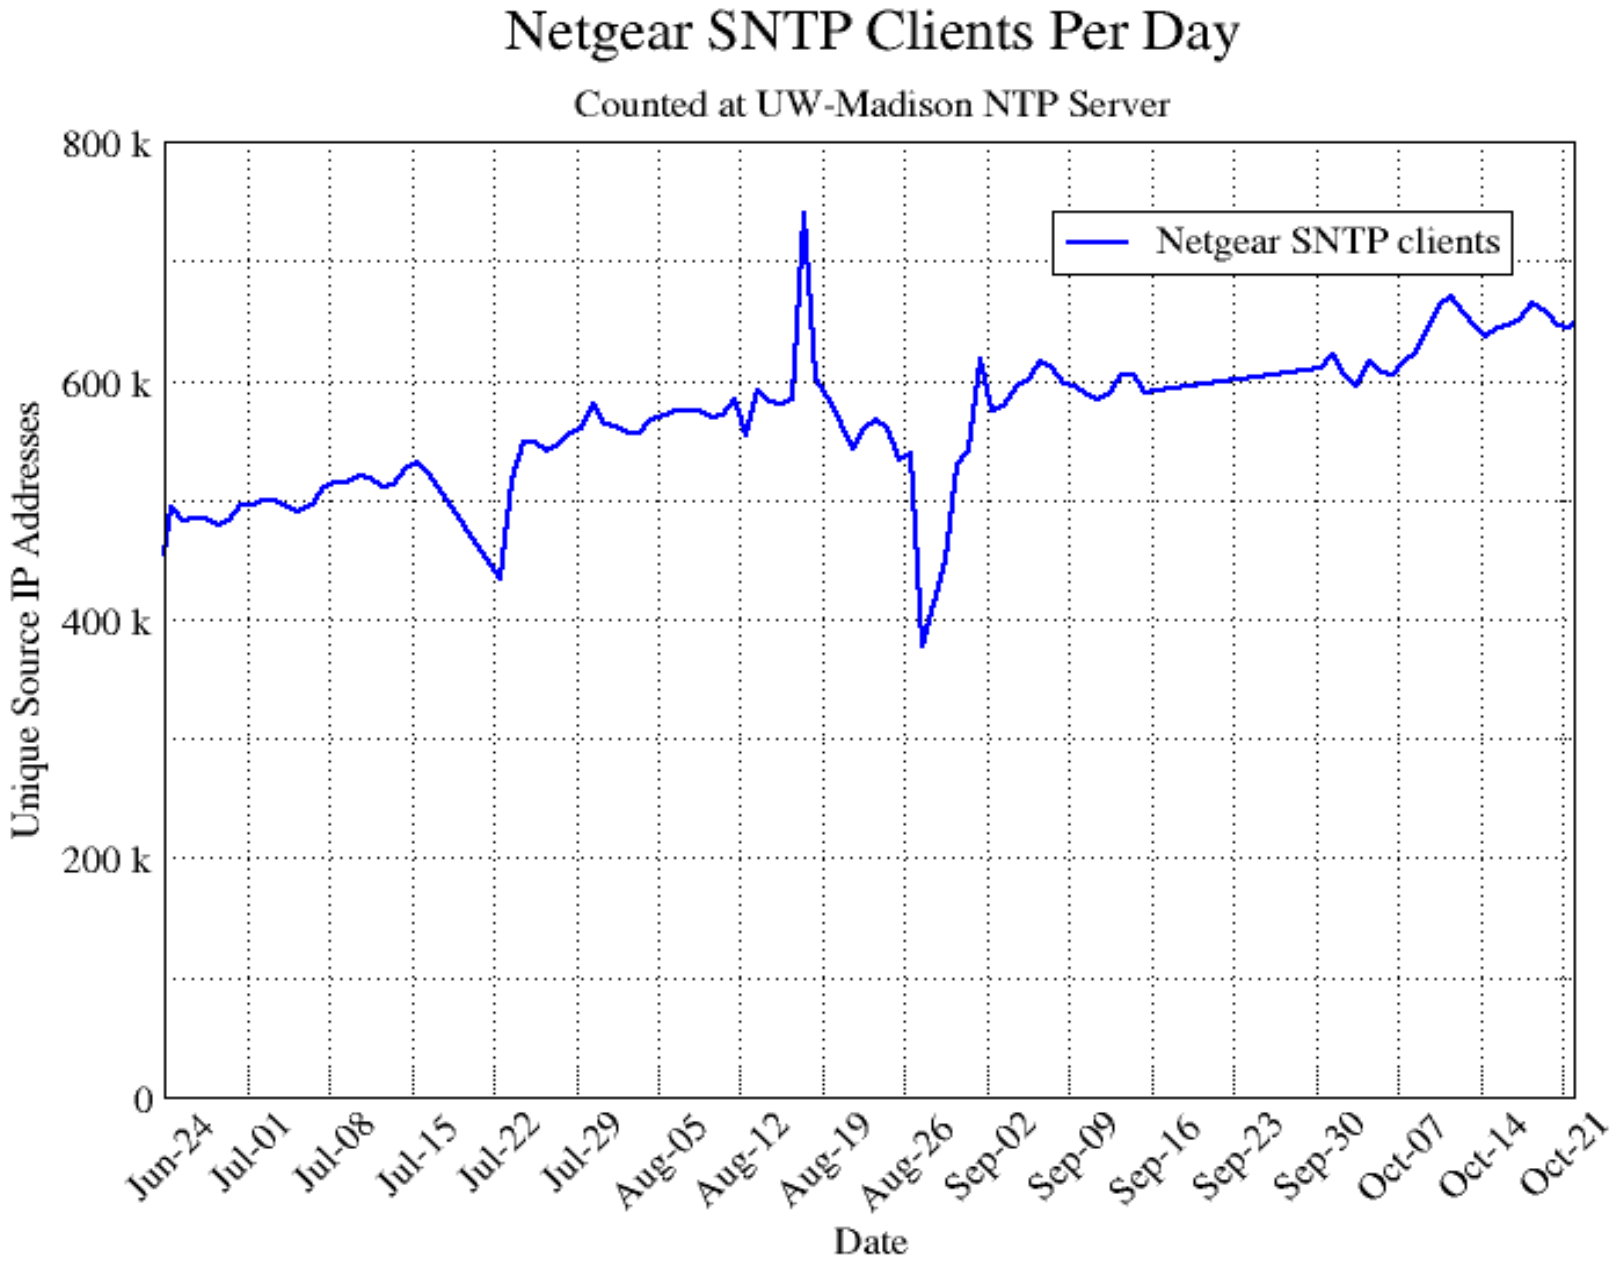

Current Status: Source Networks

- The flawed sources are within over 2400 Autonomous Systems.
- For example, on October 5, 2003:
  - Deutsche Telecom AG (3320): 76,205 clients (13%)
  - KIX (4766): 46,531 clients (8%)
  - SBIS-AS (7132): 42,172 clients (7%)
  - ATT-INTERNET4 (7018): 32,387 clients (5%)
  - DNEO-OSP1 (22909): 27,691 clients (5%)
  - others, such as U.S. broadband ISPs w/many ASNs

Network Operations Options: *To Serve or to Sever?*

- The Flawed products SNTP behavior is not easily reconfigurable.
- Both UW-Madison and Netgear believe it is not a viable strategy to rely on customers to upgrade to newer firmware.
- We've developed a multiphase plan which would likely result in one of two endgames:
  - Endgame A: Serve - Anycast Time Service
  - Endgame B: Sever - Suppress the Requests

Endgame A: Anycast Time Service

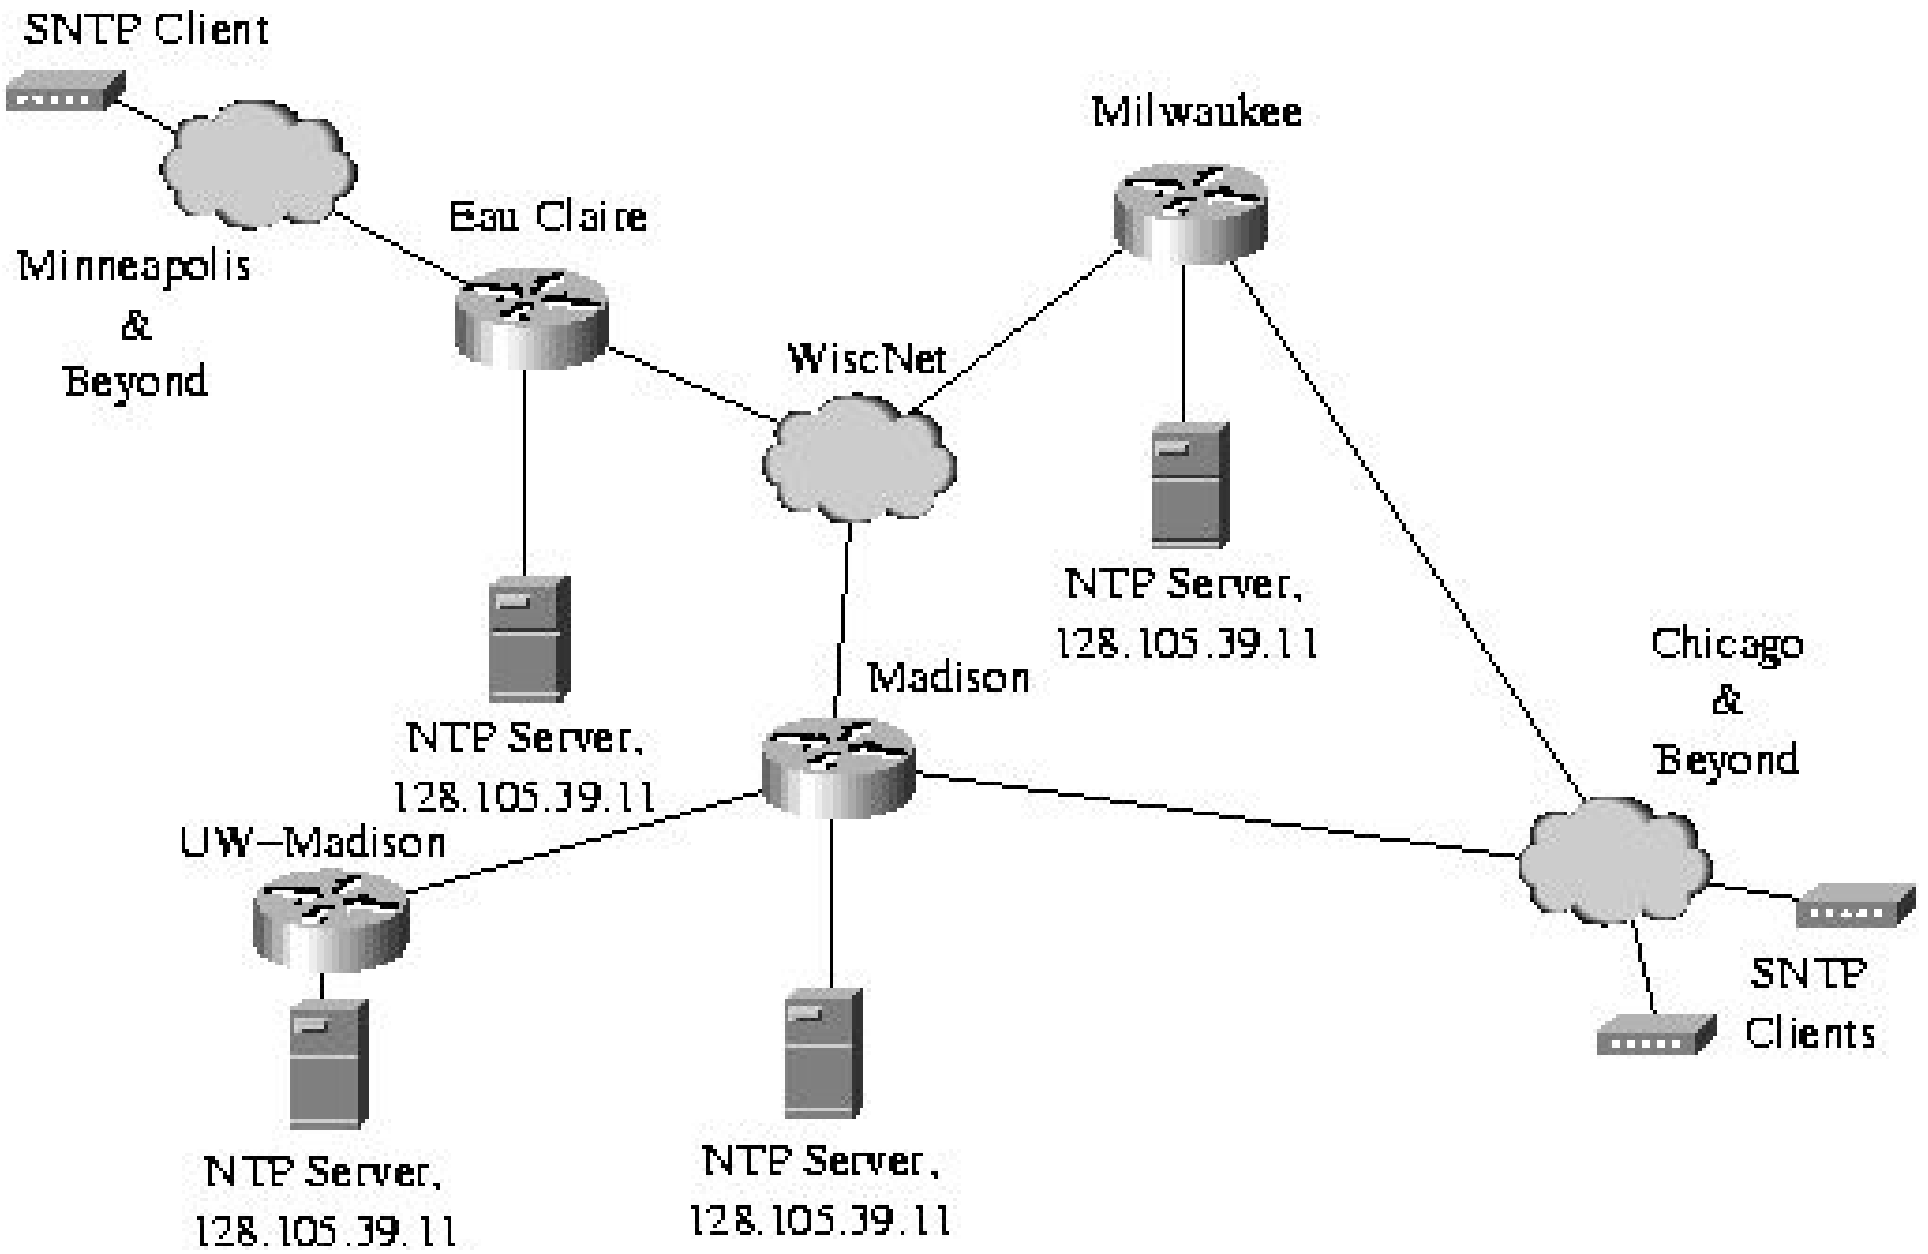

Endgame A: Requests

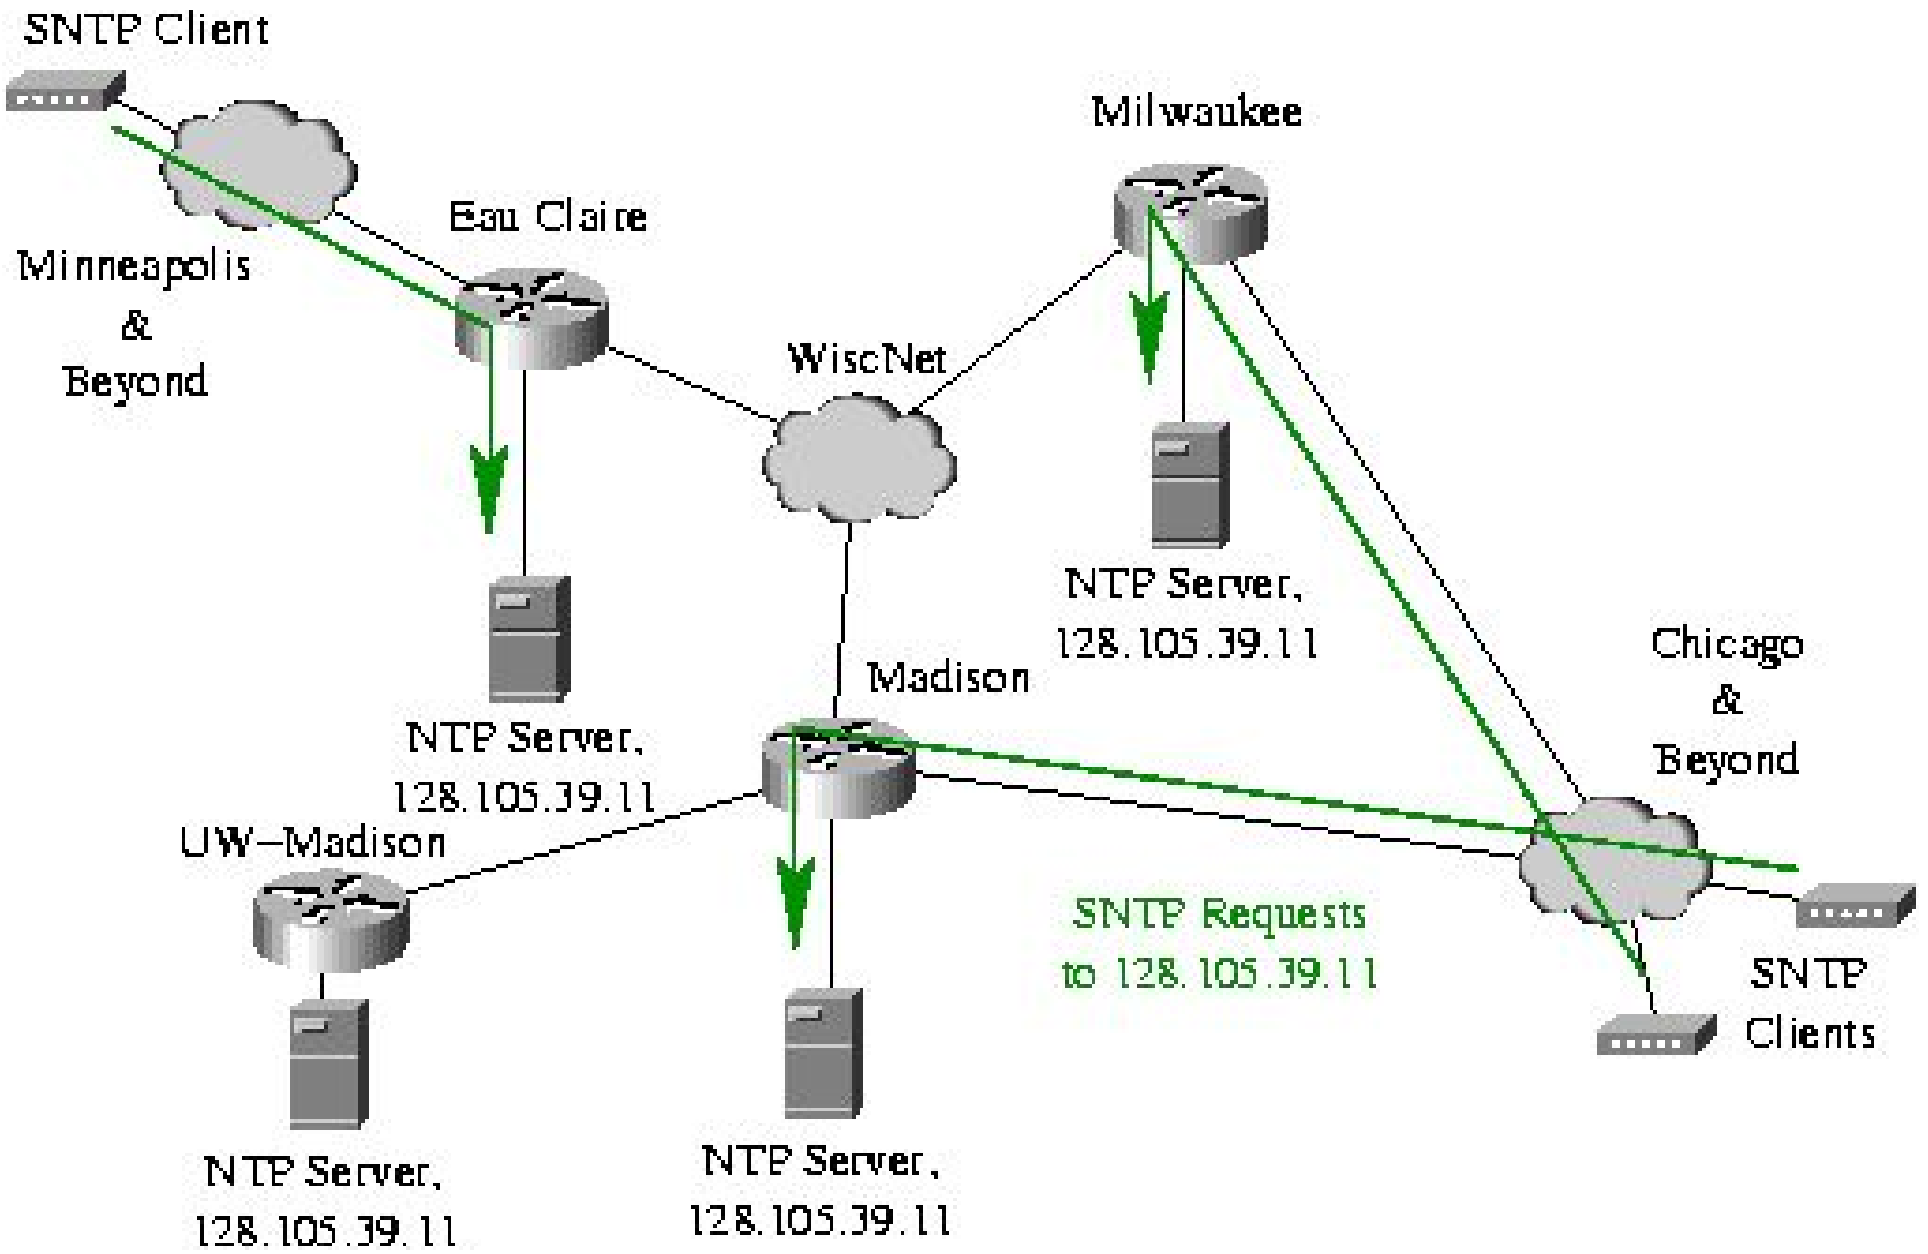

Endgame A: Anycast Responses

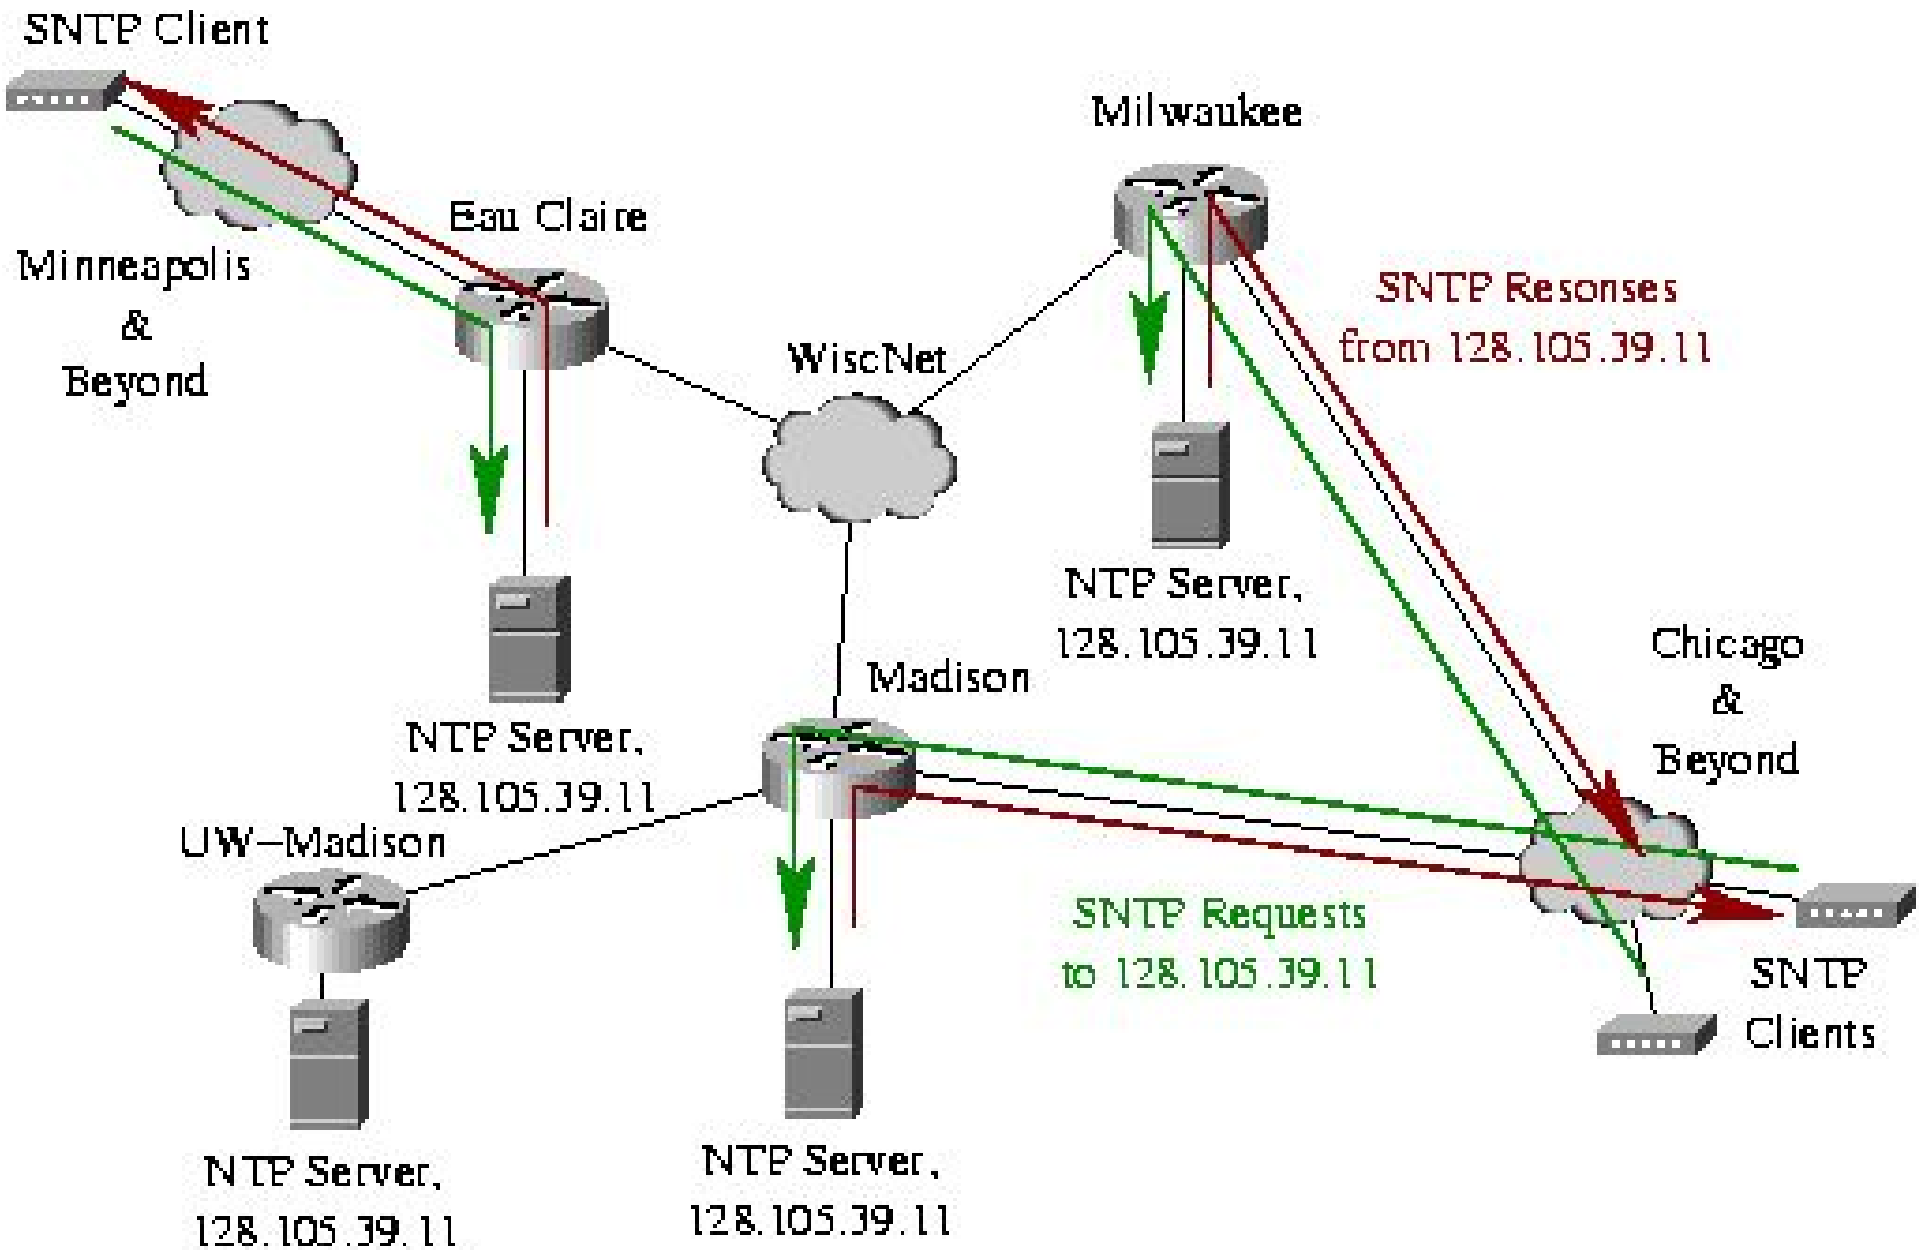

Endgame B: Suppress Requests

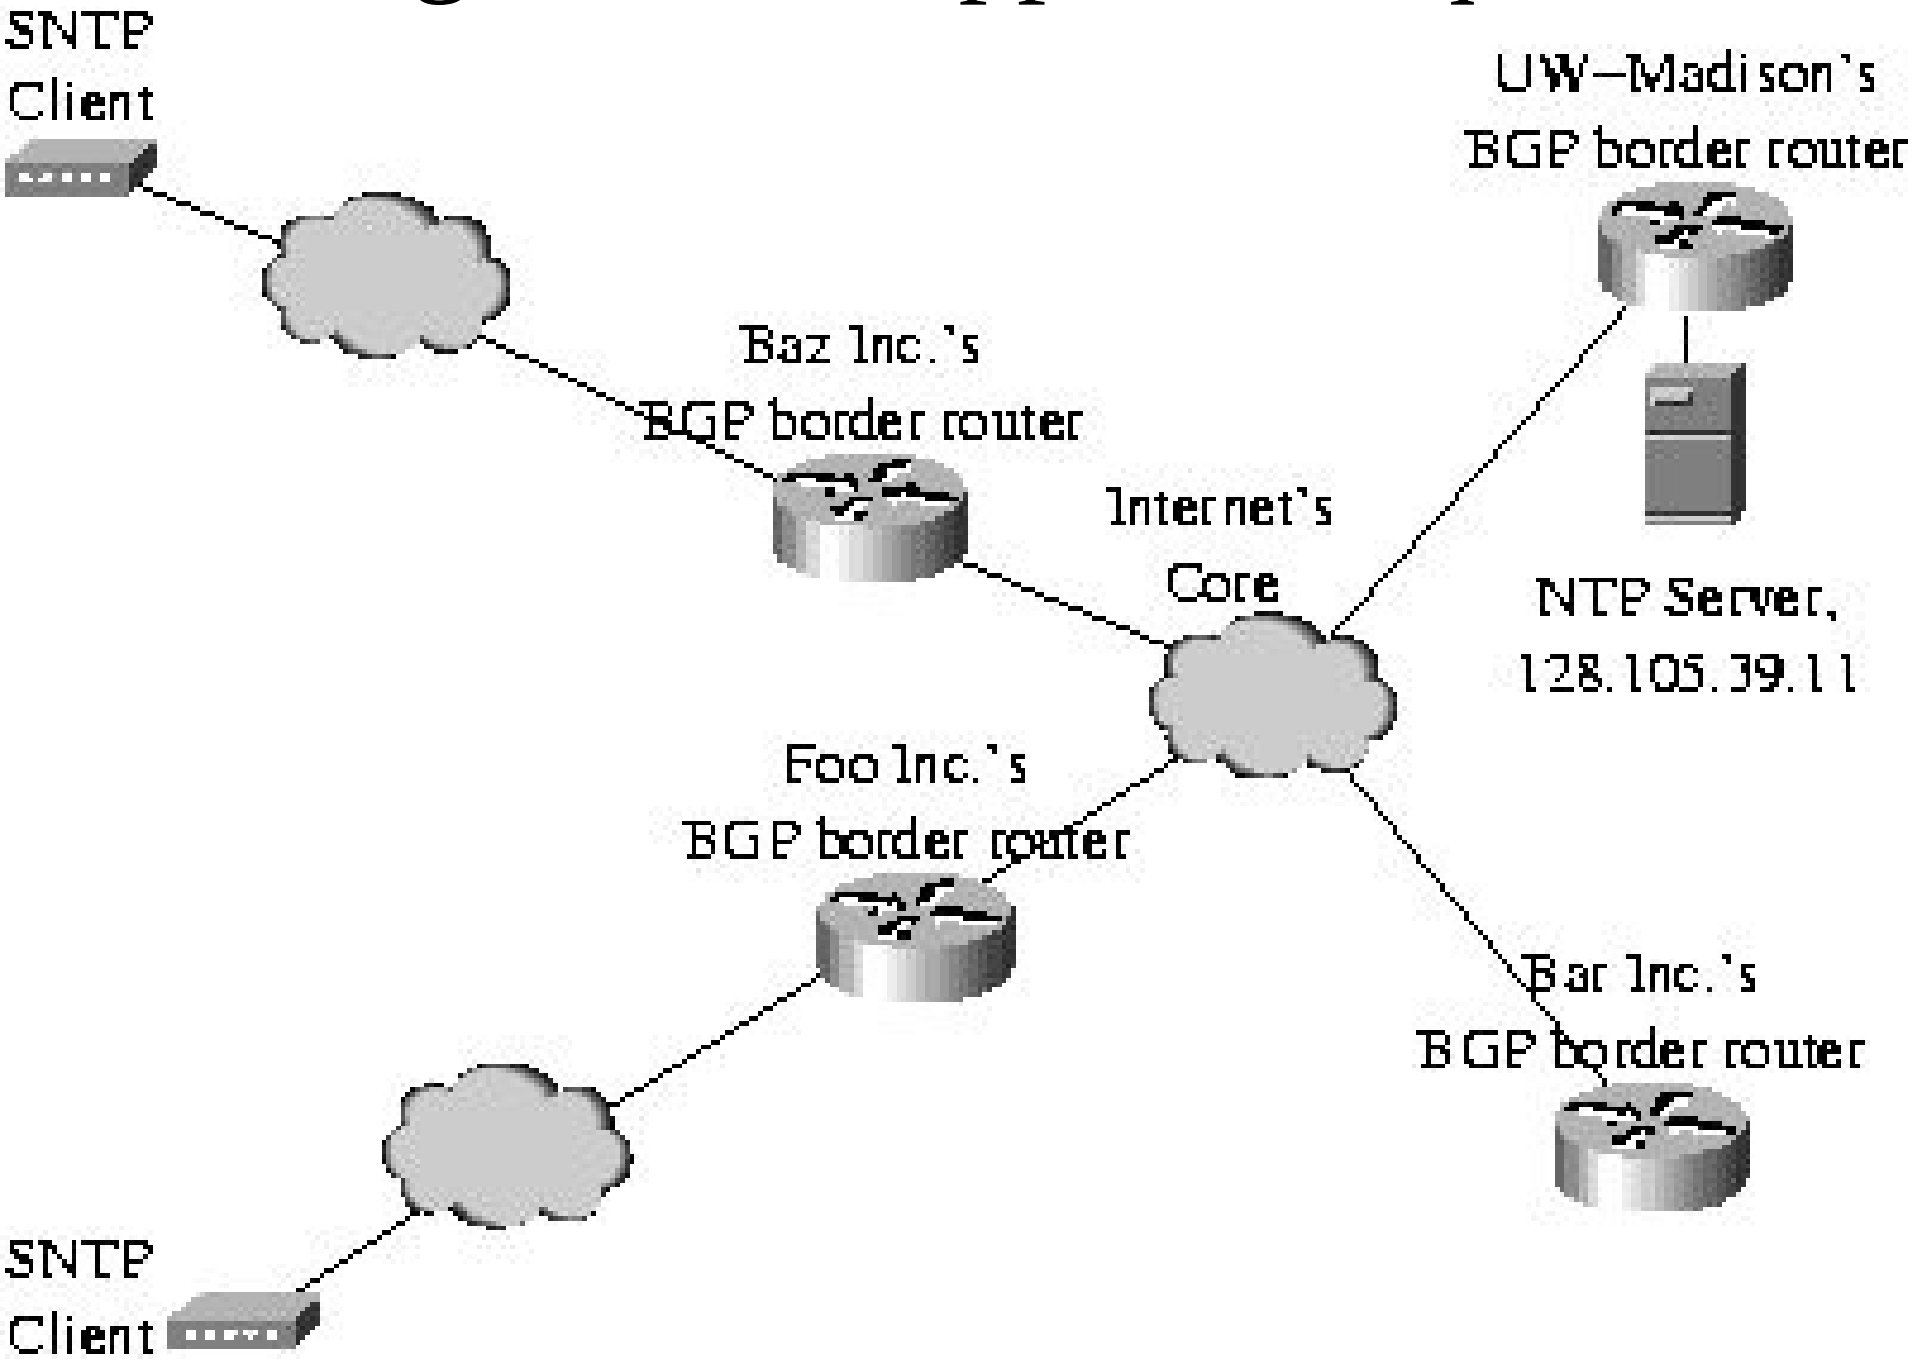

Endgame B: More-Specific Routes

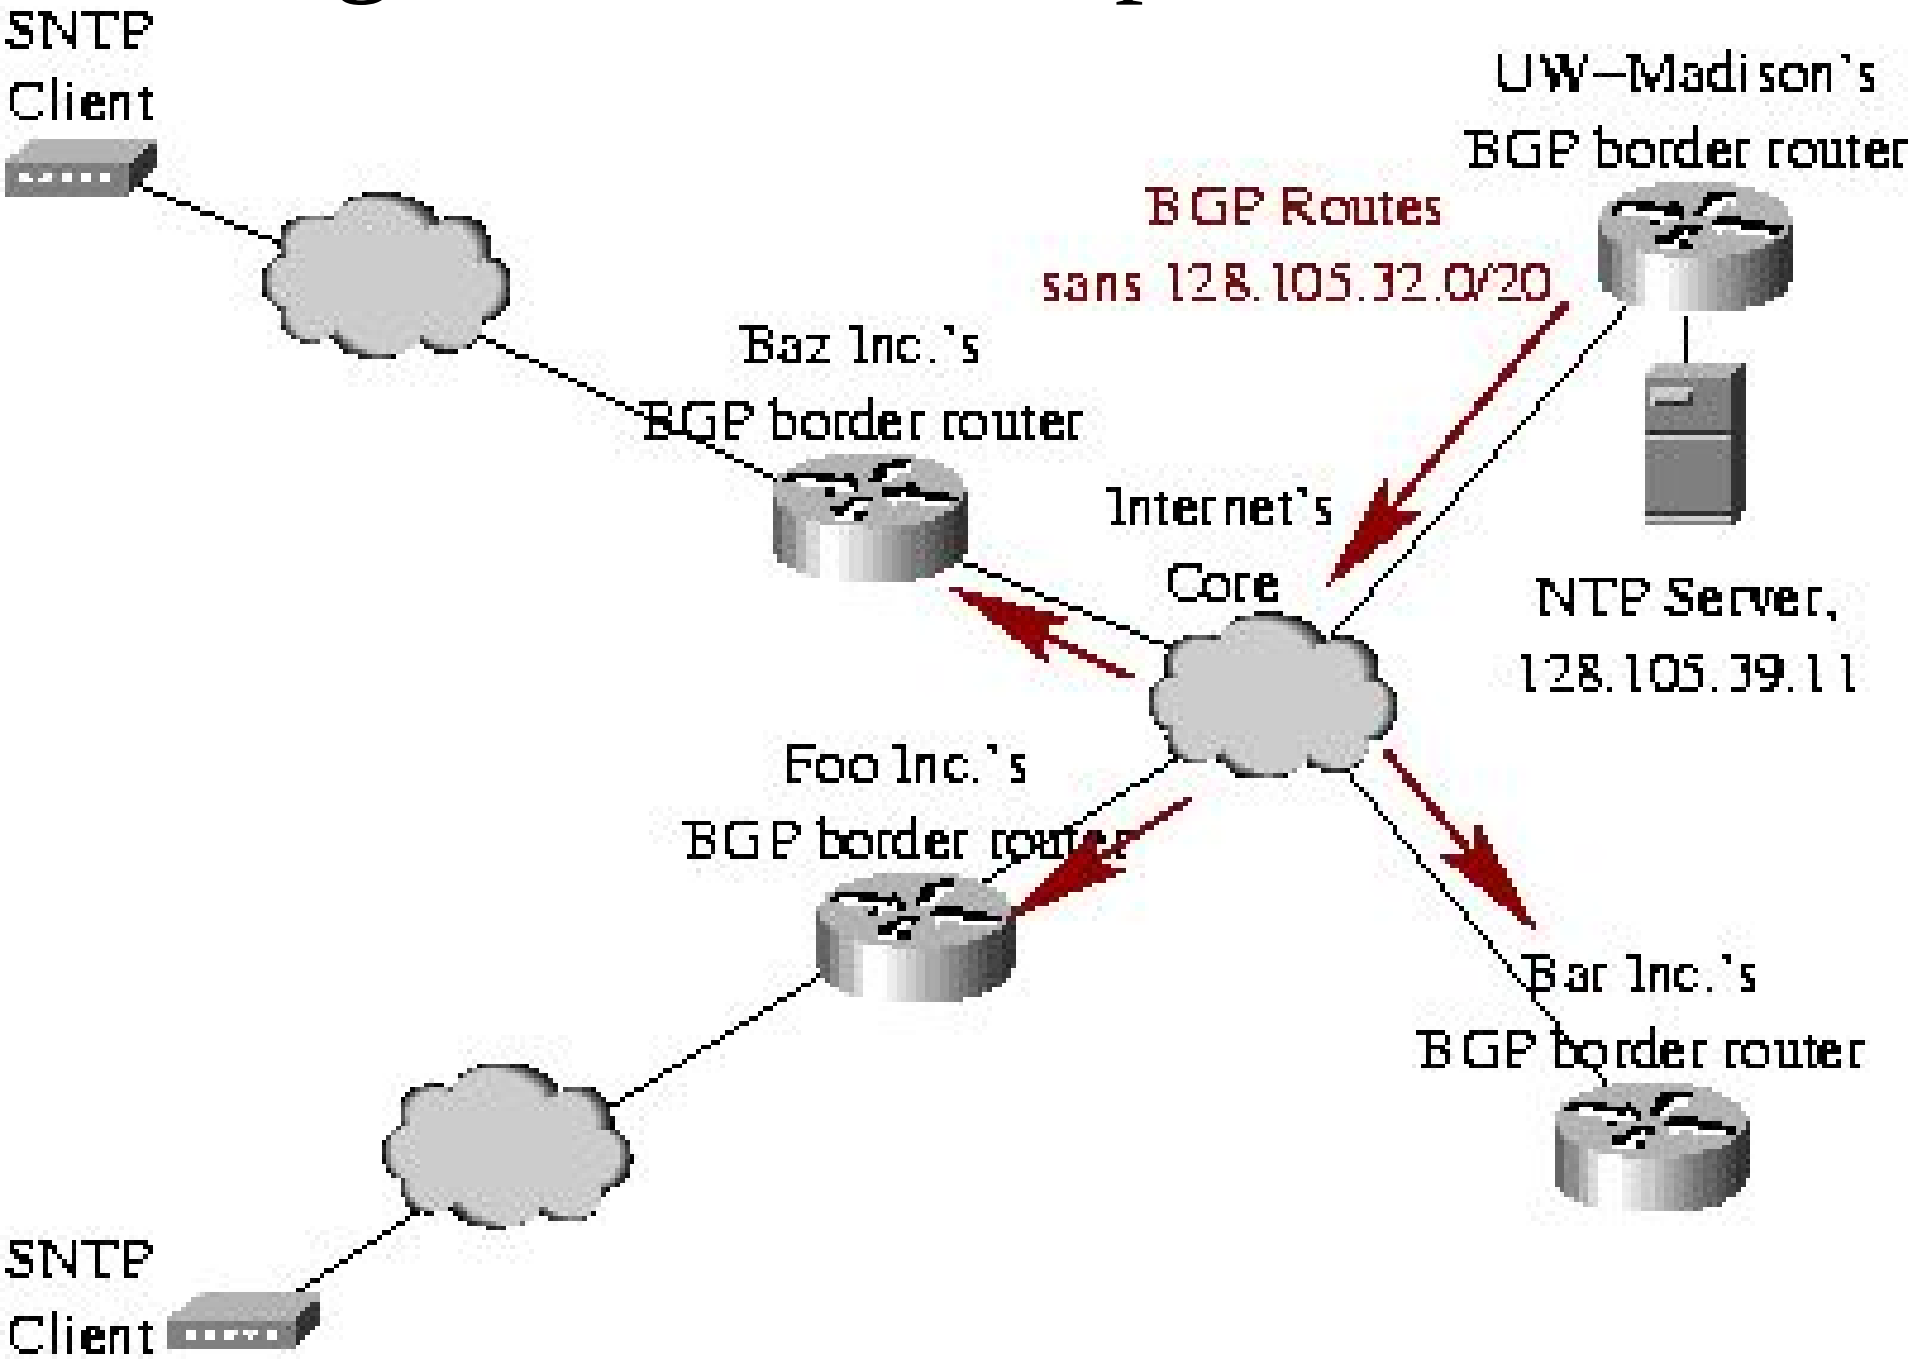

Endgame B: Requests

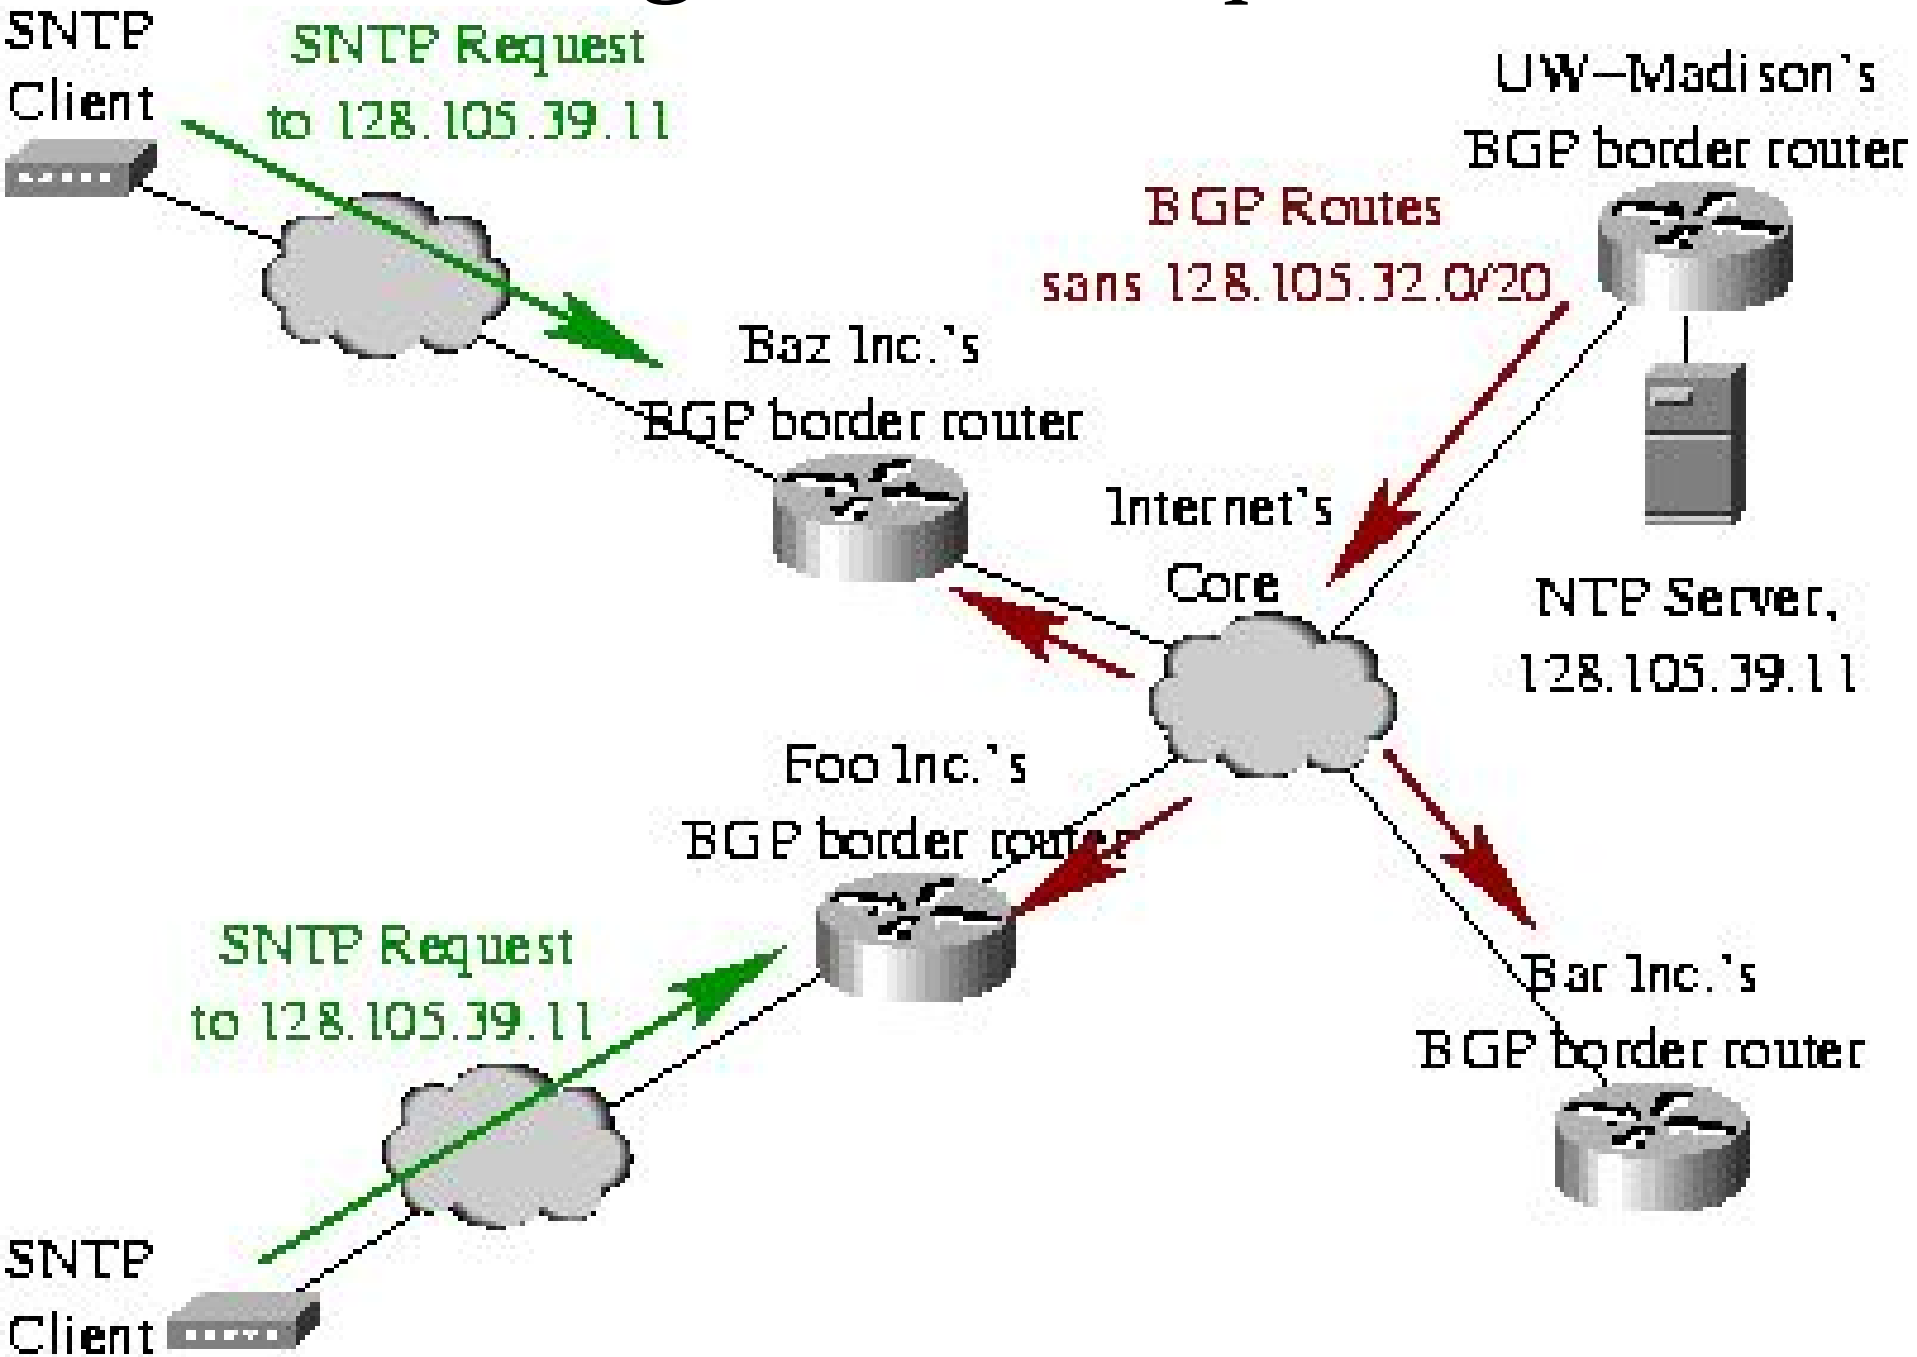

Endgame B: Unreachable "Response"

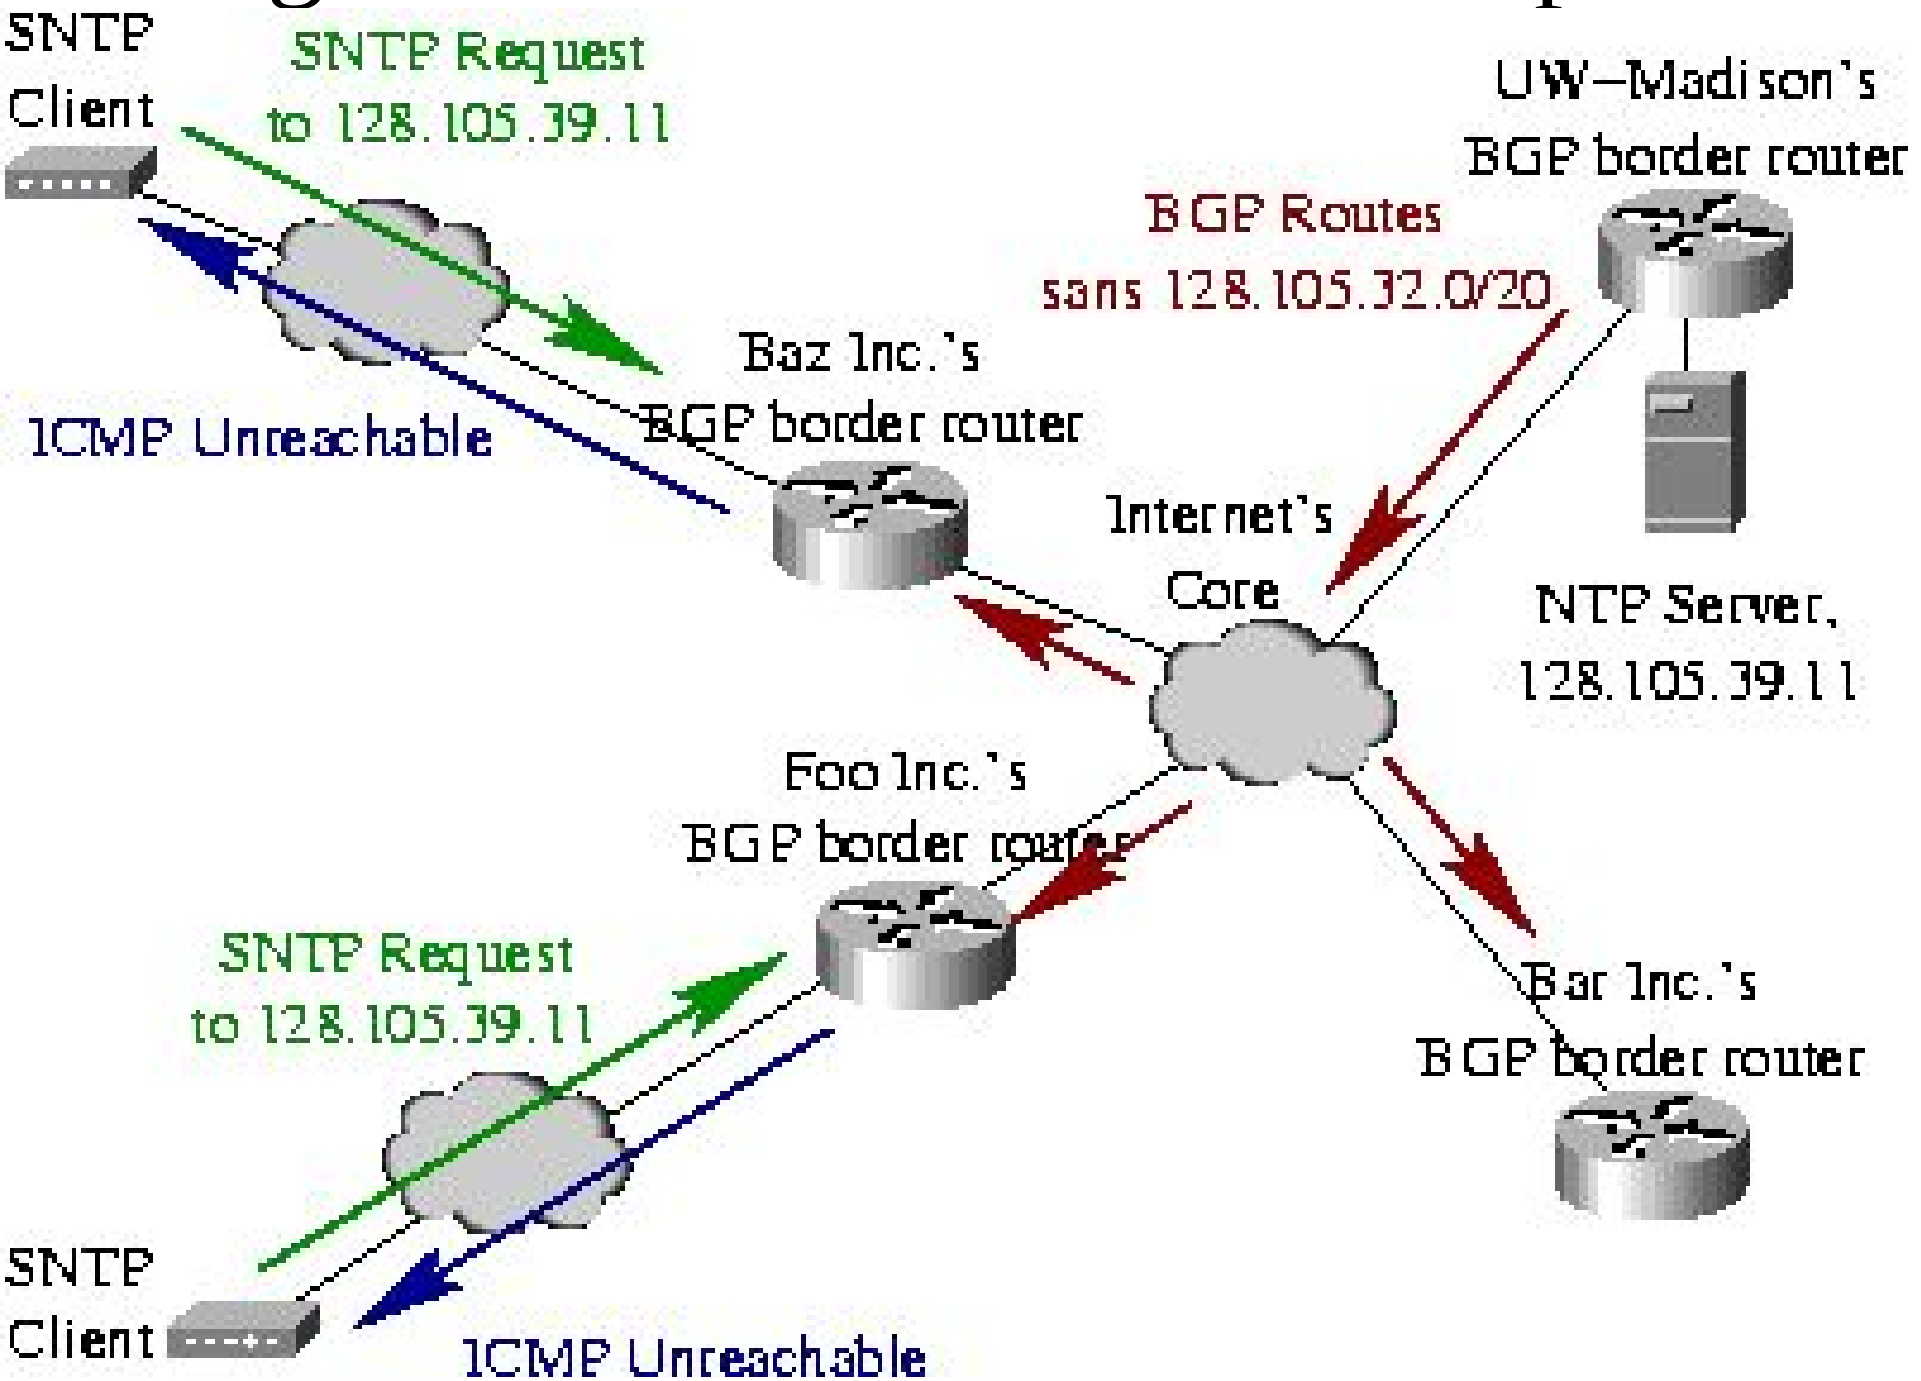

IP Addresses "Sacrificed" from 128.105 Network

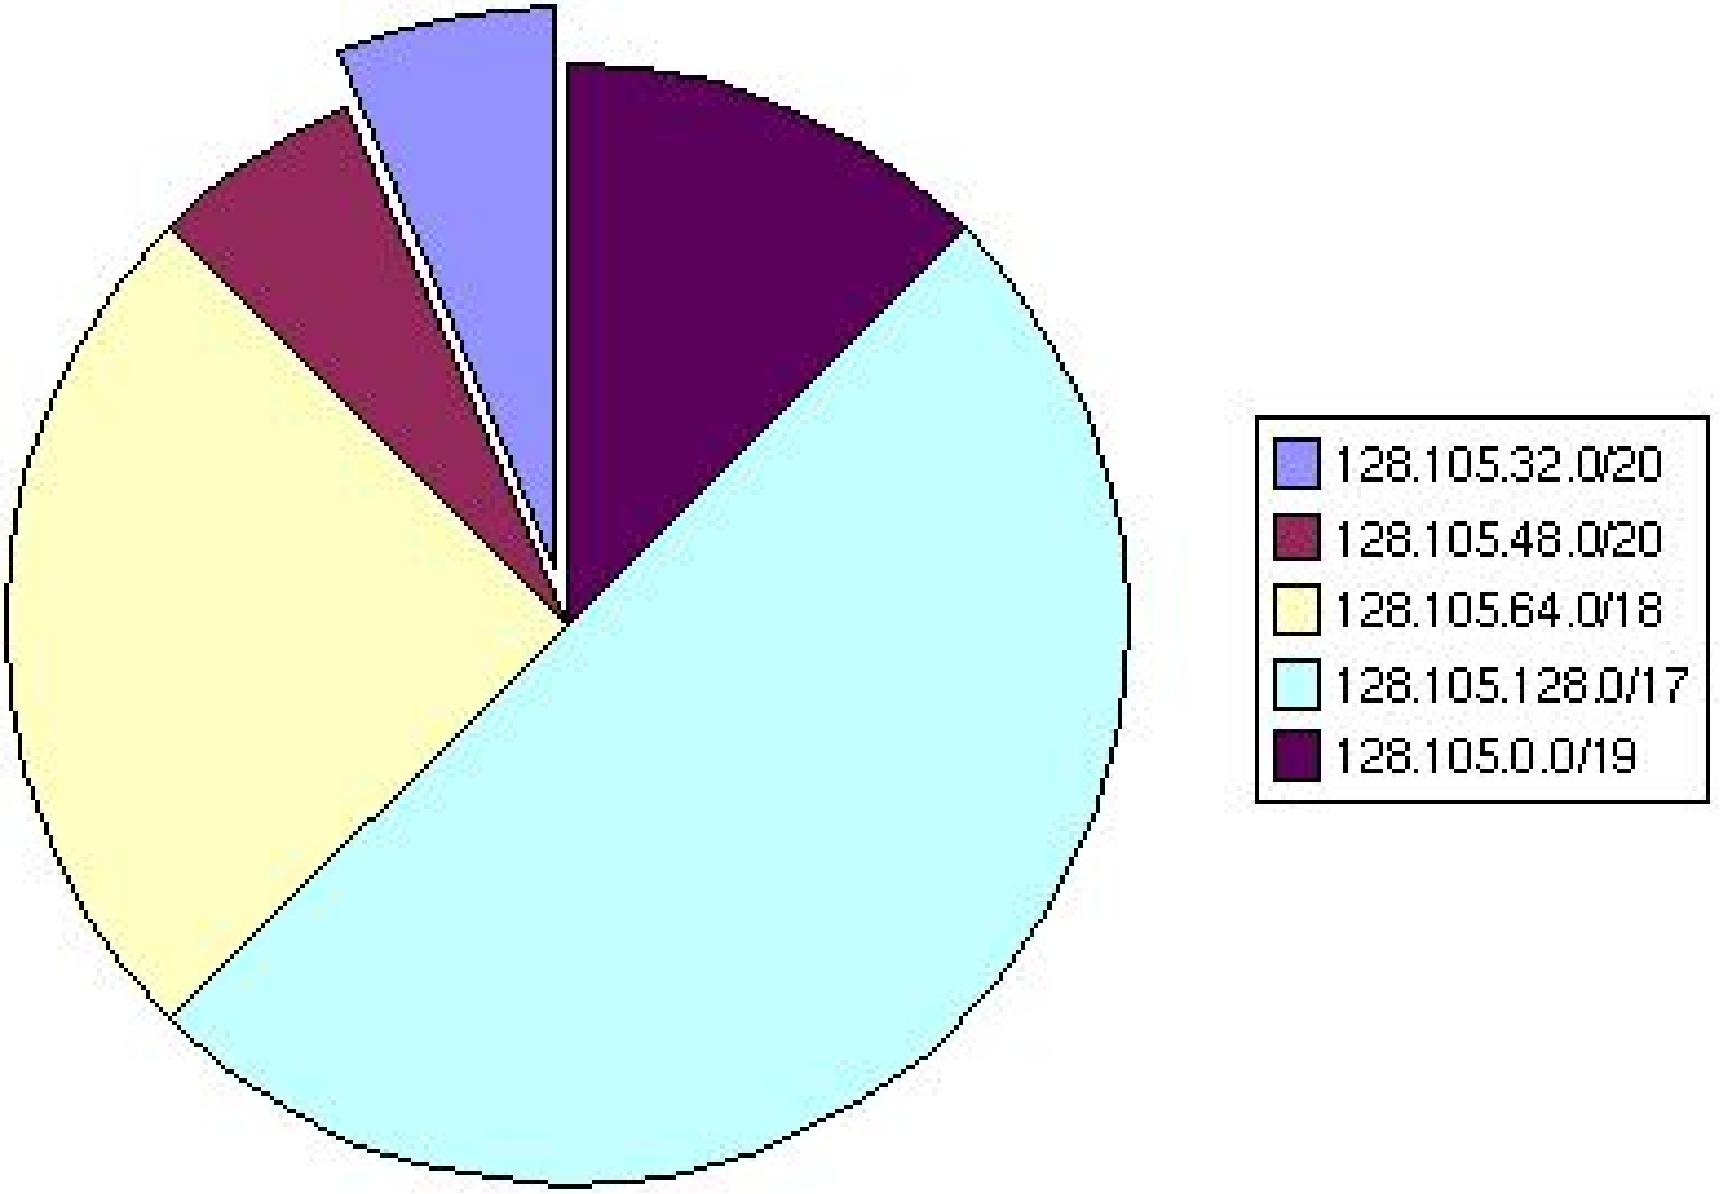

The Plan: Implement an Anycast Time Service

- Developed a detailed proposal to implement a UDP anycast SNTP service
- Also plan provide unicast NTP servers at these locations (which allows manually configuration of redundancy by NTP clients)
- Vetted proposal at NANOG 29 (October, 2003)
- Plan to monitor and measure its use to determine if any changes, improvements are necessary (such as implementing a global anycast service)

Inform the Internet Community

- Because of the scope of the problem and scale of unexpected floods, its important to inform others and continue to solicit for expert advice.
- Outreach campaign will reach a small subset of users, network operators and administrators
  - SAGE, IT Press, Universities, NANOG, USENIX LISA
- How to reach designers, manufacturers, certification labs?

Inform the Internet Community

- Can the IT press play a role in evaluating whether or not consumer products comply with Internet standards and best current practice?
- This presentation is an effort to inform the community of this flaw and the resulting floods, with the hope of minimizing the likelihood of such a mistake being repeated elsewhere.

Rogue routers cause havoc for CSIRO

- CSIRO = Commonwealth Scientific & Industrial Research Organisation
  - Like NIST, but in Australia
- 85,000 SMC routers that poll the CSIRO time server twice a minute when they don't receive a response.
- Worst case would generate about 2800 packets per second or 1.7 megabits per second

Clarify Internet Best Current Practice and Protocol Standards

- The Internet Draft titled *"**Embedding Globally Routable Internet Addresses Considered Harmful**"* denounces this practice and describes resulting problems and solutions.
- New "Simple Network Time Protocol (SNTP) Version 4" Internet Draft includes new "Best Practices" section
  - draft-mills-sntp-v4-00.txt
- With the effort and support of interested parties, we may be able to revisit NTP and SNTP as standards track protocols in the IETF.

Recommendations for Network Operators & Administrators

- Please, *do not* block nor redirect traffic destined for our NTP server's IP address - that would interfere with our operations and measurement.
- **Support the DHCP "NTP Servers" option** at the edge of your networks
- This is option 42 as defined by RFC 2132, c. 1997
- Decentralizes NTP, as it was meant to be
- This option is supported by ISC dhcpd

Recommendations for ODMs, Manufacturers, & Vendors

- Participate in and track standards organizations and the Internet operations community.
  - It is unacceptable to neglect the holistic nature of Internet engineering
- Implement a rigorous evaluation of designs that are farmed out to Original Design Manufacturers
- Consider how to remotely influence your products' behaviors, if need be, but disclose any remote control mechanisms

Recommendations for NTP Community

- Deprecate the listing of server IP addresses in the lists of public servers.
- Please continue work on the NTP protocol specifications and other implementations based on the specifications.

Afterthoughts

- What does this unintentional Denial-of-Service flood indicate about the viability of some public Internet services?
- Can the Internet routing infrastructure be improved to enable less disruptive solutions to such problems?
- Are incidents such as this a likely side-effect of ubiquitous, low-cost, perhaps even disposable Internet hosts?
- Are the manufacturer, vendor, Internet operations, and user communities willing and able to cooperate to address such problems?

Acknowledgements

- UW-Madison: Jeff Bartig, Jim Gast, Michael Hare, Adam Kunen, Dave Thompson
- U of Florida: Robert Bird, Greg Goddard
- Harvard: Greg Mazzu
- k claffy, Nevil Brownlee, George Michaelson
- Thanks to the whole review team, including those that prefer to remain anonymous.

A Case Study in Internet Pathology: Flawed Routers Flood University's Network

This slideshow:
http://net.doit.wisc.edu/~plonka/lisa/lisa2003/

Technical report publicly available at:
http://www.cs.wisc.edu/~plonka/netgear-sntp/