

On the Performance of Middleboxes

Mark Allman
ICSI Center for Internet Research
mallman@icir.org

(Work done while with BBN Technologies)

Internet Measurement Conference
October 2003

"Holly came from Miami, FLA; Hitch-hiked her way across the USA"

Middleboxes

- "Middleboxes" have cropped up all over the Internet for a variety of reasons:
 - ▶ security (firewalls, normalizers, etc.)
 - ▶ performance (PEPs, TCP snoopers, etc.)
 - ▶ address translation (NATs)
- Many have espoused the virtues and evilness of these entities.
- But, little quantitative information about their impact in real networks.
- We conducted a preliminary evaluation of one middlebox setup.

Experimental Setup


- Application measurements
 - Packet tracing and matching is future work
- Measurement period: 10/14/2002 - 1/27/2003
- Conducted in a production setting
 - A network serving thousands of users

Experimental Setup (cont.)

- Measured:
 - ▶ Transaction delay
 - ▶ Feedback time (aka "RTT")
 - ▶ Bulk transfer
 - ▶ FTP performance
 - See the paper

- Also, failures.

Experimental Setup (cont.)

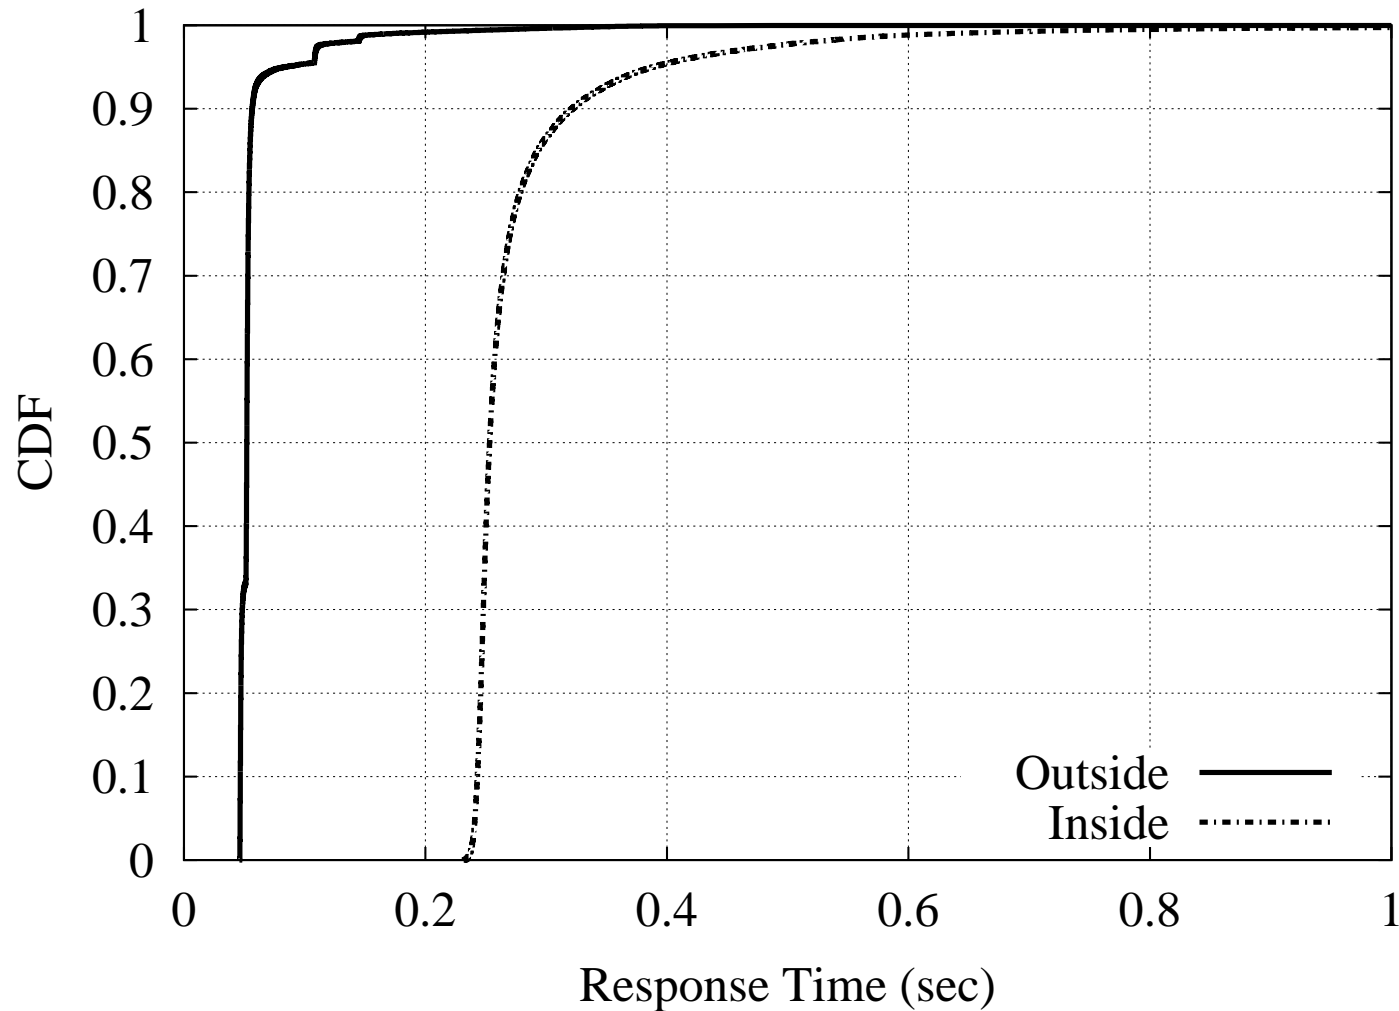


- Firewalls + Load Balancers = MBI

Transaction Delay

- How long does it take to start from nothing and run a transaction between a client and the server?
- Procedure:
 - ▶ A finger transaction between the client and server
 - ▶ Time the entire transaction at the application layer
- Conduct a transaction from each client roughly every 2 minutes.
- Over 75,000 transactions from each client.

Transaction Delay (cont.)




- 42 failures inside the MBI; 12 failures outside the MBI

Feedback Time

- Once established, how long does it take to send a message across a TCP connection?
- Procedure:
 - ▶ Open a TCP connection between the client and server
 - ▶ Send "pings" from the client; echoed by the server
 - Every (roughly) N seconds
 - ◆ We only consider $N = 30$ seconds -- others are similar
 - Until one of the pings does not come back in 20 seconds
 - Then, start a new TCP connection and start over
- Over 303,000 pings from each client.

Feedback Time (cont.)



- Failed to setup connection: 51 from inside; 46 from outside


Feedback Time (cont.)

- Connection lengths are roughly twice as long from the outside as from the inside client
 - ▶ On mean and median

Bulk Transfer


- Open a TCP connection
- Send 1 MB
 - ▶ Last 4 bytes are a random number
 - ▶ The server echos the random number back to the client
 - ▶ Measurement stops when the "ACK" arrives
- Conduct a transfer from each client roughly every 10 minutes.
- 15,000 transfers from each client

Bulk Transfer (cont.)



Bulk Transfer (cont.)

- Why the bi-model distribution?
 - Routing or provisioning changes



Bulk Transfer (cont.)

- Why the difference in performance?
 - ▶ Possibility #1: Concatenated TCP connections
 - shorter control loop
 - isolate drops
 - ▶ Possibility#2: Maybe a difference in TCP's congestion control algorithms inside and outside the MBI.

Conclusions

- Performance comparison is a muddle of contradictions
 - ▶ Bulk transfer performance is enhanced by the middleboxes
 - ▶ Transaction times increase roughly 5 times when going through the middleboxes

- Failures increase when going through the middleboxes
 - ▶ But, failures are very low in all the cases (over 99.9% across all measurements).

Future Work

- Tons
- Lots of questions can be better answered if we had packet traces from various points throughout the middlebox infrastructure.
 - ▶ Requires lots of analysis and correlation that may be non-trivial
- We can pin down why the performance is different
 - ▶ E.g., are the MBI elements getting out of sync?
 - ▶ E.g., are the firewalls dropping state?
 - ▶ Etc.
- Gather data from more locations and different kinds of middleboxes