Middlebox Discovery

Jamshid Mahdavi
Andrew Knutsen

March 23, 2010

Talk Outline

- Middlebox Discovery ID Summary and Status
- Discussion of Middlebox Needs
- Other Common Middlebox Issues of Potential Interest to IETF
- References

ID Summary

- draft-knutsen-tcpm-middlebox-discovery-03
  - Defines a new TCP Option for in-band discovery of middleboxes
  - Designed from the ground up to:
    - Consume only a single TCP Option Kind for all vendors who need this capability
    - Allow for safe proprietary use as well as future standardized use
    - Includes lessons from years of practical implementation experience
  - Incorporates numerous good suggestions from tcpm mailing list

ID Status

- Working Group has chosen not to take this up as a WG item
- Draft has been submitted for IESG approval

Evolving Internet Connectivity

- 1980's: Direct IP to IP connections
- 1990's: Firewalls and NATs become prevalent on nearly all paths
- 2000's: Increasing use of higher level middleboxes
  - Proxies (caching, security)
  - Access points
  - Acceleration devices
  - Load balancers
  - Rate shaping / TCP "enhancing" devices

What about End-to-End Arguments?

Moving away from end to end

For its first 20 years, much of the Internet's design has been guided by the end to end arguments. To a large extent, the core of the network provides a very general data transfer service, which is used by all the different applications running over it. The individual applications have been designed in different ways, but mostly in ways that are sensitive to the advantages of the end to end design approach. However, over the last few years, a number of new requirements have emerged for the Internet and its applications. To certain stakeholders, these various new requirements might best be met through the addition of new mechanism in the core of the network. This perspective has, in turn, raised concerns among those who wish to preserve the benefits of the original Internet design.

- David D. Clark, Marjory S. Blumenthal, "Rethinking the design of the Internet: The end to end arguments vs. the brave new world", August 10, 2000.
- Paper outlines many requirements that we see today

Today's Drivers

- Security
  - Cybercrime and malware are growing problems
- Performance
  - Bandwidth savings via advanced compression technologies
  - Latency savings via protocol optimizations
  - Improved goodput via TCP optimizations
- New emerging market for proxies as IPv6 transition appliances

Discovery Example

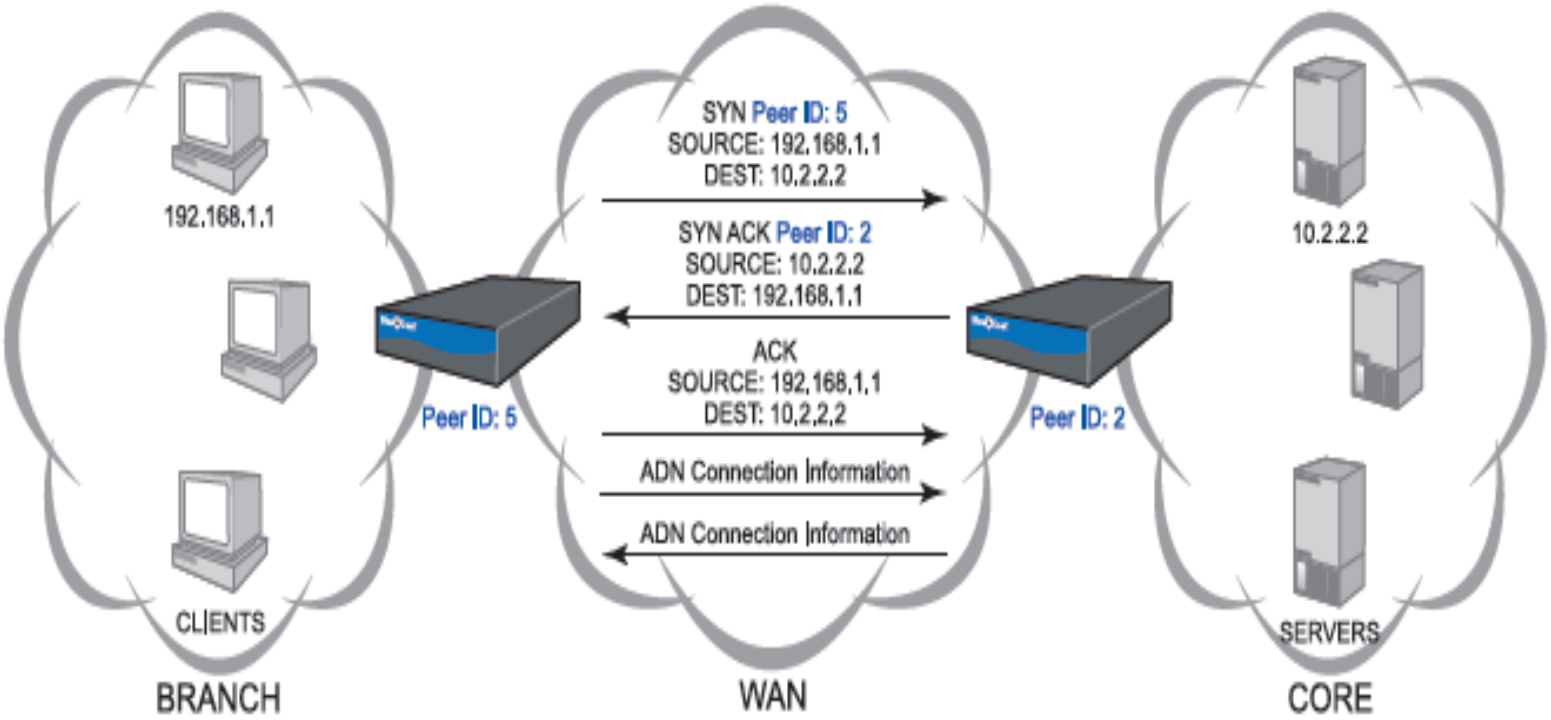

This diagram shows the setup of a Transparent ADN Tunnel

Known Problems

- There are a few problems we see all the time which the IETF could have an impact on:
  - TCP ACK storms
    - Application Networking devices often use “fail-to-wire” bridging
    - If fully transparent, when failure happens, ACK storm ensues
  - Asymmetric routing (or routing changes)
    - Often cited as a key reason transparent intercept is incompatible with Internet architecture
    - But – vendors have numerous proprietary solutions to handle this
  - Amplification of known issues
    - PMTU black holes
    - Broken support for RFC1323 and other extensions to TCP and IP

References (1/3)

- Historical references on proxies and Internet architecture:
  - Chatel, M., “Classical versus Transparent IP Proxies”, RFC1919 (1996). http://datatracker.ietf.org/doc/rfc1919/
  - Saltzer, J. H.; Reed, D. P.; Clark, D. D., “End-to-End Arguments in System Design”. (1984). http://web.mit.edu/Saltzer/www/publications/endtoend/endtoend.pdf
  - Clark, D. D.; Blumenthal, M. S., “Rethinking the design of the Internet: The end to end arguments vs. the brave new world”. (2000). http://cyberlaw.stanford.edu/e2e/papers/TPRC-Clark-Blumenthal.pdf
  - Clark, D. D.; Sollins, K.; Wroclawski, J.; Faber, T., “Addressing Reality: An Architectural Response to Real-World Demands on the Evolving Internet”. (2003). http://www.isi.edu/newarch/DOCUMENTS/Principles.FDNA03.pdf

References (2/3)

- Research publications:
  - Spring, N. T.; Wetherall, D., “A Protocol Independent Technique for Eliminating Redundant Network Traffic”. (2000). http://www.cs.umd.edu/~nspring/papers/sigcomm2000.ps.gz
  - Li, Q., “A Novel Approach to Manage Asymmetric Traffic Flows for Secure Network Proxies”. (2008). http://www.springerlink.com/content/13n10l6u011530t1/
  - Anand, A.; Gupta, A.; Akella, A.; Seshan, S.; Shenker, S., “Packet Caches on Routers: The Implications of Universal Redundant Traffic Elimination”. (2008). http://ccr.sigcomm.org/online/files/p219-anand.pdf
  - Anand, A.; Sekar, V.; Akella, A., “SmartRE: An Architecture for Coordinated Network-wide Redundancy Elimination”. (2009). http://ccr.sigcomm.org/online/files/p87.pdf

References (3/3)

- Vendor references:
  - Salchow, K. J., “Load Balancing 101: The Evolution to Application Delivery Controllers”. http://www.f5.com/pdf/white-papers/evolution-adc-wp.pdf
  - “Technology Primer: Transparent Application Delivery Networks”. http://www.bluecoat.com/doc/5276
  - Bartlett, J.; Sevcik, P., “How Network Transparency Affects Application Acceleration Deployment”. http://www.riverbed.com/docs/AnalystReport-NetForecast-Transparency.pdf