

Measuring the Dual-Stack IPv6/IPv4 Experience

Casey Deccio

Brigham Young University

NPS/CAIDA 2020 Virtual IPv6 Workshop

June 17, 2020



Motivation


- Web sites typically depend on secondary resources (e.g., images, JavaScript, CSS).
- Each of those resources might, in turn, depend on other resources.
- Questions:
 - What does an in-depth of Web-side IPv6 deployment look like?
 - Besides the main resource, what is the constituency of IPv6-enabled resources in the tree of dependencies?
 - What does the overall user experience look like over IPv4 and IPv6?
 - Is the experience consistent across different vantage points?

Methodology

- Identify (recursive) dependencies for a Web site
- Identify IPv4 and IPv6 availability for each resource
- Characterize the dependency graph, based on IPv4/IPv6 availability
- Measure the overall experience difference using IPv4/IPv6

Identifying Web Site Dependencies

- Goal: identify dynamically-loaded resources, including those invoked with scripts.
- Tool: “Page Graph” extension for Google Chrome.
- Each page loaded in by a user/researcher using the extension, output exported to json.




The above shows the graph obtained from the [IMDb](#) home page. Live demo [here](#).

<https://github.com/cyrus-and/chrome-page-graph>


Identifying IPv4/IPv6 Resources

- Create graph based on HTTP dependencies.
- Each node is marked based on IPv4/IPv6 availability.





Characterizing Dependency Graph

- Total number of nodes: 5
- Number of IPv4 vs IPv6 nodes: 5, 3
- IPv4-only (or IPv6-only) areas
- IPv6 (or IPv4) “islands”



Preliminary Results – Alexa Top Sites (top "n")



Preliminary Results – Alexa Top Sites (top “n”)

