

Paper Reviewing and Analysis

Paper: Anycast-Aware Transport for Content Delivery Networks

Familiarity: Some Knowledge

Recommendation: Likely Accept

Strengths:

- Presented a novel new solution that addressed a known issue.
- Provided a simple solution that requires minimum change to the current infrastructure.
- Improved the resiliency of the Content Delivery Networks to certain denial of service attacks.
- Provided a detailed evaluation with thorough consideration of all possible counter points.
- Focus of the paper was clearly stated and well emphasized throughout the sections.
- Well structured introduction and informative background section.
- Provided good readability and clear organization.

Weaknesses:

- Lack of details on the evaluation platform and the metrics of evaluation.
- Lacked statistics to support its evaluation conclusion, specifically the server's timeout setting.
- Couple evaluations were concluded based on speculation rather than actual testing and experiments.
- Small scale benchmarks that do not necessary imitate real networking behaviors.

Detailed Comments:

This paper was published in the International World Wide Web Conference 2009. It addresses the issue of connection disruption due to routing changes in anycast Content Delivery Networks by proposing a simple yet efficient mechanism.

Content Delivery Networks, or CDN, is a crucial element for distributing content over the Internet. Its demand and importance has dramatically increased with the growth of multimedia streaming and its related services which require large scale content distribution. The main goal of this paper is to propose a new mechanism that would address the problem of connection disruption due to routing changes in anycast CDNs. The mechanism would allow clients to continue the download of the content even after a routing change, and it requires changes mostly on the client's application layer to detect and react to a TCP connection failure.

The paper is different than other related work in the field as its proposed solution is much simpler yet efficient, rather than using more complex solution such as socket migration, the mechanism relies on the client to recover the connection failure by offloading workload from one server to another. However, while the paper provided a very specific mechanism details and implications, it focuses primarily on the theoretical details and failed to provide a comprehensive evaluation through different testing and experiments.

The performance evaluation uses small scaled benchmarks that might not necessarily represent the actual networking behaviors. In fact, not only did the testbed consist of only a single server and client but the environment was stable in comparison to the real network. However, the testing results were promising and accurately addressed the issue of providing continuing service with network interruptions. The potential contribution of this paper outweighs the shortcoming of its benchmarks.

In conclusion, while the paper failed to provide details on its evaluation metrics and environments and the small-scaled evaluating environment would not imitate the real network behavior, the paper presented a novel solution that could potential benefit and improve the overall network community. The paper was able to present a novel solution by examining the existing problem from a different perspective