

What Is in a Namespace?

Antony Rowstron*

Microsoft Research
Cambridge, UK
antr@microsoft.com

Abstract. Building large-scale distributed systems that are decentralized is challenging. Distributed Hash Tables, or structured overlays, attempted to make the process of building these distributed systems easier, by abstracting away many of the complexities of the physical topology, and provide a virtual namespace. The virtual namespace can then be exploited, for example to build or maintain other data structures like a tree to support application-level multicast or group communication. Because the namespace is independent of the underlying physical topology, it can be used to help mitigate the impact of network topology change, for example through churn. A challenge is to understand how to exploit the difference between a physical topology and a virtual topology (or namespace). In the first half of the talk, using examples from a number of systems I have worked on over the years, I will highlight how a namespace can be used and the differences between algorithms designed to work in virtual versus physical topologies.

In the second half of the talk I will then outline how recently we have begun more deeply exploring the relationship between physical and virtual topologies. The first is exploiting a virtual topology to enable disruption-tolerant routing protocols that can work on networks that are frequently and dynamically partitioning. The second is examining how we can build physical topologies that are more aligned to a virtual topology, and how this makes building servers that run in Internet-scale Data Centres easier to implement.

* Invited Speaker.