

Improving Utilization and Application Performance in Data Center Networks

Alexandru Agache

Abstract

The prevalence of cloud deployments led to a wave of research proposals and directions which aim at improving the underlying infrastructure at both hardware and software levels. Many solutions can be applied to each of the two major types of cloud infrastructures in use today: private clouds, managed and used by a single large entity, and public clouds, which are provided to customers using paradigms such as Infrastructure-as-a-Service. On the other hand, some optimizations cannot be deployed in public settings, most often because tenants, which are the main beneficiaries, do not possess a sufficient degree of information and control.

This thesis brings improvements to the performance of data center applications using a multi-layer approach. First, the hardware infrastructure is enhanced using GRIN, an overlay topology design which greatly increases the utilization of high capacity networks, which invariably face considerable underutilization. A key insight is that servers which need the most capacity are inherently limited by the use of a single network interface. GRIN enables capacity pooling by using multiple network ports to interconnect servers among themselves. The high likelihood of idle neighbours allows a single server to access much more capacity (up to as much as the number of additional ports), at a very low cost.

Next, we introduce TCP sendbuffer advertisement, a simple but powerful addition to the Linux TCP stack. It enables a finer grained control over triggering certain optimization which previously relied on heuristics, or informed guess at best. Finally, we propose CloudTalk, a middleware platform between public cloud providers and their tenants. CloudTalk allows many critical decision making instances to be offloaded by cloud customers to providers, which at the same time gain valuable insight into the future workloads of their clients. This significantly improves the status quo, where both parties are prevented from following the best course of action due to insufficient knowledge. CloudTalk empowers applications to make much better use of available resources; the lack of information leads to lackluster scheduling, which negates the benefits of solutions such as GRIN.