

Results as a (Network) Currency

christian.tschudin@unibas.ch, Oct 30, 2015

Abstract

Current computer networks don't have enough information about why their service is requested – an observation which also applies to storage and computing resources. We posit that expressing *results* (instead of URLs, remote procedure calls or memory allocation requests) and letting the network propose and orchestrate corresponding result-producing processes enables a new generation of information processing systems. Results, rather than bandwidth or CPU cycles, become the new currency which the network can offer to external users but will also use internally to optimize its own operations.

Keywords: The network is the computer (redux)

Introduction

Networks have always been task-oriented (“do this action to handle that problem”) while ultimately what matters are results. TCP does transfer bytes in some optimal way. But if the goal is to synchronize more than two file copies, and TCP is used multiple times, this is far less efficient than a solutions that use net-assisted caches and fountain codes. The problem lies not in TCP but in the lack of “language” in which users can express their needs towards the network. What matters are results instead of data, bandwidth or compute power which are mere means.

We propose to introduce “results” as the core vehicle through which end users talk to the network. Also inside the network, a similar style and set of languages (ways beyond network addresses as an example of a current “language”) should be used in order to bridge the requests with result-producing entities. These entities can be sitting at the network's edge but also inside the network: Talking “results” abstracts away how and where the network produces them. Results become the (network's) currency that ties together heterogeneous result seeking and result producing entities, including the network itself.

At an abstract level, Google is a result-producing engine in the above sense, although the languages looks very different depending on where we look. In a first place it has to produce results for the advertisers in order to survive, which is then translated in producing subresults at the network, storage and computing level (and, consistently with this terminology, turning consumers into a result (=product). However, the ambition of “Results as a Currency” is not to create another Google but to come up with the network-level mechanics that permits **to create Google-sized solutions for Science** much easier and with better elastic qualities than we can today.

Three Focal Research Areas

We see three major areas that need to be specifically pushed in order to create a “RaaC network”:

From Addresses to Result Expressions: Research initiatives like “Named Data Networking” have already moved into the direction of RaaC, for example in “named functions” where requests are expressed as recipes without reference on where and by whom it has to be produced. This leads to a network which can optimize how results are produced e.g., move the function execution closer to the data storing location, result caching, or data prefetching.

Integrated Resource Control and Choice: The disruption brought by Software Defined Networking and its plan economy style of resource handling has enlarged the spectrum how the network can do resource control. Google's way of planning bulk data transfers over time sets a new bar how much better “results” can be produced. It is urgent to import research from the fields of AI (planner, machine learning), economy (resource and “result” pricing, choice) as well as system dynamics (e.g. in chemistry) and to formulate more informed algorithms that steer the network's result production.

Modulating the Quality of Results (Security, Accuracy, Timeliness, Persistence): While the two previous areas focus on the quantitative and optimizing side, we require RaaC to also include qualitative aspects as being “nameable” and as being part of the result currency. For example, results often need approximated values while network services typically strive at perfect quality; results have different urgencies (e.g. permitting Google to shift transport tasks in time); and transactions have different needs at the security level (full anonymity vs mere result authenticity or audit capabilities). While often we know how to produce such qualities in isolation, there is an obvious lack in how requirements can be passed on “to and through” the network in order to get them produced alongside with the actual number crunching and data transfer.

RaaC is not another Middleware

The role of our RaaC proposal is to point out the importance of the *flow of results* that permeate the network: Empowering the network to reason about such flows requires that the network understands the structure of result assembly and engages in its production. Insights from the design of cloud computing centers where computation load and costs are integrated with switching decision show how shallow a middleware approach would be. We believe that “result” and “currency”, although introduced here in a metaphorical way, can and should be made real.