Distributed Systems: Controversial topics

Maarten van Steen

Introduction

Below we briefly describe a handful of controversial topics in distributed systems, each formulated in terms on advice that needs to be provided by an expert team. These descriptions are meant as a guiding principle. If you want to deviate from the advice to be given then that is, by all means, possible (but only after consultation).

1. Robustness of public peer-to-peer systems

The first wave of public peer-to-peer systems has passed and by now most are operating within the boundaries of a single organization. This shift reflects the difficulty in getting public peer-to-peer systems secured against, notably Eclipse and Sybil attacks. Nevertheless, there are still a number of systems that operate in the open, such as the DHT-based system for so-called magnet links (used by, for example, Bittorrent), as well as the Kademlia system.

Understanding the ins and outs of various types of peer-to-peer systems, you are asked to advise a new startup on developing a fully decentralized file-sharing system on whether or not they should base this system on an existing product. One particular concern they have is the robustness against churn, for which they are thinking of using a gossip-based approach which are known to quickly converge to specific overlays, yet are vulnerable to attacks.

2. Edge computing

One of your fellow students is considering setting up a startup for developing an edge-computing platform that can be easily connected to a range of cloud-service providers. One of the key issues your colleague sees is not only provisioning the basic infrastructure, but also a complete set of tools that will automatically manage resource allocation for the applications to be allocated to the edge. Many of those applications either currently run on edge devices or in the cloud.

Your expert team is aware of the hype around edge-computing systems and notably the discussions on the apparent advantages concerning security, performance, etc. You also notice that your colleague student may not be aware of the ins and outs and you decide to give a substantiated advice on whether or not she should carry on with her ideas, and if so, what the main issues are that she should consider to make the startup attractive from an end-user's perspective.

3. Virtual machines versus container technology

A small data center that operates for a largely regional, yet funtionally wide range set of companies, is considering a complete transition to container technology instead of their current use of virtual machines. Their main reason is the suspected gain in performance and scalability (so they say). Their customers rely on a diverse set of applications, requiring the need for supporting office-based operating systems such as Windows, but also a range of

different Unix machines. So far, the data center has used Debian distributions as their base operating system.

Your expert team is more than just knowledgeable when it comes to virtualization and is very much aware of the strengths and limitations of containerization. Instead of advising what to do, you offer to enhance the awareness around containerization and virtualization so that the data center operators can make a founded decision on if, where, and how to apply containers. You suspect that in the end, they may need to use both containers and virtual machines, side-by-side. Providing the proper requirements for whatever solution they decide on, is what you will be offering.

4. Scalability of publish-subscribe systems

As with so many hypes, lots of folks have started to ride the wave of publish-subscribe systems without really understanding what is going on. Fortunately, your expert team does understand what is in that wave. You have been hired by a company (called REPS) that wants to set up a nationwide competitor to funda.nl for selling and renting a range of real estate. Their unique selling point is that a customer can subscribe to new property that matches their wishes such that within only a few seconds they will receive a notification when that property comes available.

You immediately understand that you're dealing with a content-based publish-subscribe system, but with the current turnover in the housing market, you also understand that scaling may be an issue. Moreover, missing out on a notification may lead to claims by customers.

To this end, you decide to dig further into the matter and identify the tradeoffs REPS needs to consider. You ask yourself whether a simpler topic-based or channel-based pub-sub system can suffice, or perhaps a combination of the two. In any case, you will need to make clear to REPS what they are facing up to. To make matters worse, REPS has decided that it would prefer to also guarantee that matches remain anonymous, allowing them to make use of existing cloud-based services without the need to fully trust those services. If time allows, you will also advise on this additional requirement, knowing that it is not going to help keeping matters simple.

5. Scalability of blockchain systems

Blockchain technology has become a much debated topic, often for very different reasons. Part of the popularity comes from the belief that blockchains can operate without the need for a trusted third party (such as a bank) while offering scalability. When diving deeper into the technicalities, there are serious problems. One specific problem that may eventually turn out to be too difficult to solve, is that the combination of full decentralization (i.e., no trusted thrid party), high transaction processing capabilities, scalability in the number of participants, and still attaining global consensus on the commitment of transactions is practically impossible to realize. One could argue that this impossibility renders blockchains practically useless.

Imagine the situation that the CEO of a software company is considering to tender for developing a general-purpose blockchain platform that can act as a middleware solution for a range of potential applications. Aware of the controverses around blockchains, they ask your

expert team to give a well-founded advice on whether or not they should considering developing such a platform.