Universität Paderborn — Fakultät EIM
Fachgebiet Rechnernetze
http://www.upb.de/cs/cn/

Master Thesis

Consistency Models for Distributed Network Function State

Background

Network functions like firewalls, proxies, load balancers or intrusion detections systems (IDS) are essential parts of today's Internet. They are typically implemented as dedicated hardware boxes that have to be configured and integrated into the network by hand. The concept of network function virtualization (NFV) replaces these dedicated boxes by software implementations running on virtualized commodity servers in existing cloud environments. In addition, Software defined networking (SDN) is often used to dynamically forward traffic through the appropriate network functions. This allows to define a complete chain of network functions and its forwarding rules in software, giving network operators much higher flexibility. One of the main benefits of this approach is to automate the deployment of these network functions and adapt it in relation to the current load situation. Thus, a network function can be scaled-up by adding new virtual instances to the system or scaled-down by removing unneeded ones.

Unfortunately, most network functions are stateful and therefore adding or removing instances from the system does only work if the allocated state can be moved to another instance e.g. if an instance is removed. Existing solutions [1, 2] exploit the fact that most of the state is directly related to a specific network flow and typically only accessed by one of the network function instances. This makes it possible to only move these small parts of the entire state space between the instances whenever the flow assignment is changed. However, all existing approaches use a central management component that decides which state has to be moved or shared between instances. A potentially better solution may be to use distributed data stores to allocate and share this state data.

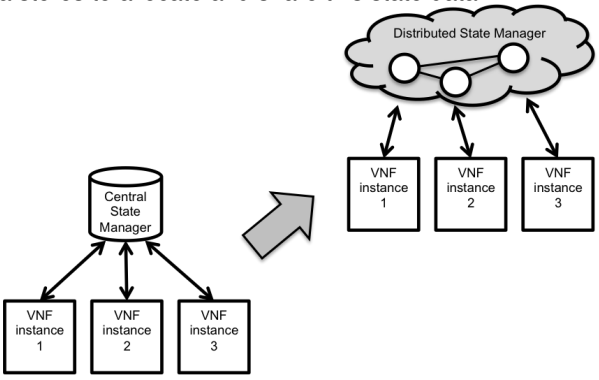

Thesis Goals

The aim of this thesis is to evaluate different consistency models for a distributed state management system. The state managed by this system is the application state of network functions. The thesis should evaluate which consistency models are needed for this kind of application state and identify problems that may occur when weaker consistency models are used, e.g. broken connections or missed intrusions.

The author of the thesis will collect background material on existing distributed memory systems (DSM) and select a DSM that can be integrated with a network function implementation, like Iptables or Bro IDS. This integration may need several changes in the implementation of the network function. Based on this, the author will perform experiments with his/her prototype to compare different consistency models. The main focus of this comparison will be the tradeoff between state consistency and reliability of the executed network function.

Milestones

- Find and select a distributed memory system (DMS)
- Integrate DMS and a network function (e.g. Bro IDS)
- Evaluate performance and reliability of the resulting distributed network function with different consistency models

Required knowledge

- Concepts of distributed systems
- Good C/C++ programming skills
- Understanding of computer networks

References

[1] A. Gember-Jacobson, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and A. Akella. OpenNF: Enabling innovation in network function control. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 163–174. ACM, 2014.

[2] S. Rajagopalan, D. Williams, H. Jamjoom, and A. Warfield. Split/Merge: System Support for Elastic Execution in Virtual Middleboxes. In *NSDI*, pages 227–240, 2013.

Contact: Manuel Peuster, manuel.peuster@upb.de, Phone: 05251 / 60 4341, Room: O3.149