

Alibaba BoF @OSDI 2018

We will introduce and discuss some key system software in Alibaba, including OS, distributed storage system and resource management. They are used to support "millions" of application use cases in large scale, especially with extreme huge workloads on promotion events.

Venue: Las Palmas Room 1

Date: October 8 , 2018

Time: 8:30pm-9:30pm

ALIBABA BOF @ OSDI 2018



OS Development and Exploration in Alibaba



Speaker: Weidong Han

Title: Senior Staff Engineer, Alibaba Group

Bio: Weidong Han is currently a senior staff engineer in OS team of Alibaba. He is leading OS development and innovation in Alibaba group. Weidong worked for Intel and Huawei before joining Alibaba. He started working on OS and virtualization development and research since 2006. His current interest includes unikernel, virtualization, secure container, etc.

Abstract: Alikernel is the OS kernel used in Alibaba. It's developed to support "millions" of application use cases in Alibaba. We will introduce some enhanced features and optimizations of Alikernel. In addition, we will also introduce our exploration on AliUK which is a unikernel and may be widely used in future.

Pangu2.0: New Distributed Storage System for Large-Scale Data Center



Speaker: Zhongjie Wu

Title: Senior Staff Engineer, Alibaba Group

Bio: Zhongjie Wu is currently a Senior Staff Engineer of Alibaba Cloud and response for Hardware & software co-design area for Pangu distributed storage system. Before joining Alibaba, he was a Senior Director in Memblaze, a startup company focus on Flash storage area and built Flash storage system department. Before that he worked in EMC and storage research center in ICT Chinese Academy. More than 12 years experience in storage system design area and issued more than 20 storage related international patents.

Abstract: In order to address large-scale data center storage issues, we proposed a new distributed storage architecture named pangu2.0 to achieve extreme low latency with high IOPS. In this presentation, the new storage architecture, core storage design ideas, fully user space storage engine and RDMA interconnection technology will be discussed in detail. Pangu2.0 has been widely deployed in large-scale data center and provide basic storage services for Alibaba. Some experience and evaluation results will be shared at last.

Sigma: Resource Management in Alibaba at Scale



Speaker: Haiyang Ding

Title: Senior Software Engineer, Alibaba Group

Bio: Haiyang Ding is a Senior Software Engineer in Cluster Resource Management and Scheduling Team in System Software Department in Alibaba Group. He joined in 2017 and is now focusing on data-driven resource management strategy and related innovation research. Prior to Alibaba, he worked at Huawei for 3 years on Cloud Computing (Kubernetes) and Blockchain (Hyperledger and general Blockchain technology). Haiyang is a Ph.D. on Automatic Control (Distributed MPC) in Grenoble University, France, supervised by Mazen Alamir and Ahmad Hably.

Abstract: Sigma is the resource management system inside Alibaba Group and it manages tens of thousands of machines distributed in several data sites, in which applications from across over 60+ Business Units are running. Besides day-to-day operation, Sigma also plays an important role in resource provisioning in the promotion events such as 11/11, aka Singles Day, in which extreme workloads (up to 100X more than daily average) take place. In this session, we will introduce some of our efforts in provisioning the resource for the promotion events and how we try to improve the overall efficiency of our fleet during the day-to-day period.