

Computer Science Seminar Series

National Capital Region

Breaking the Monolith: Rethinking Storage System Design

Speaker: Dr. Yue Cheng
George Mason University
Friday, March 30, 2018
1:00PM- 2:00PM, NVC T3

Abstract

Distributed object-based storage systems serve the growing and disparate needs of almost all of large-scale web services in use today. This is driving the proliferation of a wide variety of distributed object-based stores. However, extant, monolithic, scale-out storage system designs present unique challenges in adapting to the ever-changing storage requirements in both efficiency, and flexibility. My research takes two crucial steps on this difficult road to optimize and design better object-based stores. In this talk, I will first show that an approach to storage system design based on a simple core principle, resource partitioning, can yield systems with significantly improved performance and resource efficiency under dynamic, skewed, and multi-tenancy workloads. Then, I will discuss a fundamental challenge faced by all practitioners and developers working on scalable storage: how to implement a fast and reliable scale-out storage system with minimal engineering effort? I will present how modular design rather than the extant monolithic approaches can ease the burden of designing new storage systems, especially by enabling an innovative decoupling of the control and data plane in distributed storage.

Biography



Yue Cheng is an assistant professor in the Department of Computer Science at George Mason University. His areas of research span experimental computer systems, workload characterization, and high performance computing (HPC). His current research interests include distributed and storage systems, container virtualization, non-volatile memories, and Internet of Things (IoT). He received his Ph.D. in Computer Science in 2017 from Virginia Tech. His work has been published in premier venues in computer systems and high-performance computing, including USENIX FAST, USENIX ATC, ACM EuroSys, and ACM HPDC. He has worked and collaborated with leading storage researchers at IBM Research and Dell EMC.