

**Events**[Current Events](#)[Lectures](#)[Events Archive](#)

Other

Local Failure Local Recovery: Toward Scalable Resilient Parallel Programming Models**Keita Teranishi, Sandia National Laboratories**Digital Media Center 1034
October 02, 2019 - 11:00 am**Abstract:**

With growing scale and complexity of computational systems, HPC applications are increasingly susceptible to a wide variety of hardware and software faults. Accordingly, applications are ill-equipped to deal with the full spectrum of possible faults and often their response, particularly in synchronous programming models, is disproportionate to fault rate. Alternatively, Local Failure Local Recovery (LFLR), is based on the notion that a fault recovery that is localized around their occurrence is more scalable and efficient than a bulk response characterized by the traditional checkpoint/restart. LFLR is more amenable with an asynchronous programming model as opposed to synchronous ones. In this study, we review the existing resilient parallel programming models and then demonstrate the Efficiency and scalability of our resilient programming model for the traditional message passing and emerging asynchronous many task programming models. Also, we will discuss our recent effort to enable performance portable resilience through Kokkos.

SNL is managed and operated by NTESS under DOE NNSA contract DE-NA0003525.

Speaker's Bio:

Dr. Keita Teranishi received the BS and MS degrees from the University of Tennessee, Knoxville, in 1998 and 2000, respectively, and the PhD degree from Penn State University, in 2004. He is principal member of technical staff at Sandia National Laboratories, California. His research interests are parallel programming model, resilience and fault tolerant computation, and sparse matrix and tensor computation for high performance computing platforms.

