

# Advancing Reliability, Availability and Serviceability for High-Performance Computing

Stephen L. Scott

Senior Research Scientist  
Computer Science and Mathematics Division  
Oak Ridge National Laboratory (ORNL)  
Oak Ridge, Tennessee, USA  
scottsl@ornl.gov

Today's high performance computing systems have several reliability deficiencies resulting in noticeable availability and serviceability issues. For example, head and service nodes represent a single point of failure and control for an entire system as they render it inaccessible and unmanageable in case of a failure until repair, causing a significant downtime. Furthermore, current solutions for fault-tolerance focus on dealing with the result of a failure. However, most are unable to transparently mask runtime system configuration changes caused by failures and require a complete restart of essential system services in the case of a failure. High availability computing strives to avoid the problems of unexpected failures through preemptive measures. The overall goal of our research is to expand today's effort in high availability for high-performance computing, so that systems can be kept alive by an OS runtime environment that understands the concepts of dynamic system configuration and degraded operation mode.

This talk will present an overview of recent (and ongoing) research performed at Oak Ridge National Laboratory in collaboration with Louisiana Tech University, North Carolina State University, Tennessee Tech University, and the University of Reading in developing core technologies and proof-of-concept prototypes that improve the overall reliability, availability and serviceability of high-performance computing systems.

---

Stephen L. Scott is a Senior Research Scientist in the Network and Cluster Computing Group of the Computer Science and Mathematics Division at the Oak Ridge National Laboratory (ORNL), Oak Ridge, USA. Dr. Scott's research interest is in experimental systems with a focus on high performance distributed, heterogeneous, and parallel computing. He is a founding member of the Open Cluster Group (OCG) and Open Source Cluster Application Resources (OSCAR). Stephen is presently the OCG steering committee chair and has served as both the OSCAR release manager and working group chair. Dr. Scott is the lead principal investigator for the Modular Linux and Adaptive Runtime support for HEC OS/R research (MOLAR) research team. This multi-lab and multi-education institution research effort concentrates on adaptive, reliable, and efficient operating and runtime system solutions for ultra-scale scientific high-end computing (HEC) as part of the Forum to Address Scalable Technology for Runtime and Operating Systems (FAST-OS). He is also the ORNL lead for the Scalable Systems Software project where his group is developing technologies to scale cluster resources to 10,000's of processors. Stephen has published numerous papers on cluster and distributed computing and has both a Ph.D. and M.S. in computer science. He is also a member of ACM, IEEE Computer, and the IEEE Task Force on Cluster Computing.