

## **Deploying HIL as an automated self-service isolation tool in a Slurm HPC environment**

### **Abstract**

The MOC Hardware Isolation Layer (HIL) provides a minimal, thin API that provides a simple network isolation service for networks of clustered resources. It allows authorized users to execute a limited set of privileged compute node hardware management and network device control commands through a proxy API.

The Slurm system is a widely use tool in HPC environments for managing and scheduling compute resource allocation. In its conventional deployment Slurm provides a rich user level interface for specifying resource requirements that are then reserved for a time period for a specific "job".

Building on HIL and Slurm we have developed a set of integrated tools. These tools allow a regular user to isolate arbitrary sets of compute nodes in an HPC cluster that meet user specified criteria. These nodes can then be made available for reimaging and for other highly privileged activities, in a way that supports automated delivery of self-service bare metal work flows. This is not normally possible in an HPC environment. The resulting system allows users to use HPC resources for new sorts of workflows that in the past have not been practical, due to valid security and overall system stability concerns.

We will describe how the system works and outline some new, potentially valuable, use cases that a combined Slurm + HIL system is facilitating.

### **Presenter**

Chris Hill is MIT Director of Research Computing and a principal researcher in the MIT Earth, Atmospheric and Planetary Sciences Department. For a lot more than 20 years he has been a leader in global development activities of highly scalable open source tools for studying planetary fluid dynamics and related phenomena. He collaborates actively across many disciplines and communities, to advance the application and adoption of computational innovations that accelerate research endeavors.