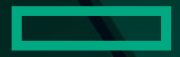

Hewlett Packard
Enterprise

StaX

HPC meta-containers from edge-to-cloud workflow orchestration

pres138

Jonathan Sparks, Technologist, Cloud Solutions

May 11, 2023

Motivation

- Uplevel the conversation in the HPC/AI container community to address workflows
- Address workflow needs and use within HPC, AI, and scientific computing
- How can workflow execute at the edge, core, and in the cloud
- Apply FAIR principles for HPC/AI workflows

StaXs are part of the conversation. Prescriptive workflows are the other

HPE Environments

Addressing the needs of …

Compute diversity

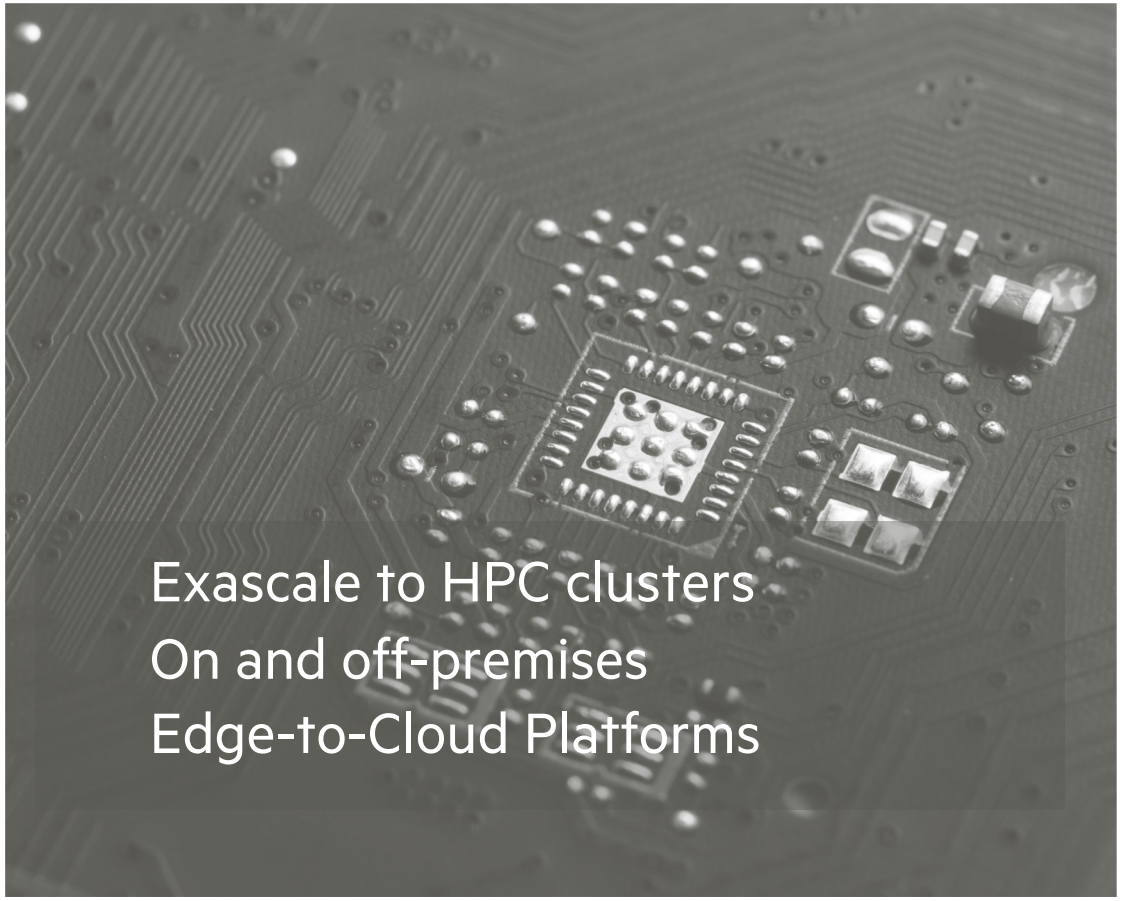

User diversity

Introduction to StaX's

- StaX's aims to make workflows easier to deploy, leverage multiple workflow runtimes, and operate with external systems.
- A StaX is a workload framework containing **executors**, **configurations**, and **applications**
- StaX's are:
  - Mobile
  - Reproducible
  - Self-contained (as much as possible)
  - Validated
  - SaaS

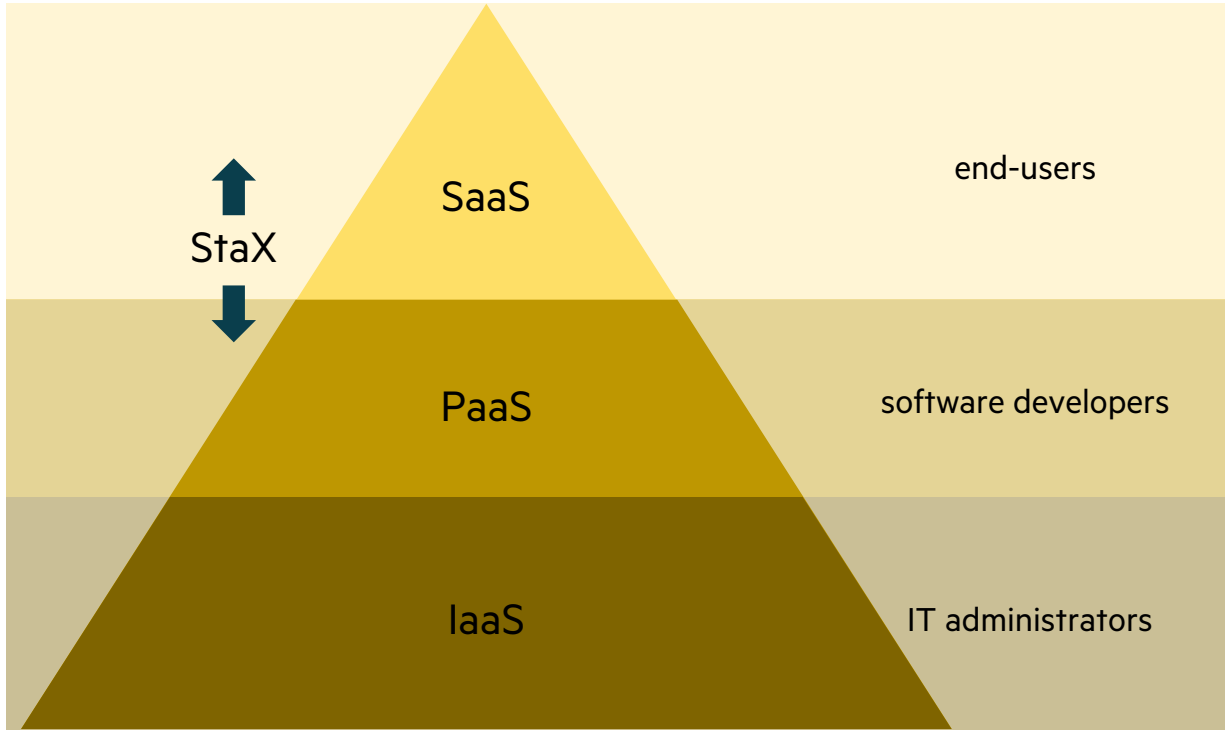

StaX Environment

- Portable environment across different platforms
- Composable software environment

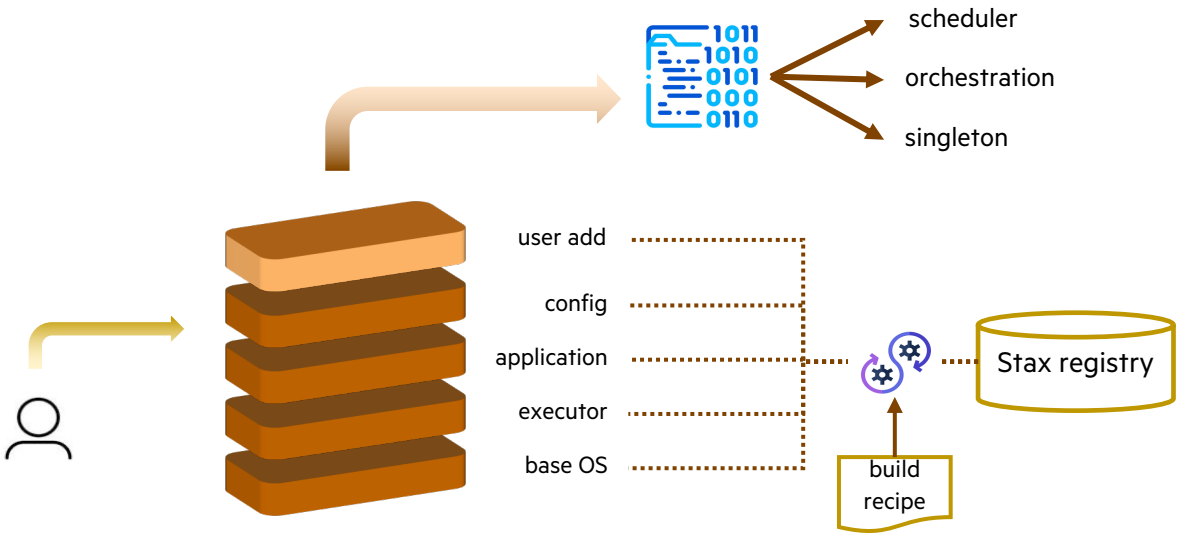

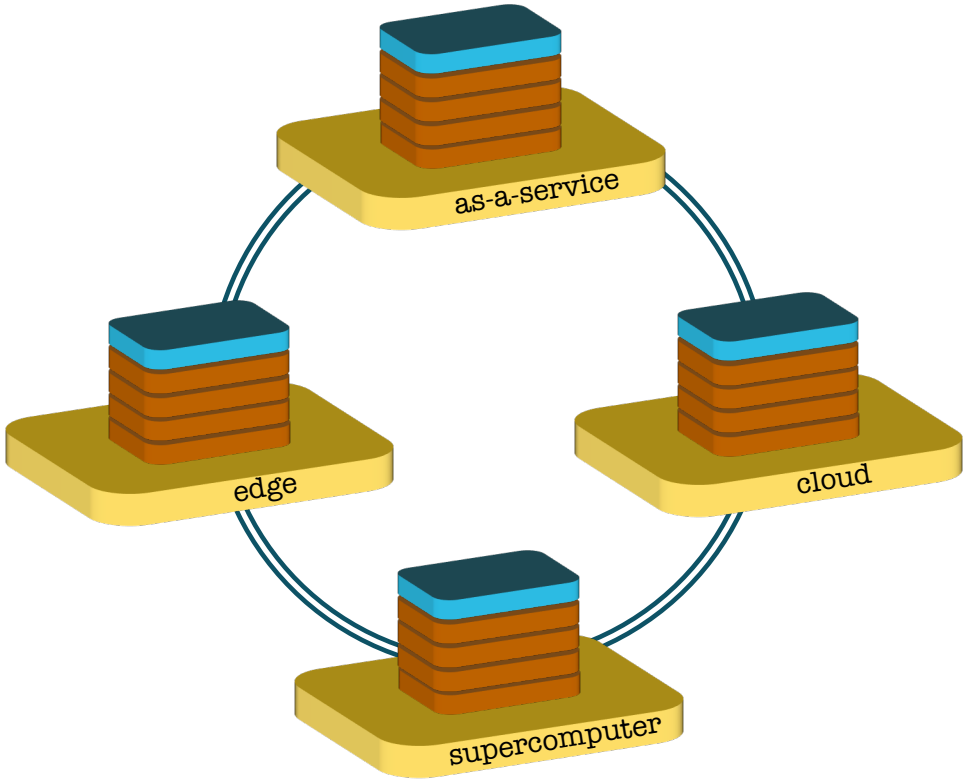

StaX Core Technologies

Workflow Engines

- Streamflow, CWL
- Nextflow
- Covalent
- ExaWorks
- …

Interface

- Common platform architecture
  - Jupyter
  - sos-notebooks
  - nteract
  - …

Platform Services

- Container Runtimes
- Image registry
- Code repository
- CI/CD/CT

Open Workflow Frameworks

Framework	Runtimes/Executors	Jupyter kernel
Nextflow	Executors (AWS, Azure, Local, PBS, …)	Bash
Covalent	Executors (local, ssh, Slurm, Dask, …)	Python
ExaWorks	PSI/J (Slurm, PBS, flux, …)	Python
Streamflow/CWL	Shell – arbitrary	Python/Bash

StaX Build Matrix

MVP StaX

framework

base

foundation

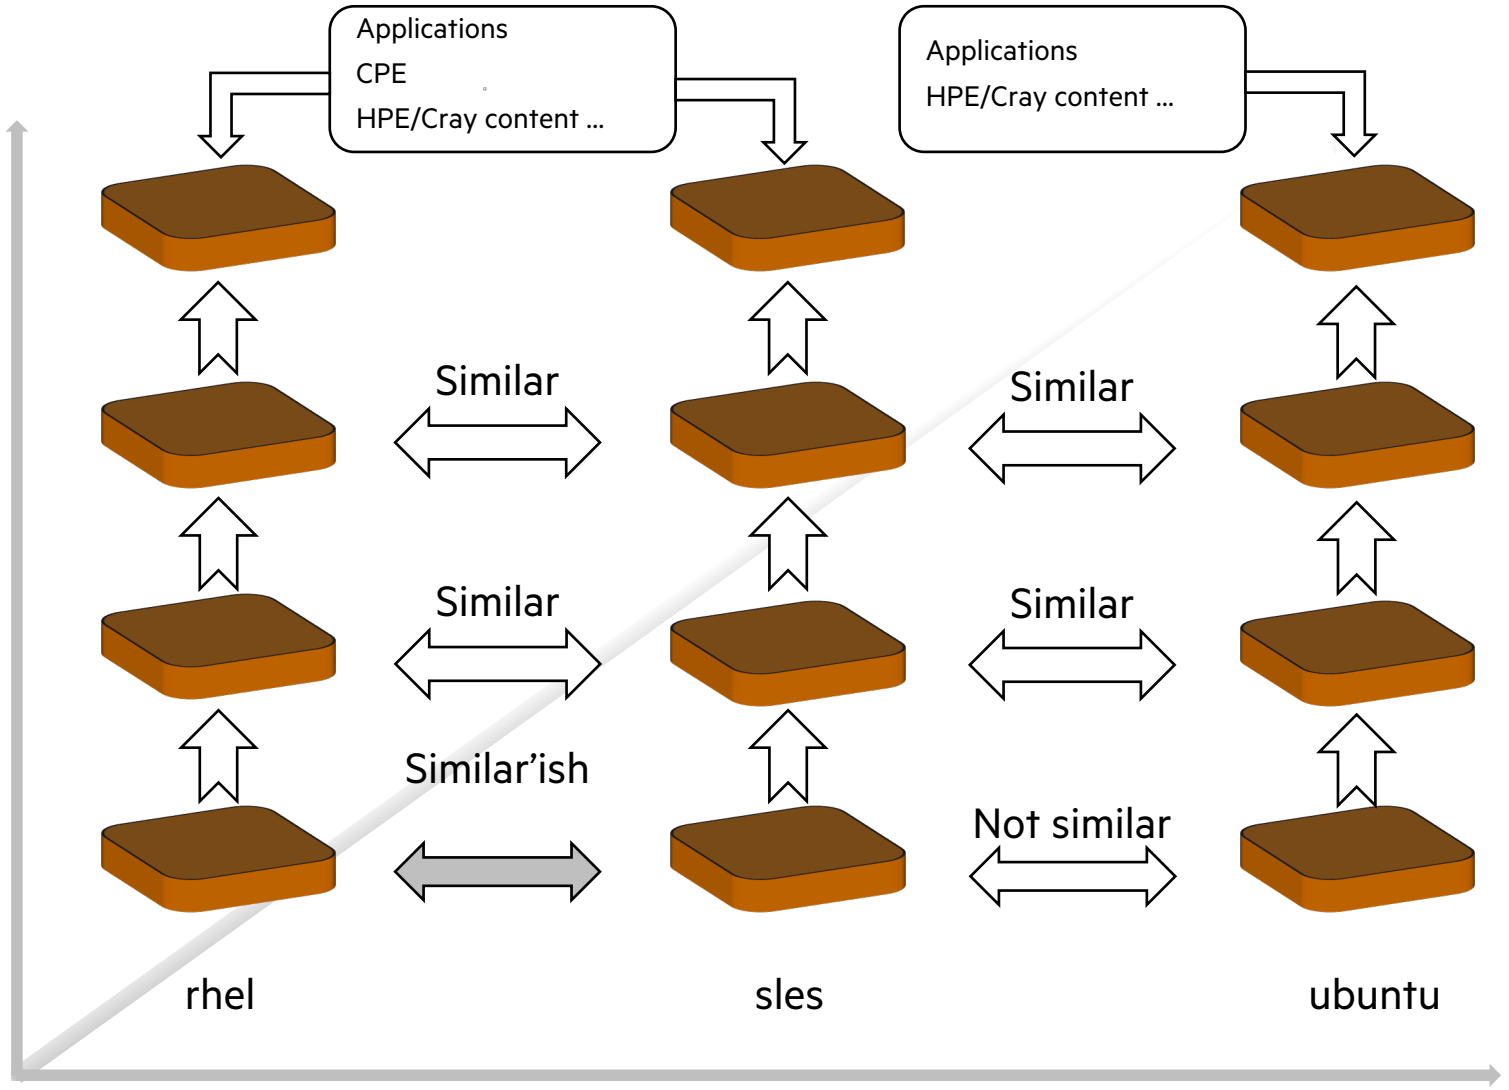

Workflow System

- **Workflow** Interface, description, DSL
- **Frameworks** provides the runtime, such as Nextflow, ExaWorks, Streamflow CWL, etc.

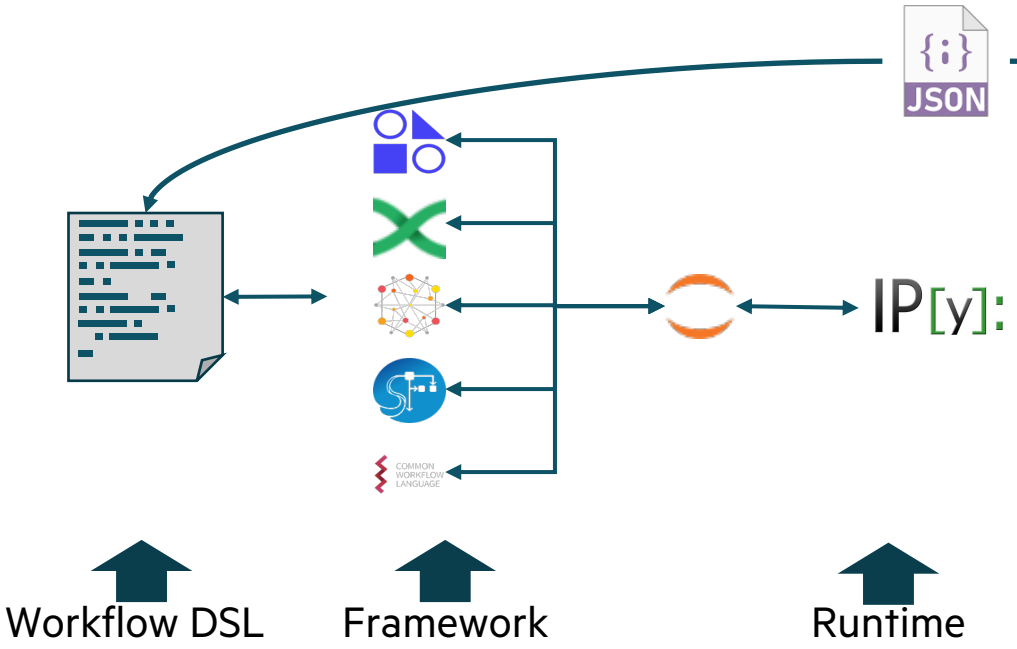


```
"org.hpe.staxs.cpu": "x86_64",
"org.hpe.staxs.framework": ”nextflow",
"org.hpe.staxs.gpu": "none",
"org.hpe.staxs.kabi": "2.28",
"org.hpe.staxs.mpi": ”mpich",
"org.hpe.staxs.network": "none",
"org.hpe.staxs.os": "Red Hat Enterprise Linux 8.7 (Ootpa)",
"org.hpe.staxs.platform": "jupyter",
```

Why Jupyter?

- Notebook framework of choice among data scientists, extensible framework
- Jupyter de facto notebook application, but by no means the only solution
  - Apache Zeppelin (Apache Foundation@2013)
  - Rodeo
  - Beaker Notebook
  - SOS Notebook
  - Rstudio
- Promote Jupyter to a first-class citizen to ssh for HPC users
- Built-in text narrative/code execution/workflow layout

Interfacing with external system services

- StaX framework calling remote/external orchestration executors

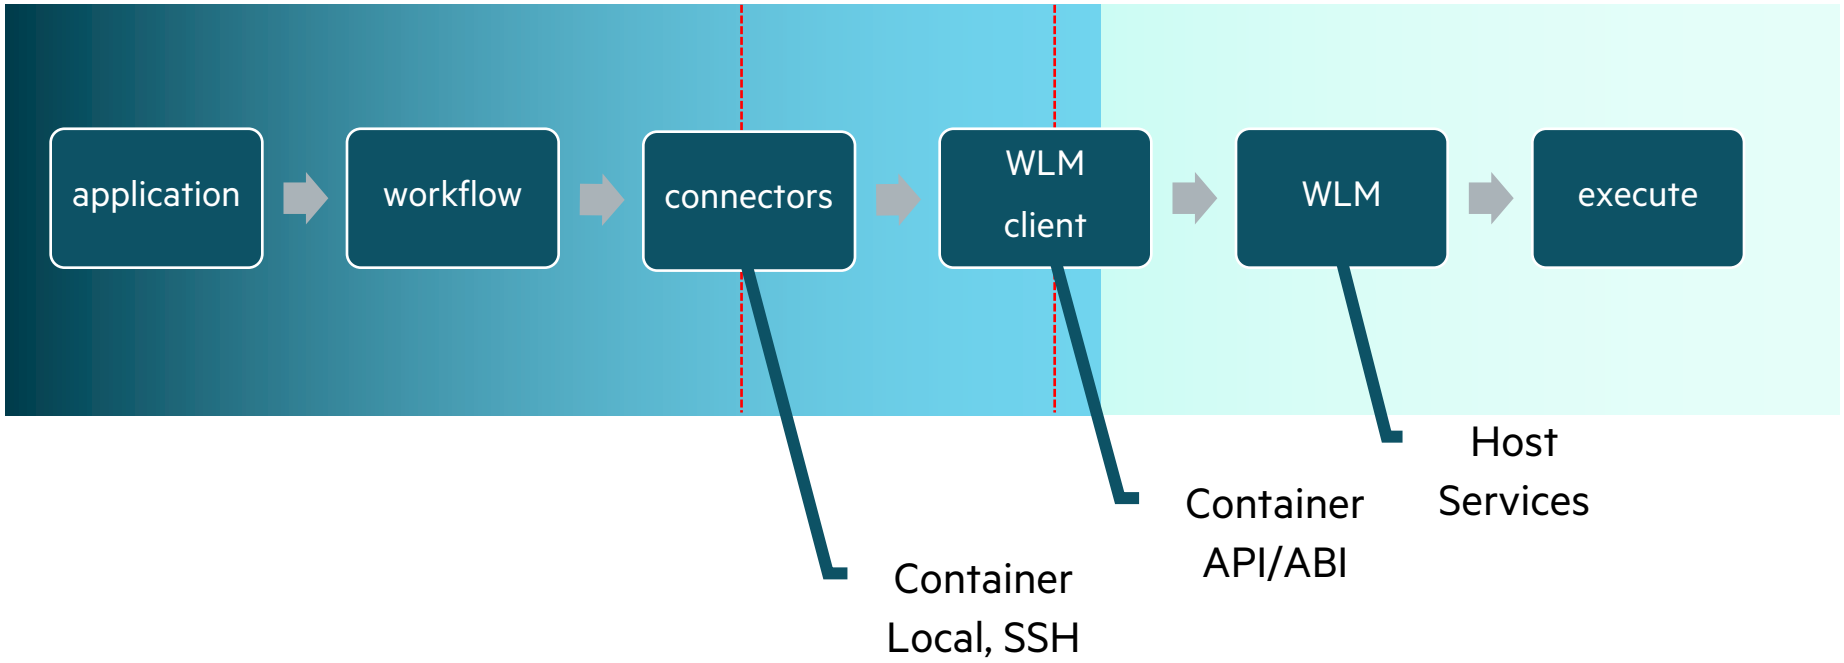

Wrap up, challenges, and next steps

- Demand for consumption of workflows and marketplaces
- Interfacing between host and StaX services
- Use cases and customer demand

Thank you

Jonathan “Bill” Sparks, jonathan.sparks@hpe.com