



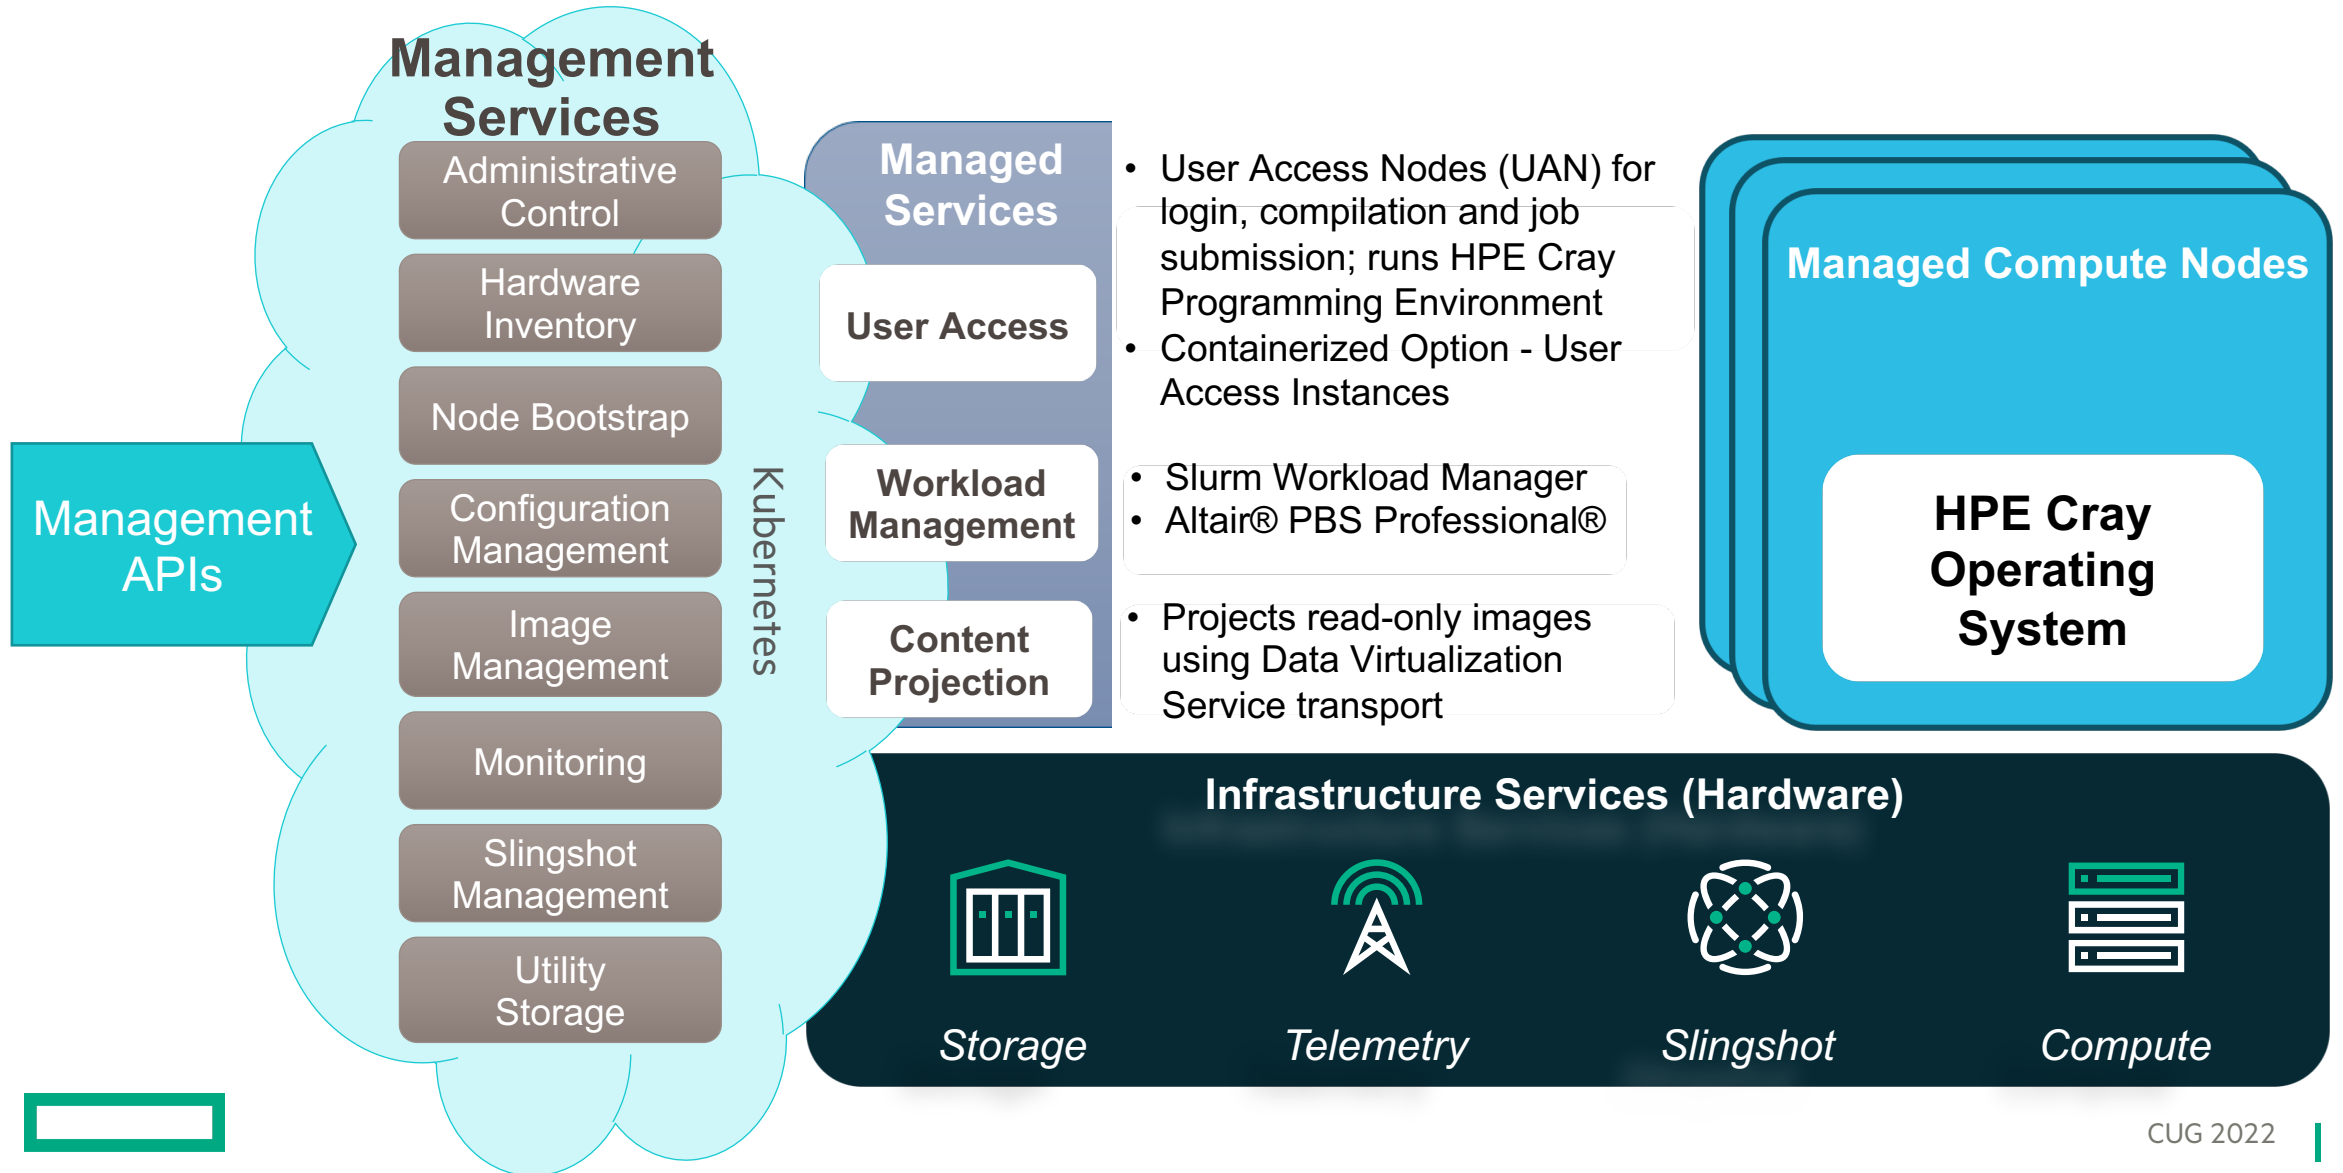
Hewlett Packard
Enterprise

HPE CRAY EX SOFTWARE RECIPE 22.03

CRAY SYSTEM MANAGEMENT OVERVIEW

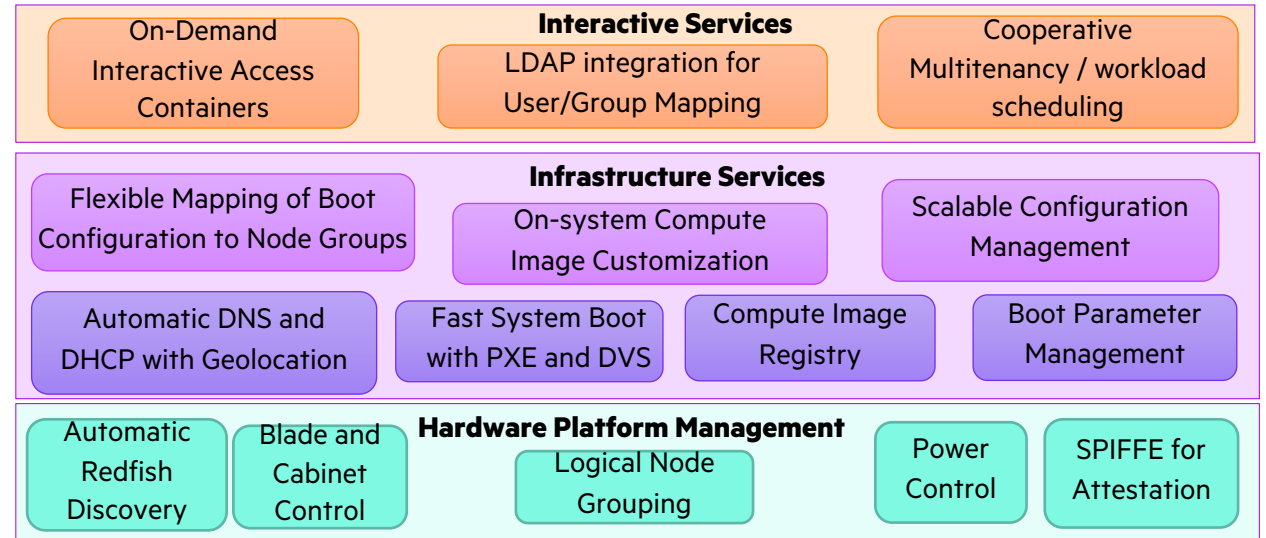
Harold Longley, CSM User Experience Solutions Architect
CUG 2022, May 2-5, 2022

HPE CRAY SYSTEM MANAGEMENT SOLUTION OVERVIEW



HPE CRAY SYSTEM MANAGEMENT IS ELASTIC AND RESILIENT

- Flexible Deployment Options
 - Management Kubernetes cluster scales with more nodes, CPUs, memory, network, and storage
 - Proven to scale from small number of nodes to more than 50 worker nodes for very large customer deployments
- Elasticity
 - Services are continuously checked and updated to match state
 - When nodes are added or subtracted or the load suddenly changes, configuration is automatically modified
 - Autoscale Horizontally and Vertically within constraints
 - When the system is under-scaled, microservices fail according to defined priorities
- Resiliency
 - Microservices are active/active HA
 - Separate gateways and individual load balancers
 - Multiple Pods
 - Rolling deployments and rollbacks
 - Managed nodes running custom app services have HA

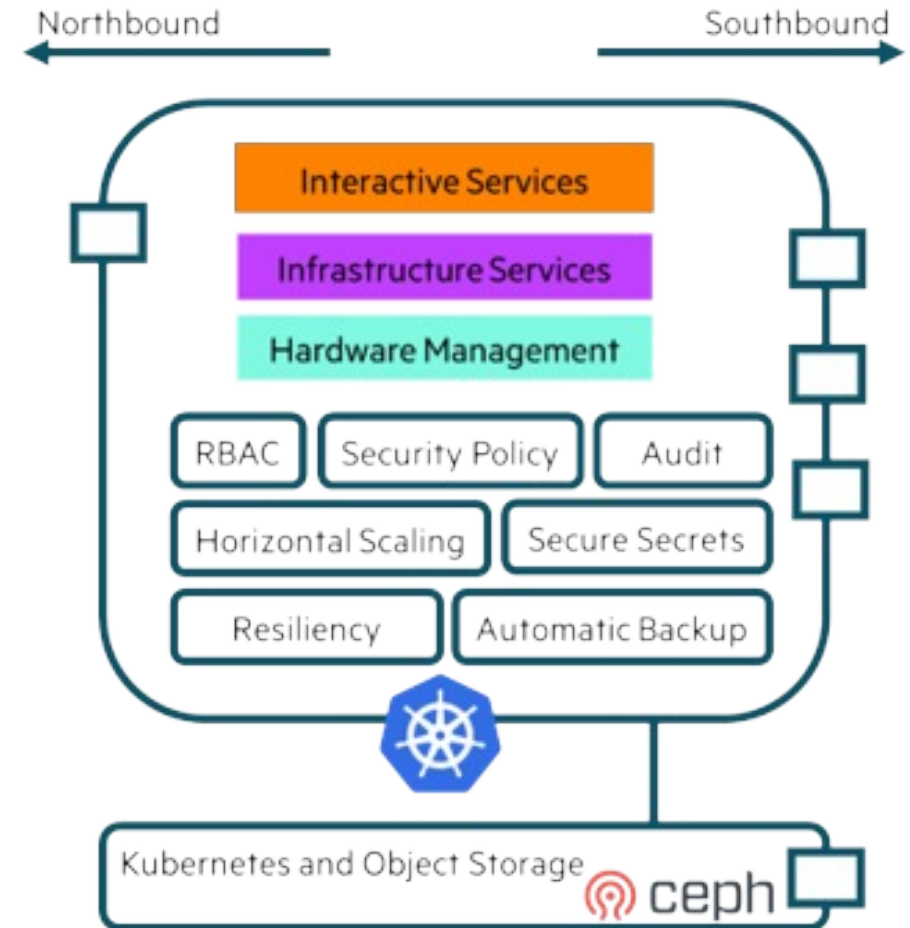


Common footprint

- 3 Kubernetes Master nodes for active failover
- 4+ K8s Worker nodes
- 3+ Utility storage nodes for state abstraction

HPE CRAY SYSTEM MANAGEMENT DESIGNED FOR AS-A-SERVICE SECURITY

- CSM supports human and non-human IAM (Identity and Access Management)
- Fully supported custom RBAC (Role Based Access Control)
 - No limits to the group or role structure, infinite customization
 - Control managed entities with a URL
 - Programmatic interface for change control after upgrades, patches, etc.
- Multiple identity providers
- Credentials management
- Certificate management
- Mesh network encryption (TLS) and access policies
- DNS and external zone transfers
- Non-root users
- User traffic isolation - necessary for multitenancy
- Node attestation



WHAT IS DIFFERENT IN THE SOFTWARE RECIPE 22.03?

- Software installation workflow improvements
- New functionality
- New hardware support
- Security
- CSM enhancements
- Deprecating features
- SAT and SDU enhancements
- CSM-diags



HPE CRAY SYSTEM MANAGEMENT SOFTWARE RECIPE 22.03

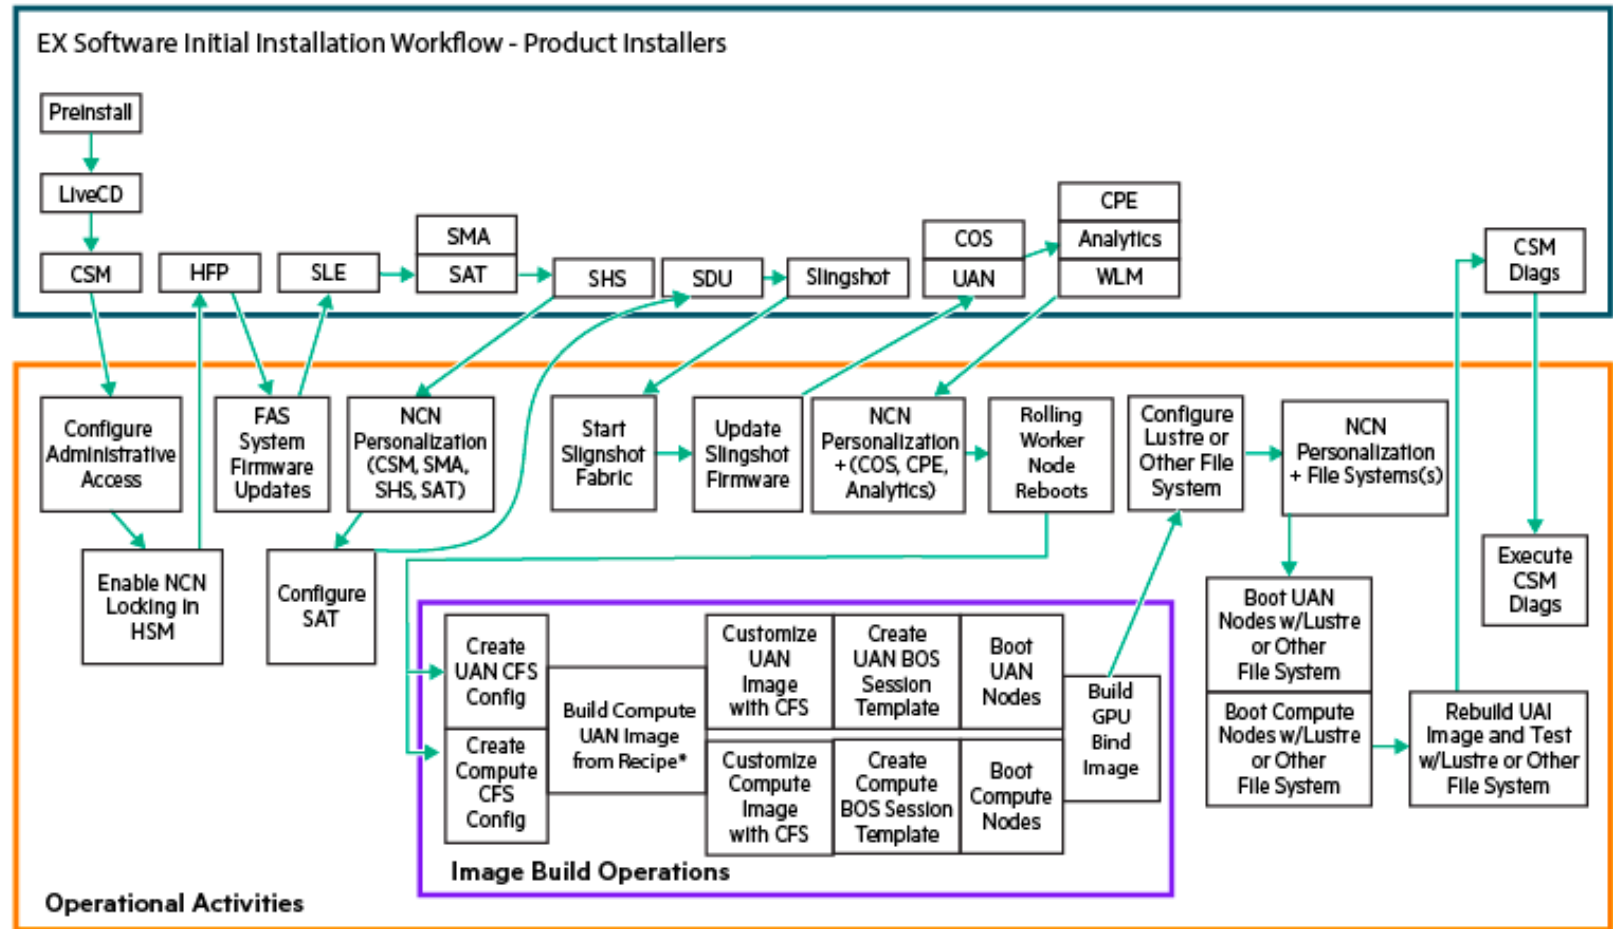
Software Product streams

- Software in recipe
 - HPE Cray System Management 1.0.11
 - HPE Cray Operating System 2.2
 - HPE Cray Programming Environment 22.03
 - HPE Slingshot 1.7.1
- Software product streams (on next slides)
 - Analytics – AI and Analytics software
 - CPE – Cray Programming Environment
 - COS – Cray Operating System for compute nodes
 - CSM – Cray System Management
 - CSM Diags – CSM Diagnostics
 - HFP – HPC Firmware Pack
 - SAT – System Admin Toolkit
 - SDU – System Diagnostic Utility
 - SHS – Slingshot Host Software
 - SLE OS – SUSE Linux Enterprise Operating System
 - Slingshot – High speed network fabric management
 - SMA – System Monitoring Application including monitoring, telemetry and log aggregation
 - UAN – User Access Nodes
 - WLM – Workload Manager (Slurm or PBS Pro)



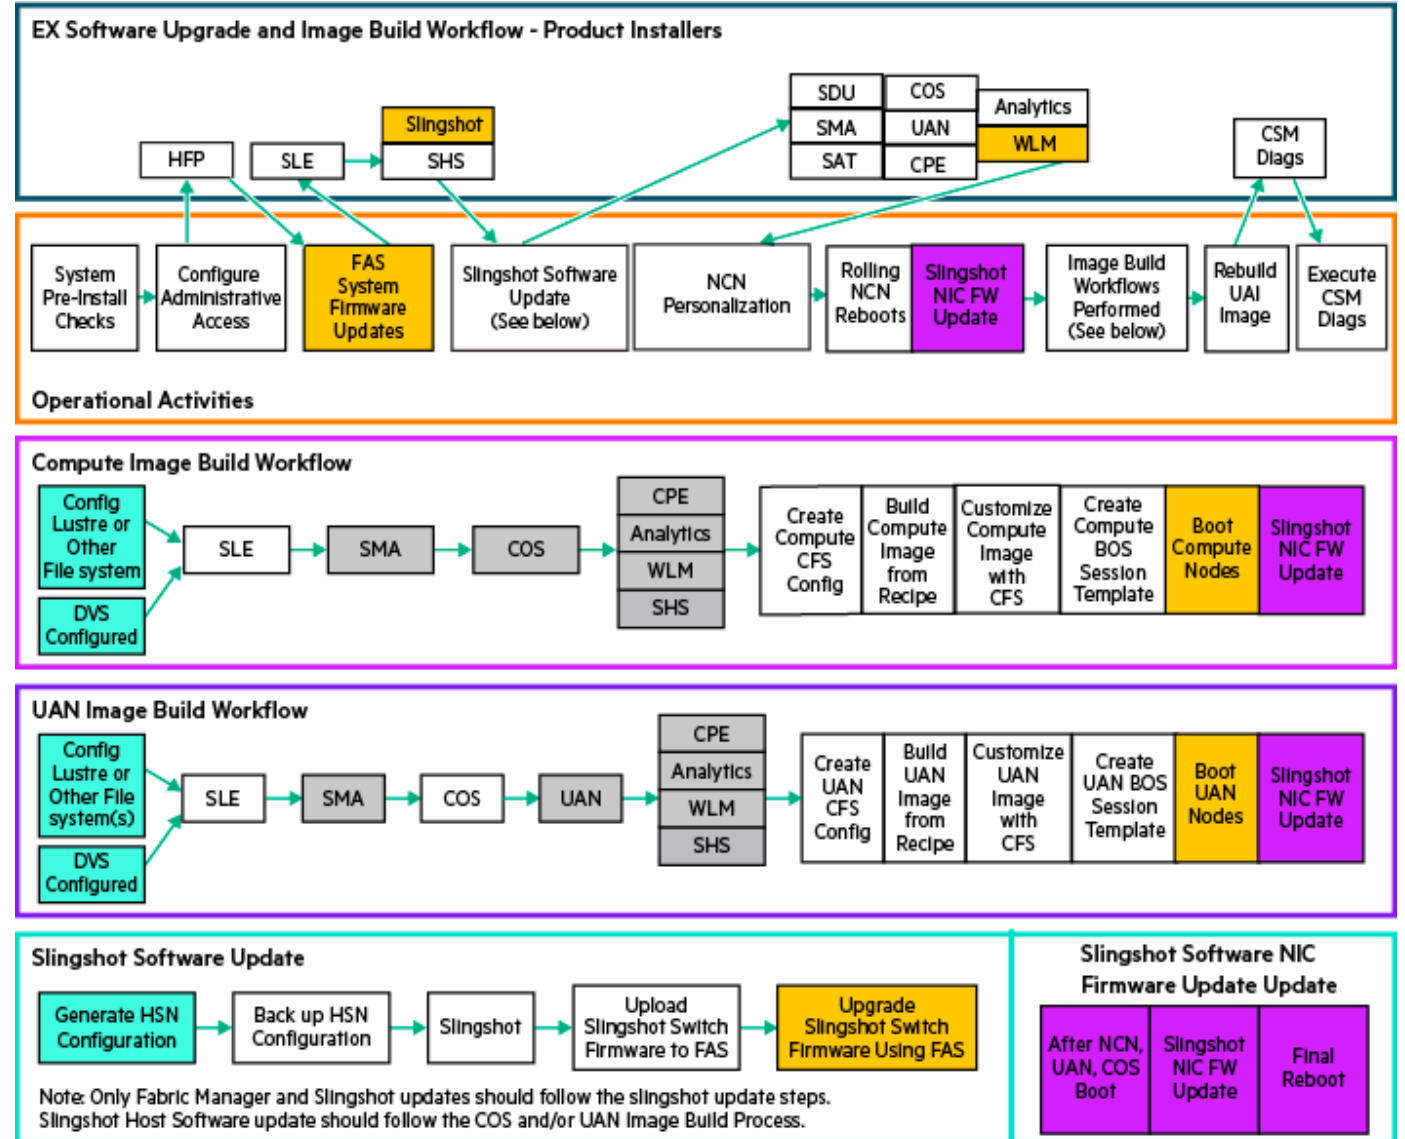
INSTALLATION OVERVIEW

- Workflow for installation of all products
 - Product Installers
 - Operational Activities
- Specific workflows
 - Compute image build workflow
 - UAN image build workflow
- Different order for some steps
 - SHS now uses CFS to configure Slingshot for nodes rather than being embedded rpms in image recipes
 - GPU image build included



UPGRADE OVERVIEW

- Workflow for upgrade of all products
 - Product Installers
 - Operational Activities
- Specific workflows
 - Compute image build workflow
 - UAN image build workflow
 - Slingshot software update
- Use rolling approach to reduce impact
 - Firmware updates
 - Slingshot firmware update
 - Booting compute nodes
 - Booting UAN nodes
- No rolling approach
 - Updating WLM may affect running jobs



NEW FUNCTIONALITY

- CSM installation and upgrade process has been improved
 - A significant number of installation-related enhancements have been integrated
 - Installation validation testing has been improved
 - Upgrades between major versions are now allowed
 - Shasta 1.4/CSM 0.9 can be upgraded to Shasta 1.5/CSM 1.0
- Documentation improvements for adding cabinets during system expansion
- Operating system upgrade
 - SLE 15SP2 for management nodes
 - SLE 15SP3 for compute nodes, application nodes (UANs)
 - UAN and compute nodes use common recipe
- UAN
 - Added kdump support for HPE DL nodes
 - Added GPU support



NEW HARDWARE

- Compute nodes
 - AMD Rome-Based HPE Apollo 6500 XL675d Gen10+ with NVIDIA 40GB A100 GPU
 - AMD Rome-Based HPE Apollo 6500 XL645d Gen10+ with NVIDIA 40GB A100 GPU
 - AMD Milan-Based HPE Apollo 6500/XL645d Gen10+ with NVIDIA 80GB A100 GPU
 - AMD Milan-Based Windom Blade with NVIDIA 40 GB A100 GPU
 - AMD Milan-Based Grizzly Peak Blade with NVIDIA 40 GB A100 GPU
 - AMD Milan-Based Grizzly Peak Blade with NVIDIA 80 GB A100 GPU
 - Bard Peak Blade
 - HPE Apollo 6500 and HPE DL385 with AMD MI100 GPUs
- User Access Nodes
 - AMD Rome-Based HPE DL385 Gen10+ with NVIDIA 40GB A100 GPU
 - AMD Rome-Based HPE DL 385 Gen10+ with AMD Mi100 GPU
 - AMD Milan-Based HPE DL 385 with NVIDIA 40 GB A100 GPU
- Management network switches
 - Aruba CX8325, 8360, and 6300M switches



SECURITY

- CSM Ansible Plays have been created to update management node OS Passwords and SSH Keys
- Significant number of security enhancements to eliminate vulnerabilities and provide security hardening
 - Removal of clear-text passwords in CSM install scripts
 - Incorporation of trusted base operating systems in containers
 - Scanning for vulnerabilities with Snyk while building software artifacts
 - Addressing many critical security CVEs
 - End-of-Life Alpine and nginx container images were removed
- SAT commands use services behind API gateway instead of Redfish directly
 - `sat sensors` rewritten to use the SMA telemetry API
 - `sat diag` rewritten to use cray-fox service to launch controller diagnostics
 - Redfish username and password command-line and config-file options removed from SAT
 - Ability to specify BMC passwords in the plain-text SAT config file was insecure
 - BMC credentials can now be set to different values on different controllers handled by APIs called by SAT



CSM ENHANCEMENTS

- CAPMC simulates reinit on hardware that does not support restart
- cray-console replaces cray-conman
 - Pods scale better to support larger systems
 - Automatically detect nodes which have been added or removed
 - Improved log rotation of conman log data inside cray-console pods
 - Dynamic autoscaling of cray-console-node pods
 - Shared filesystem for all cray-console pods to easily view log data
 - Console log data still sent to SMA for other log processing
- CFS tags enable tracing CFS sessions related to BOS sessions

```
ncn-m001# cray cfs sessions list --tags bos_session=BOS_SESSION_NAME
```
- IMS extracts kernel boot parameters from image and stores in S3 with other boot artifacts
 - BOS session template can have shorter list of kernel parameters
 - Parameters are more likely to be correct if they come from the image rather than the BOS session template
- See separate CUG 2022 talk about User Access Instances
 - UAIs Come of Age: Hosting Multiple Custom Interactive Login Experiences Without Dedicated Hardware

DEPRECATING FEATURES

- Deprecating Features

- HSM v1 REST API has been deprecated as of CSM 0.9.3

- v1 HSM APIs will be removed in the CSM version 1.3 release

- Many CAPMC v1 REST API and CLI features are being deprecated as part of CSM 1.0.1

- Full removal of the deprecated CAPMC features will happen in CSM 1.3

- Boot Orchestration Service (BOS) API in upcoming CSM 1.2.0 release:

- `cray bos --template-body` will be deprecated

- Performing a GET on the session status for a boot set currently returns a status code of 201, but it should return a status code of 200

- `/v1/session/{session_id}/status/{boot_set_name}`

- `sat linkhealth` removed

- Use Slingshot Topology Tool (STT) instead



SAT AND SDU ENHANCEMENTS

- SAT (System Admin Toolkit)

- `sat bmccreds` provides a simple interface for interacting with SCSD to set BMC Redfish credentials
- `sat bootprep` creates CFS configurations, builds IMS images, customizes IMS images with CFS configurations, and creates BOS session templates for booting compute and application nodes
- `sat hwhist` displays hardware component history by xname or Field-Replaceable Unit (FRU) ID by querying HSM
 - FRU ID was added to output of `sat hwinv`
- `sat status` added the Subrole and Aliases columns to output
 - Useful for management nodes and application nodes
- `sat slscheck` performs a check for consistency between SLS and HSM
- `sat xname2nid` can now recursively expand slot, chassis, and cabinet xnames to a list of nids in those locations

```
ncn-m# sat xname2nid x3000c0s19,x1000c0s0b0n0
nid[000001-000004,1000]
```

- SDU (System Diagnostic Utility)

- Collects data for health, inventory, and triage scenarios
- Supports optional Remote Device Access for secure asynchronous file transport to HPE
- Collection of data now includes SDU summary report
 - Metadata about the collection
 - List of all commands run
 - List of files collected
 - Exit code from all plugins

CSM DIAGS

A set of diagnostic tools to perform various node level and system wide tests on compute nodes

- Functional test suites and performance test suites with both MPI and non MPI test suites
- Tests initiated using cray-hms-badger service to submit WLM jobs on compute nodes
 - System Level Diagnostics
linpack, cwlinpack, nodeperf, stream, olcmt, oldisk, olconf, cwolconf, rank, pandora, cwhpcc
 - GPU Diagnostics
gpu-burn, xkbandwidth, xkcheck, xkdgemm, xkmemtest, xkmemtest, xkstress, dgnettest
 - OSU Benchmark
osu_startup, osu_bw_bibw, osu_single_multi_latency, osu_multiplebw_message_rate, osu_multithread_multiprocess_latency, osu_bw_latency_ops, osu_put_bibw, osu_get_acc_latency, osu_collective_blocking_barrier, osu_collective_MPI_blocking_ops, osu_collective_MPI_non_blocking_ibarrier, osu_collective_MPI_non_blocking_ops,
- sdiag_run.py using cray-hms-badger
 - Execute multiple diagnostics (MPI, NON_MPI, GPU, Slingshot) in one shot on multiple compute nodes



THANK YOU

harold.longley@hpe.com

