WHITEPAPER

SUPERCLOUD COMPOSER® ARCHITECTURAL FRAMEWORK

Your Gateway to Tomorrow's Data Center, Including Liquid-Cooling

TABLE OF CONTENTS

Executive Summary

Today's modern data centers face the growing need for operating efficiency and cost reduction in IT spending. At the same time, an organization's infrastructure needs to be agile to respond to changing business requirements. Supermicro understands that IT organizations require a management platform to span multiple generations of infrastructure technology.

Digital transformation, continuous innovation, and an ever-increasing amount of data mean that IT organizations face the rising cost of technology refresh and scale-out of systems. As a result, data center resources are underutilized to rates of 45%, data center operating efficiency is at about 50%, PUE costs are increasing, DC real estate square footage rising, and workforce hour rates climbing exponentially.

Liquid-cooling is becoming a critical technology for data centers as the TDP of CPUs and GPUs continues to increase. The ability of air to remove the heat generated by the servers is becoming more difficult and expensive. The Supermicro liquid cooling solution adds a number of additional components that must be monitored, from the CPUs to an external cooling tower.

The traditional IT paradigm resulted in a cumbersome hardware provisioning process, fixed ratio of compute, storage, accelerator resources, and a lack of one size fits all platform capable of monitoring, telemetry, analytics, and intelligent system management. Therefore, a different approach is required to meet today's business challenges. This technical whitepaper explains Supermicro's approach to software-defined and composable cloud solutions for future data centers.

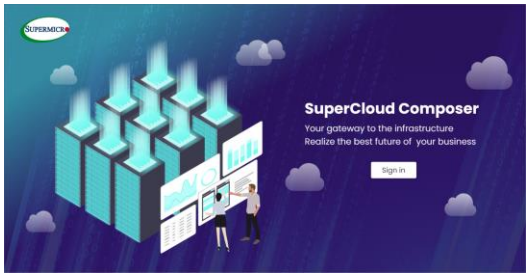

Why SuperCloud Composer?

Modern-day data centers are facing numerous challenges that SuperCloud Composer helps to solve:

- Lack of a single pane of glass platform with a streamlined, intuitive management interface
- The need for a standardized Redfish Northbound API Message Bus
- Unnecessary complexity and lack of scalability in a management platform
- The need for a unified dashboard that encompasses compute, storage, networking, and rack management
- The inability to monitor and manage resource pools in a disaggregated infrastructure
- Using platforms that don't inherently support software-defined and automated processes
- No user-based access control to support modern based data center policies
- Lack of predictive analytics, telemetry, and intelligent system management functionality
- Ability to monitor a liquid-cooling infrastructure, from cold plates to cooling towers

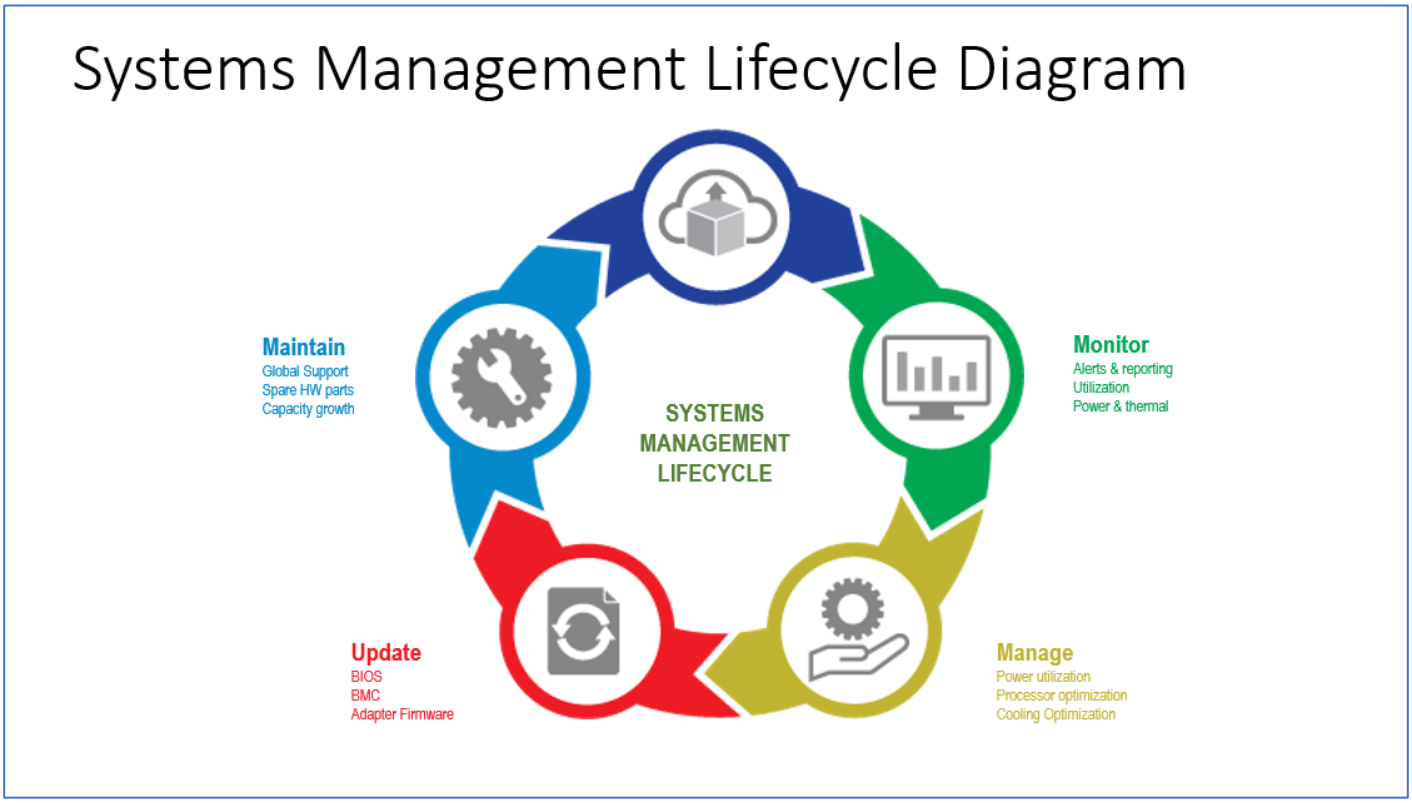

Composable Disaggregated Infrastructure

- Supermicro's SuperCloud Composer brings speed, agility, and simplicity to an IT infrastructure by integrating data center tasks into a single intelligent management solution. Our hybrid approach allows traditional IT paradigm data centers to continue to support their existing operations allowing their workloads to have the flexibility to move to a more software-defined model.

- For those more dynamic workloads, SuperCloud Composer (SCC) provides a composable cloud story that focuses on a disaggregated infrastructure methodology built on NVMe oF and PCI-E switching, utilizing the strengths of standardized Redfish API calls in a consumption-based modeling approach.

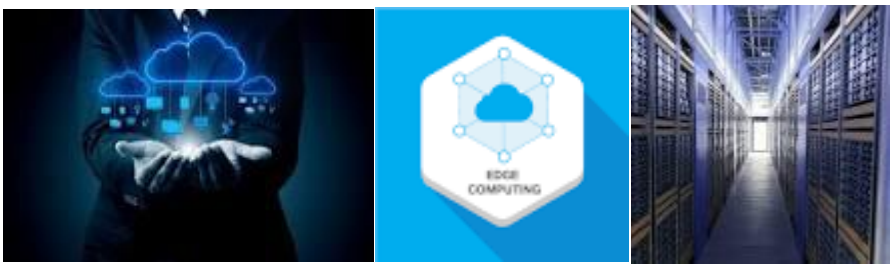

Cloud Enterprise Edge

Target Verticals

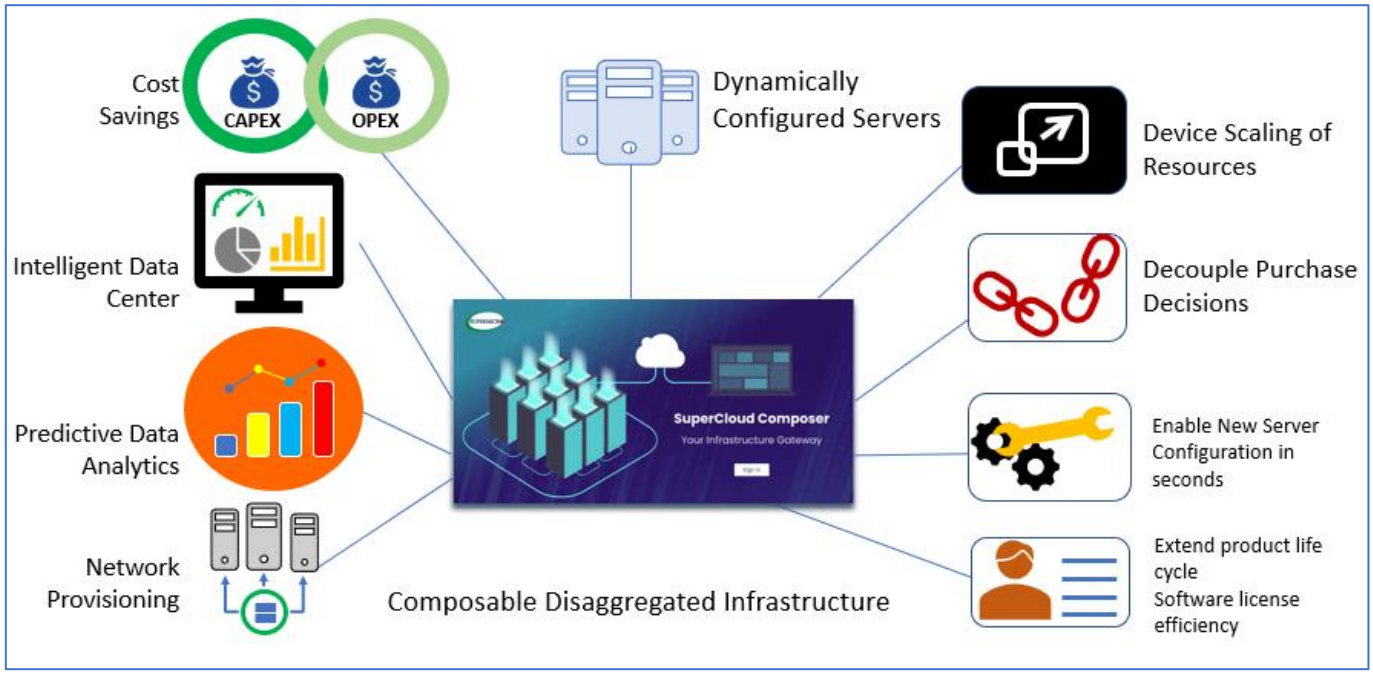

Composable Disaggregated Infrastructure – PCI-E Switching

SuperCloud Composer utilizes a cluster-level PCI-E network to deliver unprecedented performance, composability, and ease of use. This capability is provided by FabreX technology from GigaIO Networks.

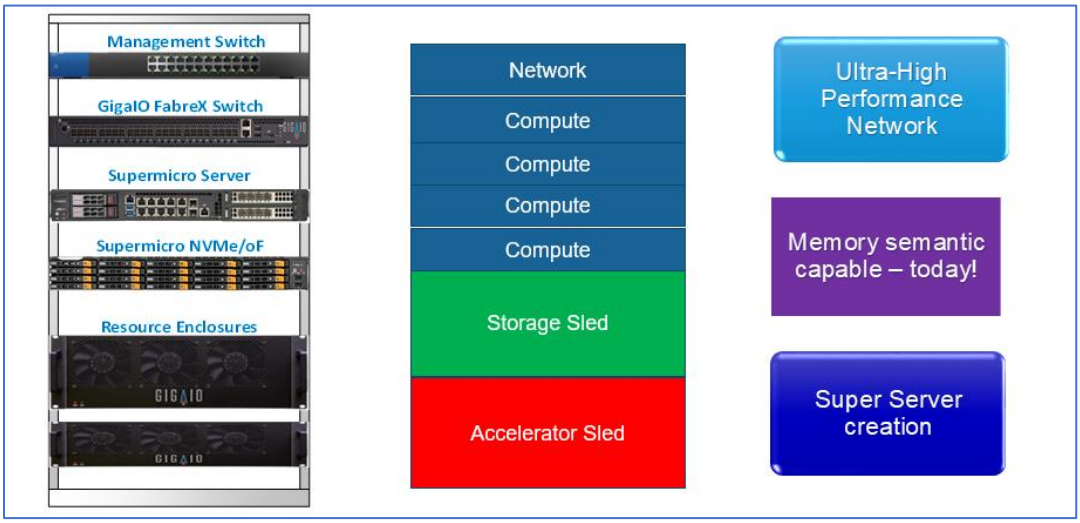

Scalable Reference Architecture Cases

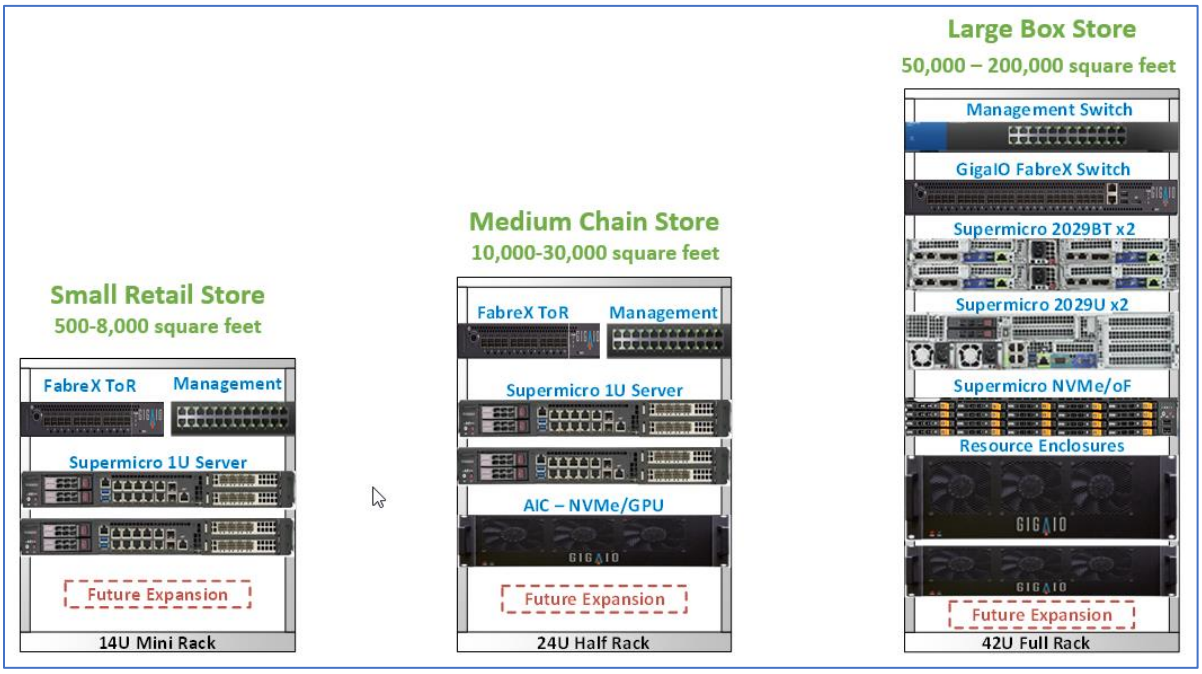

Infrastructure Management Functional Map

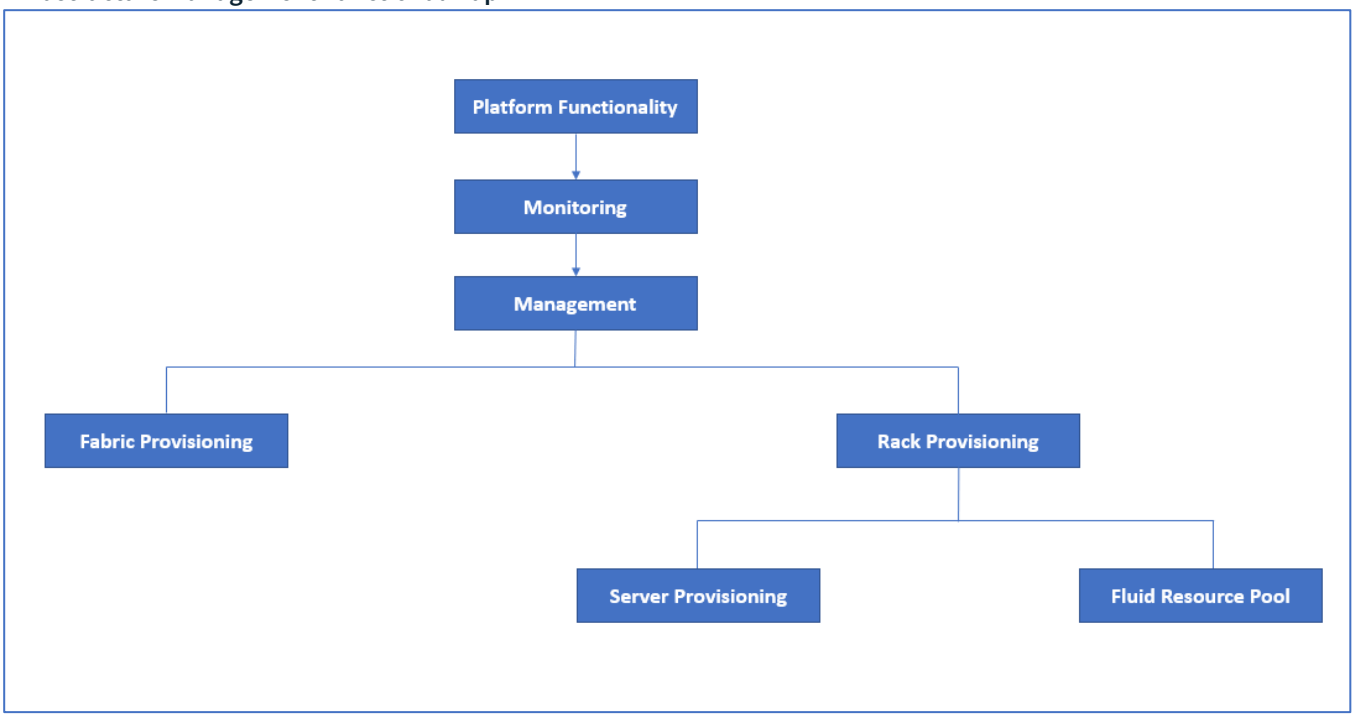

SuperCloud Composer Core Functionality and Features

Let us begin to dig deeper into the core strengths of SuperCloud Composer. To understand IT infrastructure management, it is essential to understand the critical capabilities of SuperCloud Composer, how its framework is built with standardized Open API Redfish, and how the WebUI utilizes a robust modularized angular front end.

- SuperCloud Composer streamlined install utilizes tools such as ansible playbooks for CLI based installs or an intelligent GUI based guided setup.
- SuperCloud Composer can be easily deployed as either a VM appliance or bare metal server utilizing a Linux based kernel of Ubuntu 18.04 LTS.
- SuperCloud Composer's integrated platform can improve productivity across every team member of compute, storage and networking because it focuses on robust build plans through intelligent guided wizards.
- SuperCloud Composer's architecture utilizes an open-standard based ODIM (Open Distributed Infrastructure Management) framework where cloud build plans and API plugins can be easily integrated to suit any kind of workload.
- SuperCloud Composer is tightly integrated with liquid-cooling systems to present an overall view of the entire data center.

SuperCloud Composer Core Framework

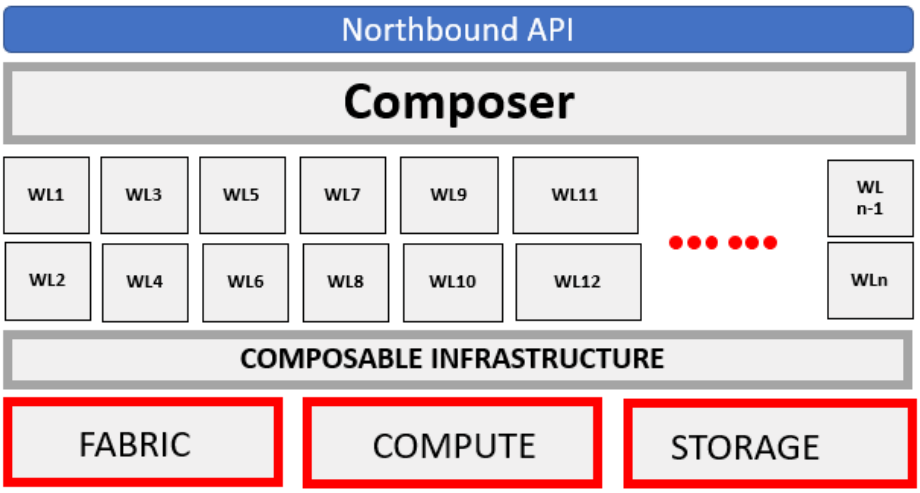

SuperCloud Composer List of Core Features

Our time-to-value intuitive web interface supports the administrator with the following features:

Dashboard	SNMP Traps	Firmware Notification
Pod View	Rack Management	Firmware Repository
Monitoring Detail	User scopes	Call home service
Composed Node	Directory Services	Analytics
Compute	License based monitoring	Fabric provisioning

Network P\provisioning	Infrastructure management	Network monitoring
Storage pools	Disaggregate infrastructure	FRU management
OS provisioning	DC management lifecycle	DMI metadata
Composable cloud	Accelerator provisioning	Hypervisors

Enhanced End-User Experience

- The SuperCloud Composer GUI functions on desktops fully embrace a web experience supporting the Chrome browser. In addition, it builds functionality around best practices every time, allowing administrators to streamline their management and configuration practices.
- Each end-user is provided with analytic capable charting represented by IOPS, network, telemetry, thermal, power, composed node status, storage allocation, and system status.
- For those DevOps environments requiring fast OS deployments, SuperCloud Composer can deploy operating systems within seconds, injecting custom metadata at the time of OS build.
- The WebUI front end is a Redfish RESTful programmatic overlay to its back end services called through standardized Redfish APIs.
- The management platform is completely redundant, utilizing an Active/Passive High availability DRDB cluster.

Industry Standard Redfish API Enhanced End-User Experience

- SuperCloud Composer supports an extensive collection of custom developed Redfish APIs supporting its back end service agents.
- Security is always an important concern of any enterprise management platform, so all Redfish API calls require HTTPS.
- For more specific standardized Redfish API schemas, visit the site: http://redfish.dmtf.org

User-Based Access Control Enhanced End-User Experience

- By default, SuperCloud Composer uses local user accounts. However, for those production environments where security is more of a concern, added directory services have been provided to support either OpenLDAP or Active Directory. A directory service comprises a fallback service that allows an end-user to be granted access to the local administrator account, preventing them from being locked out from the management platform.
- Supermicro enforces password complexity of its accounts to prevent an illegal break in attempt.
- A random retry interval will be enforced until the local account has been locked.

Local User Roles Enhanced End-User Experience

- Global Admin: Full Access
    - View, create, edit, or remove resources managed or monitored by the 1U appliance, including managing the appliance, through the UI or using REDFISH APIs.

- Viewer: Read only
    - View managed or monitored resource information.
    - Cannot add, create, edit, remove or delete resources

- Network Admin:
    - View, create, edit, or remove networks, trunk sets, VLANS, configuration parameters to TOR switch.
    - Execute Network Configurator, Network Orchestrator, and Storage Fabricator

- Architect:
    - Create and manage server profiles, server profile templates, storage volumes
    - Monitored and manage compute hosts, access the physical server through BMC console, update operating system drivers, BIOS, firmware baseline, firmware installation method, operating system deployment, and BaaS provisioning
    - Composition of a Node

Dashboard

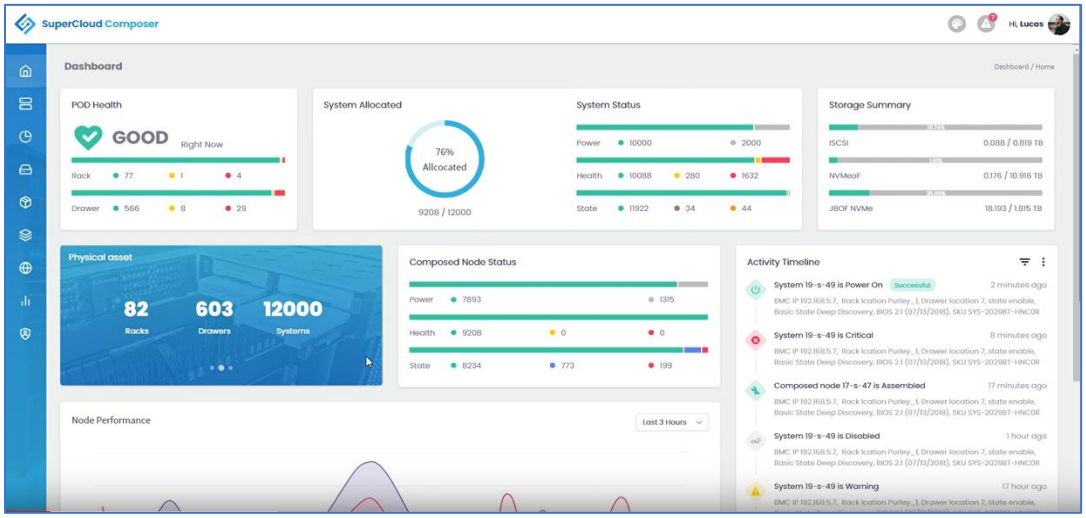

- SuperCloud Composer dashboard is an information management tool used to track, Composer health analytics through data visualizations, activity event timeline tracking utilizing standardized icon footprints, providing the administrator at a glance awareness of data center operations.
- Administrators are given the flexibility to click on each component within the dashboard to learn more detailed metadata about system status, composed node status, and storage allocated.
- In addition, the dashboard provides the end-user with the flexible charting of node performance, power, and thermal that can be represented in time intervals.

POD View

- The pod view is an intelligent engine within SuperCloud Composer that allows the administrator to organize a data center into pods that share common workloads. During Day 1, activities within a data center racks are created and logically grouped as pods.
- Devices that are physically mounted and powered within a rack are automatically discovered from the discovery pool. The entry pool consists of newly created systems, switches, JBOFs, and JBOGs.
- During rack configuration, the administrator utilizes the drawer configuration wizard to add discovered devices from the entry pool.

Monitoring Detail

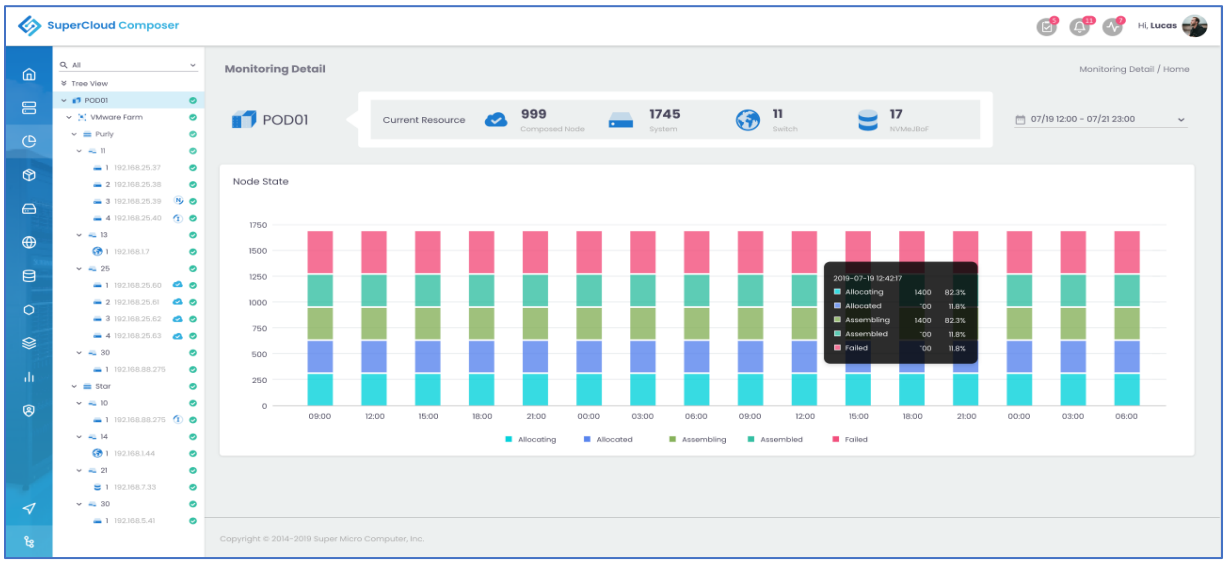

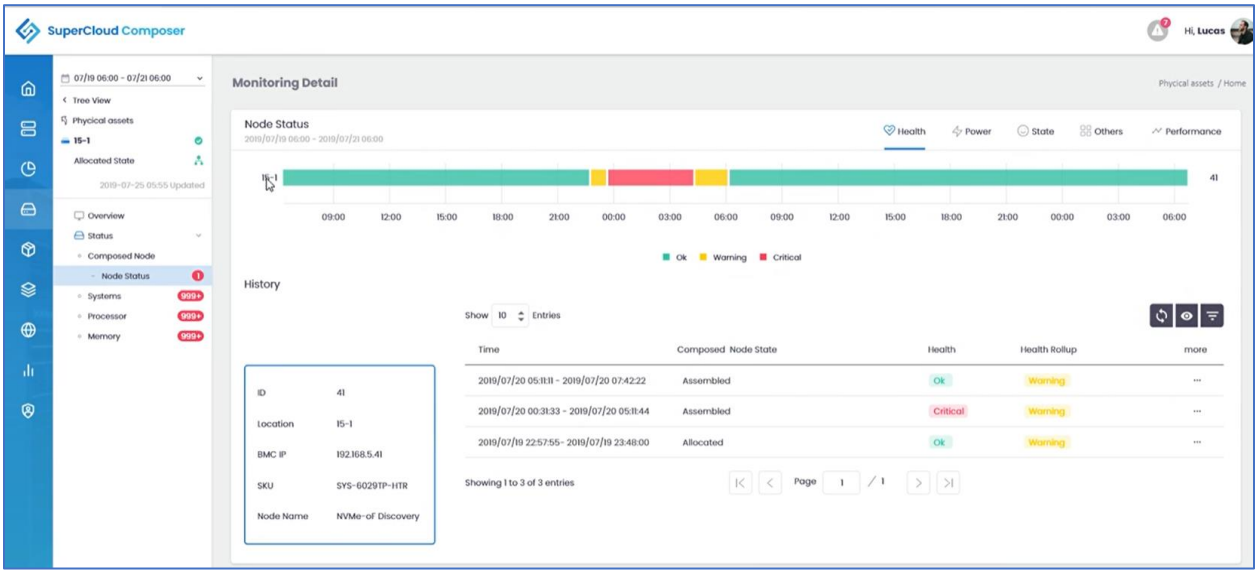

- SuperCloud Composer offers a robust analytics engine that collects historical and up-to-date analytics stored in an indexed database within its framework.
- Each data visualization is represented in charts, graphs, and tables that offer the end-user a rich granular scope providing the administrator at a glance performance metrics, telemetry, update-to-date or comprehensive historical monitoring, predictive analytics, and accurate alert notification.
- The analytics engine completely integrates with a Kubernetes cluster stack in which features such as Kibana, Elastic search, and logstash are available to provide a lookback window of historical trend data.
- Charting and graphing are available at the pod level, rack level, chassis level, and node level.
- Health log sets are captured from the BMC utilizing redfish harvesting and stored in a log repository within the Kubernetes Cluster.
- SuperCloud Composer allows the end-user the ability to make efficient searches of a logstash repository within a Kubernetes container for those data centers that would like to scale out.

Network Component

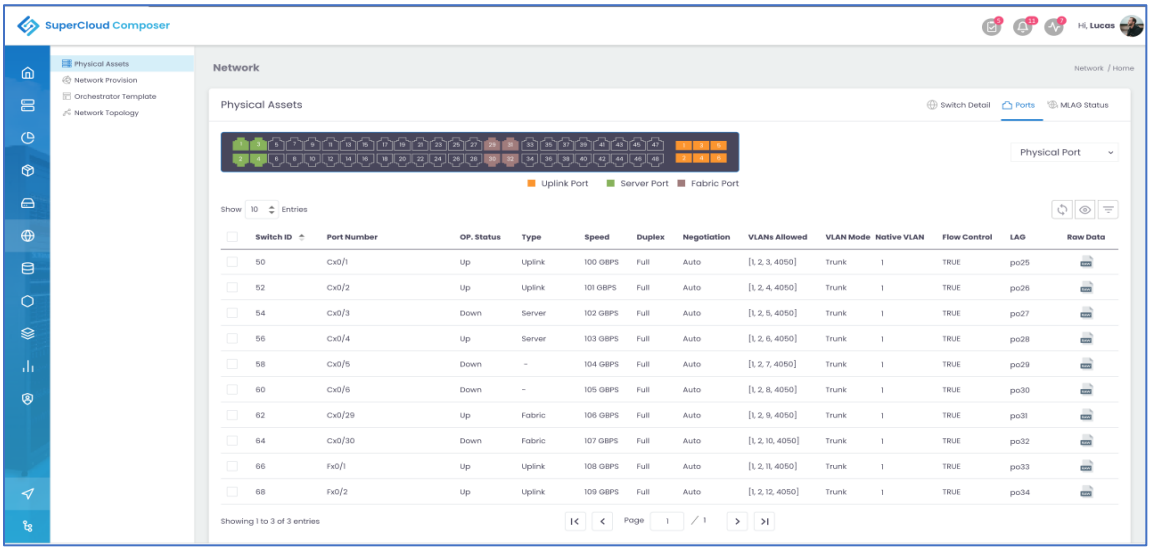

- SuperCloud Composer utilizes a rich feature called network provisioning. Build plans are pushed to data switches either as single-threaded or multithreaded operations, where multiple switches can be updated simultaneously by shared or unique build plan templates.
- Data switch build plan templates are constructed by a Network Configurator wizard formatted in JSON and pushed by a Network Orchestrator engine utilizing industry standardized API calls.
- During network management operation SuperCloud Composer offers a rich, intelligent network agent called switch sweeper to maintain configuration compliance between original build plans constructed by network configurator and operational build plans that exist within switch dynamic memory.

Storage Component

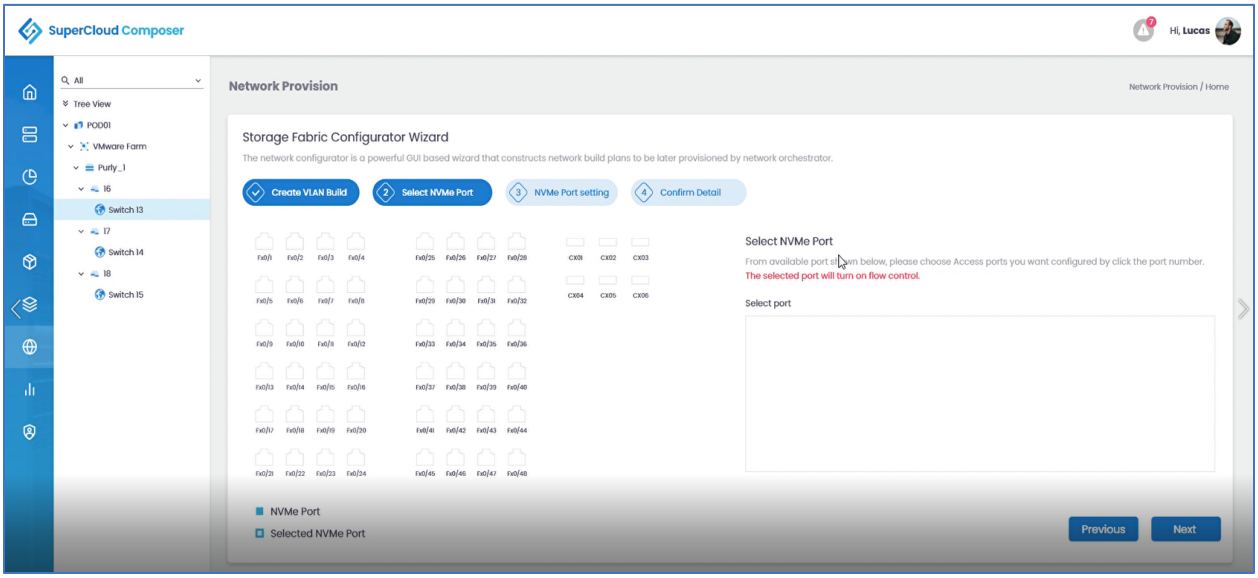

- The Storage component of SuperCloud Composer also utilizes a rich feature called storage fabric provisioning. Build plans are pushed to fabric switches either as single threaded or multithreaded operations, where multiple switches can be updated simultaneously by shared or unique build plan templates.
- Storage build plans are slightly different from traditional compute data switches because they require additional DCBX parameters to support RDMA/ROcE based storage fabrics.
- RDMA/RoCE constructs rely on ETS, flow control, and lossless queues to maintain robust communication between NVMe OF targets, initiators, and Ethernet-based switches.
- Storage switch build plan templates are constructed by a Storage Fabric Configurator wizard formatted in JSON and pushed by a Network Orchestrator engine utilizing industry standardized API calls.
- During network management operation SuperCloud Composer offers a rich, intelligent network agent called switch sweeper to maintain configuration compliance between original build plans constructed by storage fabric configurator and operational build plans that exist within switch dynamic memory.

Storage Options for SuperCloud Composer

PCI-E Switching Solution

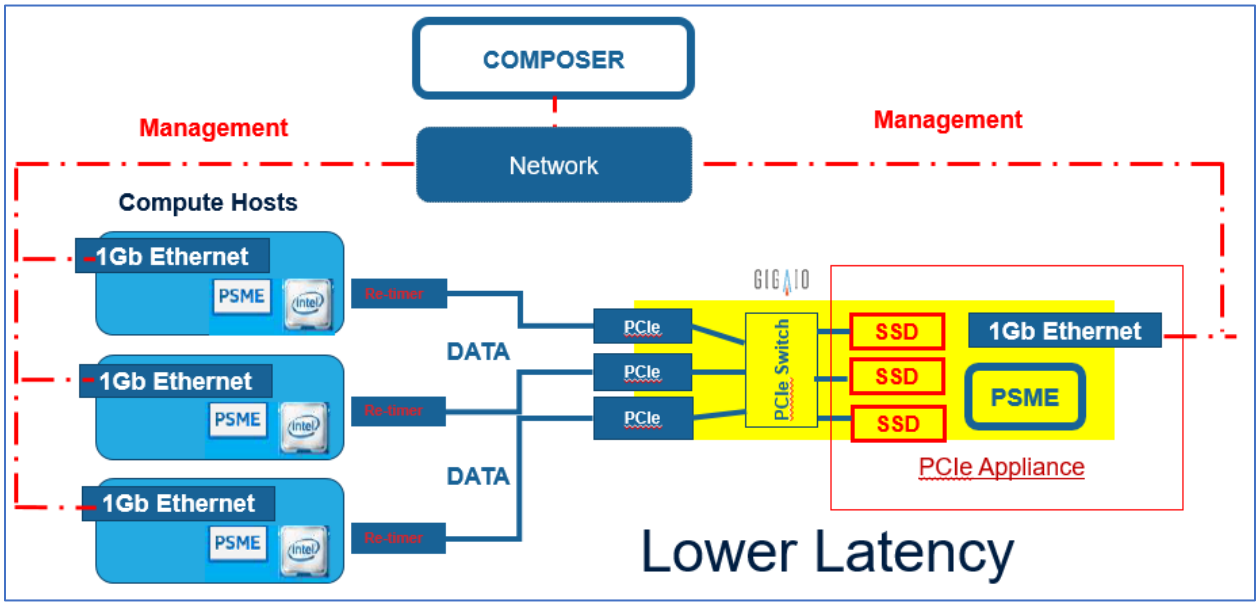

- Built on FabreX Technology from GigaIO Networks
- Network technology enabling true disaggregation and high-performance computing that support a data center and edge Infrastructure
- Realized gain in heterogeneous compute for AI, IoT, and HPC applications
- Transfers data faster, safer, cheaper utilizing a low latency PCI-E bus
- Unique PCI-E compliant network with native support for MPI, TCPIP, NVMe-oF, and GPU Direct RDMA technology
- PCI-E appliances to host industry-standard NVMe drives, FPGAs, and GPUs
- Flexibility to scale up and scale-out systems utilizing the same interconnect
- Performance advantage in latency and bandwidth

Ethernet Solution

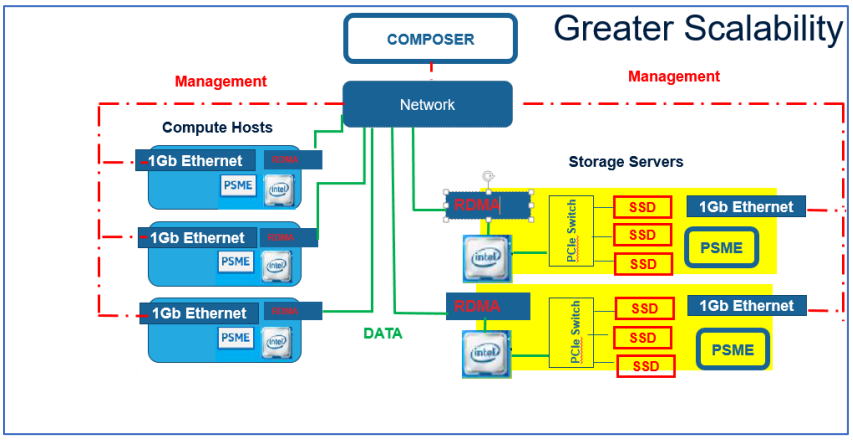

- SuperCloud Composer NVME OF technology requires RDMA/RoCE featured NICs in bare-metal servers to transfer NVMe storage commands across an Ethernet switch fabric.
- Ethernet storage fabric switches require provisioning of standardized IEEE DCBX parameters in support of NVMe OF solutions
- Flexibility to scale out utilizing the same interconnect
- Minimal integration effort required
- RDMA/RoCE are IEEE standardized features built into IT industry ethernet switches

Storage Pools

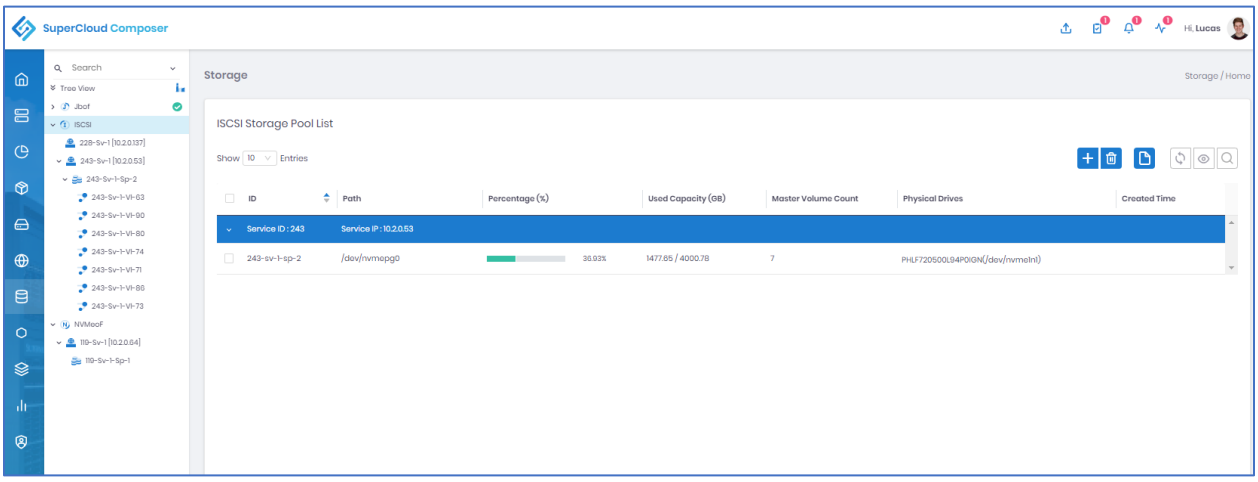

- The storage component of SuperCloud Composer technology offers two rich storage pool options i) NVMe oF storage Pools and ii) ISCSI storage pools.
- NVMe-oF targets and initiators can be configured utilizing Ethernet based or PCI-E switch fabrics.

- Dependent on the role that a bare metal NVMe fabric host will participate in, Supermicro provides tarball images and kickstart files for both Centos and Ubuntu-based architectures stored within a software repository container by default.
- For those workloads in which end-users decide to configure disk-less servers, both ISCSI storage pools and NVMe-of pools could be viable options.
- Each of the storage pools will offer raid technology in support of high availability redundant arrays.

JBOF (Just Bunch of Flash)

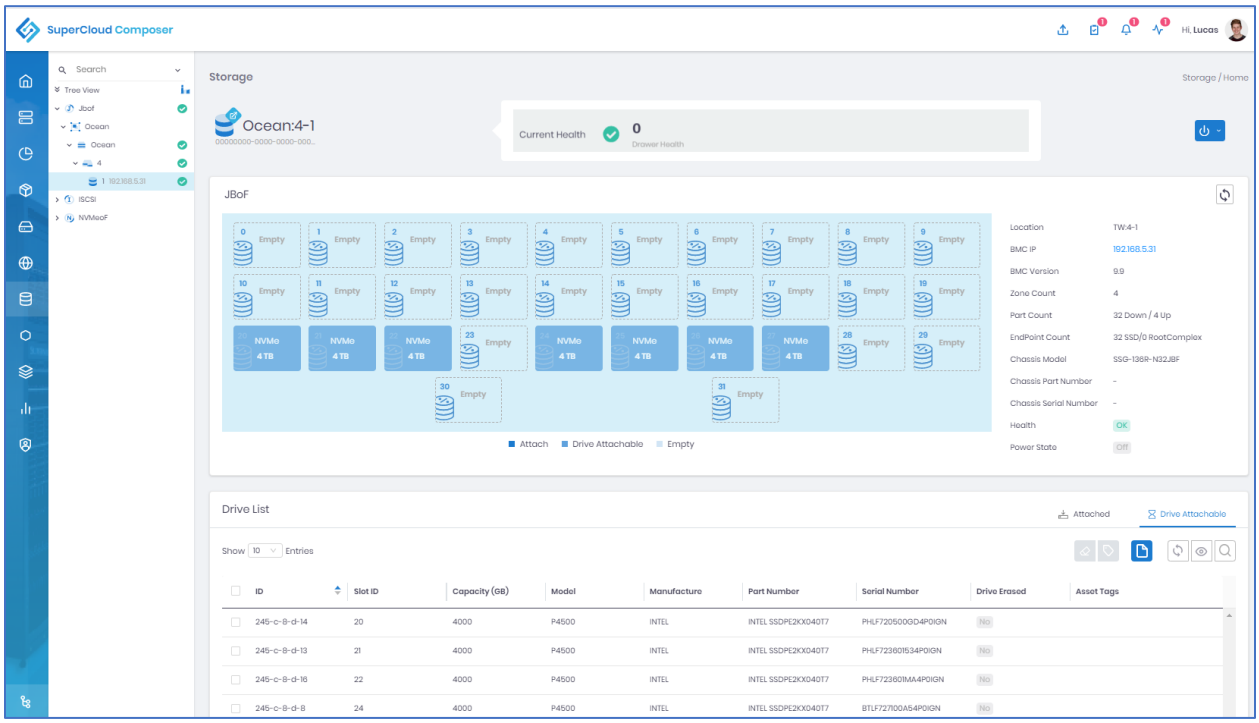

- SuperCloud Composer continues to strengthen the JBOF management experience by exposing a drive map aid to give the end-user visualization of unoccupied/occupied drive drawer collateral within the JBOF shelf configuration.
- The JBOF unit allows the end-user to attach/detach drives as needed, providing rich raid functionality and a drive erase service revealed within the JBOF unified API.
- An end-user has the flexibility to drill down to more detail for each drive end-point element within the drive list table.

OS Provisioning – Software Inventory Repository

- The software inventory repository is a folder that stores ISO distros to be utilized by a PXE boot service within SuperCloud Composer.
- The server architect is given the flexibility to either choose a file browser option or URL to upload a standard ISO image.
- Pre-defined answer files are available to be uploaded from an end-user desktop.
- Administrators are also provided the flexibility to create answer files from the WebUI GUI as well.
- By default NVMe fabric preseed files will be available for both Centos and Ubuntu to support NVMe-of workloads.

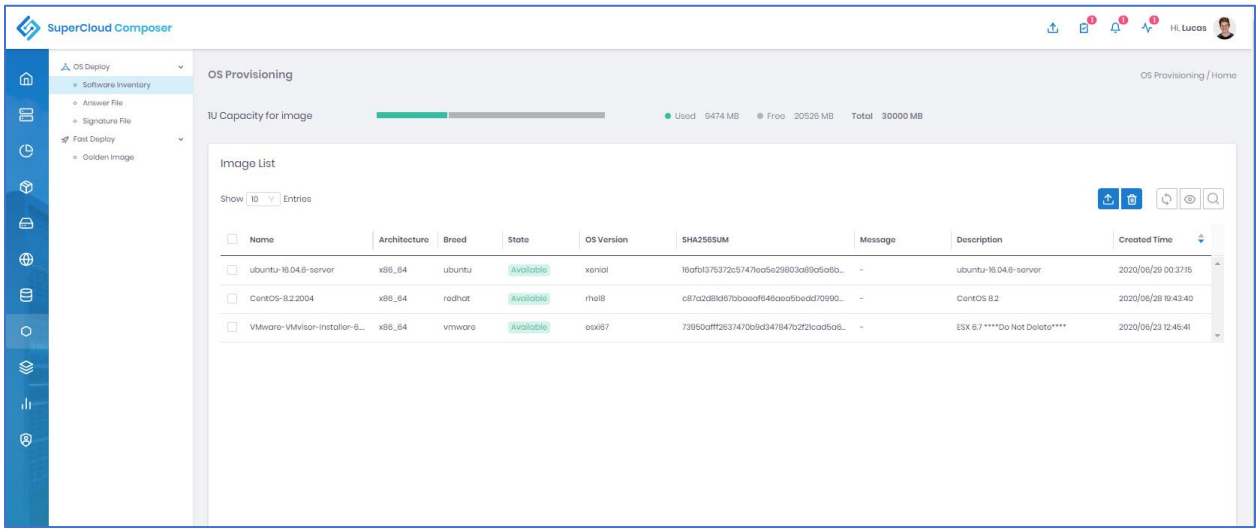

OS Provisioning – Golden Image Repository

- A golden image is a template-based OS software that has never been altered in any way within SuperCloud Composer. End-users will upload an image either as a customized base raw image or a raw image that has been built from a standard distro.
- Once an image has been uploaded to the golden image repository, it will become available to the ISCSI service pool as an available image to become either a snapshot or clone for targeted bare metal servers to utilize.
- During the provisioning phase of OS deployment, the replica image will be altered with customized metadata from the fast-deploy GUI wizard.

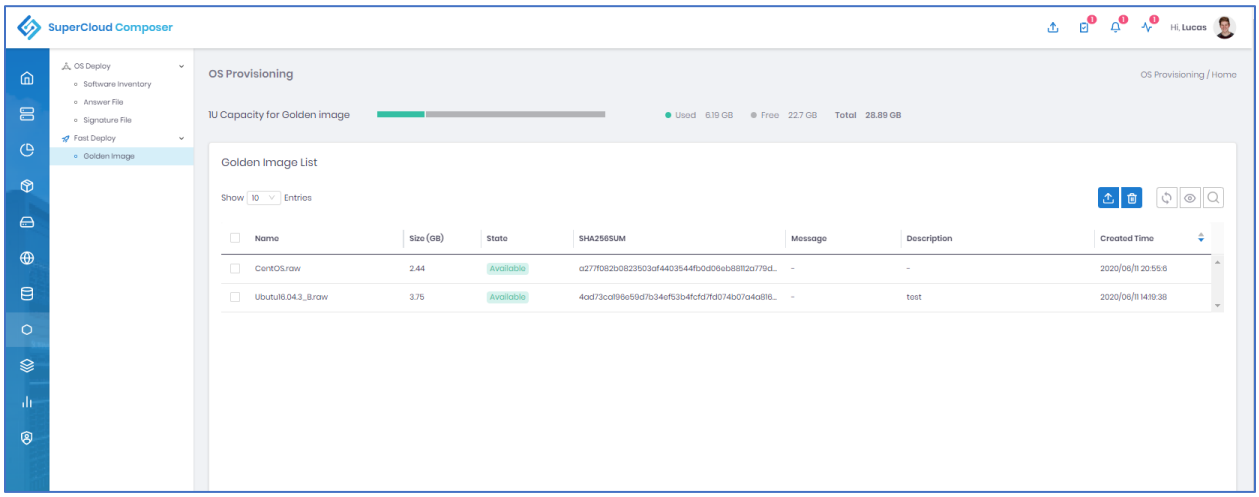

- The composition feature within SuperCloud Composer focuses on offering those Composable Platform end-users the flexibility to orchestrate and re-allocate from a fluid pool of resources. Our Rack Scale Design framework allows us to pool GPUs, FPGAs, and storage option architectures utilizing a consumption-based model approach. Use the resource when a workload requires it and when the workload has been completed, place it back in the pool.
- SuperCloud Composer offers three types of composition features combined in a single pane of glass front end GUI. Each option steps the end-user through a series of customized queries based on processor, memory, local drive, security, local NVMe storage, remote storage, and networking.
- Our intuitive, robust GUI front-end wizards allow end-users to step through a composition phase, seamlessly customizing their liquid IT servers with personalized metadata to support their dynamic workloads.
- There are three types of OS deployment build models supporting either cloud providers or traditional IT data center workload operations.

Composed Node Management

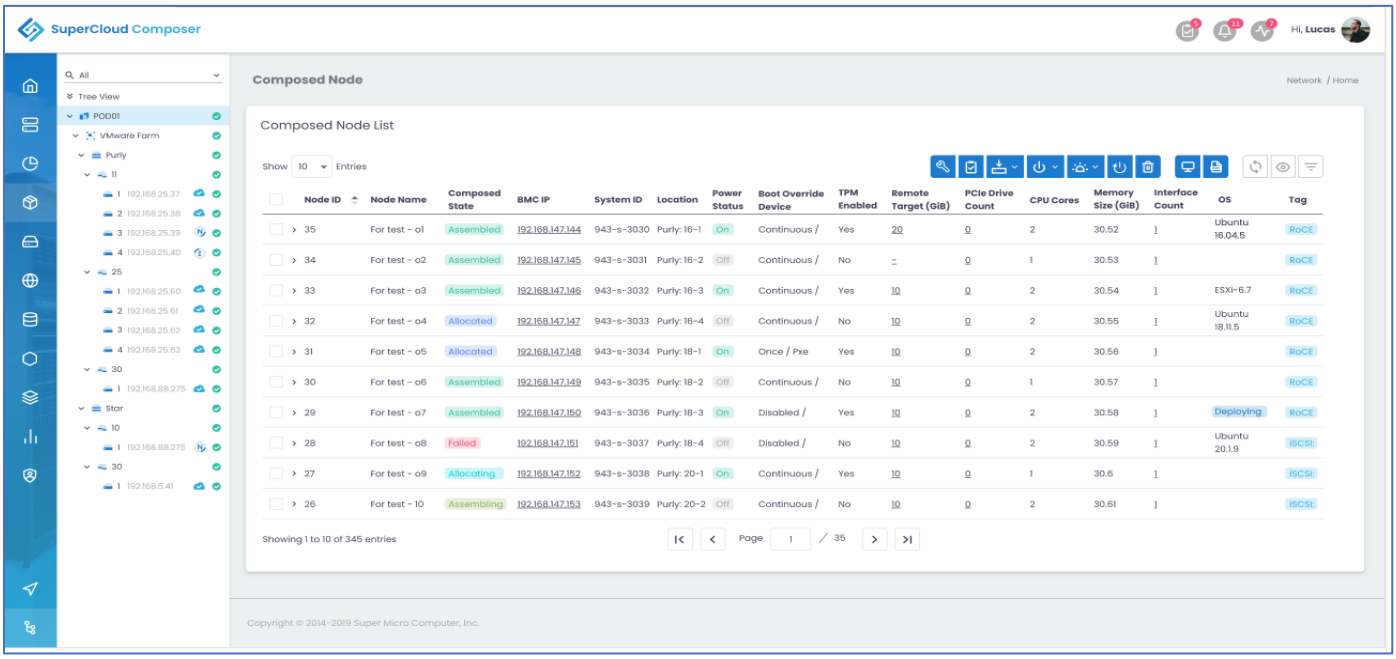

- A summarized table of assets depicted in physical inventory provides administrators a visual aid of hosts assembled and allocated. Each entry of the table emphasizes server facts allowing the administrator to troubleshoot and manage servers within a composable cloud environment.
- Administrators are given the same fast track toolbar where end-users can rapidly launch common tools as assembling a node, deploy OS, attachable resource management, power management, UID tracking, boot mode settings, and IKVM.
- Within the table list, users are encouraged to utilize the search criteria, tag management functionality, and customized table creation to quickly identify servers of interest.
- Along the left pane tree view, customized icons indicate what role a server would play, resources allocated, and general health rollout of a pod, rack, chassis, and server.
- An End-user would only be encouraged to click on the left pane of tree view to expose those servers added to a rack or POD during registration of servers within the drawer configuration wizard of PodView.

General Composed New Node – Private Cloud Model

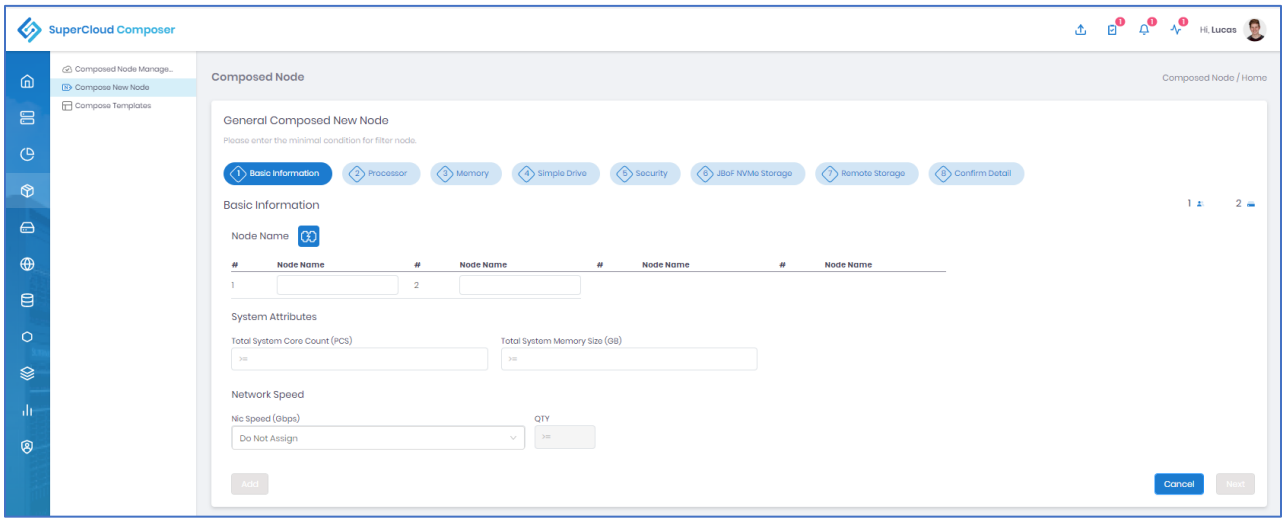

- During the composition of a host, end-users have the flexibility to build customized user-defined templates in which they can use to match certain component criteria of a system within a composable infrastructure pool. These pre-defined templates can be applied to many systems improving the overall efficiency of a data center.

Fast-Deploy – Cloud Service Provider Model

The Fast-deploy features allow an operating system to be deployed in seconds. During the composition phase, the server architect would prepare a customizable template which later would be injected within the replica.

Each fast-deploy deployment relies on a robust ISCSI service pool in which a server platform will utilize a boot from SAN operation.

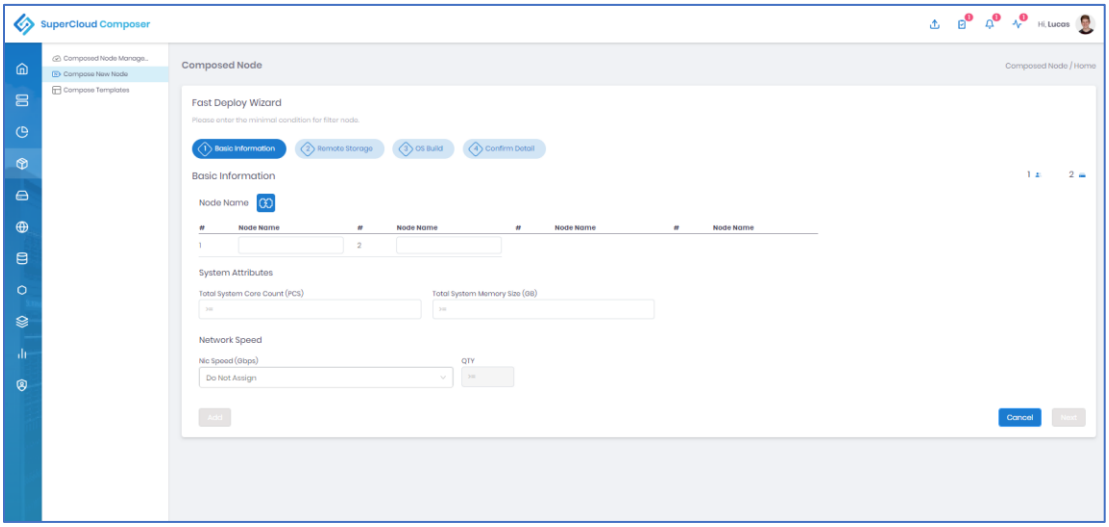

Traditional OS Deployment Model

The OS deploy option gives the end-user the ability to target specific servers instead of relying on a pool of available server resources based on selection criteria, which is utilized through user-defined templates.

A bare-metal host would send a pxe-request to deploy an operating system during the traditional OS deployment process. These deployments typically require 15-20 minutes of execution time and would not generally be used by Cloud Providers or ISPs because of some service level agreement in place.

Compute Component

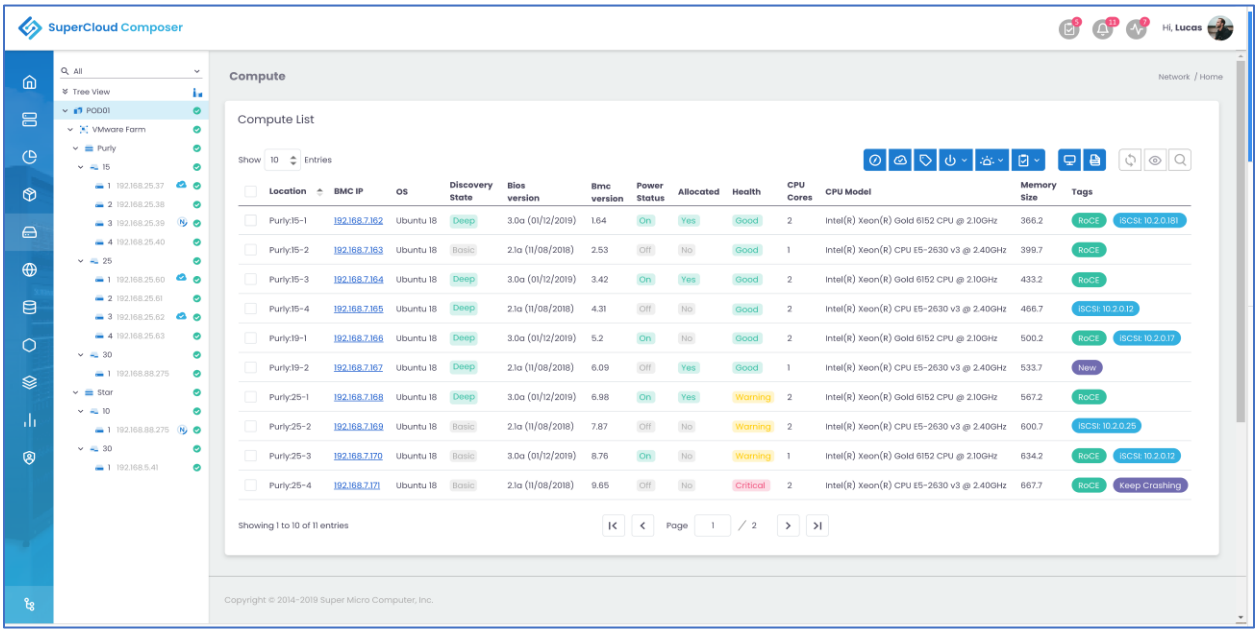

- Administrators can fast track to common tools on the top toolbar of the compute list screen without navigating to other components within SuperCloud Composer.
- Fast track options that are one click away comprise deep discovery, host allocation, tag management, UID tracking, power management, task activity, IKVM, and BMC/BIOS updates.
- All elements in a table list within SuperCloud Composer can be filtered and easily searched.
- End-users are given the flexibility to build their own customized table within fields they select from a list.

Physical Asset Collateral

- A collection of components that expose physical attributes for each bare metal server. The compute subsystem collects rich FRU contents and DMI metadata from the BMC and exposes its collection end-points of memory, CPU, local storage, remote storage, and accelerators within the Compute subsection of SuperCloud Composer.
- Without utilizing the cumbersome external SUM and IPMI utilities, SuperCloud Composer allows the administrator to manipulate FRU and DMI content from a single pane of glass.

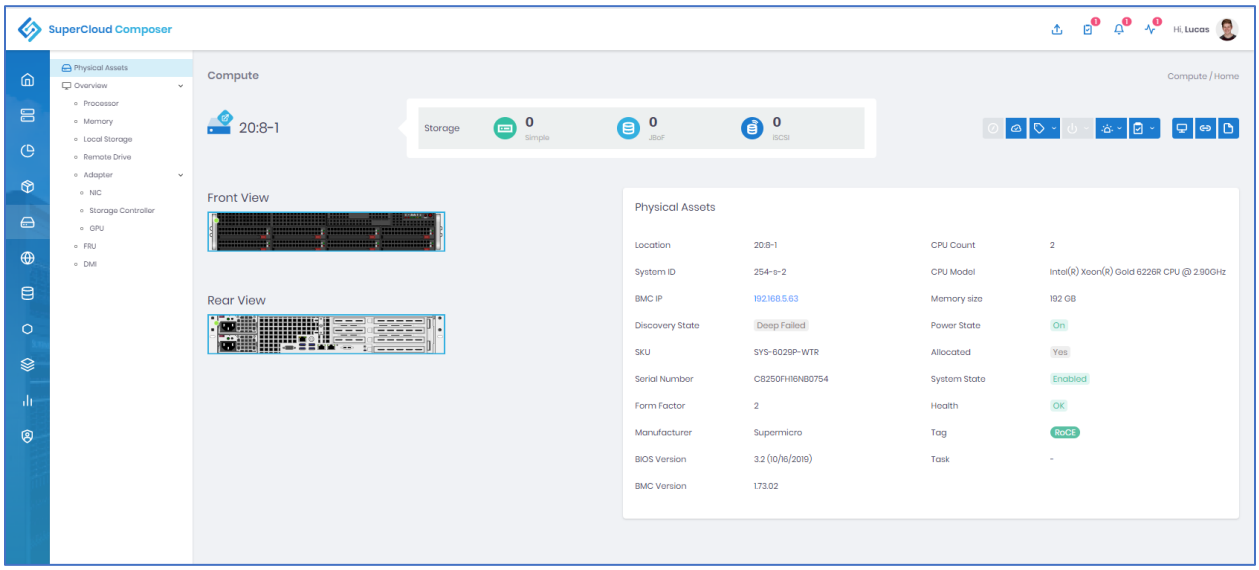

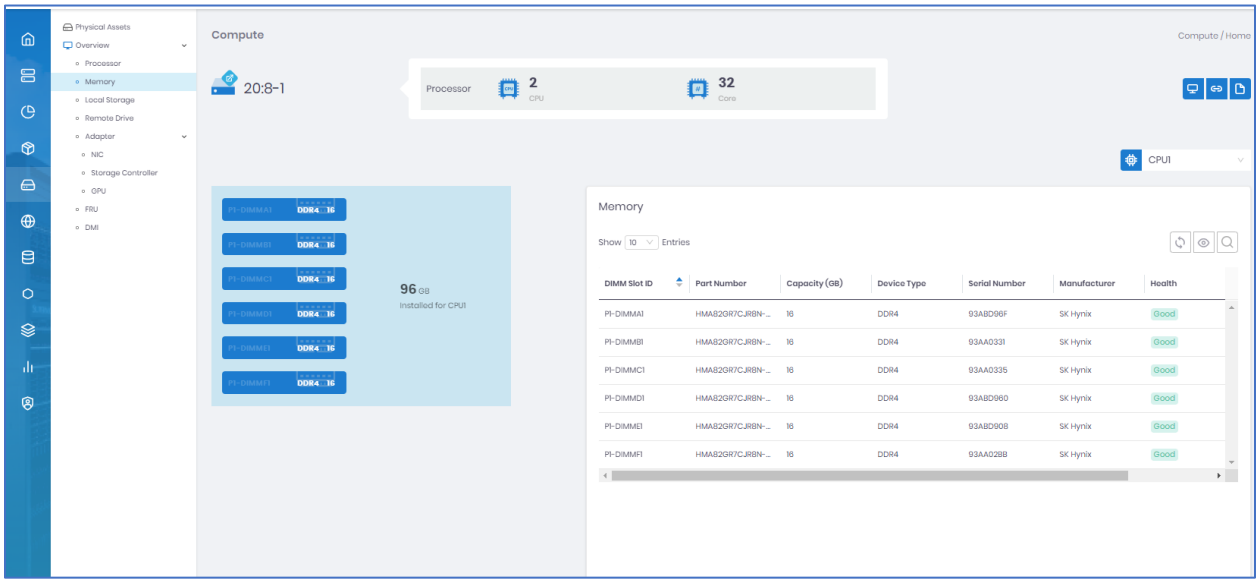

Performance Metrics

- A Grafana-like customizable analytics dashboard that provides the data center operative oversight with their GPU workload operation. Customizable widget type pulldowns that provide GPU metric metadata presented in Time Series Line Chart Sampling, Min-Max-Avg Sampling table, Gauge Meter, and Raw Data Table formats. GPU data can be represented either at the unit level or appliance level.

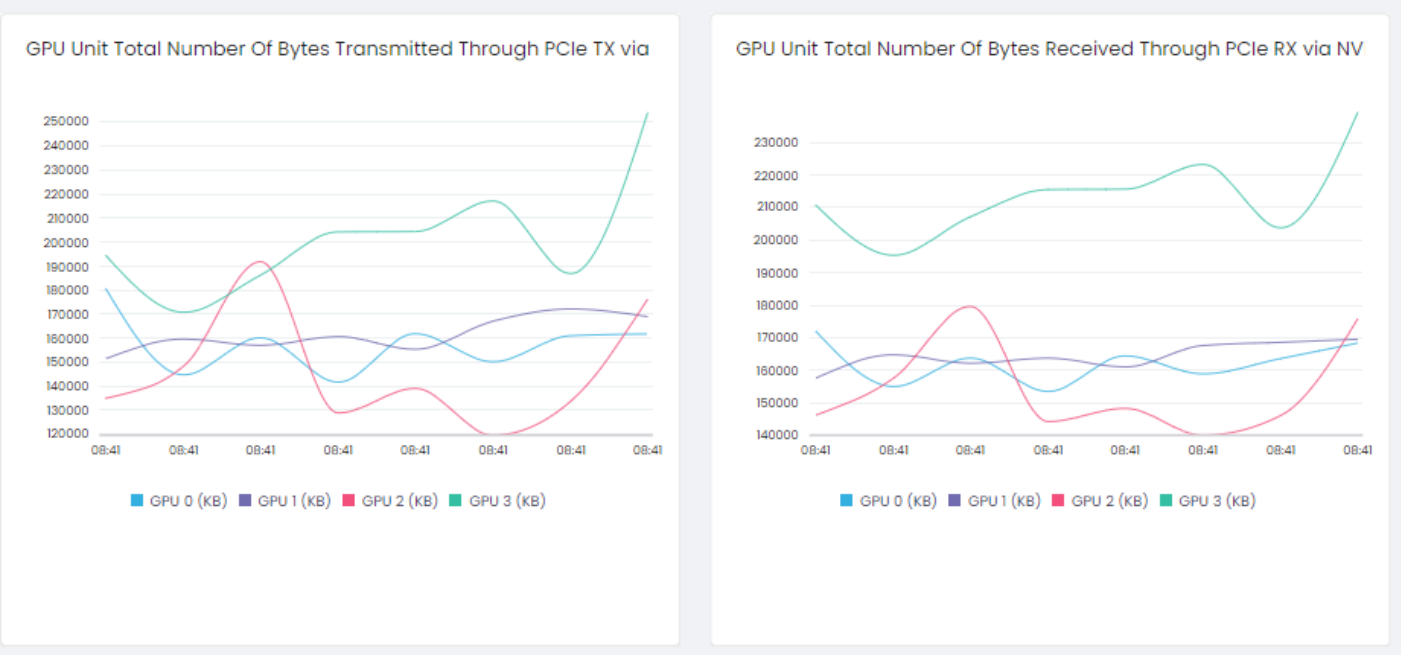

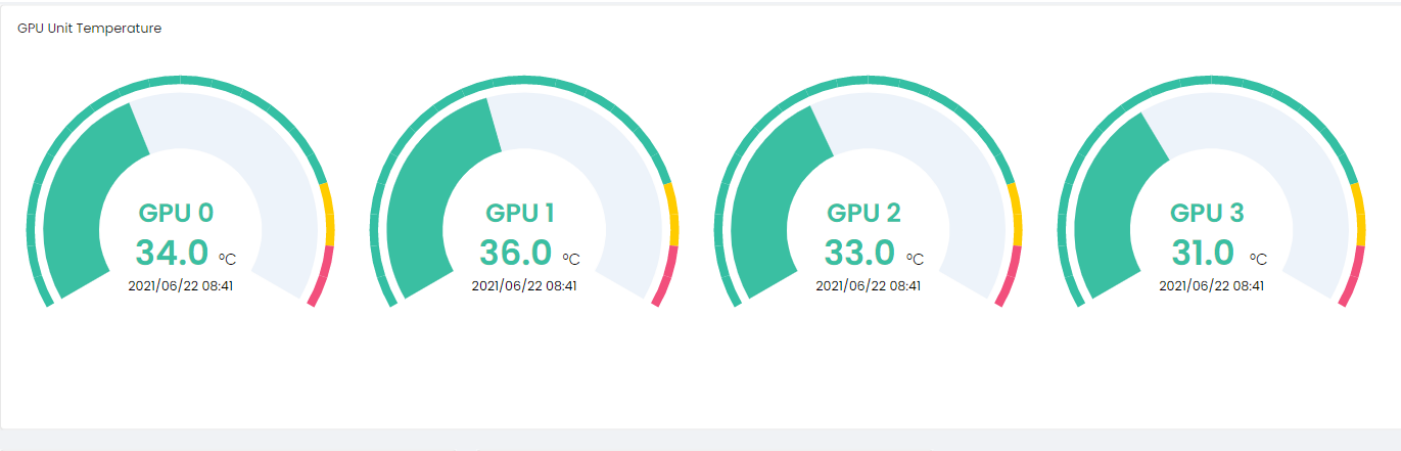

Blade Integration

- Detailed physical asset collateral for each blade within a**n** enclosure
- Oversight management of the (CMM) Computer Management Module for each enclosure
- SuperCloud presents a totally homogenous solution where its network management API collaborates with both blades and network switches within an enclosure providing the end-user seamless integration.
    - (ZTP) Zero touch provisioning of Ethernet blade switches
    - VLAN configuration
    - Physical and Logical port configuration
    - Port counter analytics
    - Detailed topology visual aid
    - Firmware updates
    - Switch detail

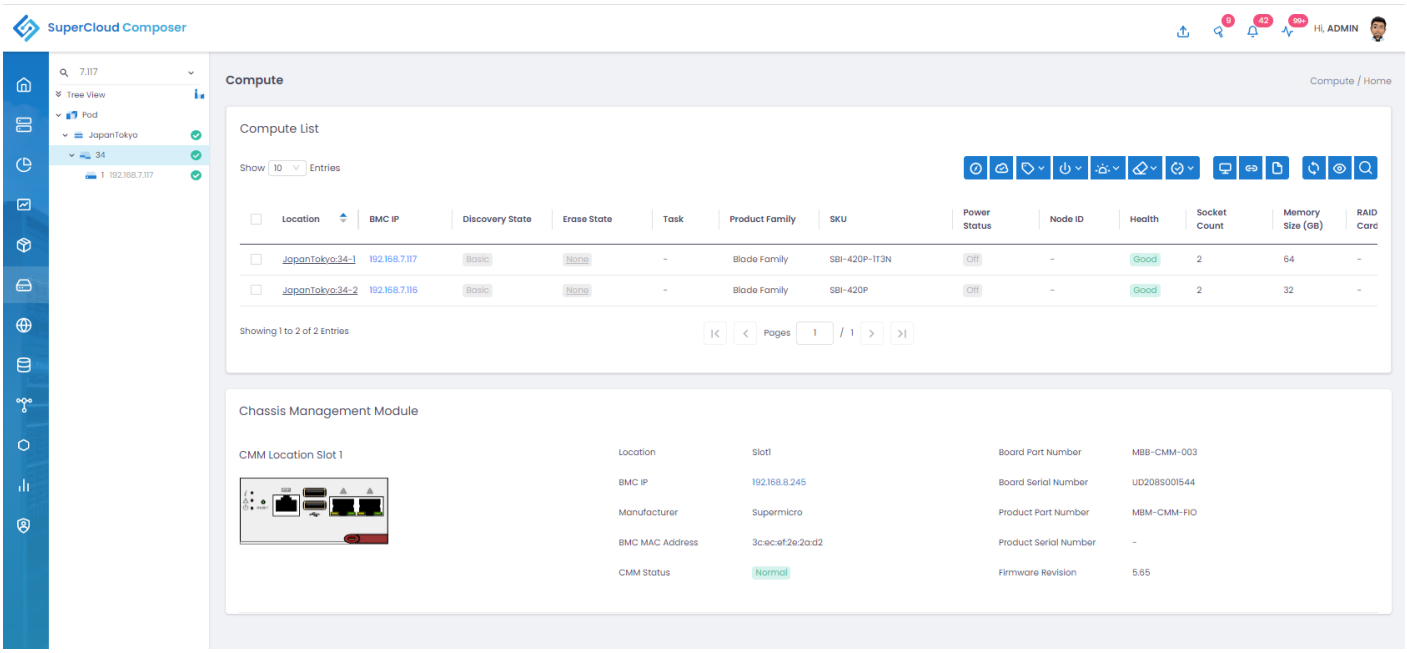

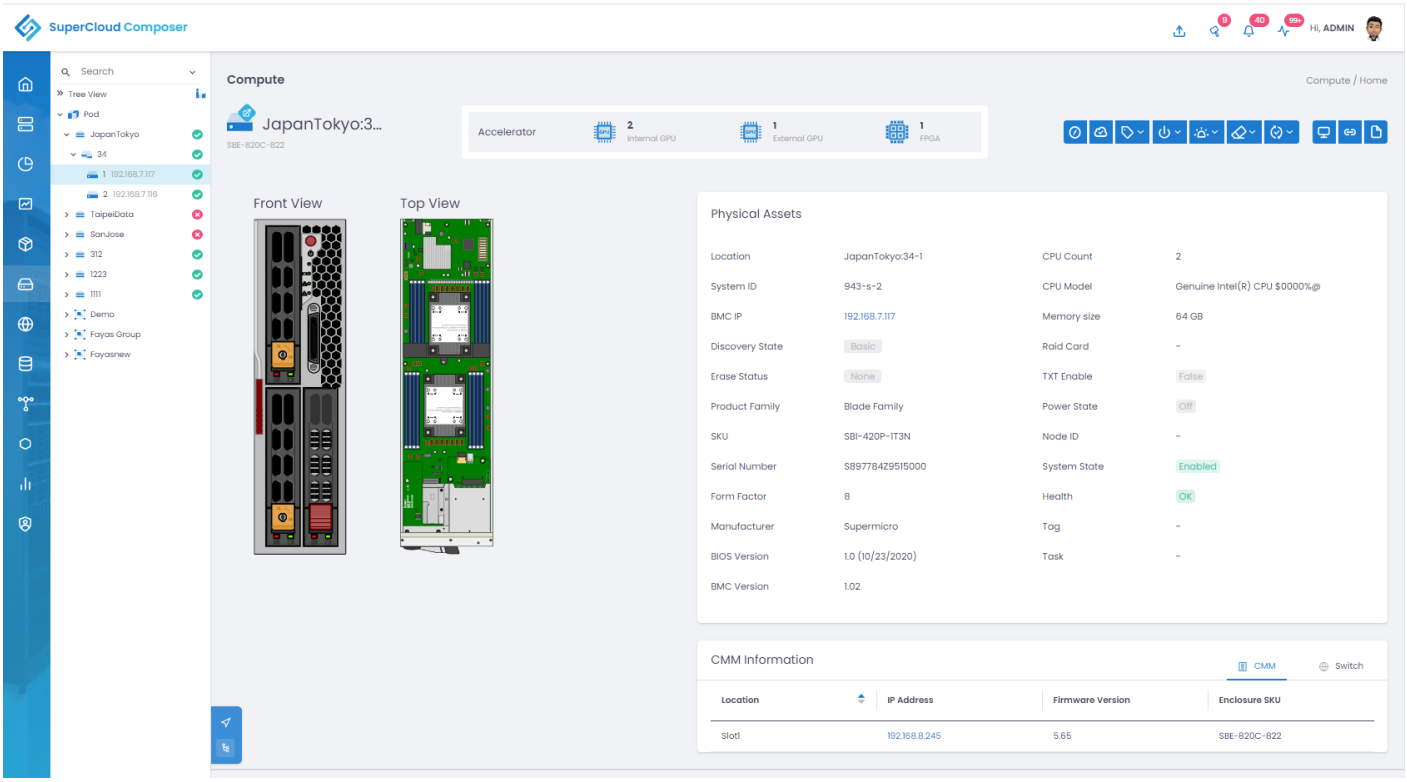

Third party device support

- System lifecycle management support featuring UID management, OS deployment, power management, and ILO/IDRAC console management
- Device asset collateral collection utilizing standardized Redfish OEM extensions support
- SuperCloud Composer is built around an ODIM (Open Distributed Infrastructure Management) framework

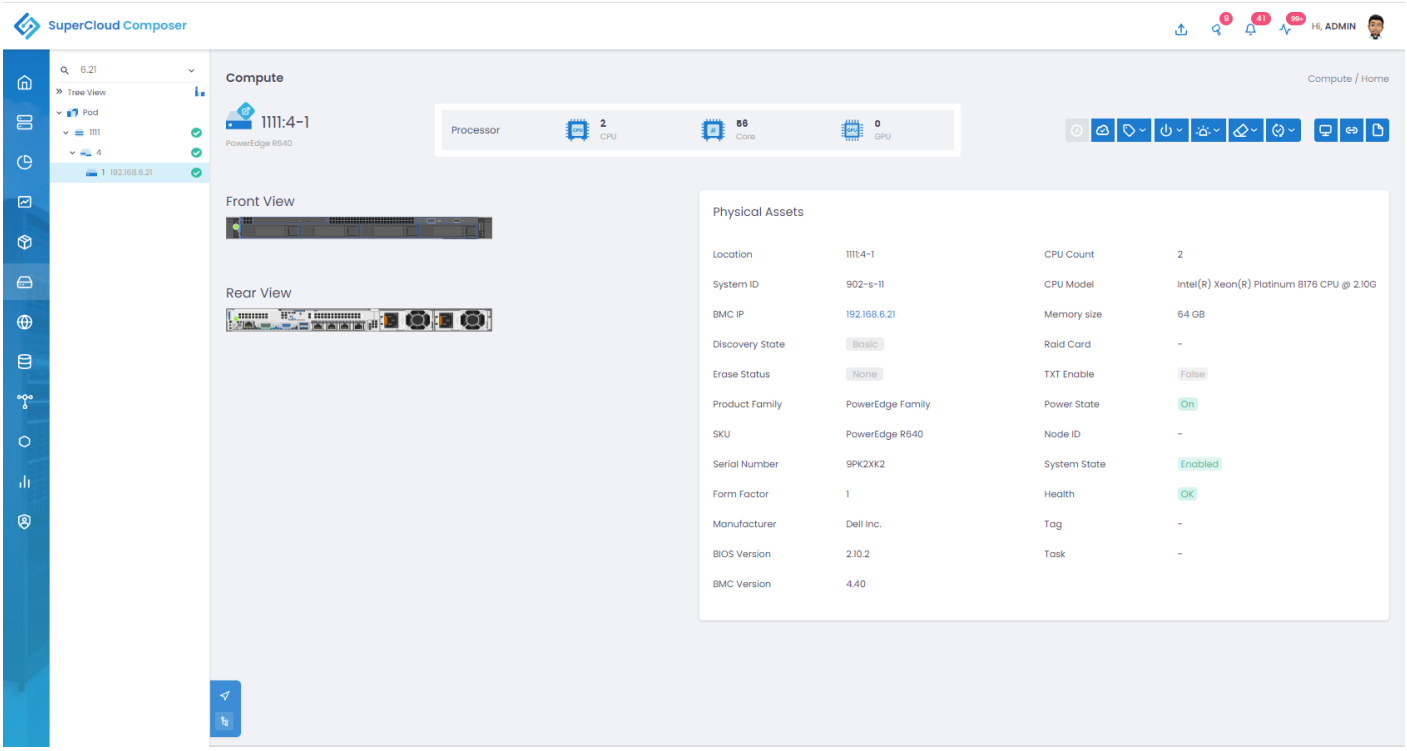

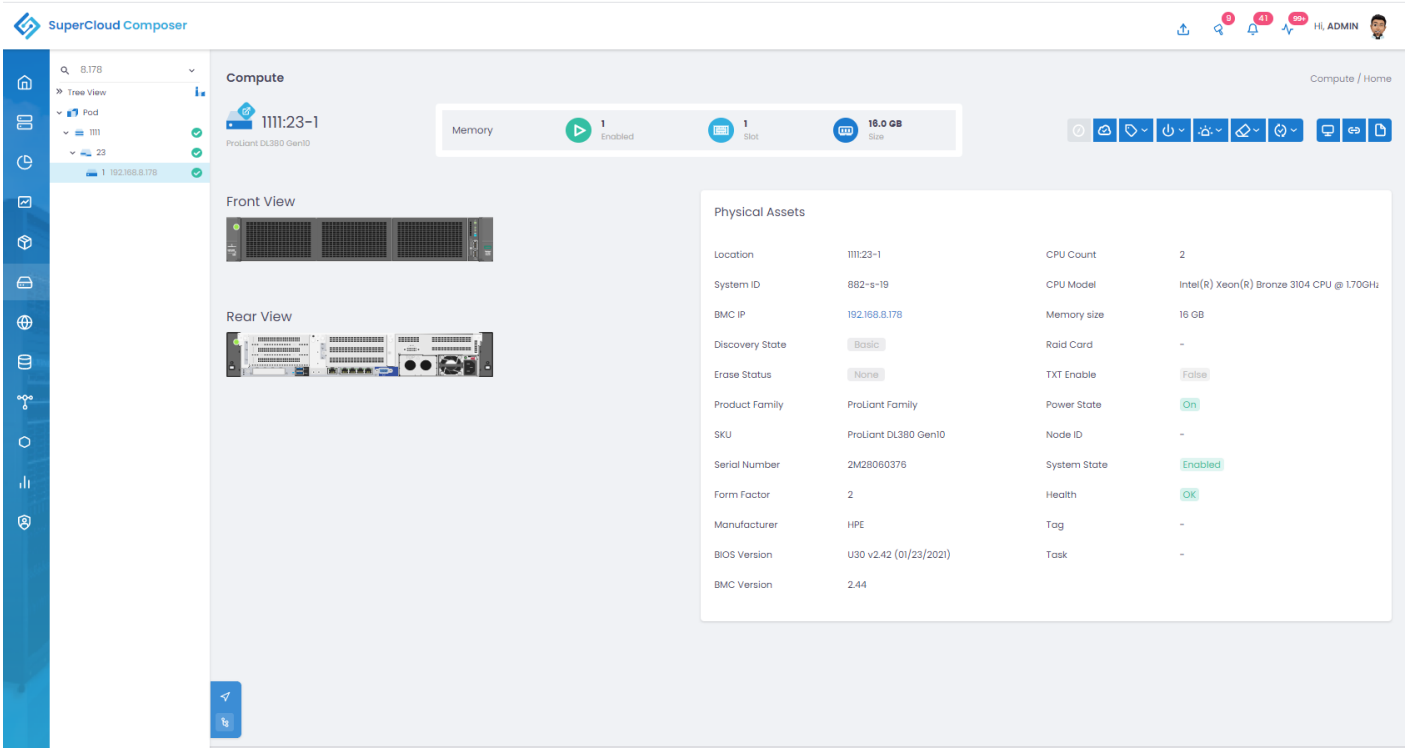

Call Alert Management

- SuperCloud composer offers a policy based alert management system where alerts can be triggered and sent to a**n** SMTP relay host, SMS mobility service**,** or a Slack workspace.
- History tracking of events triggers sent to alert management service

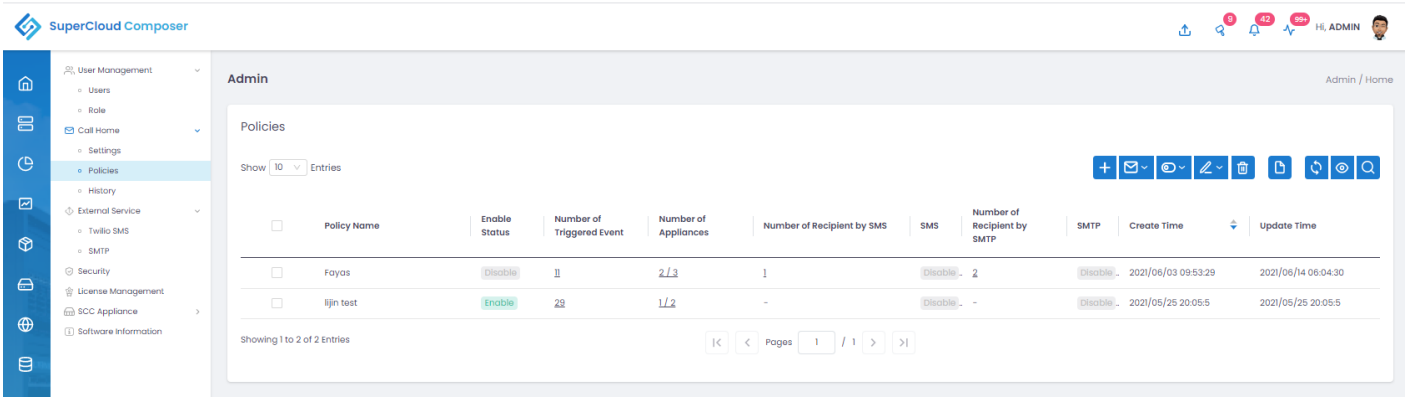

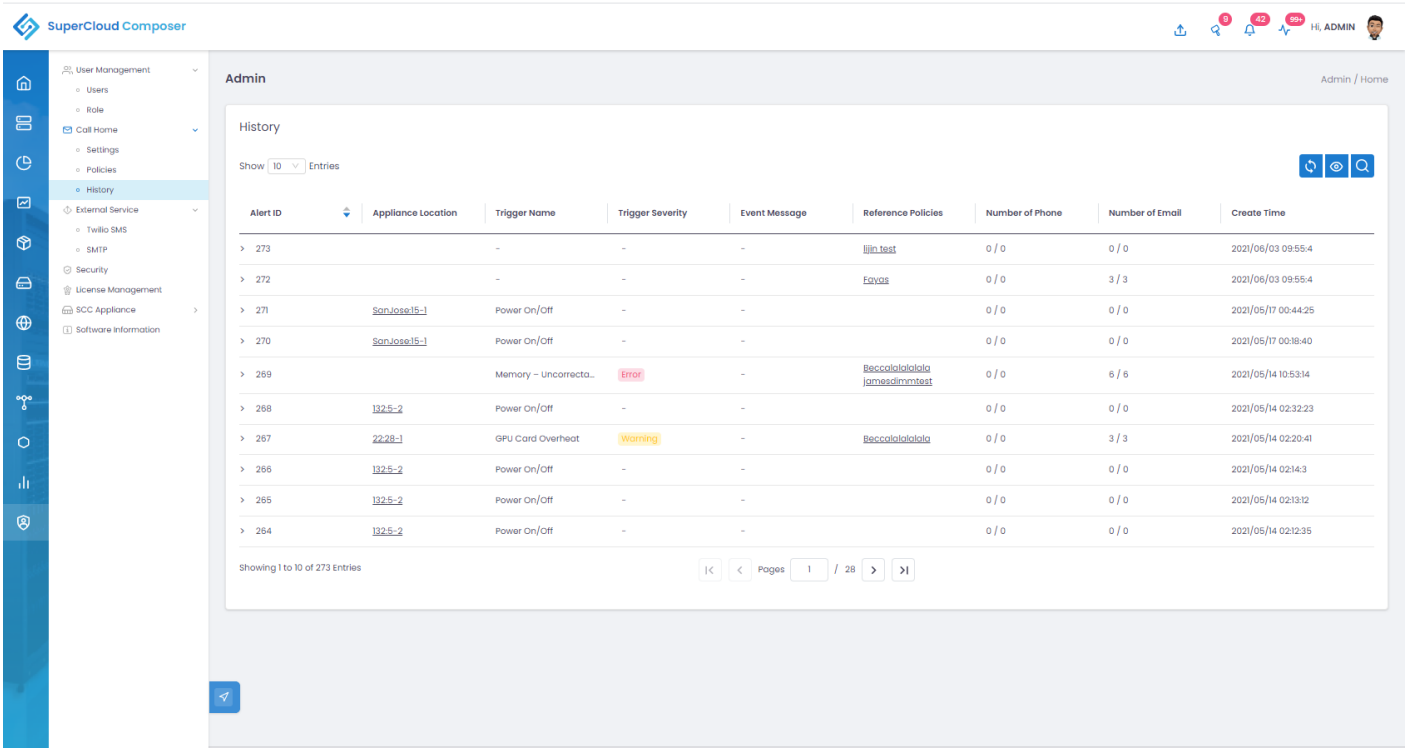

Compute List

- A summarized table of assets found in physical inventory provides administrators a visual aid of bare-metal servers currently registered within the SCC drawer configuration. Each server entry exposes detailed server attributes allowing the administrator to plan, optimize, migrate, and deploy servers within their ever-expanding data center life cycle management needs.

Firmware Compliance

- A critical requirement of Infrastructure management governance is to manage firmware, OS driver level, and SuperCloud Composer patch updates. SuperCloud Composer offers a management notification subsystem where administrators are informed of new releases from Supermicro's service portal.
- End-users can opt-in to automatic firmware downloads to a scalable repository that will host firmware bundles within the SuperCloud Composer stack.
- Software deployments will permit only one retry interval before the server is flagged as an unsuccessful firmware deployment and moved to quarantine status.
- SuperCloud Composer offers an intelligent update agent within its framework, which keeps track of successful deployments; failure rates flagged as zero percent will utilize a scale-up feature to allow those administrators to update more than one server at a given time.

Conclusion

The OS deploy option gives the end-user the ability to target specific servers instead of relying on a pool of available server resources based on selection criteria, which is utilized through user-defined templates.

SuperCloud Composer is a composable cloud management platform that provides a unified dashboard to administer software-defined data centers. Supermicro's cloud infrastructure management software brings speed, agility, and simplicity to IT administration by integrating data center tasks into a single intelligent management solution. Our robust composer engine can orchestrate cloud workloads through a streamlined industry-standard Redfish API. In addition, SuperCloud Composer monitors and manages the broad portfolio of multi-generation Supermicro servers and third-party systems through its data center lifecycle management feature set from a single unified console.

SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server and storage technology is a premier provider of advanced server Building Block Solutions® for Enterprise Data Center, Cloud Computing, Artificial Intelligence, and Edge Computing Systems worldwide. Supermicro is committed to protecting the environment through its "We Keep IT Green®" initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.