



TABLE OF CONTENTS

- 1 EXECUTIVE SUMMARY
- 2 BUSINESS PROBLEM AND BUSINESS VALUE
- 3 REQUIREMENTS
- 3 ARCHITECTURAL OVERVIEW
- 6 DEPLOYMENT
- 14 SUMMARY
- 14 ACKNOWLEDGEMENT
- 14 RESOURCES

WHITE PAPER

SUPERMICRO-SUSE KUBERNETES SOLUTION USING THE SUSE CAAS PLATFORM

EXECUTIVE SUMMARY

This white paper is intended to help an organization deploy an on-premises SUSE CaaS Platform cluster on Supermicro's BigTwin™, Ultra and SuperStorage systems with the 2nd Generation Intel® Xeon® Scalable processor to support the latest, Kubernetes compatible workloads

DESCRIPTION

As with all enterprise class IT infrastructures appropriate networking, computing and storage resources must be planned, purchased, and integrated. This document serves as a reference for the implementation of a SUSE CaaS Platform cluster for Proof of Concept (PoC) and/or non-mission critical production workloads on Supermicro SuperStorage ,Ultra,BigTwin servers and Supermicro's 10G/25G SFP + ethernet switches. Furthermore, the design and implementation suggestions offered here can be used as the starting point to design and deploy an Enterprise class, mission critical SUSE CaaS Platform cluster. While SUSE CaaS Platforms can be deployed as virtual Machines, this document focuses on bare-metal clusters for maximum performance potential. Where possible, resources required for a mission critical cluster are detailed.

The Supermicro-SUSE Kubernetes solution SUSE CaaS Platform integrates the high density compute and storage platforms to achieve the best performance and scalability. It takes advantage of Intel Xeon Scalable processors, high-performance NVMe, and 25G networking. The scale-out architecture lets users dynamically expand capacity across cloud environments to meet demand peaks without downtime or complex processes.

The deployment process for this cluster was implemented, tested and validated in a Supermicro lab.

TARGET AUDIENCE

The target audience for this document is IT professionals who design, implement, and administer Kubernetes based infrastructures. It is suggested that this document is reviewed in its entirety, along with the referenced supplemental documentation, before beginning any deployment.

BUSINESS PROBLEM AND BUSINESS VALUE

In the rapidly evolving era of DevOps, containerized workloads are becoming increasingly important. Developers need an API-centric platform to continuously update and integrate critical applications. Operational teams must implement and manage complex, software-defined infrastructure as quickly and easily as possible. Investing in an easy-to-deploy and scalable SUSE CaaS Platform allows organizations to migrate their legacy business applications as well as deploy new, cloud-native solutions onto a low-friction container-based platform.

BUSINESS PROBLEM

Supporting critical containerized workloads requires a robust container run time engine and cluster orchestrator. While there are many Kubernetes offerings that provide these features on the market, savvy business leaders realize that only a few possess the rare combination having of a reasonable cost of acquisition, ease of implementation and maintenance, and ability to protect their investment from vendor lock-in.

BUSINESS VALUE

Supermicro servers supporting the SUSE CaaS Platform offer proven value in each of these areas:

- Reasonable cost of acquisition: Supermicro and SUSE are recognized leaders in price performance in their respective market spaces. While Total Cost of Acquisition quickly increases with more expensive solutions, it is also true that performance, reliability, and value quickly decline with lower priced options. Supermicro and SUSE together offer maximum value in terms of performance and reliability.
- Ease of implementation: While cloud based solutions are often the fastest to implement, they can suffer from hidden limitations that such as difficulty/inability to work with off-platform applications and a lack of configurability that can turn the easiest option into the biggest liability. On the other end of the application spectrum, unsupported open source software can offer ample configurability, if delivered by a staff of well-paid experts. Supermicro and SUSE can strike the perfect compromise by working closely together to provide a near turn-key experience to their customers. SUSE consultants can quickly integrate Supermicro servers, preloaded with SUSE Enterprise Linux, into an Enterprise class Kubernetes platform. There is no need to invest weeks or even months of development to reach the production stage, and most importantly there is no need for a team of top-level developers to maintain the Kubernetes cluster.
- Protect against vendor lock-in: proprietary products can create unnecessary vendor lock-in. Turn-key Kubernetes offerings are often difficult or impossible to integrate with software that isn't provided by the same vendor. For that reason SUSE has consistently maintained its open source, enterprise class software closer to upstream releases. This ensures SUSE customers can interface with the latest software and hardware innovations, whether on-premises or cloud-provided. At the same time, highly scalable Supermicro Ultra, BigTwin, and Storage servers support a wide range of standards and options supporting up to 44 CPU cores, 3TB of memory, high-performance NVMe storage available on each node, and support for a wide variety of workloads. These Supermicro solutions also support

networking options such as SFP+, 10GBASE-T, 40G and InfiniBand, making them an ideal choice for the SUSE-based CaaS platform.

SUSE CAAS PLATFORM

- The SUSE CaaS Platform is an enterprise class container management solution that enables IT and DevOps professionals to easily deploy, manage, and scale container-based applications and services. It includes Kubernetes to automate lifecycle management of modern applications, and surrounding technologies that enrich Kubernetes and make the platform itself easy to operate. As a result, enterprises that use the SUSE CaaS Platform can reduce application delivery cycle times and improve business agility.
- SUSE is focused on delivering an exceptional operator experience with the SUSE CaaS Platform. With deep competencies in infrastructure, systems, process integration, platform security, lifecycle management and enterprise-grade support, SUSE aims to ensure that IT operations teams can deliver the power of Kubernetes to their users quickly, securely and efficiently.

REQUIREMENTS

Containerized workloads must be reliable, manageable and serviceable. These requirements span the multiple layers of such a solution, from the container host operating system to the container runtime engine and the container orchestration system. Such demands must be span multiple generations of IT infrastructure expectations, both forward and backward, even though the containerized workloads themselves have vastly different to through agility and resiliency.

With SUSE CaaS Platform enterprise will:

- Achieve faster time-to-value with an enterprise-ready container management platform, built from industry-leading technologies, and delivered as a complete package, with everything needed to quickly offer containerized services.
- Simplify management and control of container platforms with efficient installation, easy scaling, and updated automation.
- Maximize return on investment, with a flexible, future-proof container services solution.

ARCHITECTURAL OVERVIEW

To maximize the time to value for a Supermicro Kubernetes Proof of Concept cluster, only the minimum required components have been included in this Reference Architecture. Following this principle will ensure the quickest deployment and evaluation of the Supermicro servers and SUSE CaaS Platform. This proof of concept architecture has been designed to accommodate non-mission critical production workloads as well as fully mission critical applications utilizing the Kubernetes cluster.

As previously mentioned, the SUSE CaaS Platform is an Enterprise class Kubernetes platform that includes important software and services that differentiate from other solutions available in the marketplace:

- **Skuba:** The skuba utility is an open source software project initiated by SUSE to further automate the deployment and up-and down-scaling of a Kubernetes cluster. It leverages the well-known kubernetes project to cover virtually all of the manual and highly error prone tasks of building a Kubernetes cluster.
- **CRI-O:** While the Docker Community Edition container runtime engine enjoys a wide adoption, it is not well suited for an extremely powerful container orchestrator like Kubernetes. CRI-O is a light weight container runtime engine that was created to run all Open Container Initiative (OCI) compliant container images but unlike Docker, it is designed specifically to work with Kubernetes. CRI-O boasts a simpler codebase, lower resource overhead, and a much smaller attack surface area as compared to Docker CE.
- **Cilium:** The networking demands for scale-out container workloads are highly complex. A cluster needs, at a minimum, an overlay network that can integrate physical network segments into a cohesive cluster specific network; routing policies to allow and prevent traffic between specific containers as well between certain containers and the outside world; plus the ability to dynamically grow or contract the overlay network and policies based on real-time cluster changes. Most Kubernetes networking plugins suffer from a limited set of features or a design that lacks scalability. Cilium is a robust Container Network Interface that implements VxLANs for highly scalable overlay networking and identity aware security policies that only copies appropriate policies to the worker nodes that need them. Cilium uniquely encrypts all communication between cluster nodes and establishes encrypted communication channels, known as a Cluster Mesh, between SUSE CaaS Platform clusters, regardless of where they are deployed; on-premises or in the cloud. This allows applications in different clusters to communicate directly with the ability to load-balance incoming application requests across SUSE CaaS Platform clusters on different compute/cloud platforms.

The Host Operating System for the SUSE CaaS Platform is SUSE Linux Enterprise Server. This is the latest edition of one of the most robust Linux distributions. A variety of software packages can be added to SLES but it is not recommended to run additional applications on the SUSE CaaS Platform nodes. The Containers as a Service Platform subscription must be applied to each node and the CaaSP4.x.x repositories enabled on each SUSE CaaS Platform node.

SOLUTION ARCHITECTURE

Deployment Node

A single server is needed as the deployment platform for the SUSE CaaS Platform. Unlike prior version of SUSE CaaS Platform, this node is not considered an administration node for two reasons. This node can be used to deploy any number of SUSE CaaS Platform clusters. As well, the tools to administer a SUSE CaaS Platform cluster are not restricted to this single node.

Master Node

The Kubernetes master node provides the Kubernetes control plane as well as maintains the etcd key-value store. The control plane manages the state of the cluster nodes and pods, manages security policies, cluster-wide networking, and storage assignments. A single master node is capable of managing a high-performance non-mission critical cluster. However, for a highly available control plane, a cluster must be deployed with a properly configured load balancer and three nodes to be deployed as master nodes. A SUSE CaaS Platform cluster that has been deployed with a properly configured load balancer and a single master node can be upgraded to provide a highly available control plane with additional master nodes.

Worker Nodes

The SUSE CaaS Platform worker nodes are responsible for running customer workloads. They also run certain control plane services to help manage encryption, networking, storage services and more. The CRI-O container runtime engine works with Kubernetes to download and manage container images and run containers inside Kubernetes pods.

NETWORKING ARCHITECTURE

The following networking requirements are required for a successful on-premises deployment:

Cluster Physical Network

- A physical network with a dedicated broadcast domain is needed to support the SUSE CaaS Platform cluster. It is not recommended that services beyond the scope of the SUSE CaaS Platform cluster make use of this network since more than one CIDR block of IP addresses will be allocated in this broadcast domain. All of the SUSE CaaS Platform cluster nodes must be able to communicate on this network. This implementation leverages a single 25Gb/s network switch for simplicity and low cost. For network redundancy it is recommended that two switches be used for all data paths.
- A contiguous CIDR block of IP addresses must be available on the cluster physical network.
- Due to network isolation concerns, this implementation used a 1Gb/s network for Internet access, but this is not strictly required.
- Dual, bonded network interfaces on each node was used to communicate on the 25Gb/s network. A single network interface was used to communicate on the 1Gb/s network.

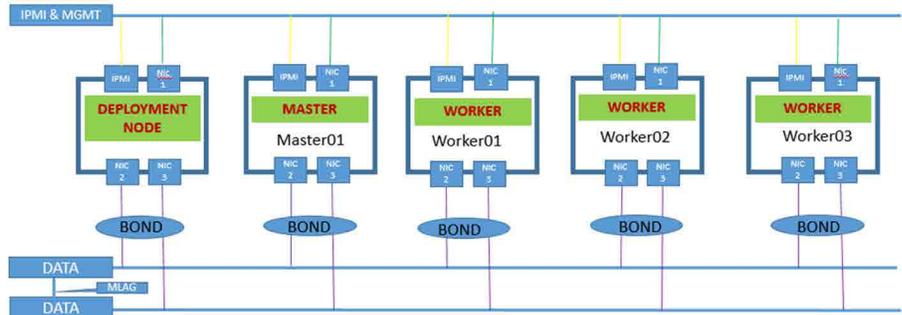
External Network Services

- An external DNS service should be used to resolve, at a minimum, the Fully Qualified Domain Name (FQDN) of the SUSE CaaS Platform master node. If a load balancer is used, the FQDN of the load balancer VIP should be resolvable to allow for TLS certification validation.
- Ensure all SUSE CaaS Platform nodes are pointed to a reliable, external NTP service.
- Ensure all SUSE CaaS Platform nodes have access to a software update repository for installation and to facilitate upgrades over time. This can be in the form of an SMT/RMT server or, with Internet access, <https://scc.suse.com>.

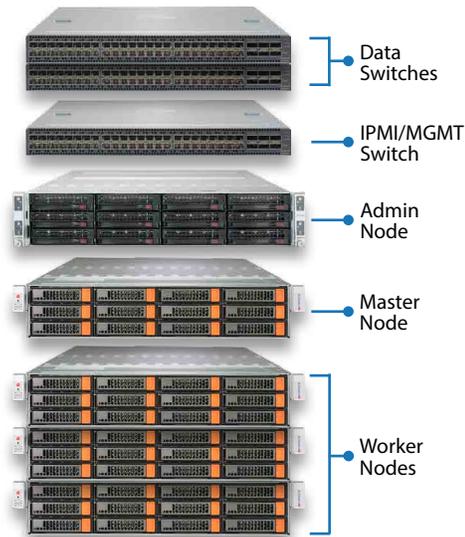
Internal networks

- Known as the Overlay and Service networks, these are used by Kubernetes and the underlying Cilium network plug-in to manage the internal cluster and container connections.
- These internal network ranges should be planned prior to deployment, are usually non-routable network ranges and cannot be changed without redeploying the entire cluster.

NETWORK ADDRESS CONFIGURATION			
FUNCTION	ROLE	HOSTNAME	IP ADDRESS
Core	Management Workstation	deployer.caaspv4.com	192.168.100.5
Cluster	Overlay Network	n/a	10.2.0.0/16
	Service Network	n/a	10.3.0.0/16
Nodes	Master	master.caaspv4.com	192.168.100.1
	Worker1	worker1.caaspv4.com	192.168.100.2
	Worker2	worker2.caaspv4.com	192.168.100.3
	Worker3	worker3.caaspv4.com	192.168.100.4



DEPLOYMENT



SUPERMICRO RACK SPECIFICATIONS (TESTED HARDWARE)		
COMPONENT TYPE	COMPONENT DESCRIPTION	QUANTITY
Rack	Standard data center rack	1
Servers	SYS-6029TP-HC1R	1
	SYS-6029-E1CR24L	4
Data switches (25G)	SSE-F3548S	2
IPMI/MGMT switch (10G)	SSE-X3348T	1

SERVER COMPONENTS FIRMWARE VERSIONS

NOTE: The versions listed below are the versions that were available at the time this Reference Architecture was developed. Ensure that the firmware on all servers, storage devices, and switches are up to date.

FIRMWARE VERSIONS	
COMPONENT	VERSION
BMC	01.46
BIOS	2.0
REDFISH	1.0.1

SUPERMICRO SERVERS SPECIFICATIONS		
BASE SERVER MODEL	SYS-6029TP-HC1R (SINGLE NODE)	SSG-6029P-E1CR24L
CPU	2x P4X-SKL4116T-SR3MQ (12C/24T 2.1G 16.5M 9.6GT 85W)	2x P4X-SKL4116T-SR3MQ (12C/24T 2.1G 16.5M 9.6GT 85W)
RAM	16x MEM-DR416L-CL07-ER26 (192GB)	16x MEM-DR416L-CL07-ER26 (192GB)
	2.5" OS SATA drive	2x HDS-S2T1-MZ7LH240HAHQ05 (240GB SATA3 2.5")
3.5" Data Disk	-	4x HDD-A14T-WUH721414AL4204 (14TB SAS, 3.5")
SIOM (Network)	AOC-MH25G-m2S2TM	AOC-MH25G-m2S2TM
AOC (Network)	1 AOC-S25G-i2S (25GbE dual-port SFP28, PCI-E 3.0 x8)	1 AOC-S25G-m2S (25GbE dual-port SFP28, PCI-E 3.0 x8)
RAID Card	1 Broadcom 3108 SAS3 controller	1 Broadcom 3008 SAS3, IT mode

RECOMMENDED HARDWARE			
ROLE	QUANTITY	DESCRIPTION	NOTES
Management Workstation	1	SYS-5019P-WTR	
Master Node	1	SYS-6019U-TN4RT	
Worker Node(s)	3	SYS-6019U-TN4RT	

RECOMMENDED SOFTWARE			
ROLE	QUANTITY	DESCRIPTION	NOTES
Software	5	SUSE CaaS Platform, x86-64, 1-2 Sockets or 1 Virtual Machine, L3-Priority Subscription, 3 year	Includes Management Workstation, Master/Worker Nodes

HARDWARE DEPLOYMENT CONFIGURATION

The following considerations for the system platforms should be followed:

- Any Supermicro hosts that have completed the [SUSE YES](#) certification for the appropriate version of SUSE Linux Enterprise Server can be used for the physical nodes of this deployment.
- Reset the BIOS setup configuration on all hosts to the default setting to have a known baseline configuration to provide consistency.
- If possible, setup RAID1 mirroring on the storage controller across a pair of drives for the operating system installation
- Ensure that all similar system devices are consistent and up-to-date with regard to BIOS/uEFI/device firmware versions to reduce potential troubleshooting issues later

SOFTWARE DEPLOYMENT CONFIGURATION

- From the [SUSE Downloads](#) website, obtain the SUSE Linux Enterprise Server installation media (DVD1)
- From the [SUSE Customer Center](#) website, obtain enough trial or purchased subscriptions for the SUSE CaaS Platform to apply to all of the cluster hosts.

MANAGEMENT WORKSTATION

- A separate SUSE Linux Enterprise host is needed as a platform to deploy the SUSE CaaS Platform cluster. This host can be used to deploy and manage one or more SUSE CaaS Platform clusters, as long as it has LAN and/or Internet access to the infrastructure that will support each cluster. A single Management Workstation can support both on-premises and cloud-based SUSE CaaS Platform clusters.

- A Management Workstation can be used as a central point of administration for any SUSE CaaS Platform clusters, even if it was not the Management Workstation that deployed a cluster.
- In this reference architecture the Management Workstation will provide the AutoYaST installation files that will be used to automatically install the worker nodes.
- The resource requirements of the Management Workstation are minimal. A 60GB boot drive, 4GB of RAM and 4 CPU cores at minimum will ensure adequate capacity and performance.
- Deploying the Management Workstation as a virtual machine is beyond the scope of this document. The additional requirements primarily involve access to the Internet and the SUSE CaaS Platform cluster network.

IMPORTANT: Installation procedures for sites without Internet access, aka “Dark Sites”, are beyond the scope of this document.

NOTE: This Reference Architecture is not intended to provide accurate deployment instructions for SUSE Linux Enterprise Server. Before attempting an installation, be sure to familiarize yourself and carefully follow the SUSE Linux Enterprise Server [Deployment Guide](#).

During the installation of the SUSE Linux Enterprise Server operating system, observe the following guidelines:

- The DVD installation media can be provided to the host either as a physical DVD or as virtual media through Supermicro IPMI.
 - When booting from the installation media, a pause will occur at the GRUB screen
 - Use the down arrow to select “Installation”, but DO NOT PRESS ENTER.
 - On the line below enter the following: `ifcfg=eth0=<IP address of Management Workstation>/<CIDR mask>,<IP address of default gateway>,<IP address of DNS server>,<DNS search domain> hostname=<hostname>.<domain>`
 - i.e.
`ifcfg=eth0=192.168.100.5/24,192.168.100.254,192.168.100.10, caaspv4.com hostname=admin.caaspv4.com`
 - Press Enter
 - The default partitioning scheme can be used, though you must disable “Enable Swap”
 - Create a regular user named “sles” and use its password as the Administrator password.
 - Apply the subscription information for SUSE CaaS Platform 4.x.x
 - Enable the following repositories:
 - SUSE CaaS Platform 4.x.x
 - This should automatically enable the Basesystem Module and Containers Module
 - Desktop Applications Module
- NOTE:** A desktop environment is not needed for this installation. If you want a desktop environment on the Management Workstation, all of the SUSE CaaS Platform nodes must will also have one. There are ways around this, but they are beyond the scope of this document.
- No software package adjustments are needed but it is recommended to disable the firewall on the final screen before initiating the install.

After the installation completes, finish configuring the Management Workstation:

- Use the YaST tool to properly configure the Management Workstation's network (i.e. enable 802.3ad network port bonding) as appropriate
 - See the [Administration Guide](#) for more information
- Create password-less sudo for the cluster user (i.e. sles):

```
echo "sles ALL=(ALL) NOPASSWD: ALL" >> /etc/sudoers
```

An AutoYaST clone will be taken of the Management Workstation and will be used as the basis for auto-installing the SUSE CaaS Platform cluster nodes:

NOTE: This Reference Architecture is only intended to offer guidance for using AutoYaST for this specific design. Additional information can be found in the [AutoYaST Guide](#).

- Create an AutoYaST clone file of the Management Workstation:

```
sudo yast2 clone_system
```
- Move the resultant autoinst.xml file to a directory dedicated to maintaining the AutoYaST templates for the SUSE CaaS Platform nodes
- Make a copy of the autoinst.xml file for the Master Node and each of the three Worker Nodes. Ensure each file contains the name of the node and ends in .xml
- Update the <hostname> entry of each file with the appropriate node's hostname
- Update the <ipaddr> entry of each file, that contains the Management Workstation's external IP address (not the 127.0.0.1 entry), with the appropriate node's external IP address

Using AutoYaST to automatically install the SUSE CaaS Platform nodes requires that their AutoYaST XML file be available during the node boot/installation. The Management Workstation will run a small webserver in a Docker container to provide this file to each SUSE CaaS Platform node. The following is a high level overview on how to set up that container. For more information see the [Docker Open Source Engine Guide](#).

- Install the Docker CE container runtime engine
- Start and enable the Docker service
- Launch the container:

```
docker run --name <name of container> -v <FQPN of directory containing the AutoYaST template files>:/usr/share/nginx/html:ro -P -d nginx:latest
```
- Take note of the network port assigned to nginx

MASTER NODE

- This reference architecture is based on all SUSE CaaS Platform nodes being configured identically. This allows using AutoYaST to install all of the SUSE CaaS Platform nodes.
- While this reference architecture relies on highly capable servers, the minimum recommended resources for the Master Node are 50GB of disk space, 4 GB of RAM, and 2 CPU cores.

During the installation of the SUSE Linux Enterprise Server operating system, observe the following guidelines:

- The DVD installation media can be provided to the host either as a physical DVD or as virtual media through Supermicro IPMI.
- When booting from the installation media, a pause will occur at the GRuB screen
 - Use the down arrow to select "Installation", but DO NOT PRESS ENTER.
 - On the line below enter the following: `autoyast=http://<FQDN of Management Workstation>:<nginx_port>/<name of Master Node AutoYaST XML file> ifcfg=eth0=<IP address of Master Node>/<CIDR mask>,<IP address of default gateway>,<IP address of DNS server>,<DNS search domain> hostname=<hostname>.<domain>`
 - i.e.
`autoyast=http://admin.caaspv4.com:32768/master.xml`
`ifcfg=eth0=192.168.100.1/24,192.168.100.254,192.168.100.10,caaspv4.com hostname=master.caaspv4.com`
 - Press Enter
- The entire installation should complete and then the host will reboot. Upon the first boot, the network bonding will be created and final installation steps completed.

WORKER NODES

- As mentioned before, this reference architecture is based on all cluster nodes being configured identically.
- While this reference architecture relies on highly capable servers, the minimum recommend resources for the Worker Nodes is primarily dependent on the containerized workload each node will sustain. Principles for calculating each resource are given in the [SUSE CaaS Platform Deployment Guide](#).

During the installation of the SUSE Linux Enterprise Server operating system, observe the following guidelines:

- The DVD installation media can be provided to the host either as a physical DVD or virtual media through Supermicro IPMI.
- When booting from the installation media, a pause will occur at the GRuB screen
 - Use the down arrow to select "Installation", but DO NOT PRESS ENTER.
 - On the line below enter the following: `autoyast=http://<FQDN of Management Workstation>:<nginx_port>/<name of Worker Node AutoYaST XML file> ifcfg=eth0=<IP address of Worker Node>/<CIDR mask>,<IP address of default gateway>,<IP address of DNS server>,<DNS search domain> hostname=<hostname>.<domain>`
 - i.e.
`autoyast=http://admin.caaspv4.com:32768/worker1.xml`
`ifcfg=eth0=192.168.100.2/24,192.168.100.254,192.168.100.10,caaspv4.com hostname=worker1.caaspv4.com`
 - Press Enter

- The entire installation should complete and then the host will reboot. Upon the first boot, the network bonding will be created and final installation steps completed.
- Repeat these steps once for each remaining worker node.
IMPORTANT: The follow section provides an overview of the process for bootstrapping the SUSE CaaS Platform cluster. Before attempting to bootstrap a SUSE CaaS Platform cluster, be sure to familiarize yourself and carefully follow the [SUSE CaaS Platform Deployment Guide](#).

PREPARE THE MANAGEMENT WORKSTATION:

- Ensure the Management Workstation has SLES, SUSE CaaS Platform 4.x.x, and the Containers Module must be registered, as per the Deployment Guide
- Install the SUSE-CaaS-Management software pattern
- Ensure the SSH Agent is running and has loaded the user's SSH key

INITIALIZE THE SUSE CAAS PLATFORM CLUSTER:

- Use the skuba command to initialize the cluster: `skuba cluster init --control-plane <FQDN of the Master Node> <name to be assigned to the SUSE CaaS Platform cluster>`
- Change directory into the cluster directory that was created

JOIN THE MASTER NODE TO THE CLUSTER:

- Use the skuba command to join the Master Node to the cluster: `skuba node bootstrap --user sles --sudo -target <FQDN of the Master Node> <name to be assigned to the Master Node>`

ADD THE WORKER NODES:

- Use the skuba command to join a Worker Node to the cluster: `skuba node bootstrap --user sles --sudo -target <FQDN of the Worker Node> <name to be assigned to the Worker Node>`
- Repeat these steps once for each remaining worker node

VERIFY THE STATUS OF THE CLUSTER:

- The status of the cluster can be verified with the skuba and/or the kubectl command:
 - `skuba cluster status`
 - `export KUBECONFIG=<FQPN of the cluster directory>/admin.conf; kubectl get nodes -o wide`

KUBERNETES CONFORMANCE TESTS

A representative set of the upstream Kubernetes conformance test suite can be found at [Heptio / Sonobuoy](#). This can be easily run via a browser on the client system that has access to the admin 'kubeconfig' file as noted in the "Getting Started" section of this site. An example run can be seen in the following screenshot:

- Log into the client system's command line, using the cluster's admin 'kubeconfig', then refer to [SUSE CaaS Platform FAQ](#) and follow the steps in the "Monitoring Stack based on Prometheus and Grafana" section.

ADDITIONAL DEPLOYMENT CONSIDERATIONS

Beyond the three distinct operational modes described in this document, some very convenient and technologically advanced features are included or can be extended:

- Fine-grained, role-based access control, relying upon a local source can be easily augmented for users with specific roles. In addition, federating to external, like LDAP/AD, authentication/authorization sources can also be accomplished. Refer to the "Managing Users and Groups" and "Role Management" section in the [SUSE CaaS Platform Administration Guide](#) to add other users. Then these configured users can login to the web interface, download their respective 'kubeconfig' files and launch containers into their created or designated namespaces with their roles and access to the specified resources.
- An additional service to consider is a repository site for container images from which to pull workloads. This can be a publicly-accessible site or can be a private collection of workload images. Other registry sites, either public or private, can be used to provide files, like Helm Charts, to deploy complete services. An option to provide more fine-grained user authorization and access to the container images is SUSE [Portus](#).
- Various types of logs are available:
 - The setup and configuration of the cluster
 - The operation of the cluster, including the containerized services and updates
 - The orchestration aspect of Kubernetes
 - More details, including locations to access and how to collect logs for external log servers are included in the "Logging" section of the [SUSE CaaS Platform Administration Guide](#).
- Increasing the node count of the cluster is as easy as joining these to the cluster with the skuba command. For more information, see the "Adding Nodes" section of the [SUSE CaaS Platform Administration Guide](#).
- The density of containerized workloads is another topic to consider. As shown in some previous sections covering around performance monitoring, one can collect the associated metrics of a given workload, kubectl top pods and the nodes, kubectl top nodes to determine how many such workloads can be accommodated. Another approach, to be executed during the launch of a manifest or helm chart, can specify the resource requirements needed. Using this method, Kubernetes will honor this request during the scheduling of the workloads to ensure, at launch, that the appropriate resources are present on the target node. It is also important to ensure that both networking and I/O of storage resources are taken into consideration.
- Other factors like certificate management, security, graceful shutdown and startup of the cluster, and troubleshooting are covered in the [SUSE CaaS Platform Administration Guide](#).



SUMMARY

Combining the features of Supermicro servers with software from the SUSE CaaS Platform yields a robust, powerful and flexible container-as-a-service infrastructure. No matter what stage of transition your organization has achieved with containerized workloads, this deployment allows industry standard compatibility coupled with industry leading support and operational ease. Any business can feel confident in the ability to rapidly and professionally address the continued growth in their container developments and usage.

ACKNOWLEDGEMENT

This document is contributed by Alex Arnoldy at SUSE and Ravi Chintala at Super Micro Computer, Inc.

RESOURCES

Supermicro

- [SSG-6029P-E1CR24L](#)
- [SYS-6029TP-HC1R](#)
- [SSE-X3348T](#)
- [SSE-F3548S](#)
- [SYS-6019U-TN4RT](#)
- [SYS-5019P-WTR](#)

SUSE

- [SUSE CaaS Platform](#)
- [Documentation](#)
- [SUSE Linux Enterprise Server](#)

ABOUT SUPER MICRO COMPUTER, INC.

Supermicro® (NASDAQ: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded 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.

www.supermicro.com

No part of this document covered by copyright may be reproduced in any form or by any means — graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system — without prior written permission of the copyright owner.

Supermicro, the Supermicro logo, Building Block Solutions, We Keep IT Green, SuperServer, Twin, BigTwin, TwinPro, TwinPro², SuperDoctor are trademarks and/or registered trademarks of Super Micro Computer, Inc.

Ultrabook, Celeron, Celeron Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Inside, Intel Inside Logo, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, vPro Inside, Xeon, Xeon Phi, and Xeon Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

Intel, the Intel logo, Intel Atom, and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

All other brands names and trademarks are the property of their respective owners.

© Copyright 2020 Super Micro Computer, Inc. All rights reserved.

Printed in USA

Please Recycle

14_SUSE-CaaS_2020_01-2

