

Network and Cloud Edge Container Bare Metal Reference System Architecture Release v22.05

Authors

Francis Cahill
Octavia Carotti
Calin Gherghie
Joel A. Gibson
Veronika Karpenko
Dana Nehama
Michael O'Reilly
Seungweon Park
Abhijit Sinha
Daniel Ugarte

1 Introduction

1.1 Purpose and Scope

The **Container Bare Metal Reference Architecture (BMRA)** is part of the Network and Cloud Edge Reference System Architectures Portfolio. The BMRA is a cloud-native, forward-looking common template platform for network implementations. It addresses the need to deploy cloud-native Kubernetes clusters optimized with Intel hardware and software innovation for diverse workloads across network locations.

Network locations (for example, On-Premises Edge, Access Edge, Remote Central Office-Forwarding) require deployment of different hardware, software, and configuration specifications due to varying workload, cost, density, and performance requirements. Configuration Profiles define prescribed sets of BMRA hardware and software components designed to optimally address these diverse deployment needs. Ansible playbooks implement the Configuration Profiles for fast, automatic deployment of needed BMRA capabilities. The result is an optimized installation of the BMRA Flavor defined by the Configuration Profile. This user guide covers implementation of BMRA using Configuration Profiles for both network-location-specific and generic deployments.

Network-location-specific Configuration Profiles covered in this document include:

- **On-Premises Edge Configuration Profile** – Typical Customer Premises deployment supporting, for example, Content Delivery Network (CDN) and Smart City scenarios.
- **Access Edge Configuration Profile** – Far edge network deployments closest to users for services that require ultra-low latency, high scalability, and high throughput, such as virtual radio access network (vRAN) and multi-access edge compute (MEC).
- **Remote Central Office-Forwarding Configuration Profile** – Near edge deployments supporting fast packet-forwarding workloads such as cable modem termination system (CMTS), user plane function (UPF), and application gateway function (AGF).
- **Regional Data Center Configuration Profile** – Central-office location typical Configuration Profile. Currently tailored exclusively for media visual processing workloads such as CDN transcoding.

Other Configuration Profiles enable flexible deployments and include the following:

- **Basic Configuration Profile** – A generic minimum BMRA Kubernetes cluster setup.
- **Full Configuration Profile** – A generic complete BMRA setup based on all software features.
- **Build-Your-Own Configuration Profile** – A BMRA Kubernetes cluster setup allowing you to select your preferred options.
- **Storage Configuration Profile** – A BMRA Flavor supporting MinIO Object Storage.

More information on Configuration Profiles is provided later in this document.

1.2 User Guide Information

This document contains step-by-step instructions on installation, configuration, and use of networking and device plug-in features for deploying the BMRA Release v22.05 per the above Configuration Profiles. Validated, open source Ansible playbooks automatically provision a Kubernetes cluster for the selected Configuration Profile, enabling users to quickly implement predictable deployments.

By following this document, it is possible to set up a Kubernetes cluster and automatically configure it using the Ansible playbooks.

The document provides the following information:

- Part 1 (Sections 2 – 5): Requirements for hardware and software to prepare for the Ansible scripts.
- Part 2 (Sections 6 – 14): Step-by-step instructions on how to build your BMRA Flavor using the Ansible scripts. If you wish to start building the BMRA right away, you may directly go to these sections and start automatically provisioning the BMRA Flavor of your choice.
- Part 3 (Section 15): BMRA application examples.
- Part 4 (Appendix A): The BMRA Release Notes.
- Part 5 (Appendix B): Abbreviations

1.3 Version 22.05 Release Information

The hardware platforms supported by BMRA v22.05 are based on 2nd, 3rd, and 4th Generation Intel® Xeon® Scalable processors and Intel® Xeon® D processors. Other advanced Intel hardware technologies supported include the Intel® Ethernet Controller, Intel® QuickAssist Technology (Intel® QAT), and Intel® Server GPU.

The supported software components comprise open-source cloud-native software delivered by Intel and its partners in open-source communities (e.g., Kubernetes, Telegraf, Istio, FD.io). The investments made for these deployments help overcome barriers to networking adoption in cloud deployments.

Following are the primary new and updated hardware ingredients and software features introduced in Release v22.05. For additional details, refer to the [BMRA Release Notes](#).

Release Highlights

Use Case Updates

- **vRAN** Distributed Units (DU)
- **5G Core** now with improved service mesh key management security

4th Generation Intel Xeon Scalable Processors (formerly codenamed Sapphire Rapids)

- Platforms: Archer City (2S-SPR XCC, MCC), Quanta S6Q SDP (2S-SPR XCC, MCC), DSG's Fox Creek Pass (2S-SPR XCC, MCC), Ruby Pass (1S-SPR XCC, MCC)
- Early access to 4th Generation Intel Xeon Scalable processor features. The following features are only available under NDA:
 - Intel® Data Streaming Accelerator (Intel® DSA)
 - Intel® Dynamic Load Balancer (Intel® DLB)
 - Intel® Quick Assist Technology (Intel® QAT)
 - Intel® Scalable I/O Virtualization (Intel® Scalable IOV)

Additional Configuration Profiles

- Access Edge Configuration Profile to support vRAN deployment
- Build-Your-Own Configuration Profile to enable developers to pick and choose hardware/software features for customized profiles

Operating Systems

- Added Rocky Linux 8.5
- Update to Ubuntu 22.04 LTS
- RHEL 8.5

Software – New and Enhancements

Power Management

- cpusets as alternative to isolcpus
- C-state configuration
- Uncore frequency scaling configuration
- Power management unit (PMU) telemetry for power C-state

Security

- Updated Key Management Reference Application (KMRA) version to 2.1

Service Mesh

- Microservices architectures with Istio service mesh (using istioctl for deployment) and

Envoy integrated with KMRA version 2.1

- Intel® Software Guard Extensions (Intel® SGX) signer for automated key management as implemented by KMRA 2.1
- Demos to showcase TLS splicing and TCP/IP bypass for EBPf
- Istio acceleration with Intel AVX-512 crypto (examples available under NDA)
- Istio acceleration and compression with Intel® QuickAssist Technology (Intel® QAT) 2.0 (examples available under NDA)

Packet Processing

- Cloud Native Data Plane (CNDP) for enabling a cloud-based framework for networking and accelerated packet processing, provisioning, orchestrating, and managing microservices to scale efficiently with Kubernetes deployment. CNDP is open-sourced.
- Updated Data Plane Development Kit (DPDK) to version 22.03

Object Storage

- Updates to MinIO tenant configurations to support multi-node cluster configurations using the whereabouts plugin

Kubernetes

- Kubernetes 1.23
- Platform Aware Scheduling update
- Node Feature Discovery update

Kubernetes Operators

- Power Manager Operator update
- Support for Ethernet operator for firmware, driver, and flow control management
- Support for Forward Error Correction (FEC) operator

Add-on Intel Software Packages

- **Telemetry Insight Reports** software, providing platform metrics and status
- **Traffic Analytics Development Kit (TADK)** offering a set of tools and libraries for accelerating applications, such as artificial intelligence and machine learning (AI/ML), while leveraging features available on Intel Xeon Scalable processors

Select advanced and validated features and capabilities of Intel's latest hardware and software are available through an NDA. Contact your Intel representative for access to the NDA material.

Experience Kits, the collaterals that explain in detail the technologies enabled in BMRA release 22.05, including benchmark information, are available in the following locations on Intel Network Builder:

- [Network Transformation Experience Kits](#)
- [Container Experience Kits](#)
- [Intel® Xeon® D-2700 and D-1700 Processor Experience Kits](#)

Table of Contents

1	Introduction	1
1.1	Purpose and Scope.....	1
1.2	User Guide Information.....	1
1.3	Version 22.05 Release Information.....	2
1.4	Document Revision History.....	7
1.5	Key Terms.....	8
1.6	Intel Investments of Capabilities.....	9
1.7	Reference Documentation	9
2	Reference Architecture Deployment.....	11
2.1	BMRA Architecture.....	11
2.2	Configuration Profiles.....	11
2.3	Reference Architecture Installation Prerequisites.....	12
2.3.1	Hardware BOM Selection and Setup for Control and Worker Nodes	12
2.3.2	BIOS Selection for Control and Worker Nodes	12
2.3.3	Operating System Selection for Ansible Host and Control and Worker Nodes.....	13
2.3.4	Network Interface Requirements for Control and Worker Nodes.....	13
2.4	Ansible Playbook.....	13
2.4.1	Ansible Playbook Building Blocks.....	13
2.4.2	Ansible Playbook Phases.....	14
2.5	Deployment Using Ansible Playbook.....	15
2.5.1	Prepare Target Servers.....	15
2.5.2	Prepare Ansible Host and Configuration Templates.....	15
2.5.3	Update Ansible Inventory File.....	16
2.5.4	Update Ansible Host and Group Variables.....	17
2.5.5	Run Ansible Cluster Deployment Playbook.....	17
2.5.6	Run Ansible Cluster Removal Playbook.....	17
3	Reference Architecture Hardware Components and BIOS	18
3.1	Hardware Components List for Control Node	18
3.2	Hardware Components List for Worker Node Base.....	19
3.3	Hardware Components List for Worker Node Plus.....	21
3.4	Hardware Components List for Storage Node.....	23
3.5	Hardware BOMs Supporting All BMRA Configuration Profiles	23
3.6	Platform BIOS	30
4	Reference Architecture Software Components.....	37
4.1	Software Components Supported	37
4.2	Software Components Compatibility Matrices.....	39
5	Post Deployment Verification Guidelines	42
5.1	Check Grafana Telemetry Visualization.....	42
5.2	Check Key Management Infrastructure with Intel SGX.....	42
6	BMRA Setup – Applicable for All Configuration Profiles.....	44
6.1	Set Up an Ansible Host.....	44
6.1.1	RHEL Version 8 as Ansible Host.....	44
6.1.2	Ubuntu 20.04 LTS as Ansible Host	44
6.2	Set Up the Control and Worker Nodes - BIOS Prerequisites.....	45
6.3	Configuration Dictionary - Group Variables.....	46
6.4	Configuration Dictionary - Host Variables.....	50
7	BMRA Basic Configuration Profile Setup	54
7.1	Step 1 - Set Up Basic Configuration Profile Hardware.....	54
7.2	Step 2 - Download Basic Configuration Profile Ansible Playbook	54
7.2.1	Basic Configuration Profile Ansible Playbook Overview.....	54
7.3	Step 3 - Set Up Basic Configuration Profile	55
7.3.1	Basic Configuration Profile Group Variables	55
7.3.2	Basic Configuration Profile Host Variables	56
7.4	Step 4 – Deploy and Validate Basic Configuration Profile Platform	56
8	BMRA Full Configuration Profile Setup	57
8.1	Step 1 - Set Up Full Configuration Profile Hardware.....	57
8.2	Step 2 - Download Full Configuration Profile Ansible Playbook	57
8.2.1	Full Configuration Profile Ansible Playbook Overview.....	57
8.3	Step 3 - Set Up Full Configuration Profile	58

8.3.1	Full Configuration Profile Group Variables.....	58
8.3.2	Full Configuration Profile Host Variables.....	59
8.4	Step 4 - Deploy and Validate Full Configuration Profile Platform.....	59
9	BMRA On-Premises Edge Configuration Profile Setup.....	60
9.1	Step 1 - Set Up On-Premises Edge Configuration Profile Hardware.....	60
9.2	Step 2 - Download On-Premises Edge Configuration Profile Ansible Playbook.....	60
9.2.1	On-Premises Edge Configuration Profile Ansible Playbook Overview.....	60
9.3	Step 3 - Set Up On-Premises Edge Configuration Profile.....	61
9.3.1	On-Premises Edge Configuration Profile Group Variables.....	61
9.3.2	On-Premises Edge Configuration Profile Host Variables.....	62
9.4	Step 4 - Deploy and Validate On-Premises Edge Configuration Profile Platform.....	62
10	BMRA Remote Central Office-Forwarding Configuration Profile Setup.....	63
10.1	Step 1 - Set Up Remote Central Office-Forwarding Configuration Profile Hardware.....	63
10.2	Step 2 - Download Remote Central Office-Forwarding Configuration Profile Ansible Playbook.....	63
10.2.1	Remote Central Office-Forwarding Configuration Profile Ansible Playbook Overview.....	63
10.3	Step 3 - Set Up Remote Central Office-Forwarding Configuration Profile.....	64
10.3.1	Remote Central Office-Forwarding Configuration Profile Group Variables.....	64
10.3.2	Remote Central Office-Forwarding Configuration Profile Host Variables.....	65
10.4	Step 4 - Deploy and Validate Remote Central Office-Forwarding Configuration Profile Platform.....	65
11	BMRA Regional Data Center Configuration Profile Setup.....	66
11.1	Step 1 - Set Up Regional Data Center Configuration Profile Hardware.....	66
11.2	Step 2 - Download Regional Data Center Configuration Profile Ansible Playbook.....	66
11.2.1	Regional Data Center Configuration Profile Ansible Playbook Overview.....	66
11.3	Step 3 - Set Up Regional Data Center Configuration Profile.....	67
11.3.1	Regional Data Center Configuration Profile Group Variables.....	67
11.3.2	Regional Data Center Configuration Profile Host Variables.....	68
11.4	Step 4 - Deploy and Validate Regional Data Center Configuration Profile Platform.....	68
12	BMRA for Storage Configuration Profile Setup.....	69
12.1	Step 1 - Set Up Storage Configuration Profile Hardware.....	69
12.2	Step 2 - Download Storage Configuration Profile Ansible Playbook.....	70
12.2.1	Storage Configuration Profile Ansible Playbook Overview.....	70
12.3	Step 3 - Set Up Storage Configuration Profile.....	71
12.3.1	Storage Configuration Profile Group Variables.....	71
12.3.2	Storage Configuration Profile Host Variables.....	71
12.4	Step 4 - Deploy and Validate Storage Configuration Profile Platform.....	71
13	BMRA Access Edge Configuration Profile Setup.....	72
13.1	Step 1 - Set Up Access Edge Configuration Profile Hardware.....	72
13.2	Step 2 - Download Access Edge Configuration Profile Ansible Playbook.....	72
13.2.1	Access Edge Configuration Profile Ansible Playbook Overview.....	72
13.3	Step 3 - Set Up Access Edge Configuration Profile.....	73
13.3.1	Access Edge Configuration Profile Group Variables.....	73
13.3.2	Access Edge Configuration Profile Host Variables.....	74
13.4	Step 4 - Deploy and Validate Access Edge Configuration Profile Platform.....	74
14	BMRA Build-Your-Own Configuration Profile Setup.....	75
14.1	Step 1 - Set Up Build-Your-Own Configuration Profile Hardware.....	75
14.2	Step 2 - Download Build-Your-Own Configuration Profile Ansible Playbook.....	75
14.2.1	Build-Your-Own Configuration Profile Ansible Playbook Overview.....	75
14.3	Step 3 - Set Up Build-Your-Own Configuration Profile.....	76
14.3.1	Build-Your-Own Configuration Profile Group Variables.....	76
14.3.2	Build-Your-Own Configuration Profile Host Variables.....	77
14.4	Step 4 - Deploy and Validate Build-Your-Own Configuration Profile Platform.....	77
15	Workloads and Application Examples.....	79
15.1	Enabling Key Management NGINX Applications.....	79
15.2	Enabling Trusted Certificate Service.....	79
15.2.1	Istio Custom CA Integration Using Kubernetes CSR.....	79
15.2.2	Remote Attestation and Manual Key Management.....	79
15.3	Service Mesh Automated Remote Attestation and Key Management with KMRA, TCS, and TCA.....	79
15.4	Istio TLS Splicing.....	80
15.5	Web Application Firewall Using Traffic Analytics Development Kit.....	80
Appendix A	BMRA Release Notes.....	82
A.1	BMRA 22.05 Notable Facts.....	82

A.2	BMRA 22.05 Bug Fixes.....	83
A.3	BMRA 22.01 Notable Facts.....	83
A.4	BMRA 22.01 Bug Fixes.....	84
A.5	BMRA 21.09 New Features.....	84
A.6	BMRA 21.09 Bug Fixes.....	85
A.7	BMRA 21.08 New Features.....	85
A.8	BMRA 21.08 Bug Fixes.....	86
A.9	Known Issues.....	86
Appendix B	Abbreviations.....	88

Figures

Figure 1.	BMRA Illustration and Applicable Elements.....	11
Figure 2.	High Level BMRA Ansible Playbooks Architecture Full Configuration Profile Example	14
Figure 3.	Basic Configuration Profile Ansible Playbook.....	55
Figure 4.	Full Configuration Profile Ansible Playbook.....	58
Figure 5.	On-Premises Edge Configuration Profile Ansible Playbook.....	61
Figure 6.	Remote Central Office-Forwarding Configuration Profile Ansible Playbook	64
Figure 7.	Regional Data Center Configuration Profile Ansible Playbook.....	67
Figure 8.	BMRA for Storage Architecture	69
Figure 9.	Storage Configuration Profile Ansible Playbook	70
Figure 10.	Access Edge Configuration Profile Ansible Playbook.....	73
Figure 11.	Build-Your-Own Configuration Profile Ansible Playbook	76

Tables

Table 1.	Terms Used.....	8
Table 2.	Hardware and Software Configuration Taxonomy	8
Table 3.	Intel Capabilities Investments and Benefits.....	9
Table 4.	Hardware Options for Control Node – 2nd Generation Intel Xeon Scalable Processor	18
Table 5.	Hardware Options for Control Node – 3rd Generation Intel Xeon Scalable Processor.....	18
Table 6.	Hardware Options for Control Node – 4th Generation Intel Xeon Scalable Processor.....	18
Table 7.	Hardware Options for Control Node – Intel® Xeon® D Processor	19
Table 8.	Hardware Components for Worker Node Base – 2nd Generation Intel Xeon Scalable Processor.....	19
Table 9.	Hardware Components for Worker Node Base – 3rd Generation Intel Xeon Scalable Processor.....	20
Table 10.	Hardware Components for Worker Node Base – 4th Generation Intel Xeon Scalable Processor.....	20
Table 11.	Hardware Components for Worker Node Base (Access Edge - vRAN) – 4th Generation Intel Xeon Scalable Processor	20
Table 12.	Hardware Components for Worker Node Base – Intel® Xeon® D Processor	21
Table 13.	Hardware Components for Worker Node Plus – 2nd Generation Intel Xeon Scalable Processor.....	21
Table 14.	Hardware Components for Worker Node Plus – 3rd Generation Intel Xeon Scalable Processor.....	21
Table 15.	Hardware Components for Worker Node Plus – 4th Generation Intel Xeon Scalable Processor.....	22
Table 16.	Hardware Components for Worker Node Plus (Access Edge - vRAN) – 4th Generation Intel Xeon Scalable Processor	22
Table 17.	Hardware Components for Worker Node Plus – Intel® Xeon® D Processor	23
Table 18.	Hardware Components for Storage Node – 3rd Generation Intel Xeon Scalable Processor.....	23
Table 19.	Control Node Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor	23
Table 20.	Control Node Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor	24
Table 21.	Control Node Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor	25
Table 22.	Control Node Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor.....	25
Table 23.	Worker Node Base Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor.....	26
Table 24.	Worker Node Plus Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor.....	26
Table 25.	Worker Node Base Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor.....	27
Table 26.	Worker Node Plus and Storage Node Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor	27
Table 27.	Worker Node Base Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor	28
Table 28.	Worker Node Plus Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor.....	29
Table 29.	Worker Node Base Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor	29
Table 30.	Worker Node Plus Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor	30
Table 31.	Platform BIOS Settings for 2nd Generation Intel® Xeon® Scalable Processor.....	31
Table 32.	Platform BIOS Settings for 3rd Generation Intel® Xeon® Scalable Processor	32
Table 33.	Platform BIOS Settings for 4th Generation Intel® Xeon® Scalable Processor	33
Table 34.	Platform BIOS Settings for Intel® Xeon® D Processor	35
Table 35.	BIOS Settings to Enable Intel SST-BF, Intel SST-TF, and Intel SST-PP	35
Table 36.	BIOS Settings to Enable Intel SGX on 2nd Generation and 3rd Generation Intel Xeon Scalable Processors.....	36
Table 37.	BIOS Settings to Enable Intel SGX on 4th Generation Intel Xeon Scalable Processor	36
Table 38.	Software Components.....	37
Table 39.	Links to Verification Guidelines on GitHub	42
Table 40.	BIOS Prerequisites for Control and Worker Nodes for Basic, Full, Storage, and Build-Your-Own Configuration Profiles.....	45

Table 41.	BIOS Prerequisites for Control and Worker Nodes for On-Premises Edge, Remote Central Office-Forwarding, Regional Data Center, and Access Edge Configuration Profiles	45
Table 42.	Configuration Dictionary – Group Variables	46
Table 43.	Configuration Dictionary – Host Variables	50
Table 44.	Hardware Setup for Basic Configuration Profile	54
Table 45.	Basic Configuration Profile – Group Variables	55
Table 46.	Basic Configuration Profile – Host Variables	56
Table 47.	Hardware Setup for Full Configuration Profile	57
Table 48.	Full Configuration Profile – Group Variables	58
Table 49.	Full Configuration Profile – Host Variables	59
Table 50.	Hardware Setup for On-Premises Edge Configuration Profile	60
Table 51.	On-Premises Edge Configuration Profile – Group Variables	61
Table 52.	On-Premises Edge Configuration Profile – Host Variables	62
Table 53.	Hardware Setup for Remote Central Office-Forwarding Configuration Profile	63
Table 54.	Remote Central Office-Forwarding Configuration Profile – Group Variables	64
Table 55.	Remote Central Office-Forwarding Configuration Profile – Host Variables	65
Table 56.	Hardware Setup for Regional Data Center Configuration Profile	66
Table 57.	Regional Data Center Configuration Profile – Group Variables	67
Table 58.	Regional Data Center Configuration Profile – Host Variables	68
Table 59.	Hardware Setup for Storage Configuration Profile	70
Table 60.	Storage Configuration Profile – Group Variables	71
Table 61.	Storage Configuration Profile – Host Variables	71
Table 62.	Hardware Setup for Access Edge Configuration Profile	72
Table 63.	Access Edge Configuration Profile – Group Variables	73
Table 64.	Access Edge Configuration Profile – Host Variables	74
Table 65.	Hardware Setup for Build-Your-Own Configuration Profile	75
Table 66.	Build-Your-Own Configuration Profile – Group Variables	76
Table 67.	Build-Your-Own Configuration Profile – Host Variables	77
Table 68.	Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 22.05	82
Table 69.	Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 22.01	84
Table 70.	Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 21.09	84

1.4 Document Revision History

Three previous editions of the BMRA document were released, starting April 2019.

- Covered 2nd Generation Intel® Xeon® Scalable processors
- Covered 2nd and 3rd Generation Intel® Xeon® Scalable processors
- Covered 2nd and 3rd Generation Intel® Xeon® Scalable processors and Intel® Xeon® D processor

REVISION	DATE	DESCRIPTION
001	February 2022	Initial release.
002	March 2022	Updated a few URLs.
003	June 2022	Covers the 4th Generation Intel® Xeon® Scalable processor (formerly codenamed Sapphire Rapids). See “Version 22.05 Release Information” for details.
004	June 2022	Changes include updates to the discussion of the BMRA for Storage Deployment Model.

1.5 Key Terms

[Table 1](#) lists the key terms used throughout the portfolio. (These terms are specific to Network and Cloud Edge Reference System Architectures Portfolio deployments.)

Table 1. Terms Used

TERM	DESCRIPTION
Experience Kits	Guidelines delivered in the form of—manuals, user guides, application notes, solution briefs, training videos—for best-practice implementation of cloud native and Kubernetes technologies to ease developments and deployments.
Network and Cloud Edge Reference System Architectures Portfolio	A templated system-level blueprint for a range of locations in enterprise and cloud infrastructure with automated deployment tools. The portfolio integrates the latest Intel platforms and cloud-native technologies for multiple deployment models to simplify and accelerate deployments of key workloads across a service infrastructure.
Deployment Model	Provides flexibility to deploy solutions according to IT needs. The portfolio offers two deployment models: <ul style="list-style-type: none"> • Container Bare Metal Reference System Architecture (BMRA) – A deployment model of a Kubernetes cluster with containers on a bare metal platform. A special version of BMRA is the BMRA for Storage that supports MinIO Object Storage. • Virtual Machine Reference System Architecture (VMRA) – A deployment model of a virtual cluster on a physical node. The virtual cluster can be a Kubernetes containers-based cluster.
Configuration Profiles	A prescribed set of components—hardware, software modules, hardware/software configuration specifications—designed for a deployment for specific workloads at a network location (such as Access Edge). Configuration Profiles define the components for optimized performance, usability, and cost per network location and workload needs. ¹ In addition, generic Configuration Profiles are available for developers' flexible deployments.
Reference Architecture Flavor	A Reference Architecture deployment using a Configuration Profile.
Ansible Playbook	A set of validated scripts that prepare, configure, and deploy a Reference Architecture Flavor by implementing a Configuration Profile.
Configuration Profile Ansible Scripts	Automates quick, repeatable, and predictive deployments using Ansible playbooks. Various Configuration Profiles and Ansible scripts allow automated installations that are application-ready, depending on the workload and network location.
Kubernetes cluster	A deployment that installs at least one worker node running containerized applications. Pods are the components of the application workload that are hosted on worker nodes. Control nodes manage the pods and worker nodes.
Intel Platforms	Prescribes Intel platforms for optimized operations. The platforms are based on 2nd, 3rd, and 4th Generation Intel® Xeon® Scalable processors plus the Intel® Xeon® D processor. These platforms include the Taylors Falls Reference Design. The platforms integrate Intel® Ethernet Controller 700 Series and 800 Series, Intel® QuickAssist Technology (Intel® QAT), Intel® Server GPU (graphics processing unit), Intel® Optane™ technology, and more. Note: This release of VMRA does not support the Intel Xeon D processor.

In addition to key terms, portfolio deployment procedures follow a hardware and software configuration taxonomy. [Table 2](#) describes the taxonomy used throughout this document.

Table 2. Hardware and Software Configuration Taxonomy

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value)

¹ [Workloads and configurations](#). Results may vary.

TERM	DESCRIPTION
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on workload. Setting does not affect the Configuration Profile or platform deployment
Software Taxonomy	
TRUE	Feature is included and enabled by default
FALSE	Feature is included but disabled by default - can be enabled and configured by user
N/A	Feature is not included and cannot be enabled or configured

1.6 Intel Investments of Capabilities

Intel investments in networking solutions are designed to help IT centers accelerate deployments, improve operational efficiencies, and lower costs. [Table 3](#) highlights Intel investments in the portfolio and their benefits.

Table 3. Intel Capabilities Investments and Benefits

CAPABILITY	BENEFIT
Performance	Intel® platform innovation and accelerators, combined with packet processing innovation for cloud-native environments, deliver superior and predictive application and network performance.
Orchestration and Automation	Implementing Kubernetes containers orchestration, including Kubernetes Operators, simplifies and manages deployments and removes barriers in Kubernetes to support networking functionality.
Observability	Collecting platform metrics by using, as an example, the collectd daemon and Telegraf server agent, publishing the data, and generating reports, enables high visibility of platform status and health.
Power Management	Leveraging Intel platform innovation, such as Intel® Speed Select Technology (Intel® SST), supports optimized platform power utilization.
Security	Intel security technologies help ensure platform and transport security. These technologies include the following: <ul style="list-style-type: none"> • Intel® Security Libraries for Data Center (Intel® SecL - DC) • Intel® QuickAssist Technology Engine for OpenSSL (Intel® QAT Engine for OpenSSL) • Intel® Software Guard Extensions (Intel® SGX) • Key Management Reference Application (KMRA) implementation
Storage	Creating a disaggregated, high-performance, scalable storage platform using MinIO Object Storage supports data-intensive applications, such as media streaming, big data analytics, AI, and machine learning.
Service Mesh	Implementing a Service Mesh architecture using Istio allows application services that can be added, connected, monitored, more secure, and load-balanced with few or no code changes. Service Mesh is integrated with the Trusted Certificate Service for Kubernetes* platform, providing secure key management.

1.7 Reference Documentation

The [Network and Cloud Edge Reference System Architectures Portfolio User Manual](#) contains a complete list of reference documents. Additionally, a virtual machine-based reference architecture (VMRA) deployment allows creation of a Kubernetes cluster for a Configuration Profile on a virtualized infrastructure. The [Network and Cloud Edge Virtual Machine Reference System Architecture User Guide](#) provides information and installation instructions for a VMRA.

Other collaterals, including technical guides and solution briefs that explain in detail the technologies enabled in BMRA release v22.05, are available in the following locations: [Network Transformation Experience Kits](#) and [Container Experience Kits](#).

Part 1:

Reference Architecture Components and
Deployment Guidelines:

Ansible Playbooks

Hardware Components

Software Ingredients

Recommended Configurations

2 Reference Architecture Deployment

This chapter explains how a BMRA Flavor is generated and deployed. The process includes installation of the hardware setup followed by system provisioning.

2.1 BMRA Architecture

The BMRA is a Kubernetes cluster that can be configured to support a flexible number of Kubernetes control nodes and worker nodes (see [Figure 1](#)). To deploy the BMRA, you deploy and configure the following elements:

- **Hardware Components:** Multiple platform hardware options are available, including a variety of 4th, 3rd, and 2nd Generation Intel® Xeon® Scalable processor SKUs, Intel® Xeon® D processor SKUs, Intel® Ethernet Network Adapters, Intel® QAT, and Intel® Server GPU. BIOS options are listed elsewhere in this guide. Deployment engineers should refer to [Section 3.6](#) during deployment to select and configure optimal BIOS values before cluster provisioning.
- **Software Capabilities:** The software capabilities are based on open source software delivered by cloud-native and CNCF communities driving Kubernetes, Istio, observability, DDPK, FD.io. OVS, OVS-DPDK, and through Intel GitHub. Options for RHEL and Ubuntu Linux operating systems are available. The container environment is based on Docker, containerd, or CRI-O container runtimes.
- **Configuration Profiles:** Specific hardware and software configurations are provided in the Configuration Profiles based on Intel assessment and verification. Hardware configurations address two performance capabilities: base and plus.
- **Installation Playbooks:** Ansible playbooks implement the Configuration Profiles for best-practice, reliable, and accelerated BMRA Flavor deployment.

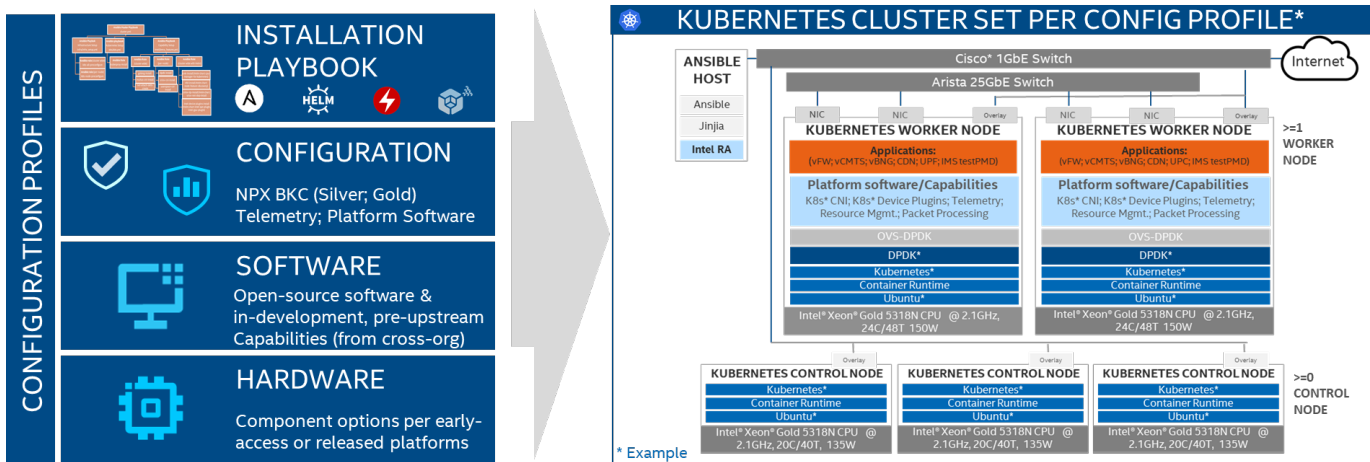


Figure 1. BMRA Illustration and Applicable Elements

2.2 Configuration Profiles

A Configuration Profile describes specific hardware and software bills of material (BOM) and configurations, applicable for a specific deployment. Configuration Profiles consider the best-known configuration (BKC) validated by Intel for optimized performance.²

Installation scripts are available to deploy the required components for a BMRA Flavor. Each BMRA is built on the following:

- **Intel Platform foundation** with Intel processors and technologies.
- **Hardware BOM** optimized for delivering an application at a specific location using a deployment model. For example, to support a UPF workload at the Remote CO, the BMRA deployment is populated with the maximum available network interface cards (NICs).
- **Software BOM** leverages the Intel platform and enables cloud-native adoption.
- **Installation (Ansible) Playbook** automates the installation of a BMRA Flavor per a Configuration Profile specification.

The following Reference Architecture Configuration Profiles are network location-specific:

- **On-Premises Edge Configuration Profile** – Small cluster of stationary or mobile server platforms, ranging from one to four servers. Usage scenarios include data collection from sensors, local (edge) processing, and upstream data transmission. Sample locations are hospitals, factory floors, law enforcement, media, cargo transportation, and power utilities. This Configuration Profile recommends a Kubernetes cluster hardware configuration, software capabilities, and specific hardware and software configurations that typically support enterprise edge workloads used in Smart City deployments, CDN, and Ad-insertion.
- **Access Edge Configuration Profile** – This Configuration Profile is designed for a small cluster of one to four servers for data

² [Workloads and configurations](#). Results may vary.

collection, aggregation, and, in some cases, data processing for high-speed and low-latency services. These nodes represent a large portion of the total number of network elements for an operator with the equipment possibly being in the outside plant in harsh, minimally controlled temperature cabinets. Use cases are: IoT, virtual Radio Access Networks (vRAN) for 5G distribution unit (DU), central unit (CU), vBNG, Multi-access Edge Control (MEC), video security and Smart City.

- **Remote Central Office-Forwarding Configuration Profile** – Clusters ranging from a half rack to a few racks of servers, typically in a pre-existing, repurposed, unmanned structure. The usage scenarios include running latency-sensitive applications near the user (for example, real-time gaming, stock trading, video conferencing). This Configuration Profile addresses a Kubernetes cluster hardware, software capabilities, and configurations that enable high performance for packet forwarding packets. In this category, you can find workloads such as UPF, vBNG, vCMTS, and vCDN.
- **Regional Data Center Configuration Profile** – The Regional Data Center consists of a management domain with many racks of servers, typically managed and orchestrated by a single instance of resource orchestration. Usage scenarios include services such as content delivery, media, mobile connectivity, and cloud services. This Configuration Profile is tailored exclusively and defined for Media Visual Processing workloads such as CDN Transcoding.

Additional Reference Architecture Configuration Profiles are not location-specific and enable flexible deployments per need:

- **Basic Configuration Profile** – A minimum set of software features where network acceleration is the only concern.
- **Full Configuration Profile** – A complete set of all available software features targeted at developers and deployers that are looking to evaluate, control, and configure all the software and hardware ingredients and dependencies simultaneously. This profile is targeting developers and deployers that are looking to evaluate, control, and configure the software and hardware ingredients and dependencies.
- **Build-Your-Own Configuration Profile** – A complete set of all available software features targeted at developers and deployers that are looking to evaluate, control, and configure all the software and hardware ingredients and dependencies individually.

Note: The BMRA with Storage model is implemented using the Storage Configuration Profile, which is a Kubernetes-native high-performance object store that supports deploying MinIO tenants onto private and public cloud infrastructures (“Hybrid” Cloud). The API is also compatible with Amazon Web Services Simple Storage Service (S3).

2.3 Reference Architecture Installation Prerequisites

This section helps you get ready for running the Ansible scripts. Before the Ansible playbook can begin, you must identify the required hardware components, ensure hardware connectivity, and complete the initial configuration, for example BIOS setup. This section describes the minimal system prerequisites needed for the Ansible host and Kubernetes control and worker nodes. It also lists the steps required to prepare hosts for successful deployment. Detailed instructions are provided in relative sections, which are referred to in this section. Steps include:

- Hardware BOM selection and setup
- Required BIOS/UEFI configuration, including virtualization and hyper-threading settings
- Network topology requirements – a list of necessary network connections between the nodes
- Installation of software dependencies needed to execute Ansible playbooks
- Generation and distribution of SSH keys that are used for authentication between the Ansible host and Kubernetes cluster target servers

After satisfying these prerequisites, Ansible playbooks for 2nd, 3rd, and 4th Generation Intel Xeon Scalable processors, and Intel Xeon D processors can be downloaded directly from the dedicated GitHub page ([Container Experience Kits Releases](#)) or cloned using the Git. Request access to the NDA Ansible playbooks for 4th Generation Intel Xeon Scalable processor from your regional Intel representative.

2.3.1 Hardware BOM Selection and Setup for Control and Worker Nodes

Before software deployment and configuration, deploy the physical hardware infrastructure for the site. To obtain ideal performance and latency characteristics for a given network location, Intel recommends the hardware BOMs and configurations described in the following sections:

- Control nodes – Review [Section 3.1](#) for recommended control node assembly.
- Worker nodes – Refer to the following sections for recommended worker node assembly:
 - Base worker node – Review [Section 3.2](#) to satisfy base performance characteristics.
 - Plus worker node – Review [Section 3.3](#) to satisfy plus performance characteristics.
- Configuration Profile BOM – See Sections 7 through 14 for details about hardware BOM selection and setup for your chosen Configuration Profile.

2.3.2 BIOS Selection for Control and Worker Nodes

Enter the UEFI or BIOS menu and update the configuration as listed in [Section 6](#) and [Table 31](#), which describe the BIOS selection in detail.

2.3.3 Operating System Selection for Ansible Host and Control and Worker Nodes

The following Linux operating systems are supported for Control and Worker Nodes:

- RHEL for x86_64 Version 8 (8.5)
- Rocky Linux 8.5
- Ubuntu 20.04.4 LTS
- Ubuntu 21.10
- Ubuntu 22.04
- Ubuntu 22.04 RT

For all supported distributions, the base operating system installation images are sufficient when built using the "Minimal" option during installation. In addition, the following must be met:

- The Control and Worker Nodes must have network connectivity to the Ansible host.
- All systems must have public internet connectivity.
- SSH connections are required. If needed, on Ubuntu, install SSH Server with the following commands (internet access is required):

```
# sudo apt update
# sudo apt install openssh-server
```

2.3.4 Network Interface Requirements for Control and Worker Nodes

The following list provides a brief description of different networks and network interfaces needed for deployment.

- Internet network
 - Ansible host accessible
 - Capable of downloading packages from the internet
 - Can be configured for Dynamic Host Configuration Protocol (DHCP) or with static IP address
- Management network and Calico pod network interface (This can be a shared interface with the internet network)
 - Kubernetes control and worker node inter-node communications
 - Calico pod network runs over this network
 - Configured to use a private static address
- Tenant data networks
 - Dedicated networks for traffic
 - Single Root Input/Output Virtualization (SR-IOV) enabled
 - Virtual function (VF) can be DPDK bound in pod

2.4 Ansible Playbook

This section describes how the Ansible playbooks allow for an automated deployment of a fully functional BMRA cluster, including initial system configuration, Kubernetes deployment, and setup of capabilities as described in [Section 2.5](#).

2.4.1 Ansible Playbook Building Blocks

The following components make up the BMRA Ansible playbooks.

Note: Ansible playbooks for 2nd, 3rd, and 4th Generation Intel Xeon Scalable processors and Intel® Xeon® D processors are open source and available [here](#).

Ansible playbooks for the NDA 4th Generation Intel® Xeon® Scalable processor are available under NDA. To obtain these Ansible playbooks, contact your regional Intel representative.

Configuration Files provide examples of cluster-wide and host-specific configuration options for each of the Configuration Profiles. With minimal changes, these configuration files can be used directly with their corresponding playbooks. The path to these configuration files is:

- inventory.ini
- group_vars/all.yml
- host_vars/node1.yml

For default values in these files, refer to the Configuration Profile-specific sections for BMRA installations: [BMRA Basic Configuration Profile Setup](#); [BMRA Full Configuration Profile Setup](#); [BMRA On-Premises Edge Configuration Profile Setup](#); [BMRA Remote Central Office-Forwarding Configuration Profile Setup](#); [BMRA Regional Data Center Configuration Profile Setup](#), [BMRA for Storage Configuration Profile Setup](#); [BMRA Access Edge Configuration Profile Setup](#); and [BMRA Build-Your-Own Configuration Profile Setup](#).

Ansible Playbooks act as a user entry point and include all relevant Ansible roles and Helm charts. Top-level Ansible playbooks exist for each Configuration Profile, allowing lean use-case-oriented cluster deployments. Each playbook includes only the Ansible roles and configuration files that are relevant for a given use case. See High Level Ansible Playbooks in [Figure 2](#).

- playbooks/basic.yml

- playbooks/full_nfv.yml
- playbooks/on_prem.yml
- playbooks/remote_fp.yml
- playbooks/regional_dc.yml
- playbooks/storage.yml
- playbooks/access.yml
- playbooks/build_your_own.yml

Additionally, an optional Cluster Removal Playbook exists to remove an existing cluster, which is useful to try different deployment models.

- playbooks/redeploy_cleanup.yml

Each of these playbooks encompasses **Ansible Roles** grouped into three main execution phases, which are depicted in [Figure 2](#) and further explained in the next section:

- Infrastructure Setup
- Kubernetes Deployment
- Capabilities Setup

Note that several capabilities setup roles include nested Helm charts for easier deployment and lifecycle management of deployed applications as well as a group of Common Utility Roles that provide reusable functionality across the playbooks.

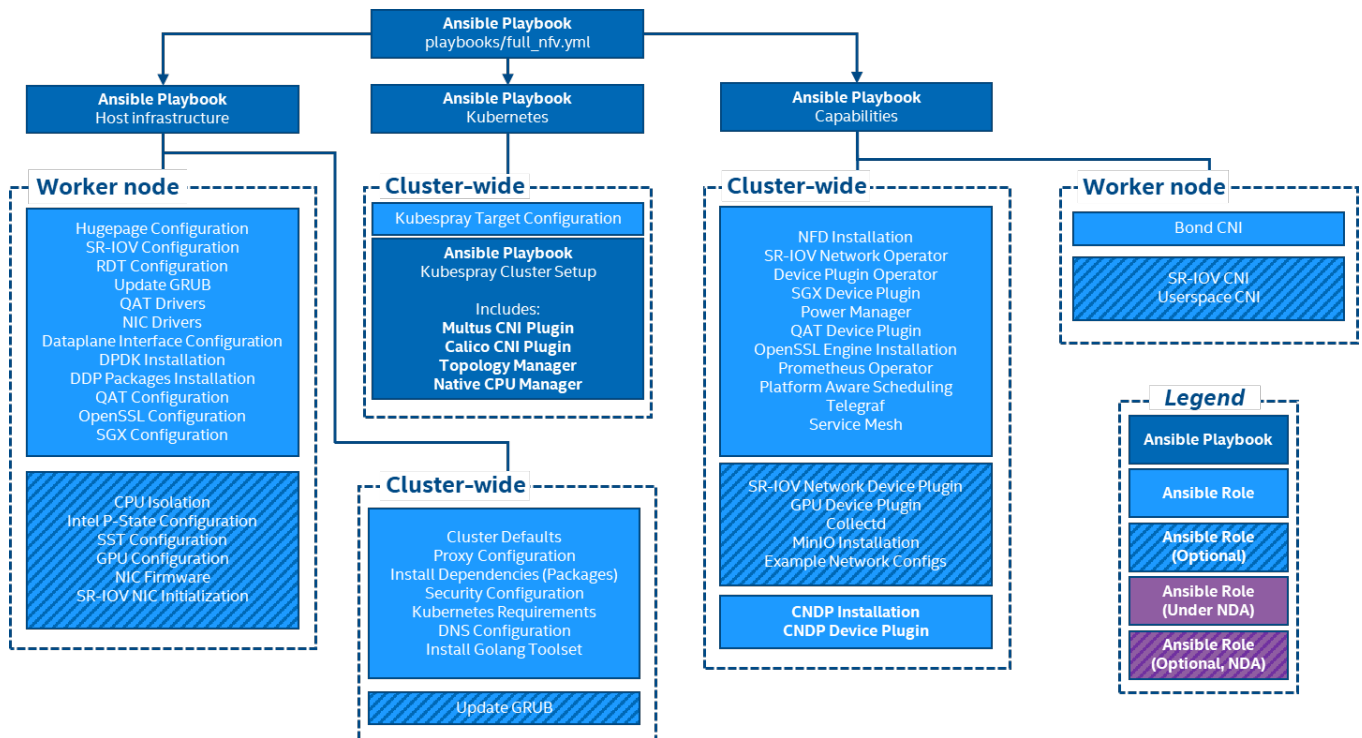


Figure 2. High Level BMRA Ansible Playbooks Architecture³ Full Configuration Profile Example

2.4.2 Ansible Playbook Phases

Regardless of the selected Configuration Profile, the installation process always consists of three main phases:

1. Infrastructure Setup (sub-playbooks in `playbooks/infra/` directory)

These playbooks modify kernel boot parameters and apply the initial system configuration for the cluster nodes. Depending on the selected Configuration Profile, Infrastructure Setup includes:

- Generic host OS preparation, e.g., installation of required packages, Linux kernel configuration, proxy and DNS configuration, and modification of SELinux policies and firewall rules.
- Configuration of the kernel boot parameters according to the user-provided configuration in order to configure CPU isolation, SR-IOV related settings such as IOMMU, hugepages, or explicitly enable/disable Intel P-state technology.
- Configuration of SR-IOV capable network cards and QAT devices. This includes the creation of virtual functions and binding to appropriate Linux kernel modules.

³ Refer to <https://software.intel.com/articles/optimization-notice> for more information regarding performance and optimization choices in Intel software products.

- Network Adapter drivers and firmware updates, which help ensure that all latest capabilities such as Dynamic Device Personalization (DDP) profiles are enabled.
 - Intel® Speed Select Technology (Intel® SST) configuration, which provides control over base frequency.
 - Installation of DDP profiles, which can increase packet throughput, help reduce latency, and lower CPU usage by offloading packet classification and load balancing to the network adapter.
2. **Kubernetes Setup** (in `playbooks/k8s/` directory)
This playbook deploys a high availability (HA) Kubernetes (K8s) cluster using Kubespray, which is a project under the Kubernetes community that deploys production-ready Kubernetes clusters. The Multus container network interface (CNI) plugin, which is specifically designed to support multiple networking interfaces in a Kubernetes environment, is deployed by Kubespray along with Calico and Helm. Preferred security practices are used in the default configuration. On top of Kubespray, there is also a container registry instance deployed to store images of various control-plane Kubernetes applications, such as Telemetry Aware Scheduling (TAS), CPU Manager for Kubernetes (CMK), or device plugins.
 3. **BMRA System Capabilities Setup** (sub-playbooks in the `playbooks/intel` directory)
Advanced networking technologies, enhanced platform awareness, and device plugin features are deployed by this playbook using operators or Helm charts as part of the BMRA. The following capabilities are deployed:
 - Device plugins that allow using, for example, SR-IOV, QAT, and GPU devices in workloads running on top of Kubernetes.
 - SR-IOV CNI plugin, Bond CNI plugin, and Userspace CNI plugin, which allow Kubernetes pods to be attached directly to accelerated and highly available hardware and software network interfaces.
 - Native CPU Manager for Kubernetes (replacement for CMK), which performs a variety of operations to enable core pinning and isolation on a container or a thread level.
 - Node Feature Discovery (NFD), which is a Kubernetes add-on to detect and advertise hardware and software capabilities of a platform that can, in turn, be used to facilitate intelligent scheduling of a workload.
 - Telemetry Aware Scheduling (TAS), which allows scheduling workloads based on telemetry data.
 - Full Telemetry Stack consisting of collectd, Kube-Prometheus, and Grafana, which provides cluster and workload monitoring capabilities and acts as a source of metrics that can be used in TAS to orchestrate scheduling decisions.
 - MinIO operator/console, which supports deploying MinIO tenants onto private and public cloud infrastructures ("Hybrid" Cloud).

2.5 Deployment Using Ansible Playbook

This section describes common steps to obtain the BMRA Ansible Playbooks source code, prepare target servers, configure inventory and variable files, and deploy the BMRA Kubernetes cluster.

2.5.1 Prepare Target Servers

For each target server that will act as a control or worker node, you must make sure that it meets the following requirements:

- Install Python 3. The following example assumes that the host is running RHEL. Other operating systems may have slightly different installation steps:

```
yum install python3
```
- Internet access on all target servers is mandatory. Proxies are supported and can be configured in the Ansible vars.
- BIOS configuration matching the desired state is applied. For details, refer to the specific Configuration Profile section below for your profile: [BMRA Basic Configuration Profile Setup](#); [BMRA Full Configuration Profile Setup](#); [BMRA On-Premises Edge Configuration Profile Setup](#); [BMRA Remote Central Office-Forwarding Configuration Profile Setup](#); [BMRA Regional Data Center Configuration Profile Setup](#); [BMRA for Storage Configuration Profile Setup](#); [BMRA Access Edge Configuration Profile Setup](#); and [BMRA Build-Your-Own Configuration Profile Setup](#).

For detailed steps on how to build the Ansible host, refer to [Section 6.1](#).

2.5.2 Prepare Ansible Host and Configuration Templates

Perform the following steps:

1. Log in to your Ansible host (the one that you will run these Ansible playbooks from).
2. Install packages on Ansible host. The following example assumes that the host is running RHEL. Other operating systems may have slightly different installation steps:

```
yum install python3
pip3 install --upgrade pip
```
3. Enable passwordless login between all nodes in the cluster.
Create authentication SSH-Keygen keys on Ansible host:

```
ssh-keygen
```
4. SSH is used by the Ansible host to communicate with each target node. Configure the same SSH keys on each machine. Copy your generated public keys to all the nodes from the Ansible host:

```
ssh-copy-id root@<target_server_address>
```
5. Clone the source code and change working directory.

```
git clone https://github.com/intel/container-experience-kits/
cd container-experience-kits
```


Check out the latest version of the playbooks – using the tag from [Table 38](#), for example:

```
git checkout v22.05
```

Note: Alternatively go to [Container Experience Kits Releases](#), download the latest release tarball, and unarchive it:

```
wget https://github.com/intel/container-experience-kits/archive/v22.05.tar.gz
tar xzf v22.05.tar.gz
cd container-experience-kits-22.05
```

6. Initialize Git submodules to download Kubespray code.

```
git submodule update --init
```

7. Decide which Configuration Profile that you want to deploy and export the environmental variable.

For Kubernetes **Basic** Configuration Profile deployment:

```
export PROFILE=basic
```

For Kubernetes **Regional Data Center** Configuration Profile deployment:

```
export PROFILE=regional_dc
```

For Kubernetes **Remote Central Office-Forwarding** Configuration Profile deployment:

```
export PROFILE=remote_fp
```

For Kubernetes **On-Premises Edge** Configuration Profile deployment:

```
export PROFILE=on_prem
```

For Kubernetes **Full** Configuration Profile deployment:

```
export PROFILE=full_nfv
```

For Kubernetes **Storage** Configuration Profile deployment:

```
export PROFILE=storage
```

For Kubernetes **Access Edge** Configuration Profile deployment:

```
export PROFILE=access
```

For Kubernetes **Build-Your-Own** Configuration Profile deployment:

```
export PROFILE=build_your_own
```

8. Install requirements needed by deployment scripts.

```
pip3 install -r requirements.txt
```

9. Generate example profiles. Be aware of the machine's architecture and data plane network before generating profiles. Example machine architectures (ARCH) are `spr`, `icx`, and `clx` and data plane networks (NIC) are `fv1` and `cv1`.

```
make k8s-profile PROFILE=$PROFILE ARCH=spr NIC=cv1
```

2.5.3 Update Ansible Inventory File

Perform the following steps:

1. Edit the `inventory.ini` file generated in the previous steps.
2. In the section `[all]`, specify all your target servers. Use their actual hostnames and Management IP addresses. Also update `ansible_user` and `ansible_password` to match the SSH configuration of the target servers. If any of the servers are configured with passwordless SSH, the `ansible_password` host variable can be removed.

```
[all]
controller1 ansible_host=10.0.0.1 ip=10.0.0.1 ansible_user=USER ansible_password=XXXX
controller2 ansible_host=10.0.0.2 ip=10.0.0.2 ansible_user=USER ansible_password=XXXX
controller3 ansible_host=10.0.0.3 ip=10.0.0.3 ansible_user=USER ansible_password=XXXX
node1 ansible_host=10.0.0.4 ip=10.0.0.4 ansible_user=USER ansible_password=XXXX
node2 ansible_host=10.0.0.5 ip=10.0.0.5 ansible_user=USER ansible_password=XXXX
localhost ansible_connection=local ansible_python_interpreter=/usr/bin/python3
```

```
[vm_host]
```

```
[kube_control_plane]
```

```
controller1
controller2
controller3
```

```
[etcd]
```

```
controller1
controller2
controller3
```

```
[kube_node]
```

```
node1
node2
```

```
[k8s_cluster:children]
```

```
kube_control_plane
kube_node

[all:vars]
ansible_python_interpreter=/usr/bin/python3
```

2.5.4 Update Ansible Host and Group Variables

Perform the following steps.

1. Create `host_vars/<hostname>.yaml` files for all worker nodes, matching their hostnames from the inventory file. The provided `host_vars/node1.yaml` file can be used as a template.
2. Edit all `host_vars/<hostname>.yaml` and `group_vars/all.yaml` files to match your desired configuration. Each Configuration Profile uses its own set of variables. Refer to the specific Configuration Profile section for your profile to get a full list of variables and their documentation: [Section 7, BMRA Basic Configuration Profile Setup](#); [Section 8, BMRA Full Configuration Profile Setup](#); [Section 9, BMRA On-Premises Edge Configuration Profile Setup](#); [Section 10, BMRA Remote Central Office-Forwarding Configuration Profile Setup](#); [Section 11, BMRA Regional Data Center Configuration Profile Setup](#); [Section 12, BMRA for Storage Configuration Profile Setup](#); [Section 13, BMRA Access Edge Configuration Profile Setup](#); and [Section 14, BMRA Build-Your-Own Configuration Profile Setup](#).

2.5.5 Run Ansible Cluster Deployment Playbook

After the inventory and vars are configured, you can run the provided playbooks from the root directory of the project.

It is recommended that you check dependencies of components enabled in `group_vars` and `host_vars` with the packaged dependency checker:

```
ansible-playbook -i inventory.ini playbooks/preflight.yaml
```

If you are deploying an RHEL 8 cluster, you need to patch Kubespray:

```
ansible-playbook -i inventory.ini playbooks/k8s/patch_kubespray.yaml
```

Otherwise, you can skip directly to your chosen Configuration Profile playbook:

```
ansible-playbook -i inventory.ini playbooks/${PROFILE}.yaml
```

Pay attention to logs and messages displayed on the screen. Depending on the selected Configuration Profile, network bandwidth, storage speed, and other similar factors, the execution may take up to 30-40 minutes.

After the playbook finishes without any “Failed” tasks, you can proceed with the deployment validation described in [Section 5, Post Deployment Verification Guidelines](#).

Note: Additional information can be found in the Ansible project root directory readme.

2.5.6 Run Ansible Cluster Removal Playbook

If the playbook fails or if you want to clean up the environment to run a new deployment, you can optionally use the provided Cluster Removal Playbook (`redeploy_cleanup.yaml`) to remove any previously installed Kubernetes and related plugins.

```
ansible-playbook -i inventory.ini playbooks/redeploy_cleanup.yaml
```

After successful removal of Kubernetes components, you can repeat [Section 2.5.5](#).

Note: Any OS and/or hardware configurations (for example, proxies, drivers, kernel parameters) are not reset by the cleanup playbook.

3 Reference Architecture Hardware Components and BIOS

For all Configuration Profiles, this section provides a menu of all possible hardware components for control node and worker node as well as the BIOS components available.

3.1 Hardware Components List for Control Node

The following tables list the hardware options for control nodes.

Table 4. Hardware Options for Control Node – 2nd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
2nd Generation Intel® Xeon® Scalable processors	Intel® Xeon® Gold 5218 or 5218N processor at 2.3 GHz, 16 C/32 T, 125 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	DRAM only configuration: 192 GB (12 x 16 GB DDR4 2666 MHz)	Required
Network Adapter	Option 1: Dual Port 25 GbE Intel® Ethernet Network Adapter XXV710-DA2 SFP28+, or	Required
	Option 2: Dual Port 10 GbE Intel® Ethernet Converged Network Adapter X710-DA2 SFP+, or	
	Option 3: Dual Port 10 GbE Intel® Ethernet Converged Network Adapter X520-DA2 SFP+	
Intel® QAT	Intel® QuickAssist Adapter 8970 (PCIe) AIC or equivalent third-party Intel® C620 Series Chipset Intel® QAT enabled PCIe AIC, with minimum 8 lanes of PCIe connectivity	Recommended
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® NVMe SSD DC P4510 Series at 2 TB or equivalent (Recommended NUMA Aligned)	Recommended
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 5. Hardware Options for Control Node – 3rd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
3rd Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 5318N processor at 2.1 GHz, 20 C/40 T, 135 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	256 GB DRAM (16 x 16 GB DDR4, 2666 MHz)	Required
Network Adapter	Dual Port 100 GbE Intel® Ethernet Network Adapter E810-CQDA2 QSFP28	Required
Intel® QAT	Intel® QuickAssist Adapter 8960 or 8970 (PCIe*) AIC or equivalent third-party Intel® C620 Series Chipset	Recommended
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Recommended
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 6. Hardware Options for Control Node – 4th Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
4th Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 5418N processor at 2.0 GHz, 24 C/48 T, 165 W	Required
Memory	DRAM only configuration: 256 GB DRAM (16 x 16 GB DDR5)	Required
Network Adapter	Intel® Ethernet Network Adapter E810-CQDA2 or E810-XXVDA2	Required

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Intel® QAT	Integrated in the processor	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Recommended
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 7. Hardware Options for Control Node – Intel® Xeon® D Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Intel® Xeon® D processors	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher	Required
Memory	DRAM only configuration: 16 GB DDR4 2933 MHz	Required
Network Adapter	2 x 10 GbE integrated Ethernet ports	Required
Intel® QAT	20 G Intel® QAT	Recommended
Storage (Boot Drive)	Intel® SSD 256 GB 2.5" internal SSD/M.2	Required
Additional Plug-in cards	N/A	

3.2 Hardware Components List for Worker Node Base

The following tables list the hardware options for worker nodes in the “base” configuration. If your configuration needs improved processing, you may choose to use the “plus” configuration instead ([Section 3.3](#)).

Table 8. Hardware Components for Worker Node Base – 2nd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
2nd Generation Intel® Xeon® Scalable processors	Intel® Xeon® Gold 6230 processor @ 2.1 GHz or 6230N CPU @ 2.3 GHz 20 C/40 T, 125 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	Option 1: DRAM only configuration: 384 GB (12 x 32 GB DDR4 2666 MHz)	Required
	Option 2: DRAM only configuration: 384 GB (24 x 16 GB DDR4 2666 MHz)	
	Option 3: DRAM + Intel® Optane™ Persistent Memory DRAM: 192 GB (12x 16 GB DDR4, 2666 MHz)	
Intel® Optane™ Persistent Memory	512 GB (4x 128 GB Intel® Optane™ persistent memory in 2-1-1 Topology)	Recommended
Network Adapter	Option 1: Dual Port 100 GbE Intel® Ethernet Network Adapter E810-CQDA2 QSFP28	Required
	Option 2: Intel® Ethernet Network Adapter XXV710-DA2 QSFP28	
Intel® QAT	Intel® QuickAssist Adapter 8970 (PCIe) AIC or equivalent third-party Intel® C620 Series Chipset Intel® QAT enabled with minimum 8 lanes of PCIe connectivity	Required
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® NVMe SSD DC P4510 Series P4510 at 2 TB or equivalent (Recommended NUMA Aligned)	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 9. Hardware Components for Worker Node Base – 3rd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
3rd Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 5318N processor at 2.1 GHz, 24 C/48 T, 150 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	Option 1: DRAM only configuration: 256 GB (8 x 32 GB DDR4, 2666 MHz) Option 2: DRAM only configuration: 256 GB (16 x 16 GB DDR4, 2666 MHz)	Required
Intel® Optane™ Persistent Memory	512 GB (4x 128 GB Intel® Optane™ persistent memory in 2-1-1 Topology)	Recommended
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2 Option 2: Intel® Ethernet Network Adapter E810-XXVDA-2	Required
Intel® QAT	Intel® QuickAssist Adapter 8960 or 8970 (PCIe*) AIC or equivalent third-party Intel® C620 Series Chipset	Required
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 10. Hardware Components for Worker Node Base – 4th Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
4th Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 5418N processor at 2.0 GHz, 24 C/48 T, 165 W	Required
Memory	DRAM only configuration: 256 GB DRAM (16 x 16 GB DDR5)	Required
Intel® Optane™ Persistent Memory	512 GB (4 x 128 GB Intel® Optane™ persistent memory in 2-1-1 topology)	Recommended
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2 Option 2: Intel® Ethernet Network Adapter E810-2CQDA2	Required
Intel® QAT	Integrated in the processor	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 11. Hardware Components for Worker Node Base (Access Edge - vRAN) – 4th Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
4th Generation Intel Xeon Scalable processors	Intel® Xeon®-SP 5411N 24 C/48 T 1.9 GHz 165 W	Required
Memory	DRAM only configuration: 128 GB DRAM (8 x 16 GB DDR5)	Required
Intel® Optane™ Persistent Memory	512 GB (4 x 128 GB Intel® Optane™ persistent memory in 2-1-1 topology)	Recommended
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2 Option 2: Intel® Ethernet Network Adapter E810-XXVAM-DA4	Required
Intel® QAT	Integrated in the processor	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Additional Plug-in cards	Intel® vRAN Accelerator ACC100/200 Adapter	Required

Table 12. Hardware Components for Worker Node Base – Intel® Xeon® D Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Intel® Xeon® D processors	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or Intel® Xeon® D-1700 processor, 10 core LCC, or Intel Xeon D-2733 NT processor, 8 cores HCC, 80 W	Required
Memory	DRAM only configuration: 32 GB DDR4 2667 MHz	Required
Network Adapter	2 x 10/25 GbE integrated Ethernet ports OR Intel® Ethernet Network Adapter E810-CQDA2	Required
Intel® QAT	Intel® QuickAssist Adapter 8960 or 8970 (PCIe*) AIC or equivalent third-party Intel® C620 Series Chipset	Recommended
Storage (Boot Drive)	Intel® SSD 256 GB 2.5" internal SSD/M.2	Required
Additional Plug-in cards	N/A	

3.3 Hardware Components List for Worker Node Plus

The following tables list the hardware options for worker nodes in the “plus” configuration, which helps improve the processing capability due to more powerful CPU, more memory, more disk space, and a faster network.

Table 13. Hardware Components for Worker Node Plus – 2nd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
2nd Generation Intel® Xeon® Scalable processors	Intel® Xeon® Gold 6252 processor @ 2.1 GHz or 6252N processor @ 2.3 GHz 24 C/48 T, 150 W, or higher number Intel® Xeon® Gold/Platinum CPU SKU	Required
Memory	Option 1: DRAM only configuration: 384 GB (12 x 32 GB DDR4 2666 MHz) Option 2: DRAM only configuration: 384 GB (24 x 16 GB DDR4 2666 MHz) Option 3: DRAM + Intel® Optane™ persistent memory DRAM: 192 GB (12 x 16 GB DDR4 2666 MHz)	Required
Intel® QAT	Intel® C620 Series Chipset integrated on base board Intel® C627/C628 Chipset, integrated with NUMA connectivity to each CPU or minimum 16 Peripheral Component Interconnect Express (PCIe) lane connectivity to one CPU	Required
Intel® Optane™ Persistent Memory	Option 1: 1 TB (8x 128 GB Intel® Optane™ persistent memory in 2-2-1 Topology) Option 2: 1.5 TB (12x 128 GB Intel® Optane™ persistent memory in 2-2-2 Topology)	Recommended
Network Adapter	Option 1: Dual Port 25 GbE Intel® Ethernet Network Adapter X710-DA4 SFP+, or Option 2: Dual Port 100 GbE Intel® Ethernet Network Adapter E810-CQDA2 QSFP28	Required
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® NVMe SSD DC P4510 Series at 2 TB or equivalent (Recommended NUMA Aligned)	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	N/A	

Table 14. Hardware Components for Worker Node Plus – 3rd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
3rd Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 6338N CPU @ 2.2 GHz 32 C/64 T, 185 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	Option 1: DRAM only configuration: 512 GB (16 x 32 GB DDR4, 2666 MHz)	Required

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Intel® QAT	Option 2: DRAM only configuration: 512 GB (3 2x 16 GB DDR4, 2666 MHz)	Required
	Intel® C620 Series Chipset integrated on base board Intel® C627/C628 Chipset, integrated with NUMA connectivity to each CPU or minimum 16 Peripheral Component Interconnect Express (PCIe) lane connectivity to one CPU	
Intel® Optane™ Persistent Memory	Option 1: 1 TB (8 x 128 GB Intel® Optane™ persistent memory in 8+4 Topology)	Recommended
	Option 2: 2 TB (16 x 128 GB Intel® Optane™ persistent memory in 8+8 Topology)	
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2	Required
	Option 2: Intel® Ethernet Network Adapter E810-2CQDA2	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 4 TB or equivalent drive (recommended NUMA aligned)	Recommended
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	Intel® Server Graphics 1 card	Optional

Table 15. Hardware Components for Worker Node Plus – 4th Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
4th Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 6428N processor at 1.8GHz, 32 C/64 T, 185 W	Required
Memory	Option 1: DRAM only configuration: 512 GB (16 x 32 GB DDR5)	Required
	Option 2: DRAM only configuration: 512 GB (32 x 16 GB DDR5)	
Intel® QAT	Integrated in the processor	Required
Intel® Optane™ Persistent Memory	Option 1: 1 TB (8 x 128 GB Intel® Optane™ persistent memory in 8+4 Topology)	Recommended
	Option 2: 2 TB (16 x 128 GB Intel® Optane™ persistent memory in 8+8 Topology)	
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2	Required
	Option 2: Intel® Ethernet Network Adapter E810-2CQDA2	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 4 TB or equivalent drive (recommended NUMA aligned)	Recommended
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	Intel® Server Graphics 1 card	Optional

Table 16. Hardware Components for Worker Node Plus (Access Edge - vRAN) – 4th Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
4th Generation Intel Xeon Scalable processors	Intel® Xeon®-SP Gold 6421N 32 C/ 64 T 1.8 GHz 185 W	Required
Memory	DRAM only configuration: 128 GB DRAM (8 x 16 GB DDR5)	Required
Intel® Optane™ Persistent Memory	1 TB (8 x 128 GB Intel® Optane™ persistent memory in 2-1-1 topology)	Recommended
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2	Required
	Option 2: Intel® Ethernet Network Adapter E810-2CQDA2	
	Option 3: Intel® Ethernet Network Adapter E810-XXVAM-DA4	
Intel® QAT	Integrated in the processor	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Intel® SSD D7-P5510 Series at 3.84 TB or equivalent drive (recommended NUMA aligned)	Required

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required
Additional Plug-in cards	Intel® vRAN Accelerator ACC100/200 Adapter	Required

Table 17. Hardware Components for Worker Node Plus – Intel® Xeon® D Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
Intel® Xeon® D processors	Intel® Xeon® D-2766NT processor 2.1 GHz, 14 core HCC, 97 W, or higher	Required
Memory	DRAM only configuration: 64 GB DDR4 2667 MHz	Required
Network Adapter	4 x 10/25 GbE integrated Ethernet ports	Required
	Intel® Ethernet Network Adapter E810-CQDA2	
Intel® QAT	Intel® QuickAssist Adapter 8960 or 8970 (PCIe*) AIC or equivalent third-party Intel® C620 Series Chipset	Recommended
Storage (Boot Drive)	Intel® SSD 512 GB 2.5" internal SSD/M.2	Required
Additional Plug-in cards	N/A	

3.4 Hardware Components List for Storage Node

Table 18. Hardware Components for Storage Node – 3rd Generation Intel Xeon Scalable Processor

INGREDIENT	REQUIREMENT	REQUIRED/ RECOMMENDED
3rd Generation Intel Xeon Scalable processors	Intel® Xeon® Gold 6338N CPU @ 2.2 GHz 32 C/64 T, 185 W, or higher number Intel® Xeon® Gold or Platinum CPU SKU	Required
Memory	Option 1: DRAM only configuration: 512 GB (16 x 32 GB DDR4, 2666 MHz)	Required
	Option 2: DRAM only configuration: 512 GB (32 x 16 GB DDR4, 2666 MHz)	
Intel® QAT	Intel® C620 Series Chipset integrated on base board Intel® C627/C628 Chipset, integrated with NUMA connectivity to each CPU or minimum 16 Peripheral Component Interconnect Express (PCIe) lane connectivity to one CPU	Required
Intel® Optane™ Persistent Memory	512 GB (4 x 128 GB Intel® Optane™ persistent memory in 2-1-1 topology)	Recommended
Network Adapter	Option 1: Intel® Ethernet Network Adapter E810-CQDA2	Required
	Option 2: Intel® Ethernet Network Adapter E810-2CQDA2	
Storage (Boot Drive)	Intel® SATA Solid State Drive D3 S4510 at 480 GB or equivalent boot drive	Required
Storage (Capacity)	Kioxia CM6 3.2 TB NVMePCIe4x4 2.5"15mm S1E 3DWPDP - KCM6XVUL3T20	Required
LAN on Motherboard (LOM)	10 Gbps or 25 Gbps port for Preboot Execution Environment (PXE) and Operation, Administration, and Management (OAM)	Required
	1/10 Gbps port for Management Network Adapter	Required

3.5 Hardware BOMs Supporting All BMRA Configuration Profiles

The following tables list the hardware BOMs for control nodes, worker node base, and worker node plus.

Choose your controller profile from the three available profiles (Controller_xGen_1, Controller_xGen_2, or Controller_xGen_3) based on your BIOS profile (Profiles available: Energy Balance, Deterministic, or Max Performance, respectively).

The profiles for Worker Nodes vary with respect to network interface card, Intel® QuickAssist Technology, and BIOS profiles. You may choose based on the requirements for the workloads to be run on the worker nodes.

Table 19. Control Node Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor

NAME	Controller_2ndGen_1	Controller_2ndGen_2	Controller_2ndGen_3
Platform	S2600WFQ	S2600WFQ	S2600WFQ
CPU/node	2x 5218 or 2x 5218N	2x 5218 or 2x 5218N	2x 5218 or 2x 5218N

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

Mem	192 GB	192 GB	192 GB
Intel Optane Persistent Memory	Recommended	Recommended	Recommended
Network Adapter	2x XXV710-DA2 or 2x X710-DA2 or 2x X520-DA2	2x XXV710-DA2	2x XXV710-DA2
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Recommended - 2x (1 per NUMA)	Recommended - 2x (1 per NUMA)	Recommended - 2x (1 per NUMA)
LOM	No	No	No
Intel® QAT	Recommended	N/A	N/A

BIOS Configuration

Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	No	Yes	Yes
Intel® VT-d enabled	No	Yes	Yes
BIOS Profile	Energy Balance	Deterministic	Max Performance
Virtualization enabled	No	Yes	Yes

Table 20. Control Node Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor

NAME	Controller_3rdGen_1	Controller_3rdGen_2	Controller_3rdGen_3
Platform	M50CYP	M50CYP	M50CYP
CPU/node	2x 5318N 20c	2x 5318N 20c	2x 5318N 20c
Mem	256 GB	256 GB	256 GB
Intel Optane Persistent Memory	Recommended	Recommended	Recommended
Network Adapter	2x E810-CQDA2	2x E810-CQDA2	2x E810-CQDA2
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Recommended - 2x (1 per NUMA)	Recommended - 2x (1 per NUMA)	Recommended - 2x (1 per NUMA)
LOM	No	No	No
Intel® QAT	Recommended	N/A	N/A

BIOS Configuration

Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	No	Yes	Yes
Intel® VT-d enabled	No	Yes	Yes

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

BIOS Profile	Energy Balance	Deterministic	Max Performance
Virtualization enabled	No	Yes	Yes

Table 21. Control Node Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor

NAME	Controller_4thGen_1	Controller_4thGen_2	Controller_4thGen_3
Platform	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q
CPU/node	2x 5418N	2x 5418N	2x 6428N
Mem	256 GB	256 GB	512 GB
Intel Optane Persistent Memory	Recommended	Recommended	Recommended – 2 TB
Network Adapter	2x E810-CQDA2 or E810-XXVDA2	2x E810-CQDA2 or E810-XXVDA2	2x E810-2CQDA2 or 4x E810-CQDA2
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Recommended - 2x (1 per NUMA)	Recommended - 2x (1 per NUMA)	Required- 4x (2 per NUMA)
LOM	No	No	Yes
Intel® QAT	Integrated in the processor	Integrated in the processor	Integrated in the processor
Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	No	Yes	Yes
Intel® VT-d enabled	No	Yes	Yes
BIOS Profile	Energy Efficiency Turbo	Energy Efficiency Turbo	Energy Balance Turbo
Virtualization enabled	No	Yes	Yes

Table 22. Control Node Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor

NAME	Controller_Xeon_D_1	Controller_Xeon_D_2	Controller_Xeon_D_3
Platform	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101
CPU/node	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher
Mem	16 GB DDR4 2933 MHz	16 GB DDR4 2933 MHz	16 GB DDR4 2933 MHz
Network Adapter	2 x 10 GbE integrated Ethernet ports	2 x 10 GbE integrated Ethernet ports	2 x 10 GbE integrated Ethernet ports

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

Storage (Boot Media)	Intel® SSD 256 GB 2.5" internal SSD/M.2	Intel® SSD 256 GB 2.5" internal SSD/M.2	Intel® SSD 256 GB 2.5" internal SSD/M.2
LOM	No	No	No
Intel® QAT AIC	Recommended	N/A	N/A

BIOS Configuration

Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	No	Yes	Yes
Intel® VT-d enabled	No	Yes	Yes
BIOS Profile	Energy Balance	Deterministic	Max Performance
Virtualization enabled	No	Yes	Yes

Table 23. Worker Node Base Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor

NAME	Worker_2ndGen_Base_1	Worker_2ndGen_Base_2	Worker_2ndGen_Base_3
Platform	S2600WFQ	S2600WFQ	S2600WFQ
CPU/node	2x 6230 or 6230N	2x 6230 or 6230N	2x 6230 or 6230N
Mem	384 GB	384 GB	384 GB
Intel Optane Persistent Memory	Recommended – 512 GB	Recommended – 512 GB	Recommended – 512 GB
Network Adapter	2x XXV710-DA2 or 2x E810-CQDA2	2x XXV710-DA2	2x XXV710-DA2
Storage (Boot Media)	Required – 2x	Required – 2x	Required – 2x
Storage (Capacity)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)
LOM	No	No	Yes
Intel® QAT	No	Optional	Yes
Additional Plug-in cards	No	No	No
BIOS Configuration			
Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes
BIOS Profile	Energy Balance	Deterministic	Max Performance
Virtualization enabled	No	Yes	Yes

Table 24. Worker Node Plus Hardware Setup for all Configuration Profiles – 2nd Generation Intel Xeon Scalable Processor

NAME	Worker_2ndGen_Plus_1	Worker_2ndGen_Plus_2
Platform	S2600WFQ	S2600WFQ
CPU/node	2x 6252 or 6252N	2x 6252 or 6252N
Mem	384 GB	384 GB
Intel Optane Persistent Memory	Recommended – 1 TB/1.5 TB	Recommended – 1 TB/1.5 TB
Network Adapter	2x E810-CQDA2	2x E810-CQDA2
Storage (Boot Media)	Required – 256 GB	Required – 256 GB
LOM	No	Yes

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

NAME	Worker_2ndGen_Plus_1	Worker_2ndGen_Plus_2
Intel® QAT	No	Yes
Additional Plug-in cards	No	No
BIOS Configuration		
Intel® HT Technology enabled	Yes	Yes
Intel® VT-x enabled	Yes	Yes
Intel® VT-d enabled	Yes	Yes
BIOS Profile	Deterministic	Max Performance
Virtualization enabled	Yes	Yes

Table 25. Worker Node Base Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor

NAME	Worker_3rdGen_Base_1	Worker_3rdGen_Base_2	Worker_3rdGen_Base_3
Platform	M50CYP	M50CYP	M50CYP
CPU/node	2x 5318N 24c	2x 5318N 24c	2x 5318N 24c
Mem	512 GB	512 GB	512 GB
Intel Optane Persistent Memory	Recommended – 512 GB	Recommended – 512 GB	Recommended – 512 GB
Network Adapter	2x E810-CQDA2	2x E810-CQDA2	2x E810-2CQDA2 or 4x E810-CQDA2
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)
LOM	No	Yes	No
Intel® QAT	No	Yes	Optional
Additional Plug-in cards	No	No	No
BIOS Configuration			
Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes
BIOS Profile	Energy Balance	Max Performance	Deterministic
Virtualization enabled	No	Yes	Yes

Table 26. Worker Node Plus and Storage Node Hardware Setup for all Configuration Profiles – 3rd Generation Intel Xeon Scalable Processor

NAME	Worker_3rdGen_Plus_1	Worker_3rdGen_Plus_2	Worker_3rdGen_Plus_3	Storage_3rdGen_1
Platform	M50CYP	M50CYP	M50CYP	M50CYP
CPU/node	2x 6338N 32c	2x 6338N 32c	2x 6338N 32c	2x 6338N 32c
Mem	512 GB	512 GB	512 GB	512 GB
Intel Optane Persistent Memory	Recommended – 512 GB	Recommended – 512 GB	Recommended – 512 GB	Recommended – 512 GB
Network Adapter	2x E810-2CQDA2 or 4x E810-CQDA2	2x E810-2CQDA2	2x E810-2CQDA2 or 4x E810-CQDA2	2x E810-2CQDA2 or 4x E810-CQDA2

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

NAME	Worker_3rdGen_Plus_1	Worker_3rdGen_Plus_2	Worker_3rdGen_Plus_3	Storage_3rdGen_1
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Required- 4x (2 per NUMA)	Required- 4x (2 per NUMA)	Required- 4x (2 per NUMA)	Required Kioxia 3.2TB - 8x
LOM	Yes	Yes	No	Yes
Intel® QAT	Yes	No	Optional	Yes
Additional Plug-in cards	No	Intel Server GPU	No	No
BIOS Configuration				
Intel® HT Technology enabled	Yes	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes	Yes
BIOS Profile	Max Performance	Max Performance	Deterministic	Max Performance
Virtualization enabled	Yes	Yes	Yes	Yes

Table 27. Worker Node Base Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor

NAME	Worker_4thGen_Base_1	Worker_4thGen_Base_2	Worker_4thGen_Base_3 (Access Edge - vRAN)
Platform	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q / Ruby Pass
CPU/node	2x 5418N	2x 5418N	1x 5411N
Mem	512 GB	512 GB	128 GB
Intel Optane Persistent Memory	Recommended – 512 GB	Recommended – 512 GB	Recommended – 512 GB
Network Adapter	2x E810-CQDA2	2x E810-CQDA2	2x E810-CQDA2 or 8x E810-XXVAM-DA4
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)	Required- 2x (1 per NUMA)
LOM	No	Yes	Yes
Intel® QAT	Integrated in the processor	Integrated in the processor	Integrated in the processor
Additional Plug-in cards	No	No	Intel® vRAN Accelerator ACC100/200
BIOS Configuration			
Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes
BIOS Profile	Energy Efficiency Turbo	Max Performance Turbo	Low Latency
Virtualization Enable	No	Yes	Yes

Table 28. Worker Node Plus Hardware Setup for all Configuration Profiles – 4th Generation Intel Xeon Scalable Processor

NAME	Worker_4thGen_Plus_1	Worker_4thGen_Plus_2	Worker_4thGen_Plus_3	Worker_4thGen_Plus_4 (Access Edge - vRAN)
Platform	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q	Archer City / Quanta - S6Q / Ruby Pass
CPU/node	2x 6428N	2x 6428N	2x 6428N	1x 6421N
Mem	512 GB	512 GB	512 GB	128 GB
Intel Optane Persistent Memory	Recommended – 2 TB	Recommended – 1 TB	Recommended – 2 TB	Recommended – 1 TB
Network Adapter	2x E810-2CQDA2 or 4x E810-CQDA2	2x E810-2CQDA2 or 8x E810-XXVAM-DA4	2x E810-2CQDA2 or 4x E810-CQDA2	1x E810-2CQDA2 or 2x E810-CQDA2 8x E810-XXVAM-DA4
Storage (Boot Media)	Required - 2x	Required - 2x	Required - 2x	Required - 2x
Storage (Capacity)	Required- 4x (2 per NUMA)	Required- 4x (2 per NUMA)	Required- 4x (2 per NUMA)	Required- 4x (1 per NUMA)
LOM	Yes	Yes	No	Yes
Intel® QAT	Integrated in the processor	Integrated in the processor	Integrated in the processor	Integrated in the processor
Additional Plug-in cards	No	Intel Server GPU	No	Intel® vRAN Accelerator ACC100/200
BIOS Configuration				
Intel® HT Technology enabled	Yes	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes	Yes
BIOS Profile	Energy Balance Turbo	Energy Balance Turbo	Max Performance Turbo	Low Latency
Virtualization enabled	Yes	Yes	Yes	Yes

Table 29. Worker Node Base Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor

NAME	Worker_Xeon_D_Base_1	Worker_Xeon_D_Base_2	Worker_Xeon_D_Base_3
Platform	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101 or Taylors Falls Reference Design	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101 or Taylors Falls Reference Design	Intel® SoC Based Server Reference Platform Board (Codename Brighton City) K97971-101 or Taylors Falls Reference Design
CPU/node	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher	Intel® Xeon® D-1700 processor, 4 core LCC, 45 W, or higher
Mem	16 GB DDR4 2933 MHz	16 GB DDR4 2933 MHz	16 GB DDR4 2933 MHz
Network Adapter	Intel® Ethernet Network Adapter E810-CQDA2	Intel® Ethernet Network Adapter E810-CQDA2	Intel® Ethernet Network Adapter E810-CQDA2
Storage (Boot Media)	Required – 256 GB	Required – 256 GB	Required – 256 GB
LOM	No	Yes	No
Intel® QAT	No	Yes	Optional
Additional Plug-in cards	No	No	No

BIOS Configuration

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

Intel® HT Technology enabled	Yes	Yes	Yes
Intel® VT-x enabled	Yes	Yes	Yes
Intel® VT-d enabled	Yes	Yes	Yes
BIOS Profile	Max Performance	Deterministic	Max Performance
Virtualization enabled	Yes	Yes	Yes

Table 30. Worker Node Plus Hardware Setup for all Configuration Profiles – Intel® Xeon® D Processor

NAME	Worker_Xeon_D_Plus_1	Worker_Xeon_D_Plus_2
Platform	Intel® SoC Based Server Reference Platform Board (Codename Moro City)	Intel® SoC Based Server Reference Platform Board (Codename Moro City)
CPU/node	Intel Xeon D-2700 processor, 16 core HCC, 105 W, or higher	Intel Xeon D-2700 processor, 16 core HCC, 105 W, or higher
Mem	64 GB DDR4 2933 MHz	64 GB DDR4 2933 MHz
Network Adapter	Intel® Ethernet Network Adapter E810-CQDA2	Intel® Ethernet Network Adapter E810-CQDA2
Storage (Boot Media)	Required – 512 GB	Required – 512 GB
LOM	Yes	No
Intel® QAT	Yes	Optional
Additional Plug-in cards	No	No
BIOS Configuration		
Intel® HT Technology enabled	Yes	Yes
Intel® VT-x enabled	Yes	Yes
Intel® VT-d enabled	Yes	Yes
BIOS Profile	Max Performance	Deterministic
Virtualization enabled	Yes	Yes

3.6 Platform BIOS

This section provides BIOS Configuration Profiles for each of the BMRA Configuration Profiles. For details on how the BIOS configuration should be set per each Configuration Profile, see the tables in [Section 3.5](#).

For more information about BIOS settings, visit [Intel BIOS Setup Utility User Guide](#).

Table 31. Platform BIOS Settings for 2nd Generation Intel® Xeon® Scalable Processor

MENU (ADVANCED)	PATH TO BIOS SETTING	BIOS SETTINGS	ENERGY BALANCE	MAX PERFORMANCE	DETERMINISTIC
Advanced	Processor Configuration	Intel® Hyper-Threading Tech	Enabled	Enabled	Enabled
		Intel® Virtualization Technology	Enabled	Enabled	Enabled
	Integrated IO Configuration	Intel® VT for Directed I/O	Enabled	Enabled	Enabled
Advanced/Power Configuration	Power and Performance	CPU Power and Performance Policy	Balanced Performance	Performance	Performance
		Workload Configuration	I/O sensitive	I/O sensitive	I/O sensitive
	CPU P-state control	Enhanced Intel SpeedStep® Technology	Enabled	Enabled	Disabled* [Read footnote]
		Activate PBF	Disabled	Enabled	Enabled
		Configure PBF	Disabled	Disabled	Disabled
		Intel® Turbo Boost Technology	Enabled	Enabled	Disabled* [Read footnote]
		Energy Efficient Turbo	Enabled	Disabled	N/A
		Intel Configurable TDP	Disabled	Disabled	Disabled
	Hardware P-states	Hardware P-states	Native Mode with no legacy Support	Disabled** [Read footnote]	Disabled** [Read footnote]
		EPP Enable	Enabled	Enabled	Enabled
		RAPL Prioritization	Disabled	Enabled	Enabled
	CPU C-state Control	Package C-state	C6 Retention	C6 Retention	C0/C1 State
		C1E	Enabled	Enabled	Disabled
		Processor C6	Enabled	Enabled	Disabled
	Uncore Power Management	Uncore Frequency scaling	Enabled	Disabled	Disabled
		Performance P-limit	Enabled	Disabled	Disabled
Advanced	Memory Configuration	IMC Interleaving	2-way interleave	2-way Interleave	2-way Interleave
	System Acoustic and Performance Configuration	Set Fan Profile	Acoustic	Performance	Performance
GPU	GPU Fz	Lock 900Mhz	Optional	Optional	Optional

* Enabled in the case where Intel® SST-BF is enabled to allow for configuration of individual core speeds.

** "Native Mode with No Legacy Support" where Intel® SST-BF need to be enabled

Table 32. Platform BIOS Settings for 3rd Generation Intel® Xeon® Scalable Processor

MENU (ADVANCED)	PATH TO BIOS SETTING	BIOS SETTING	ENERGY BALANCE	MAX PERFORMANCE WITH TURBO	DETERMINISTIC
Socket Configuration	Processor Configuration	Hyper-Threading	Enable	Enable	Enable
		XAPIC	Enable	Enable	Enable
		VMX	Enable	Enable	Enable
		Uncore frequency scaling	Enable	Enable	Disable
		Uncore frequency	800-2400	1.8MHz (hex 0x12)	2400
Power Configuration	Power and Performance	CPU Power and Performance Policy	Balance Performance	Performance	Performance
		Workload Configuration	I/O sensitive	I/O sensitive	I/O sensitive
	CPU P-state Control	EIST PSD Function	HW_ALL	HW_ALL	HW_ALL
		Boot Performance Mode	Max. Performance	Max. Performance	Max. Performance
		AVX License Pre-Grant	Disable	Disable	Disable
		AVX ICCP Pre Grant Level	NA	NA	NA
		AVX P1	Nominal	Nominal	Nominal
		Energy Efficient Turbo	Enable	Enable	Disable
		WFR Uncore GV rate Reduction	Enable	Enable	Enable
		GPSS timer	500us	0us	0us
		Intel Turbo Boost Technology	Enable	Enable	Disable
		Intel SpeedStep® Technology (P-states)	Enable	Enable	Disable
	Frequency Prioritization	RAPL Prioritization	Enable	Disable	Disable
	Hardware PM State Control	Hardware P-states	Native Mode with no legacy Support	Native Mode with no legacy Support	Disable
		EPP enable	Enable	Disable	Disable
	CPU C-state Control	Enable Monitor Mwait	Enable	Enable	Enable
		CPU C1 Auto Demotion	Enable	Disable	Disable
		CPU C1 Auto unDemotion	Enable	Disable	Disable
		CPU C6 Report	Enable	Enable	Disable
		Processor C6	Enable	Enable	Disable
		Enhanced Halt State (C1E)	Enable	Enable	Disable
		OS ACPI Cx	ACPI C2	ACPI C2	ACPI C2
	Energy Performance Bias	Power Performance Tuning	OS Controls EPB	OS Controls EPB	OS Controls EPB

		ENERGY_PERF_BIAS_CFG mode	Performance	Performance	Performance
		Workload Configuration	I/O Sensitive	I/O Sensitive	I/O Sensitive
	Package C-state Control	Package C-state	C6 Retention	C0/C1 State	C0/C1 State
		Dynamic L1	Enable	Disable	Disable
		Package C-state Latency Negotiation	Disable	Disable	Disable
		PKG_CSA_PS_CRITERIA	Disable	Disable	Disable
	Memory Configuration		Memory Configuration	2-way interleave	2-way interleave
Enforce POR			Enable	Enable	Enable
Platform Configuration	Miscellaneous Configuration	Serial Debug Message Level	Minimum	Minimum	Minimum
	PCI Express* Configuration	PCIe* ASPM Support	Per Port	Per Port	Per Port
	PCI Express* Configuration	PCIe* ASPM	Enable	Disable	Disable
	PCI Express* Configuration	ECRC generation and checking	Enable	Enable	Enable
Server Management		Resume on AC Power Loss	Power On	Power On	Power On
System Acoustic and Performance Configuration		Set Fan Profile	Acoustic	Performance	Performance

Table 33. Platform BIOS Settings for 4th Generation Intel® Xeon® Scalable Processor

MENU (ADVANCED)	PATH TO BIOS SETTING	BIOS SETTING	LOW LATENCY	MAX PERFORMANCE WITH TURBO	ENERGY BALANCE TURBO
Socket Configuration	Processor Configuration	Hyper-Threading	Enable	Enable	Enable
		X2APIC	Enable	Enable	Enable
		VMX	Enable	Enable	Enable
		Homeless Prefetch	Enable	Disable (default)	Disable (default)
		LLC Prefetch	Enable	Enable	Enable
		SNC	Disable	Disable	Disable
		Uncore RAPL	Disable	Disable	Enable
		Uncore frequency scaling	Disable	Disable	Enable
		Uncore frequency	1.8GHz (hex 0x12)	1.6MHz (hex 0x10)	800MHz to 2.5GHz
Power Configuration	CPU P-state Control	EIST PSD Function	HW_ALL	HW_ALL	HW_ALL
		Boot Performance Mode	Max. Performance	Max. Performance	Max. Performance
		AVX License Pre-Grant	Enable	Disable	Disable
		AVX ICCP Pre Grant Level	Level 5	NA	NA
		AVX P1 (ConfigTDP)	Level 2	Nominal (default)	Nominal
		Energy Efficient Turbo	Disable	Disable	Enable
		GPSS timer	0us	0us	0us

		Turbo	Enable	Enable	Enable
		Intel® SpeedSte p® Technology	Enable	Enable	Enable
	Frequency Prioritization	RAPL Prioritization	Disable	Disable	Disable
	Common Ref Code	UMA-Based Clustering	Quadrant	Quadrant	Quadrant
	Hardware PM State Control	Hardware P- states	Native with no Legacy Support	Native with no Legacy Support	Native with no Legacy Support
		EPP enable	Disable	Disable	Disable
	CPU C-state Control	Enable Monitor Mwait	Enable	Enable	Enable
		CPU C1 Auto Demotion	Disable	Disable	Disable
		CPU C1 Auto unDemoti on	Disable	Disable	Disable
		Processor C6 or CPU C6 Report	Enable	Enable	Enable
		Enhanced Halt State (C1E)	Enable (per Core Level)	Enable	Enable
		OS ACPI Cx	ACPI C2	ACPI C2	ACPI C2
	Energy Performance Bias	Power Performance Tuning	OS Control EPB	OS Controls EPB	OS Controls EPB
		Workload Configuration	I/O Sensitive	I/O Sensitive	Balanced
	Package C-state Control	Package C- state	C6 Retention	C0/C1 State	C0/C1 State
		Dynamic L1	Enable	Disable	Disable
Memory Configuration		Memory Configuration	8-way interleave	8-way interleave	8-way interleave
		Enforce POR / Memory Patrol Scrub	Enable/Disa ble	Enable/Enable	Enable/Enable
		Memory DIMM Refresh Rate	1x	1x	2x
Platform Configuration	Miscellaneous Configuration	Serial Debug Message Level	Minimum	Minimum	Minimum
	PCI Express* Configuration	PCIe* ASPM	Disable	Enable	Enable
		ECRC generation and checking	Disable	Enable	Enable
Server Management		Resume on AC Power Loss	Power On	Power On	Power On
System Acoustic and Performance Configuration		Set Fan Profile	Performance	Acoustic	Acoustic

Table 34. Platform BIOS Settings for Intel® Xeon® D Processor

MENU (ADVANCED)	PATH TO BIOS SETTING	BIOS SETTINGS	ENERGY BALANCE	MAX PERFORMANCE	DETERMINISTIC
Power Configuration	Power and Performance	CPU Power and Performance Policy	Balanced Performance	Performance	Performance
		Workload Configuration	I/O sensitive	I/O sensitive	I/O sensitive
		Turbo	Disabled	Enabled	Disabled
	CPU P-state control	Enhanced Intel SpeedStep® Technology	Enabled	Enabled	Disabled
		GPSS timer	500 µs	0 µs	0 µs
	Hardware P- states	Hardware P- states	Native Mode with no legacy Support	Disabled]	Disabled
	CPU C-state Control	Package C-state	C6 Retention	C6 Retention	C0/C1 State
		C1E	Enabled	Enabled	Disabled
		Processor C6	Enabled	Enabled	Disabled
	Uncore Power Management	Uncore Frequency scaling	Enabled	Disabled	Disabled
		Performance P- limit	Enabled	Disabled	Disabled
Memory Configuration	Memory Configuration	IMC Interleaving	2-way interleave	2-way interleave	2-way interleave
Thermal Configuration	System Acoustic and Performance Configuration	Set Fan Profile	Acoustic	Performance	Performance
GPU	GPU Fz	Lock 900 MHz	Optional	Optional	Optional

Use the following table to configure the BIOS settings to use Intel SST-BF, Intel SST-TF, and Intel SST-PP in 3rd Generation and 4th Generation Intel Xeon Scalable processor systems.

Table 35. BIOS Settings to Enable Intel SST-BF, Intel SST-TF, and Intel SST-PP

BIOS SETTING	STATUS
Hardware PM State Control	
Scalability	Disable
Hardware PM Interrupt	Disable
CPU P-state	
Dynamic SST-PP	Enable
Speed Step (P-states)	Enable
Activate SST-BF	Enable
Configure SST-BF	Enable
EIST PSD Function	HW_All
Turbo	Enable
Energy Efficient Turbo	Enable
Boot Performance	Max
Freq: Prioritization AC	

BIOS SETTING	STATUS
SST-CP	Enable

In BIOS, the configuration paths might be slightly different, depending on platform, but the key settings are as follows and must be performed in order.

Table 36. BIOS Settings to Enable Intel SGX on 2nd Generation and 3rd Generation Intel Xeon Scalable Processors

BIOS SETTING	STATUS
Socket Configuration > Processor Configuration > Total Memory Encryption (TME)	Enable
Socket Configuration > Common RefCode Configuration > UMA-Based Clustering	Disable (All2All)
Socket Configuration > Processor Configuration > SW Guard Extensions (SGX)	Enable
Socket Configuration > Processor Configuration > Enable/Disable SGX Auto MP Registration Agent	Enable

Table 37. BIOS Settings to Enable Intel SGX on 4th Generation Intel Xeon Scalable Processor

BIOS SETTING	STATUS
Advanced > Processor Configuration > Total Memory Encryption (TME)	Enable
Advanced > Memory Configuration > Memory RAS and Performance Configuration > UMA-Based Clustering	Disable (All2All)
Advanced > Processor Configuration > SW Guard Extensions (SGX)	Enable
Advanced > Processor Configuration > Enable/Disable SGX Auto MP Registration Agent	Enable

4 Reference Architecture Software Components

4.1 Software Components Supported

Table 38 lists the software components automatically deployed per Configuration Profile in a BMRA and their sources.

Table 38. Software Components





SOFTWARE FUNCTION	SOFTWARE COMPONENT	LOCATION
	Ubuntu 20.04.4 Kernel version: 5.4.0-100-generic	https://www.ubuntu.com
	Ubuntu 21.10 Kernel version: 5.13.0-19-generic	
	Ubuntu 22.04 Kernel version: 5.15.0-25-generic	
	Ubuntu 22.04 RT Kernel version: 5.15.0-25-realtime	
	RHEL 8.5 Kernel version: 4.18.0-348.el8.x86_64	https://www.redhat.com/
	Rocky 8.5 Kernel version: 4.18.0-348.el8.x86_64	https://rockylinux.org/
Data Plane Development Kit	DPDK 22.03	https://core.dpdk.org/download/
Open vSwitch with DPDK	OVS-DPDK v2.17.1	https://github.com/openvswitch/ovs
Vector Packet Processing	VPP 21.10	https://docs.fd.io/vpp/
Telegraf	1.1	https://github.com/intel/observability-telegraf
collectd	a7cea43d9d2f67c38fbf0407786edbe660ee9072945f7bb272b55fd255e8eaca	https://www.collectd.org/
Grafana	8.5.3	https://www.grafana.com/
Prometheus	2.35.0	quay.io/prometheus/prometheus
Prometheus nginx image	1.21.6-alpine	docker.io/library/nginx:1.21.6-alpine
Ansible	4.10.0	https://www.ansible.com/
BMRA Ansible Playbook	v22.05	https://github.com/intel/container-experience-kits
Python	Python 3.6.x for RHEL 8 Python 3.8.x for Ubuntu 20.04 and Python 3.9.x for Ubuntu 21.10	https://www.python.org/
Kubespray	2.18	https://github.com/kubernetes-sigs/kubespray
Docker	20.10	https://www.docker.com/
containerd	1.5.11	Dependency of other software - not downloaded independently
CRI-O	1.22	Dependency of other software - not downloaded independently
Container orchestration engine	Kubernetes v1.23.4	https://github.com/kubernetes/kubernetes
	Kubernetes v1.22.3	
	Kubernetes v1.21.5	
CPU Manager (native to Kubernetes)	Available natively in Kubernetes	N/A

User Guide | Network and Cloud Edge Container Bare Metal Reference System Architecture

SOFTWARE FUNCTION	SOFTWARE COMPONENT	LOCATION
Platform Aware Scheduling (TAS)	TAS 0.2.0	https://github.com/intel/platform-aware-scheduling
Platform Aware Scheduling (GAS)	GAS 0.3.0	https://github.com/intel/platform-aware-scheduling
k8s-prometheus-adapter	0.8.4	Dependency of other software - not downloaded independently
K8s node-exporter	1.3.1	Dependency of other software - not downloaded independently
K8s prometheus-operator	0.50.0	https://github.com/prometheus-operator/kube-prometheus
K8s kube-rbac-proxy	0.11.0	Dependency of other software - not downloaded independently
Node Feature Discovery	0.11.0	https://github.com/kubernetes-sigs/node-feature-discovery
Multus CNI	3.8	https://github.com/intel/multus-cni
SR-IOV CNI	2.6.2	https://github.com/intel/sriov-cni
SR-IOV network device plugin	3.4.0	https://github.com/intel/sriov-network-device-plugin
SR-IOV Network Operator	3f1b2e2d0792e96f2c6031a8bc8e42a4a093e573844f15f6f6f2c086646b19ba	https://github.com/k8snetworkplumbingwg/sriov-network-operator
Whereabouts Service	05cc22a9c8165c5cba875bebfa58d1b504a2e6c9	https://github.com/k8snetworkplumbingwg/helm-charts.git
Device Plugins Operator	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
QAT device plugin	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
GPU device plugin	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
Intel® SGX device plugin	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
Intel DLB device plugin	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
Intel DSA device plugin	0.23	https://github.com/intel/intel-device-plugins-for-kubernetes
Userspace CNI	1.3	https://github.com/intel/userspace-cni-network-plugin
Bond CNI plugin	1.0	https://github.com/intel/bond-cni
Intel® Ethernet Drivers	i40e v2.19.3 ice v1.8.8 iavf v4.4.2.1	https://sourceforge.net/projects/e1000/files/i40e%20stable/2.19.3/ https://sourceforge.net/projects/e1000/files/ice%20stable/1.8.8/ https://sourceforge.net/projects/e1000/files/iavf%20stable/4.4.2.1/
Intel® Ethernet NVM Update Package 700 Series	8.70	https://www.intel.com/content/www/us/en/download/18190/non-volatile-memory-nvm-update-utility-for-intel-ethernet-network-adapter-700-series.html
Intel® Ethernet NVM Update Package 800 Series	3.20	https://www.intel.com/content/www/us/en/download/19626/non-volatile-memory-nvm-update-utility-for-intel-ethernet-network-adapters-e810-series-linux.html
DDP Profiles	Dynamic Device Personalization for Intel® Ethernet 700 Series Version 25.4	https://downloadmirror.intel.com/28940/eng/mpslogreudp.zip https://downloadmirror.intel.com/28040/eng/ppp-oe-ol2tpv2.zip https://downloadmirror.intel.com/29446/eng/esp-ah.zip https://downloadmirror.intel.com/29780/eng/ecpri.zip
	Intel® Ethernet 800 Series DDP specific package for Comms 1.3.35.0	https://www.intel.com/content/www/us/en/download/19660/intel-ethernet-800-series-telecommunication-comms-dynamic-device-personalization-ddp-package.html
Intel® Ethernet Operator	22.04	https://github.com/intel/intel-ethernet-operator.git
Intel® Ethernet Operator SDK	1.18.1	https://github.com/operator-framework/operator-sdk.git
Intel® Ethernet UFT	22.03	https://github.com/intel/UFT.git
Intel® QAT Drivers	QAT.L.4.18.0-0008	Dependency of other software - not downloaded independently
Intel® QAT Drivers (NDA)	QAT20.L.0.8.0-00071	Dependency of other software - not downloaded independently
OpenSSL	openssl-3.0.3	https://github.com/openssl/openssl https://www.openssl.org/source/

SOFTWARE FUNCTION	SOFTWARE COMPONENT	LOCATION
OpenSSL (NDA)	openssl-1.1.1j	Dependency of other software - not downloaded independently
OpenSSL QAT Engine	0.6.12	https://github.com/intel/QAT_Engine
OpenSSL QAT Engine (NDA)	0.6.10_INT	Dependency of other software - not downloaded independently
Intel ipsec-mb	1.2	https://github.com/intel/intel-ipsec-mb
Intel® SGX DCAP Drivers	1.41	https://download.01.org/intel-sgx/sgx-dcap/1.10.3/linux/
Intel® SGX SDK	2.16.100.4	https://download.01.org/intel-sgx/sgx-dcap/1.10.3/linux/
Intel® KMRA	2.1	https://01.org/key-management-reference-application-kmra
Intel® KMRA AppHSM	2.1	https://hub.docker.com/r/intel/apphsm
Intel® KMRA CTK	2.1	https://hub.docker.com/r/intel/ctk_loadkey
Intel® KMRA PCCS	2.1	https://hub.docker.com/r/intel/pccs
Istio operator	1.13.1	https://github.com/istio/istio/releases/download/
Istio operator (NDA)	1.10.1	https://github.com/intel-innersource/applications.services.cloud.istio.deployment
istio-intel/pilot-cryptomb (NDA)	21.06.2	https://github.com/intel-innersource/applications.services.cloud.istio.deployment
istio-intel/proxyv2-cryptomb (NDA)	21.06.2	https://github.com/intel-innersource/applications.services.cloud.istio.deployment
istio-intel/proxyv2-openssl-qat (NDA)	21.09	https://github.com/intel-innersource/applications.services.cloud.istio.deployment
istio-intel/proxyv2-openssl (NDA)	1.9.4	https://github.com/intel-innersource/applications.services.cloud.istio.deployment
istio-intel/trusted-certificate-issuer	0.1.0	https://github.com/intel/trusted-certificate-issuer
istio-intel/trusted-attestation-controller	1.0	https://github.com/intel/trusted-attestation-controller
istio-intel/tcpip-bypass-ebpf	1.0	https://hub.docker.com/r/intel/istio-tcpip-bypass
CNDP DP	0.0.1	https://github.com/intel/afxdp-plugins-for-kubernetes.git
CNDP CNI	22.04.0	https://github.com/CloudNativeDataPlane/cndp.git http://www.github.com/otcs-hare/CNDP
MinIO operator	4.4.2	https://github.com/minio/operator
MinIO console	0.13.2	https://github.com/minio/operator
Power Manager Operator	1.0.2	https://hub.docker.com/r/intel/power-operator
Power Node Agent Operator	1.0.2	https://hub.docker.com/r/intel/power-node-agent
Intel® RDT	4.3.0	https://github.com/intel/intel-cmt-cat
FEC Operator	2.2.0	https://github.com/smart-edge-open/sriov-fec-operator
FEC Operator SDK	1.20.0	https://github.com/operator-framework/operator-sdk.git
Operator Package Manager	1.22.0	https://github.com/operator-framework/operator-registry/releases/

4.2 Software Components Compatibility Matrices

Legend for the tables in this section	
	Indicates that the combination is unsupported
	Indicates that the combination is supported and tested
	Indicates that the combination is expected to work but untested
	Indicates that the combination is not applicable

Note: Features that are not listed have been verified to work for all combinations

Feature / Platform Compatibility Limitations				
	2nd Generation Intel® Xeon® Scalable Processor	3rd Generation Intel® Xeon® Scalable Processor	4th Generation Intel® Xeon® Scalable Processor	Intel® Xeon® D Processor
dsa			NDA only	
dlb			NDA only	
sst-bf				
sst-cp				
sst-tf				
sst-pp				
sst-operator				
minio / nvme				
sgx				
kmra				
gpu				
vm				
qat off-chip				
qat on-chip			NDA only	
siov			NDA only	

Profile / Platform Validation Matrix				
	2nd Generation Intel® Xeon® Scalable Processor	3rd Generation Intel® Xeon® Scalable Processor	4th Generation Intel® Xeon® Scalable Processor	Intel® Xeon® D Processor
Access				
Basic				
Full				
On-Premises Edge				
Remote Central Office-Forwarding				
Regional Data Center				
Build-Your-Own				
Storage				

Feature / OS Compatibility Limitations					
	Ubuntu 20.04.4 5.4.0-100- generic	Ubuntu 21.10 5.13.0-19- generic	Ubuntu 22.04 5.15.0-25- generic	RHEL 8.5 4.18.0- 348.el8.x86_64	Rocky Linux 8.5 4.18.0- 348.el8.x86_64
dsa		NDA only	NDA only		
dlb		NDA only	NDA only		
cndp					

Feature / Feature Compatibility Limitations											
	DPDK 21.11	DPDK 20.08	OVS 2.16.2	VPP	collectd	Telegraf	SST-BF	SST-CP	SST-TF	SST-PP	Power Manager
DPDK 21.11											
DPDK 20.08											
OVS 2.16.2											
VPP											
collectd											
Telegraf											
SST-BF											
SST-CP											
SST-TF											
SST-PP											
Power Manager											

5 Post Deployment Verification Guidelines

This section describes a set of processes that you can use to verify the components deployed by the scripts. The processes are not Configuration Profile-specific but relate to individual components that may not be available in all profiles. Details for each of the Configuration Profiles are described in Sections 7 through 14.

Many verification guidelines and output examples can be found on GitHub, as listed in [Table 39](#), and others are described after the table.

Table 39. Links to Verification Guidelines on GitHub

VERIFICATION STEP
Check the Kubernetes Cluster
Check Intel SST-BF Configuration on 2nd Generation Intel Xeon Scalable Processor
Check Intel SST-BF and Intel SST-CP on 3rd Generation Intel Xeon Scalable Processor
Check Intel SST-PP with Intel SST-TF on 3rd and 4th Generation Intel Xeon Scalable Processors
Check DDP Profiles on Intel® Ethernet 700 and 800 Series Network Adapters
Check Node Feature Discovery
Check Topology Manager
Check SR-IOV Network Operator
Check SR-IOV Device Plugin
Check QAT Device Plugin
Check SGX Device Plugin
Check DSA Device Plugin
Check GPU Device Plugin
Check Multus CNI Plugin
Check SR-IOV CNI Plugin
Check Userspace CNI Plugin
Check Bond CNI Plugin
Check Telemetry Aware Scheduling
Check Intel® Server GPU Device and Driver
Check Intel QAT Engine with OpenSSL
Check MinIO Operator/Console and Tenant
Check Intel Power Manager (Balance Performance Power-Profile & Sample Power-Pods)

5.1 Check Grafana Telemetry Visualization

BMRA deploys Grafana for telemetry visualization. It is available on every cluster node on port 30000. Due to security reasons, this port is not exposed outside the cluster by default. Default credentials are `admin/admin` and you should change the default password after first login.

The Grafana TLS certificate is signed by the cluster certificate authority (CA) and it is available in `/etc/kubernetes/ssl/ca.crt`

Visit Grafana at `https://<node-ip>:30000/`

BMRA comes with a set of dashboards from the kube-prometheus project ([kube-prometheus](#)). Dashboards are available in the Dashboards > Manage menu.

5.2 Check Key Management Infrastructure with Intel SGX

To verify the Key Management infrastructure with SGX and use the private keys provisioned to Intel SGX enclaves, see [Section 15.1](#) for step-by-step instructions to set up and run the NGINX workload.

Part 2: Building a BMRA Step-by-Step

6 BMRA Setup – Applicable for All Configuration Profiles

This section is relevant for generating BMRA Flavors based on their Configuration Profiles. It provides the prerequisites for system setup and includes information that enables you to review BIOS prerequisites and software BOMs at a glance. The information is presented in multi-column tables to provide an easy way to compare and assess the differences between the BMRA Flavors that are available.

After setting up the Kubernetes system, refer to the specific section from the following list to build the BMRA Flavors:

- [Section 7, BMRA Basic Configuration Profile Setup](#)
- [Section 8, BMRA Full Configuration Profile Setup](#)
- [Section 9, BMRA On-Premises Edge Configuration Profile Setup](#)
- [Section 10, BMRA Remote Central Office-Forwarding Configuration Profile Setup](#)
- [Section 11, BMRA Regional Data Center Configuration Profile Setup](#)
- [Section 12, BMRA for Storage Configuration Profile Setup](#)
- [Section 13, BMRA Access Edge Configuration Profile Setup](#)
- [Section 14, BMRA Build-Your-Own Configuration Profile Setup](#)

6.1 Set Up an Ansible Host

BMRA Kubernetes clusters require an Ansible host that stores information about all managed remote nodes. In general, any machine running a recent Linux distribution can be used as Ansible host for any of the supported BMRA deployments (regardless of target OS on the control and worker nodes), as long as it meets the following basic requirements:

- Network connectivity to the control and worker nodes, including SSH
- Internet connection (using proxy if necessary)
- Git utility installed
- Python 3 installed
- Ansible version 4.10.0 installed (Ansible-base at 2.11.3)

Step-by-step instructions for building the Ansible host are provided below for the same list of operating systems that are supported for the control and worker nodes (see [Section 2.3.3](#)).

6.1.1 RHEL Version 8 as Ansible Host

1. Install the Linux OS. If using the iso image, choose the Minimal iso version, or select the "Minimal Install" (Basic functionality) option under Software Selection.
2. Make the proper configuration during installation for the following key elements: Network (Ethernet) port IP Address, Host Name, Proxies (if necessary), and Network Time Protocol (NTP).
3. After the installation completes and the machine reboots, log in as root and confirm that it has a valid IP address and can connect (ping) to the control and worker nodes.
4. Make sure that the HTTP and HTTPS proxies are set, if necessary, for internet access. The configuration can be completed with the `export` command or by including the following lines in the `/etc/environment` file:

```
http_proxy=http://proxy.example.com:1080
https_proxy=http://proxy.example.com:1080
```

Then, load the proxies configuration in the current environment:

```
# source /etc/environment
```

5. Install Git:


```
# yum install -y git
```
6. Install Python 3:


```
# yum -y install python3
```
7. Install Ansible:


```
# pip install ansible-base==2.11.3
```

The Ansible host box is now ready to deploy the Container BMRA. Follow the instructions in [Section 2.5](#).

6.1.2 Ubuntu 20.04 LTS as Ansible Host

1. Install the OS using any method supported by the vendor (Canonical Ltd.). Either the Desktop or Server distribution can be used. Select the "Minimal installation" option under "Updates and Other software".
2. Follow steps 2, 3, and 4 as described above for RHEL.
3. Update the installation:


```
# sudo apt update
```
4. Install SSH utilities:


```
# sudo apt install openssh-server
```

5. Install Git:

```
# sudo apt install -y git
```

6. Install Python 3-pip:

```
# sudo apt install -y python3-pip
```

7. Install Ansible:

```
# sudo pip install ansible-base==2.11.3
```

The Ansible host box is now ready to deploy the Container BMRA. Follow the instructions in [Section 2.5](#).

6.2 Set Up the Control and Worker Nodes - BIOS Prerequisites

This section is applicable for all **Configuration Profiles**.

Enter the UEFI or BIOS menu and update the configuration as shown in [Table 40](#) and [Table 41](#).

Note: The method for accessing the UEFI or BIOS menu is vendor-specific, for example: [How to boot into the BIOS or the Lifecycle Controller on your PowerEdge Server](#)

The BIOS profile referenced in these tables consists of configurations in the power management, thermal management, and configuration for Intel® platform technologies such as Intel® Virtualization Technology, Intel® Hyper-Threading Technology, Intel SpeedStep® technology, and Intel® Turbo Boost Technology.

The table provides four different BIOS profiles.

- Energy Balance
- Max Performance
- Deterministic
- Low Latency (4th Generation Intel® Xeon® Scalable processor)

The configuration and values set per each BIOS profile are defined in [Table 31](#), [Table 32](#), and [Table 33](#).

Table 40. BIOS Prerequisites for Control and Worker Nodes for Basic, Full, Storage, and Build-Your-Own Configuration Profiles

PROFILES	BASIC CONFIGURATION PROFILE	FULL CONFIGURATION PROFILE	STORAGE CONFIGURATION PROFILE	BUILD-YOUR-OWN CONFIGURATION PROFILE
Configuration				
BIOS Profile	Energy Balance	Max Performance	Max Performance	Any
Grub Command Line (values are set by Ansible)				
Isolcpus	Optional	Yes	No	Optional
Hugepages	Optional	Yes	No	Optional
P-state=disable	Optional	Yes, No-SST-BF	No	Optional
Limit C-state	Optional	Yes	No	Optional

Table 41. BIOS Prerequisites for Control and Worker Nodes for On-Premises Edge, Remote Central Office-Forwarding, Regional Data Center, and Access Edge Configuration Profiles

PROFILES	ON-PREMISES EDGE CONFIGURATION PROFILE	REMOTE CENTRAL OFFICE-FORWARDING CONFIGURATION PROFILE	REGIONAL DATA CENTER CONFIGURATION PROFILE	ACCESS EDGE CONFIGURATION PROFILE
Configuration				
BIOS Profile	Max Performance	Deterministic / Energy Balance	Max Performance	Low Latency
Grub Command Line (values are set by Ansible)				
Isolcpus	Yes	Yes	Optional	Yes
Hugepages	Yes	Yes	Optional	Yes
P-state=disable	No	Yes, No-SST-BF	Optional	No
Limit C-state	No	Yes	Optional	Yes

Note: The above values are the recommended configuration options on the Intel S2600WFQ and Intel M50CYP server boards. Some server boards may not provide the same options that are documented in this table. Vendors typically provide options for max performance configuration with virtualization.

6.3 Configuration Dictionary - Group Variables

Table 42 lists the parameters available as group variables with their type (for example, Boolean, string, URL, list, integer), possible values, and descriptions. The variables in **bold** must be updated to match the target environment. The variables with blue highlight must be updated according to your infrastructure. Refer to the section that describes your Configuration Profile to see the parameters enabled for that Configuration Profile.

Table 42. Configuration Dictionary – Group Variables

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
Deployment Configuration				
	profile_name	String	""	See chapters on supported profiles
	configured_arch	String	""	See chapters on supported hardware
	preflight_enabled	Boolean	true/false	Verify user configuration is accurate and compatible across feature sets
Common Cluster Configuration				
Kubernetes		Boolean	true/false	Specifies whether to deploy Kubernetes
	kube_version	String	v1.23.4	Kubernetes version
	container_runtime	String	docker, cri-o, containerd	Container runtime to use as base engine for cluster deployment
	container_runtime_only_deployment	Boolean	true/false	Deploy only the container runtime without Kubernetes
	docker_version	String	20.10	Docker version
	containerd_version	String	1.5.11	Containerd version
	crio_version	String	1.22	CRI-O version
	update_all_packages	Boolean	false	Runs system-wide package update (apt dist-upgrade, yum update, ...). Tip: Can be set using host_vars for more granular control.
	http_proxy	URL	http://proxy.example.com:1080	HTTP proxy address. Comment out if your cluster is not behind proxy.
	https_proxy	URL	http://proxy.example.com:1080	HTTPS proxy address. Comment out if your cluster is not behind proxy.
	additional_no_proxy	Comma-separated list of addresses	.example.com	Additional URLs that are not behind proxy, for example your corporate intra network DNS domain, e.g., ".intel.com". Note: Kubernetes nodes addresses, pod network, and the like are added to no_proxy automatically.
	kube_network_plugin_multus	Boolean	True	Specifies whether to use the network plugin Multus
	multus_version	String	V3.8	Multus version
	kube_network_plugin	String	calico/flannel	Specifies networking CNI to use
	kube_pods_subnet	CIDR	10.244.0.0/16	Kubernetes pod subnet. Make sure that it matches your CNI plugin requirements (Calico by default) and doesn't overlap with your corporate LAN.
	kube_service_addresses	CIDR	10.233.0.0/18	Kubernetes service subnet. Make sure that it matches your CNI plugin requirements (Calico by default) and doesn't overlap with your corporate LAN.
	kube_proxy_mode	String	iptables	Instructs kube_proxy how to set up NAT and load balancing functions
	kube_proxy_nodeport_addresses_cidr	CIDR	127.0.0.0/8	Kubernetes service subnet
	cluster_name	DNS domain	cluster.local	Name of the cluster
	registry_local_addresses	String	"localhost:30500"	Container registry address IP and port

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
	psp_enabled	Boolean	true/false	Enable pod security policy admission controller and create minimal set of rules
	always_pull_enabled	Boolean	true/false	Set image pull policy to Always. Pulls images before starting containers. Valid credentials must be configured.
Node Feature Discovery				
	nfd_enabled	Boolean	true/false	Specifies whether to deploy Node Feature Discovery
	nfd_version	String	0.11.0	NFD version
	nfd_build_image_locally	Boolean	false	Builds NFD image locally instead of using the one from public registry.
	nfd_namespace	String	kube-system	Kubernetes namespace used for NFD deployment
	nfd_sleep_interval	String	60s	Defines how often NFD queries node status and update node labels
Native Built-in Kubernetes CPU Manager				
	native_cpu_manager_enabled	Boolean	true/false	Kubernetes CPU manager controls CPU management policies on the nodes. Setting this option as "true" enables the "static" policy; otherwise the default "none" policy is used.
	native_cpu_manager_system_reserved_cpus	Kubernetes millicores	2000m	Number of CPU cores to be reserved for housekeeping (2000m = 2000 millicores = 2 cores)
	native_cpu_manager_kube_reserved_cpus	Kubernetes millicores	1000m	Number of CPU cores to be reserved for Kubelet
	native_cpu_manager_reserved_cpus	Comma-separated list of integers or integer ranges	0,1,2	Explicit list of the CPUs reserved from pods scheduling. Note: Supported only with kube_version 1.17 and newer, overrides two previous options.
Topology Manager (Kubernetes Built-in)⁴				
	topology_manager_enabled	Boolean	true/false	Enables Kubernetes built-in Topology Manager
	topology_manager_policy	String, options: none, best-effort, restricted, single-numa-node	best-effort	Topology Manager policy
Intel SR-IOV Network Device Plugin				
	sriov_network_operator_enabled	Boolean	true/false	Enables SR-IOV Network Operator
	sriov_network_operator_namespace	String	sriov-network-operator	Kubernetes namespace used to deploy SR-IOV network operator
	sriov_net_dp_enabled	Boolean	true/false	Enables SR-IOV network device plugin
	sriov_net_dp_namespace	String	kube-system	Kubernetes namespace used to deploy SR-IOV network device plugin
	sriov_net_dp_build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
	sriovdp_config_data	Multi-line string in JSON format	Two resource pools for kernel stack and DPDK-based networking respectively	SR-IOV network device plugin configuration. For more information on supported configurations, refer to Configurations

Intel Device Plugins for Kubernetes

⁴ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
Intel_dp_namespace		String	kube-system	Kubernetes namespace used to deploy Intel device plugin operator
dsa_dp_enabled		Boolean	true/false	Enables Intel DSA device plugin
	configure_dsa_devices	Boolean	true/false	Specifies whether to configure Intel DSA devices
	dsa_devices	List	[]	Intel DSA devices for which to configure worker queues
dlb_dp_enabled				Enables Intel Dynamic Load Balancing device plugin
	configure_dlb_devices	Boolean	true/false	Enables configuration of DLB on worker node
	dlb_dp_build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
	dlb_dp_verbosity	Integer	4	
qat_dp_enabled		Boolean	true/false	Enables Intel QAT device plugin
	qat_dp_namespace	String	kube-system	Namespace used for Intel QAT device plugin
sgx_dp_enabled		Boolean	true/false	Enables Intel SGX device plugin
	sgx_dp_build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
	sgx_aesmd_namespace	String	kube-system	Kubernetes namespace used to deploy SGX device plugin
	sgx_dp_provision_limit	Integer	20	
	sgx_dp_enclave_limit	Integer	20	
gpu_dp_enabled		Boolean	true	Enables Intel GPU device plugin
	gpu_dp_namespace	String	kube-system	Namespace used for Intel GPU device plugin
	gpu_dp_shared_devices	Integer	10	Number of containers (min. 1) that can be shared across the same GPU device
	gpu_dp_monitor_resources	Boolean	true/false	Enable monitoring all GPU resources on the node
	gpu_dp_fractional_manager	Boolean	true/false	Enable handling of fractional resources for multi-GPU nodes
	gpu_dp_preferred_allocation	String	none	Policy to apply to GPU nodes: 'balanced', 'packed', 'none'
	gpu_dp_max_memory	String	8 GB	Maximum memory per card
Intel Key Management Reference Application				
kmra_enabled		Boolean	true/false	Enables Intel Key Management Reference Application
	kmra_pccs_api_key	String	"fffff....."	API key obtained from Intel's Provisioning Certificate Service
	kmra_deploy_demo_workload	Boolean	true/false	Enable to deploy a KMRA demo workload (NGINX Server)
Service Mesh				
service_mesh		Boolean	true/false	Enables Istio service mesh for Kubernetes
	profile	String	custom-ca or default	
	tcpip_bypass_ebpf	Boolean	true/false	Enable TCP/IP ebpf bypass demo
	tls_splicing	Boolean	true/false	Enable TLS splicing demo
	sgx_signer	Boolean	true/false	Enable automated key management
tcs		Boolean	true/false	Enable trusted certificate issuer
tca		Boolean	true/false	Enable trusted certificate attestation controller

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
Intel Platform Aware Scheduling				
pas_namespace		String	kube-system	Kubernetes namespace used for TAS deployment
tas_enabled		Boolean	true/false	Enables Intel Telemetry Aware Scheduling
	tas_build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
	tas_enable_demo_policy	Boolean	true/false	Creates demo TAS policy
gas_enabled		Boolean	true/false	
	gas_build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
Telemetry Configuration				
prometheus_operator		Boolean	true/false	Enable Prometheus for telemetry management
collectd_enabled		Boolean	true/false	Gather platform metrics with collectd
	collectd_scrap_interval	Integer	30	Duration to gather metrics using collectd
telegraf_enabled		Boolean	true/false	Gather platform metrics with Telegraf
	telegraf_scrap_interval	Integer	30	Duration to gather metrics using Telegraf
Example Network Attachment Definitions (Ready to Use Examples of Custom CNI Plugin Configuration)				
example_net_attach_defs		List of dictionaries	[]	Example network attached definition objects to create
	userspace_ovs_dpdk	Boolean	true/false	Example net-attach-def for userspace CNI with OVS-DPDK
	userspace_vpp	Boolean	true/false	Example net-attach-def for userspace CNI with VPP
	sriov_net_dp	Boolean	true/false	Example net-attach-def for SR-IOV Net DP and SR-IOV CNI
MinIO Configuration				
minio_enabled		Boolean	true/false	Enables MinIO operator/console
	minio_tenant_enabled	Boolean	true/false	Specifies whether to install sample MinIO tenant
	minio_tenant_servers	Integer	4	The number of MinIO tenant nodes
	minio_tenant_volumes_per_server	Integer	4	The number of volumes per server
	minio_deploy_test_mode	Boolean	true/false	<ul style="list-style-type: none"> true (Test Mode) – use a file as a loop device when creating storage called "virtual block device", which is useful for test or automation purpose false (Performance Mode) – use an actual NVME or SSD device when creating storage
CNDP Configuration				
cndp_dp_enabled		Boolean	true/false	Enable Cloud Native Data Plane device plugin
	cndp_net_attach_def_enabled	Boolean	true/false	Creates a network attach definition resource
	cndp_net_attach_def_conf	List of dicts	[]	Network configuration for CNDP
Ethernet Operator Configuration				
	intel_ethernet_operator_enabled	Boolean	true/false	Enable firmware and driver management of network resources with Ethernet Operator
	intel_ethernet_operator_flow_config_enabled	Boolean	true/false	Enable flow management of network resources with Ethernet Operator

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
Power Manager Configuration				
Intel_power_manager		List of dictionaries		Power manager settings to configure
	enabled	Boolean	true/false	Enable Intel power manager
	power_profiles	List	[performance, balance-performance, balance-power]	Power profiles to be available on the nodes
	power_nodes	List	[]	Nodes power operator should manage
	build_image_locally	Boolean	true/false	Build and store image locally or use one from public external registry
	deploy_example_pods	Boolean	true/false	Deploy example pods to use power features
	global_shared_profile_enabled	Boolean	true/false	Deploy custom power profile with user defined frequencies
	max_shared_frequency	Integer	1500	Max frequency to be applied for cores by shared workload
	min_shared_frequency	Integer	1000	Min frequency to be applied for cores by shared workload
FEC Operator Configuration				
	intel_sriov_fec_operator_enabled	Boolean	true/false	Enables operator deployment to manage forward error correction HW

6.4 Configuration Dictionary - Host Variables

[Table 43](#) lists the parameters available as host variables with their type (for example, Boolean, string, URL, list, integer), possible values, and descriptions. The variables in **bold** must be updated to match the target environment. The variables with blue highlight must be updated according to your infrastructure. Refer to the section that describes your Configuration Profile to see the parameters enabled for that Configuration Profile.

Table 43. Configuration Dictionary – Host Variables

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
SR-IOV and Network Devices Configuration				
iommu_enabled		Boolean	true/false	Sets up SR-IOV-related kernel parameters and enables further SR-IOV configuration
dataplane_interfaces		List of dictionaries	n/a	SR-IOV-related NIC configuration using per-port approach
	dataplane_interfaces[*].name	String	enp24s0f0, enp24s0f1	Name of the interface representing PF port
	dataplane_interfaces[*].bus_info	String (PCI address)	18:00.0, 18:00.1	PCI address of the PF port
	dataplane_interfaces[*].pf_driver	String	ice	PF driver, "i40e", "ice"
	dataplane_interfaces[*].sriov_numvfs	Integer	6, 4	Number of VFs to be created, associated with the PF
	dataplane_interfaces[*].minio_vf	Boolean	true/false	Set MinIO tenant pods use the SR-IOV as additional network interfaces
	dataplane_interfaces[*].default_vf_driver	String, options: "i40evf", "iavf", "vfio-pci", "igb_uio"	vfio-pci for DPDK, iavf for kernel network stack	Default driver module name that the VFs are bound to

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
	<code>dataplane_interfaces[*].sriov_vfs[*]</code>	List of dictionaries	n/a	List of VFs to create with specific driver (non-default)
	<code>dataplane_interfaces[*].ddp_profile</code>	String, optional	<code>gtp.pkggo</code>	Name of the DDP package to be loaded onto the Network Adapter. Note: Use only for the port 0 of the Network Adapter (PCI address ending with :00.0)
<code>update_nic_drivers</code>		Boolean	true/false	Set to 'true' to update Linux kernel drivers for Intel Network Adapters
<code>update_nic_firmware</code>		Boolean	true/false	Set 'true' to update Network Adapter firmware
<code>install_ddp_packages</code>		Boolean	true/false	Install Intel X700 and X800 series Network Adapters DDP packages. Required if DDP packages configured in <code>dataplane_interfaces</code> .
<code>install_dpdk</code>		Boolean	true/false	DPDK installation is required for <code>sriov_cni_enabled:true</code>
	<code>dpdk_version</code>	String	22.03	DPDK version to install
	<code>dpdk_local_patches_dir</code>	String	Empty	Path to user-supplied patches to apply against the specified version of DPDK
SR-IOV and Bond CNI Plugins				
<code>sriov_cni_enabled</code>		Boolean	true/false	Installs SR-IOV CNI plugin binary on the node
<code>bond_cni_enabled</code>		Boolean	true/false	Installs Bond CNI plugin binary on the node
Userspace Networking Plugins and Accelerated Virtual Switches				
<code>userspace_cni_enabled</code>		Boolean	true/false	Installs userspace CNI plugin binary on the node
<code>ovs_dpdk_enabled</code>		Boolean	true/false	Installs OVS-DPDK on the node
	<code>ovs_dpdk_lcore_mask</code>	Hex integer	0x1	CPU mask for OVS-DPDK PMD threads
	<code>ovs_dpdk_socket_memory</code>	Integer or comma-separated list of integers	256,0	Amount of memory per NUMA node allocated to OVS-DPDK PMD threads
<code>vpp_enabled</code>		Boolean	true/false	Installs FD.io VPP
Hugepages/Memory Configuration				
<code>hugepages_enabled</code>		Boolean	true/false	Enables hugepages support
	<code>default_hugepage_size</code>	String, options: 2M, 1G	1G	Default hugepages size
	<code>number_of_hugepages</code>	Integer	4	Sets how many hugepages should be created
CPU Configuration				
<code>isolcpus_enabled</code>		Boolean	true/false	Enables CPU cores isolation from Linux scheduler
	<code>isolcpus</code>	Comma-separated list of CPU cores/ranges	4-11	CPU cores isolated from Linux scheduler
<code>intel_pstate</code>		String	<code>hwp_only</code>	Enables Intel P-state scaling driver. Available parameters: <code>disable</code> , <code>passive</code> , <code>force</code> , <code>no_hwp</code> , <code>hwp_only</code> , <code>support_aci_pcc</code> , <code>per_cpu_perf_limites</code>
	<code>turbo_boost_enabled</code>	Boolean	true/false	Enables Turbo Boost for P-state attribute
<code>sst_pp_configuration_enabled</code>		Boolean	true/false	Enables Intel SST Performance Profiles for flexible configuration of SST-BF, SST-CP, and SST-TF

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
	sst_pp_config_list	List of dictionaries	sst_bf: enable/disable sst_cp: enable/disable sst_tf: enable/disable	Enables configuration of SST features through SST-PP
	online_cpus_range	String	auto	Specifies automatic configuration of online CPUs versus manual configuration of each SST feature
sst_bf_configuration_enabled		Boolean	true/false	Enables Intel SST Base Frequency technology. Support of SST-BF requires 'intel_pstate' to be 'enabled'
	clx_sst_bf_mode	Character, options: s, m, r	s	Configure SST-BF mode for 2nd Generation Intel® Xeon® [s] Set SST-BF config (set min/max to 2700/2700 and 2100/2100) [m] Set P1 on all cores (set min/max to 2300/2300) [r] Revert cores to min/Turbo (set min/max to 800/3900)
	icx_sst_bf_enabled	Boolean	true/false	Enables Intel SST Base Frequency technology. 3rd Generation Intel® Xeon® support of SST-BF requires 'intel_pstate' to be 'enabled'.
	icx_sst_bf_with_core_priority	Boolean	true/false	Prioritize (SST-CP) power flow to high frequency cores
sst_cp_configuration_enabled		Boolean	true/false	Enables Intel SST Core Power technology on 3rd Generation Intel® Xeon®. SST-CP overrides any 'SST-BF configuration'.
	sst_cp_priority_type	Integer	1	0 – proportional 1 - ordered
	sst_cp_clos_groups	List of dictionaries	[]	Allows for configuration of up to 4 CLOS groups including id, frequency_weight, min_MHz, max_MHz
	sst_cp_cpu_clos	List of dictionaries	[]	Allows for definition of CPU cores per close group
sst_tf_configuration_enabled		Boolean	true/false	Enables Intel SST Turbo Frequency
Miscellaneous				
dns_disable_stub_listener	dns_disable_stub_listener	Boolean	true/false	(Ubuntu only) Disables DNS stub listener from the systemd-resolved service, which is known to cause problems with DNS and Docker containers on Ubuntu
install_real_time_package	install_real_time_package	Boolean	true/false	Installs real-time Linux kernel packages.
QAT Configuration				
update_qat_drivers		Boolean	true/false	Install QAT drivers and services
qat_devices		List of dictionaries	[]	SR-IOV related QAT configuration using per-port approach
	qat_devices[*].qat_id	String (PCI address)	0000:ab:00.0, 0000:xy:00.0, 0000:yz:00.0	PCI address of the PF port
	qat_devices[*].qat_sriov_numvfs	Integer	10	Number of VFs to be created per QAT device physical function
openssl_install		Boolean	true/false	Install OpenSSL for use with QAT engine
FEC Configuration				
	fec_acc	String	0000:ab:00.0	PCI address of the FEC device
MinIO Configuration				
minio_pv		List of dictionaries	[]	PV related MinIO configuration
	minio_pv[*].name	String	mnt-data-1	PV identifier followed by node name for creating PVs
	minio_pv[*].storageClassName	String	local-storage	Storage class name to match with PVC

COMPONENT	COMPONENT PARAMETER	TYPE	VALUE	DESCRIPTION/COMMENT
	minio_pv[*].accessMode	String	ReadWriteOnce	Access mode when mounting a volume: ReadWriteOnce, ReadOnlyMany, ReadWriteMany, ReadWriteOncePod
	minio_pv[*].persistentVolumeReclaimPolicy	String	Retain	Reclaim policy when a volume is released once it's bound: Retain, Recycle, Delete
	minio_pv(*).mountPath	String	/mnt/data0	Mount path of a volume
	minio_pv(*).device	String	/dev/nvme0n1	Target storage device name when creating a volume. When group_var: minio_deploy_test_mode == true, use files (/tmp/diskimage[*]) as a loop device (/dev/loop[*]) for storage. Otherwise, use an actual NVME or SSD device for storage on the device name for storage.
	minio_pv(*).capacity	String	1GiB	Volume capacity when creating a partition on the target device. Supports units: GiB, TiB

7 BMRA Basic Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA Basic Flavor.

To use the Basic Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 7.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 7.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 7.3](#) for details.
4. Deploy the platform. Refer to [Section 7.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

7.1 Step 1 - Set Up Basic Configuration Profile Hardware

The table in this section lists the hardware BOM for the Basic Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 44. Hardware Setup for Basic Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_1	Controller_3rdGen_1	Controller_4thGen_1	Controller_Xeon_D_1
Worker node options	Worker_2ndGen_Base_1	Worker_3rdGen_Base_1	Worker_4thGen_Base_1	Worker_Xeon_D_Base_1

7.2 Step 2 - Download Basic Configuration Profile Ansible Playbook

This section contains details for downloading the Basic Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Basic Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

7.2.1 Basic Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Basic Configuration Profile allows you to provision a production-ready Kubernetes cluster. Every capability included in the Basic Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Basic Configuration Profile.

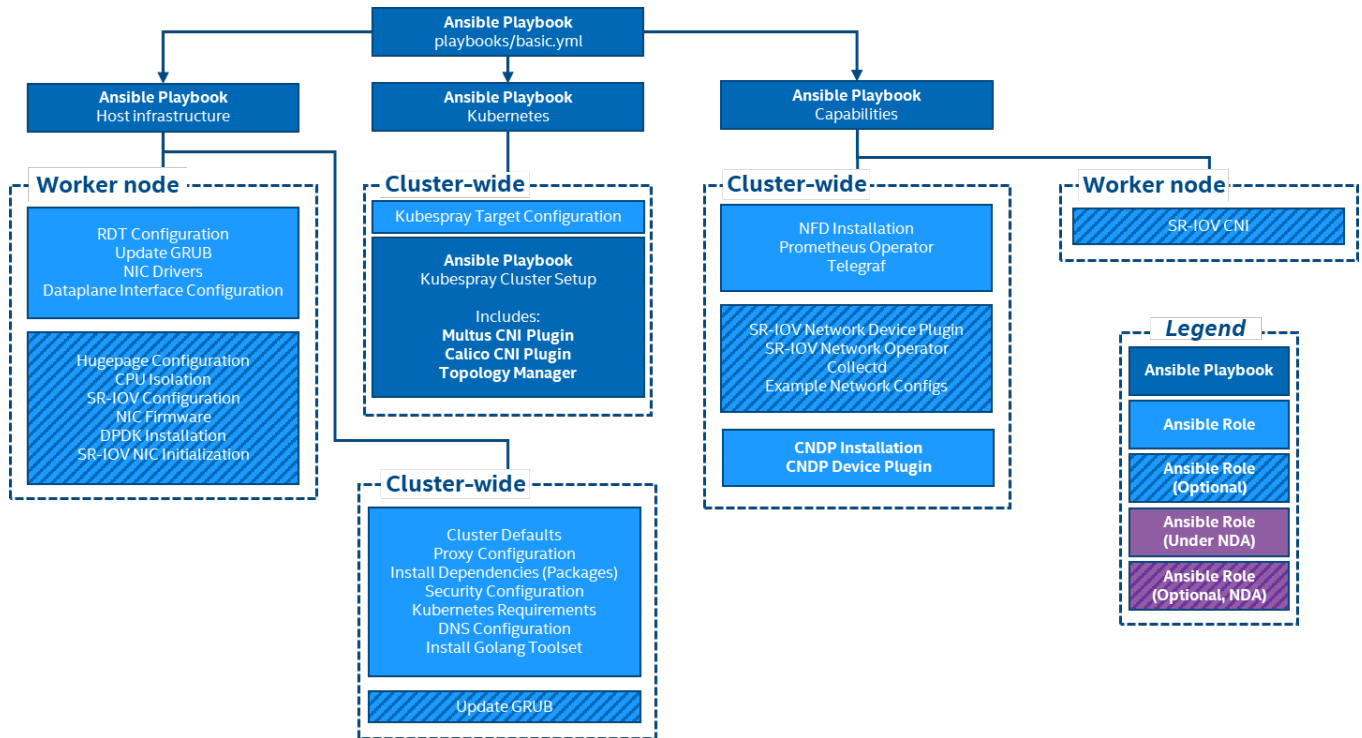


Figure 3. Basic Configuration Profile Ansible Playbook

7.3 Step 3 - Set Up Basic Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 7.3.1](#) and [Section 7.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

7.3.1 Basic Configuration Profile Group Variables

Table 45. Basic Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	For the list of all configurable properties, see Section 6.3
nfd_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	false	
sriov_net_dp_enabled	false	
example_net_attach_defs	false	
collectd_enabled	false	
telegraf_enabled	true	

7.3.2 Basic Configuration Profile Host Variables⁵

Table 46. Basic Configuration Profile – Host Variables

COMPONENT	VALUE	
iommu_enabled	false	For the list of all configurable properties, see Section 6.4
sriov_cni_enabled	false	
install_dpdk	false	
isolcpus_enabled	false	
dataplane_interfaces	[]	

7.4 Step 4 – Deploy and Validate Basic Configuration Profile Platform

Deploy the Basic Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

⁵ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

8 BMRA Full Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA Full Flavor.

To use the BMRA Full Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 8.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 8.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 8.3](#) for details.
4. Deploy the platform. Refer to [Section 8.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value.)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

8.1 Step 1 - Set Up Full Configuration Profile Hardware

The table in this section lists the hardware BOM for the Full Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least three control nodes and two worker nodes.

Table 47. Hardware Setup for Full Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_3	Controller_3rdGen_3	Controller_4thGen_3	Controller_Xeon_D_3
Worker node options	Worker_2ndGen_Plus_1	Worker_3rdGen_Plus_1	Worker_4thGen_Plus_1	Worker_Xeon_D_Plus_1

8.2 Step 2 - Download Full Configuration Profile Ansible Playbook

This section contains details for downloading the Full Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Full Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

8.2.1 Full Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Full Configuration Profile allows you to provision a production-ready Kubernetes. It also applies any additional requirements, such as host OS configuration or Network Adapter drivers and firmware updates. Full Configuration Profile playbook includes all features available through BMRA Ansible Playbook and provides one of the highest degrees of configurability. Every capability included in the Full Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Full Configuration Profile.

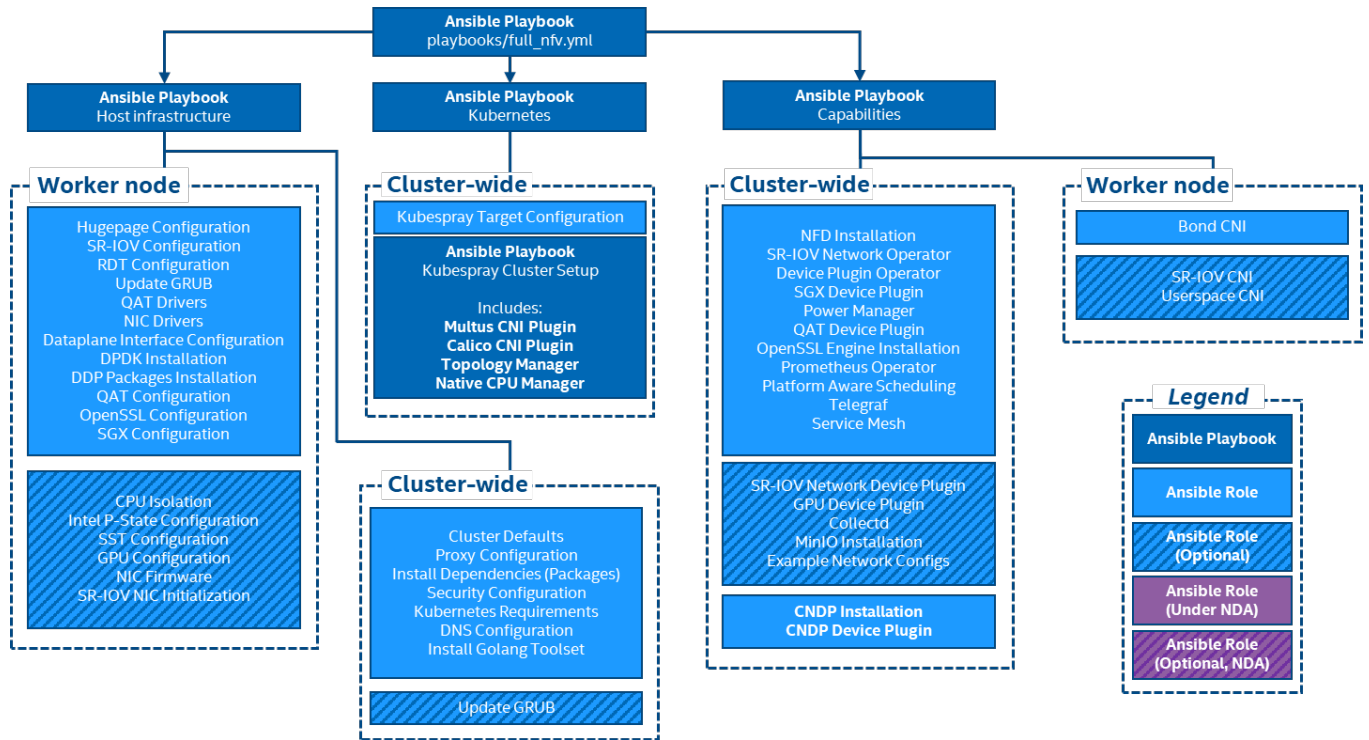


Figure 4. Full Configuration Profile Ansible Playbook

8.3 Step 3 - Set Up Full Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 8.3.1](#) and [Section 8.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

8.3.1 Full Configuration Profile Group Variables

Table 48. Full Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	For the list of all configurable properties, see Section 6.3
nfd_enabled	true	
native_cpu_manager_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	true	
sriov_net_dp_enabled	false	
sgx_dp_enabled	true	
gpu_dp_enabled	false	
qat_dp_enabled	true	
openssl_engine_enabled	true	
kmra_enabled	true	
tas_enabled	true	

COMPONENT	VALUE
gas_enabled	false
example_net_attach_defs	false
collectd_enabled	false
telegraf_enabled	true
service_mesh	true
power_manager	false
minio_enabled	false
kube_network_plugin_multus	true

8.3.2 Full Configuration Profile Host Variables⁶

Table 49. Full Configuration Profile – Host Variables

COMPONENT	VALUE	
iommu_enabled	true	For the list of all configurable properties, see Section 6.4
sriov_cni_enabled	false	
bond_cni_enabled	true	
ddp_enabled	true	
sst_pp_configuration_enabled	false	
userspace_cni_enabled	false	
hugepages_enabled	true	
isolcpus_enabled	false	
install_dpdk	true	
install_ddp_packages	true	
qat_devices	[]	
dataplane_interfaces	[]	
minio_pv	[]	

8.4 Step 4 - Deploy and Validate Full Configuration Profile Platform

Deploy the Full Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

⁶ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

9 BMRA On-Premises Edge Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA On-Premises Edge Flavor.

To use the On-Premises Edge Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 9.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 9.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 9.3](#) for details.
4. Deploy the platform. Refer to [Section 9.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value.)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

9.1 Step 1 - Set Up On-Premises Edge Configuration Profile Hardware

The table in this section lists the hardware BOM for the On-Premises Edge Configuration Profile, including Control Node, Worker Node Base, and Worker Node Plus. We recommend that you set up at least one control node and one worker node.

Table 50. Hardware Setup for On-Premises Edge Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_1	Controller_3rdGen_1	Controller_4thGen_1	Controller_Xeon_D_1
Worker node options	Worker_2ndGen_Base_2 or Worker_2ndGen_Plus_1	Worker_3rdGen_Base_2 or Worker_3rdGen_Plus_1	Worker_4thGen_Base_2 or Worker_4thGen_Plus_1	Worker_Xeon_D_Base_2 or Worker_Xeon_D_Plus_1

9.2 Step 2 - Download On-Premises Edge Configuration Profile Ansible Playbook

This section contains details for downloading the On-Premises Edge Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the On-Premises Edge Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

9.2.1 On-Premises Edge Configuration Profile Ansible Playbook Overview

The Ansible playbook for the On-Premises Edge Configuration Profile allows you to provision a production-ready Kubernetes cluster. It also applies any additional requirements, such as host OS configuration or Network Adapter drivers and firmware updates. Every capability included in the On-Premises Edge Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the On-Premises Edge Configuration Profile.

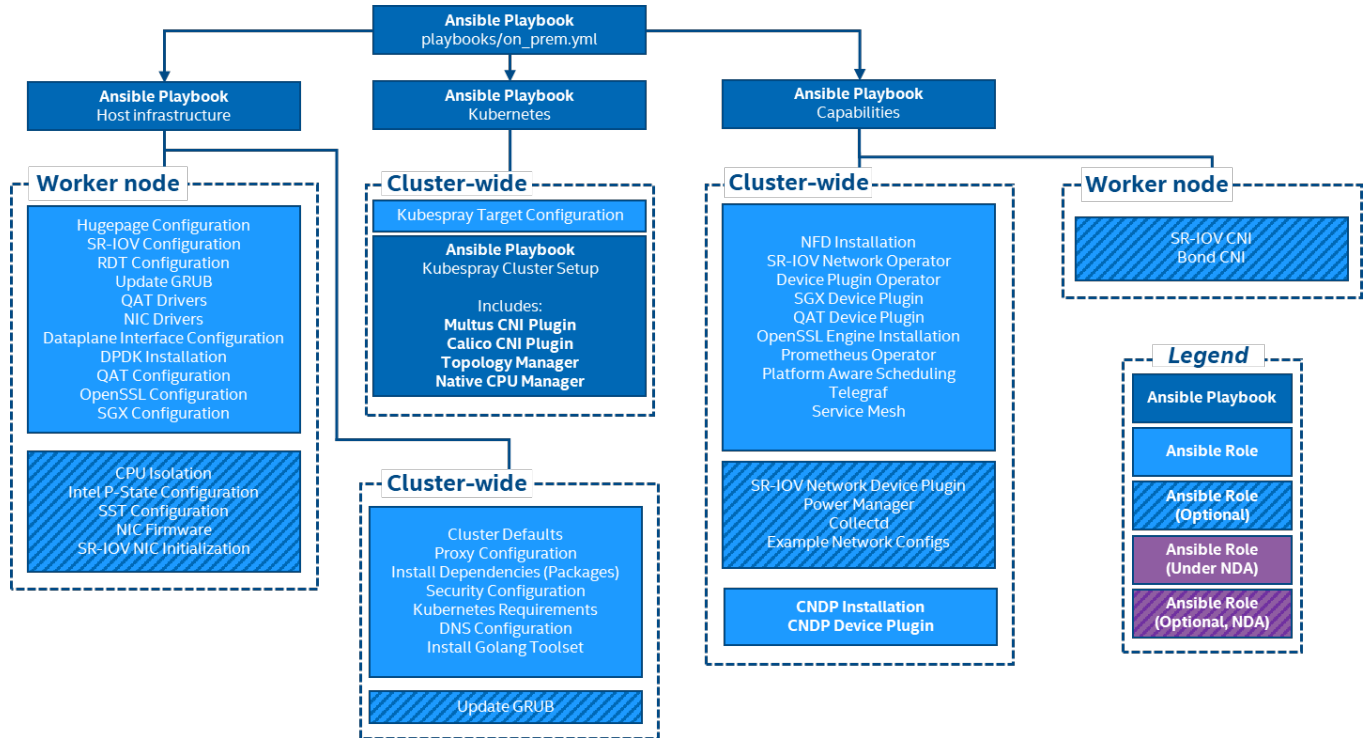


Figure 5. On-Premises Edge Configuration Profile Ansible Playbook

9.3 Step 3 - Set Up On-Premises Edge Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 9.3.1](#) and [Section 9.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

9.3.1 On-Premises Edge Configuration Profile Group Variables

Table 51. On-Premises Edge Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	
nfd_enabled	true	
native_cpu_manager_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	true	
sriov_net_dp_enabled	false	
sgx_dp_enabled	true	
qat_dp_enabled	true	
openssl_engine_enabled	true	
kmra_enabled	true	
tas_enabled	true	
example_net_attach_defs	false	

For the list of all configurable properties, see [Section 6.3](#)

COMPONENT	VALUE
collectd_enabled	false
telegraf_enabled	true
service_mesh	true
power_manager	false

9.3.2 On-Premises Edge Configuration Profile Host Variables⁷

Table 52. On-Premises Edge Configuration Profile – Host Variables

COMPONENT	VALUE
iommu_enabled	true
sriov_cni_enabled	false
bond_cni_enabled	false
hugepages_enabled	true
isolcpus_enabled	false
sst_pp_configuration_enabled	false
install_dpdk	true
qat_devices	[]
dataplane_interfaces	[]

For the list of all configurable properties, see [Section 6.4](#)

9.4 Step 4 - Deploy and Validate On-Premises Edge Configuration Profile Platform

Deploy the On-Premises Edge Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

⁷ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

10 BMRA Remote Central Office-Forwarding Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA Remote Central Office-Forwarding Flavor.

To use the Remote Central Office-Forwarding Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 10.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 10.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 10.3](#) for details.
4. Deploy the platform. Refer to [Section 10.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value.)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

10.1 Step 1 - Set Up Remote Central Office-Forwarding Configuration Profile Hardware

The table in this section lists the hardware BOM for the Remote Central Office-Forwarding Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 53. Hardware Setup for Remote Central Office-Forwarding Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_2	Controller_3rdGen_2	Controller_4thGen_2	Controller_Xeon_D_2
Worker node options	Worker_2ndGen_Base_3 or Worker_2ndGen_Plus_2	Worker_3rdGen_Base_3 or Worker_3rdGen_Plus_2	Worker_4thGen_Base_3 or Worker_4thGen_Plus_2	Worker_Xeon_D_Base_3 or Worker_Xeon_D_Plus_2

10.2 Step 2 - Download Remote Central Office-Forwarding Configuration Profile Ansible Playbook

This section contains details for downloading the Remote Central Office-Forwarding Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Remote Central Office-Forwarding Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

10.2.1 Remote Central Office-Forwarding Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Remote Central Office-Forwarding Configuration Profile allows you to provision a production-ready Kubernetes cluster. It also applies any additional requirements, such as host OS configuration or Network Adapter drivers and firmware updates. Every capability included in the Remote Central Office-Forwarding Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Remote Central Office-Forwarding Configuration Profile.

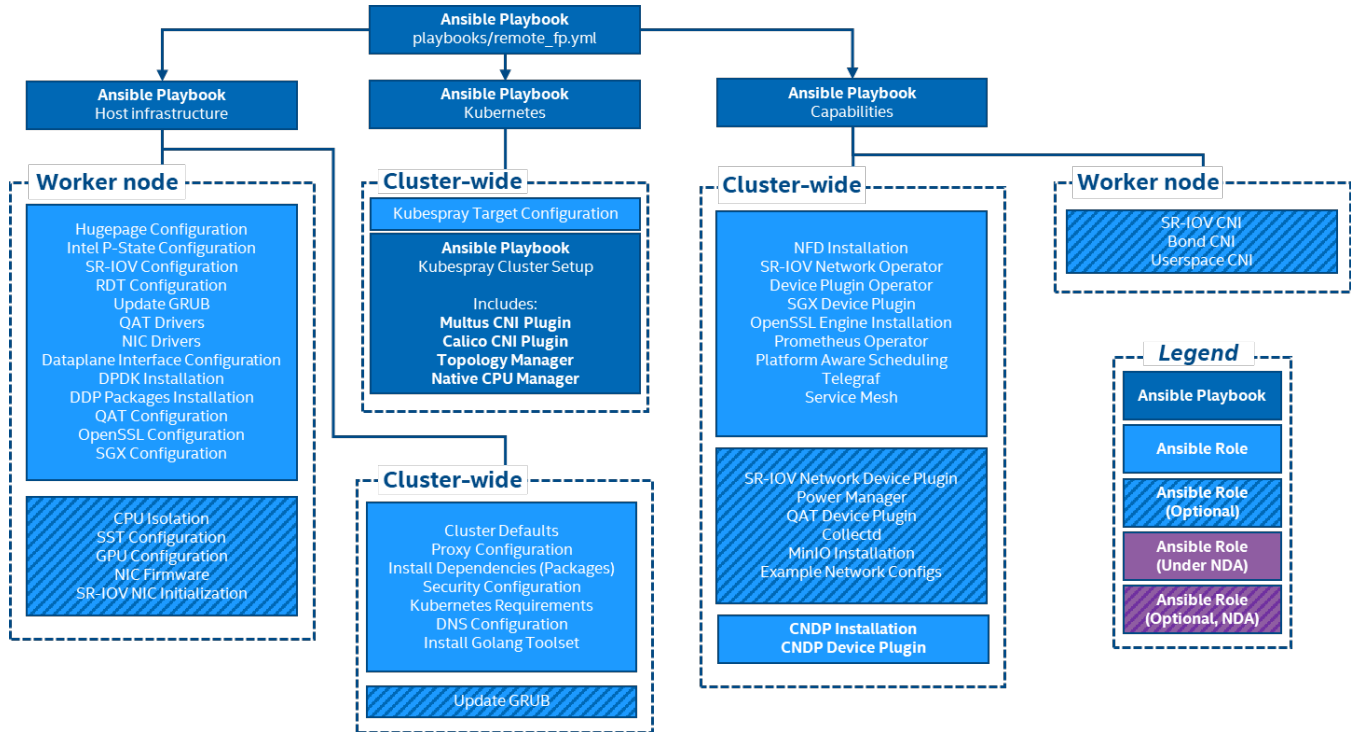


Figure 6. Remote Central Office-Forwarding Configuration Profile Ansible Playbook

10.3 Step 3 - Set Up Remote Central Office-Forwarding Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.3.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 10.3.1](#) and [Section 10.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

10.3.1 Remote Central Office-Forwarding Configuration Profile Group Variables

Table 54. Remote Central Office-Forwarding Configuration Profile – Group Variables

COMPONENT	VALUE
Kubernetes	true
nfd_enabled	true
native_cpu_manager_enabled	true
topology_manager_enabled	true
sriov_network_operator_enabled	true
sriov_net_dp_enabled	false
sgx_dp_enabled	true
qat_dp_enabled	false
openssl_engine_enabled	true
kmra_enabled	true
tas_enabled	true
example_net_attach_defs	false

For the list of all configurable properties, see [Section 6.3](#)

COMPONENT	VALUE
collectd_enabled	false
telegraf_enabled	true
service_mesh	true
power_manager	false

10.3.2 Remote Central Office-Forwarding Configuration Profile Host Variables⁸

Table 55. Remote Central Office-Forwarding Configuration Profile – Host Variables

COMPONENT	VALUE
iommu_enabled	true
sriov_cni_enabled	false
bond_cni_enabled	false
ddp_enabled	true
userspace_cni_enabled	false
hugepages_enabled	true
isolcpus_enabled	false
sst_pp_configuration_enabled	false
install_dpdk	true
install_ddp_packages	true
qat_devices	[]
dataplane_interfaces	[]

For the list of all configurable properties, see [Section 6.4](#)

10.4 Step 4 - Deploy and Validate Remote Central Office-Forwarding Configuration Profile Platform

Deploy the Remote Central Office-Forwarding Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

⁸ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

11 BMRA Regional Data Center Configuration Profile Setup

This section contains a step-by-step description of how to set up your BMRA Regional Data Center Flavor.

To use the Regional Data Center Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 11.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 11.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 11.3](#) for details.
4. Deploy the platform. Refer to [Section 11.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value.)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

11.1 Step 1 - Set Up Regional Data Center Configuration Profile Hardware

The table in this section lists the hardware BOM for the Regional Data Center Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 56. Hardware Setup for Regional Data Center Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	N/A*	Controller_3rdGen_3	N/A*	N/A*
Worker node options	N/A*	Worker_3rdGen_Plus_3	N/A*	N/A*

*Configuration Profile only tested with 3rd Generation Intel Xeon Scalable processor

11.2 Step 2 - Download Regional Data Center Configuration Profile Ansible Playbook

This section contains details for downloading the Regional Data Center Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Regional Data Center Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

11.2.1 Regional Data Center Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Regional Data Center Configuration Profile allows you to provision a production-ready Kubernetes cluster. It also applies any additional requirements, such as host OS configuration or Network Adapter drivers and firmware updates. Every capability included in the Regional Data Center Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host vars tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Regional Data Center Configuration Profile.

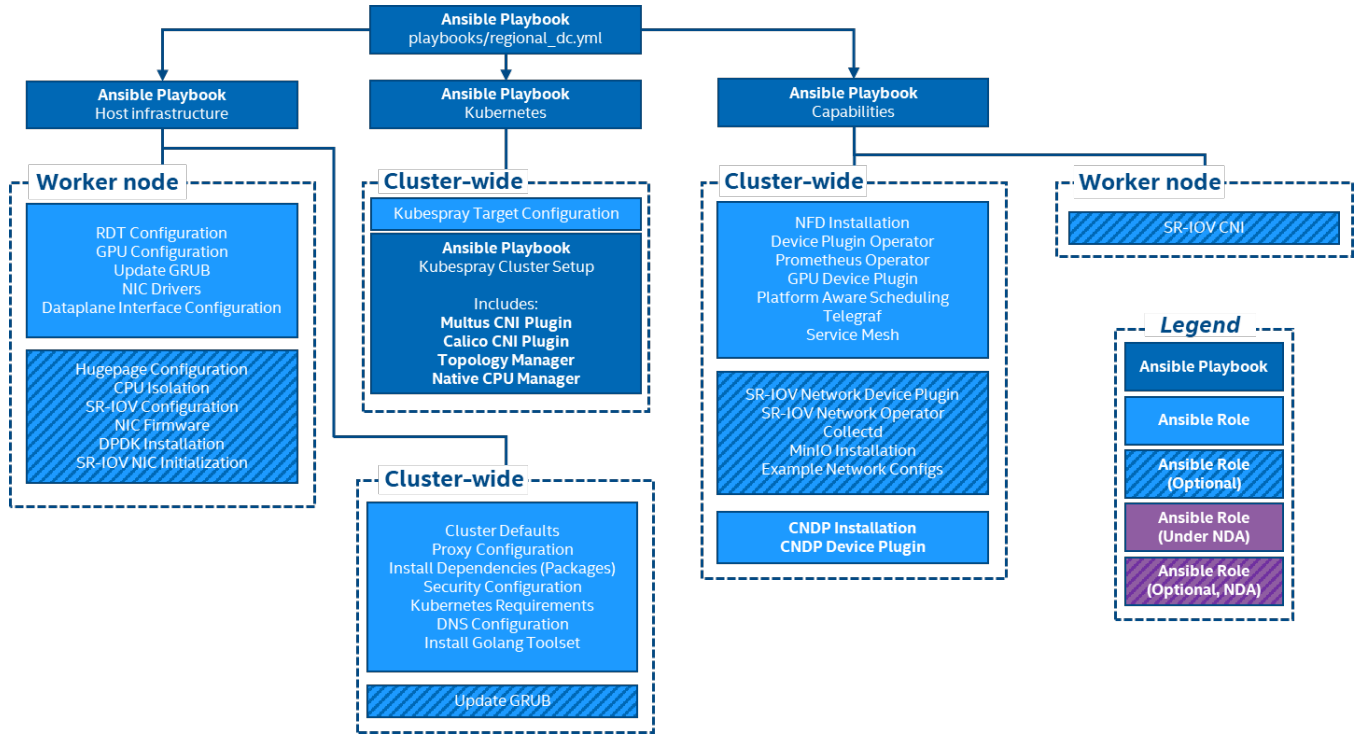


Figure 7. Regional Data Center Configuration Profile Ansible Playbook

11.3 Step 3 - Set Up Regional Data Center Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 11.3.1](#) and [Section 11.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

11.3.1 Regional Data Center Configuration Profile Group Variables

Table 57. Regional Data Center Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	For the list of all configurable properties, see Section 6.3
nfd_enabled	true	
native_cpu_manager_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	false	
sriov_net_dp_enabled	false	
gpu_dp_enabled	true	
tas_enabled	true	
gas_enabled	true	
example_net_attach_defs	false	
collectd_enabled	false	
telegraf_enabled	true	

COMPONENT	VALUE
service_mesh	true

11.3.2 Regional Data Center Configuration Profile Host Variables⁹

Table 58. Regional Data Center Configuration Profile – Host Variables

COMPONENT	VALUE
iommu_enabled	false
sriov_cni_enabled	false
hugepages_enabled	false
isolcpus_enabled	false
install_dpdk	false
dataplane_interfaces	[]

For the list of all configurable properties, see [Section 6.4](#)

11.4 Step 4 – Deploy and Validate Regional Data Center Configuration Profile Platform

Deploy the Regional Data Center Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

⁹ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

12 BMRA for Storage Configuration Profile Setup

The BMRA for Storage focuses exclusively on object storage with the use of MinIO, where the singular workload for the storage solution is the MinIO application, as shown in [Figure 8](#). MinIO has no requirements for telco-specific infrastructure such as DPDK. MinIO's only architectural requirement is the basic OS support for it to service clients and use connected storage devices such as HDDs and SSDs.

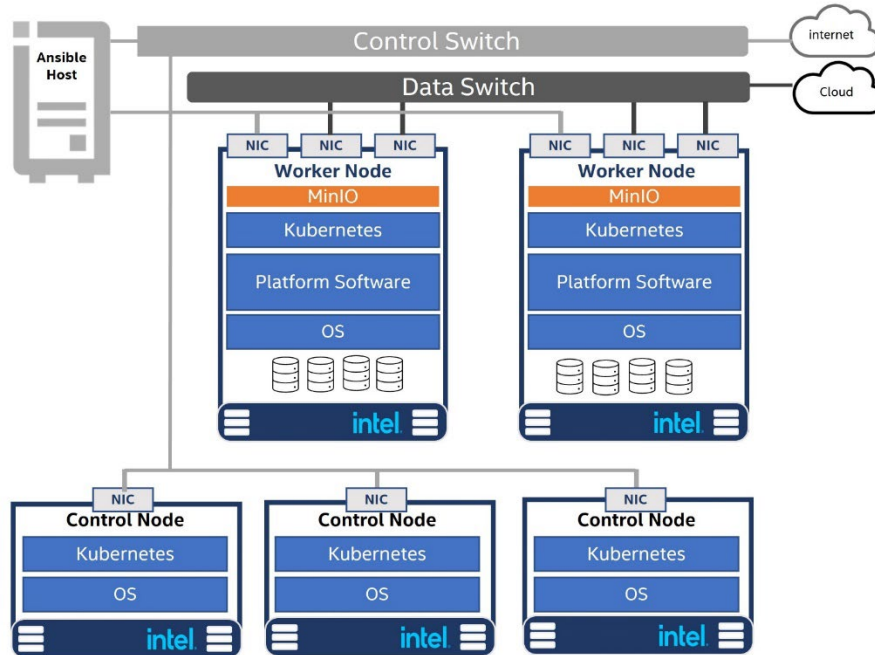


Figure 8. BMRA for Storage Architecture

This section contains a step-by-step description of how to set up your BMRA for Storage using the Storage Configuration Profile.

To use the Storage Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 12.1](#) for details. You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 12.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 12.3](#) for details.
4. Deploy the platform. Refer to [Section 12.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value.)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

12.1 Step 1 - Set Up Storage Configuration Profile Hardware

The table in this section lists the hardware BOM for the Storage Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 59. Hardware Setup for Storage Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	N/A*	Controller_3rdGen_3	N/A*	N/A*
Worker node options	N/A*	Storage_3rdGen_1	N/A*	N/A*

*Configuration Profile only tested with 3rd Generation Intel Xeon Scalable processor

12.2 Step 2 - Download Storage Configuration Profile Ansible Playbook

This section contains details for downloading the Storage Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Storage Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

12.2.1 Storage Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Storage Configuration Profile allows you to provision a production-ready Kubernetes cluster. It also applies any additional requirements, such as host OS configuration or Network Adapter drivers and firmware updates. Every capability included in the Storage Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host vars tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Storage Configuration Profile.

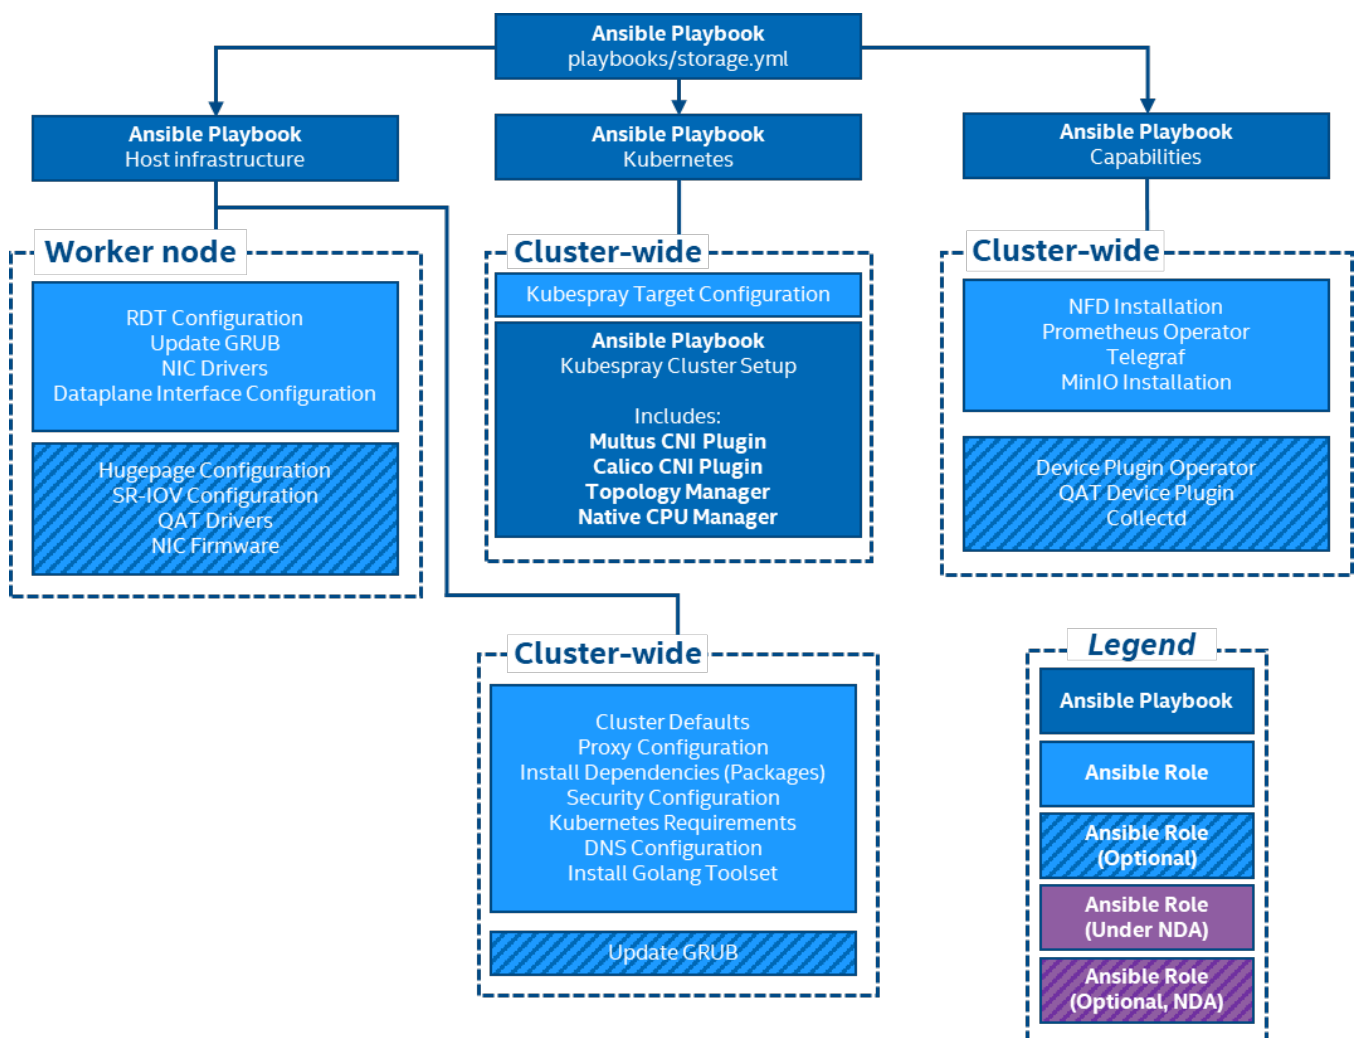


Figure 9. Storage Configuration Profile Ansible Playbook

12.3 Step 3 - Set Up Storage Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 12.3.1](#) and [Section 12.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

12.3.1 Storage Configuration Profile Group Variables

Table 60. Storage Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	For the list of all configurable properties, see Section 6.3
nfd_enabled	true	
native_cpu_manager_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	true	
qat_dp_enabled	false	
tas_enabled	true	
minio_enabled	true	
kube_network_plugin_multus	true	
collectd_enabled	false	
telegraf_enabled	true	

12.3.2 Storage Configuration Profile Host Variables¹⁰

Table 61. Storage Configuration Profile – Host Variables

COMPONENT	VALUE	
iommu_enabled	false	For the list of all configurable properties, see Section 6.4
sriov_cni_enabled	false	
hugepages_enabled	true	
isolcpus_enabled	false	
qat_devices	[]	
dataplane_interfaces	[]	
minio_pv	[]	

12.4 Step 4 - Deploy and Validate Storage Configuration Profile Platform

Deploy the Storage Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

¹⁰ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

13 BMRA Access Edge Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA Access Edge Flavor.

To use the Access Edge Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 13.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 13.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 13.3](#) for details.
4. Deploy the platform. Refer to [Section 13.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

13.1 Step 1 - Set Up Access Edge Configuration Profile Hardware

The table in this section lists the hardware BOM for the Access Edge Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 62. Hardware Setup for Access Edge Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_1	Controller_3rdGen_1	Controller_4thGen_1	Controller_Xeon_D_1
Worker node options	Worker_2ndGen_Base_1	Worker_3rdGen_Base_1	Worker_4thGen_Base_1	Worker_Xeon_D_Base_1

13.2 Step 2 - Download Access Edge Configuration Profile Ansible Playbook

This section contains details for downloading the Access Edge Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Access Edge Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

13.2.1 Access Edge Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Access Edge Configuration Profile allows you to provision a production-ready Kubernetes cluster. Every capability included in the Access Edge Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Access Edge Configuration Profile.

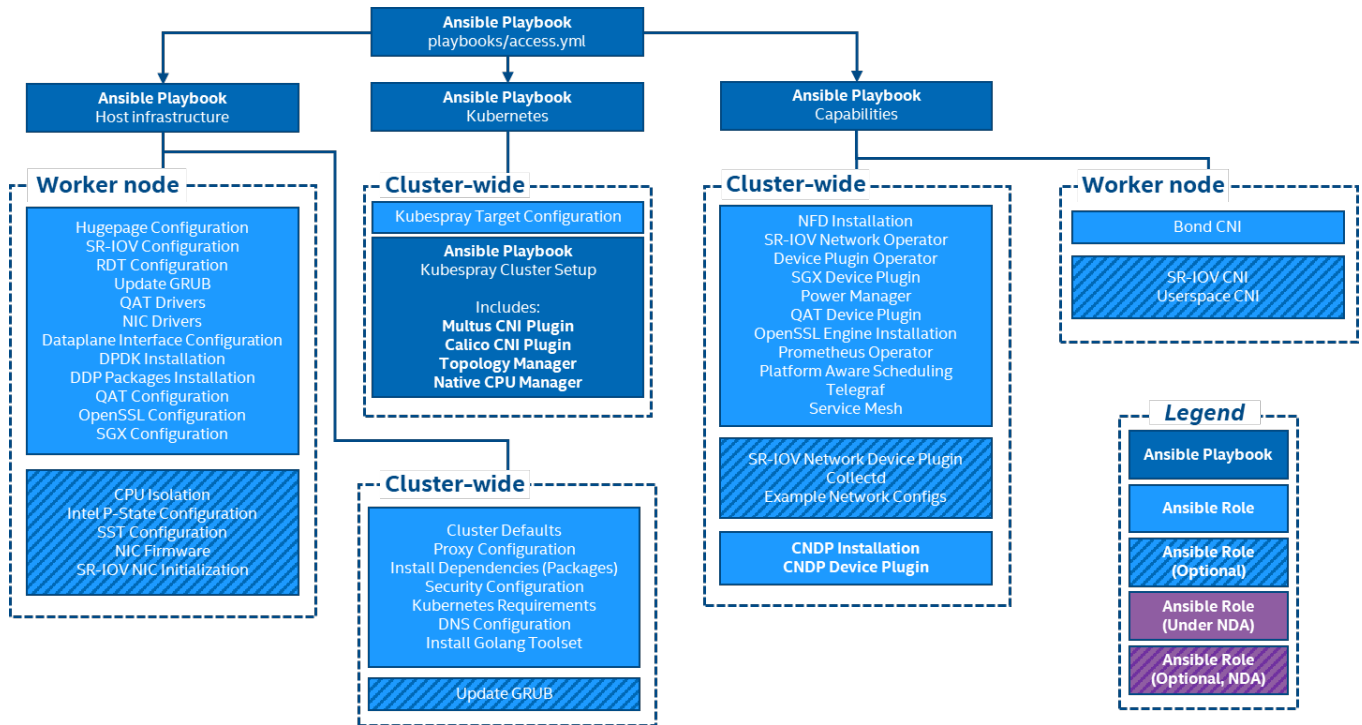


Figure 10. Access Edge Configuration Profile Ansible Playbook

13.3 Step 3 - Set Up Access Edge Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 13.3.1](#) and [Section 13.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

13.3.1 Access Edge Configuration Profile Group Variables

Table 63. Access Edge Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	true	For the list of all configurable properties, see Section 6.3
nfd_enabled	true	
topology_manager_enabled	true	
sriov_network_operator_enabled	true	
sriov_net_dp_enabled	false	
example_net_attach_defs	false	
collectd_enabled	false	
telegraf_enabled	true	

13.3.2 Access Edge Configuration Profile Host Variables¹¹

Table 64. Access Edge Configuration Profile – Host Variables

COMPONENT	VALUE	
iommu_enabled	true	For the list of all configurable properties, see Section 6.4
sriov_cni_enabled	false	
install_dpdk	true	
isolcpus_enabled	true	
dataplane_interfaces	[]	

13.4 Step 4 - Deploy and Validate Access Edge Configuration Profile Platform

Deploy the Access Edge Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

¹¹ See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

14 BMRA Build-Your-Own Configuration Profile Setup

This section contains a step-by-step description of how to set up a BMRA Build-Your-Own Flavor.

To use the Build-Your-Own Configuration Profile, perform the following steps:

1. Choose your hardware, set it up, and configure the BIOS. Refer to [Section 14.1](#) for details.
You also need to build your Kubernetes cluster.
2. Download the Ansible playbook for your Configuration Profile. Refer to [Section 14.2](#) for details.
3. Configure the optional Ansible parameters using the information in the Configuration Profile tables. Refer to [Section 14.3](#) for details.
4. Deploy the platform. Refer to [Section 14.4](#) for details.
5. Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

Be aware of the definitions of terminology used in tables in this section.

TERM	DESCRIPTION
Hardware Taxonomy	
ENABLED	Setting must be enabled in the BIOS (configured as Enabled, Yes, True, or similar value)
DISABLED	Setting must be disabled in the BIOS (configured as Disabled, No, False, or any other value with this meaning.)
OPTIONAL	Setting can be either disabled or enabled, depending on user's workload. Setting does not affect the Configuration Profile or platform deployment.
Software Taxonomy	
TRUE	Feature is included and enabled by default.
FALSE	Feature is included but disabled by default - can be enabled and configured by user.
N/A	Feature is not included and cannot be enabled or configured.

14.1 Step 1 - Set Up Build-Your-Own Configuration Profile Hardware

The table in this section lists the hardware BOM for the Build-Your-Own Configuration Profile, including control node, worker node base, and worker node plus. We recommend that you set up at least one control node and one worker node.

Table 65. Hardware Setup for Build-Your-Own Configuration Profile

NODE OPTIONS	2ND GENERATION INTEL XEON SCALABLE PROCESSOR	3RD GENERATION INTEL XEON SCALABLE PROCESSOR	4TH GENERATION INTEL XEON SCALABLE PROCESSOR	INTEL XEON D PROCESSOR
Control node options	Controller_2ndGen_1	Controller_3rdGen_1	Controller_4thGen_1	Controller_Xeon_D_1
Worker node options	Worker_2ndGen_Base_1	Worker_3rdGen_Base_1	Worker_4thGen_Base_1	Worker_Xeon_D_Base_1

14.2 Step 2 - Download Build-Your-Own Configuration Profile Ansible Playbook

This section contains details for downloading the Build-Your-Own Configuration Profile Ansible playbook. It also provides an overview of the Ansible playbook and lists the software that is automatically installed when the playbook is deployed.

Download the Build-Your-Own Configuration Profile Ansible playbook using the steps described in [Section 2.5](#).

14.2.1 Build-Your-Own Configuration Profile Ansible Playbook Overview

The Ansible playbook for the Build-Your-Own Configuration Profile allows you to provision a production-ready Kubernetes cluster. Every capability included in the Build-Your-Own Configuration Profile playbook can be disabled or enabled. Refer to the diagram and group and host variables tables below to see which Ansible roles are included and executed by default.

The diagram shows the architecture of the Ansible playbooks and roles that are included in the Build-Your-Own Configuration Profile.

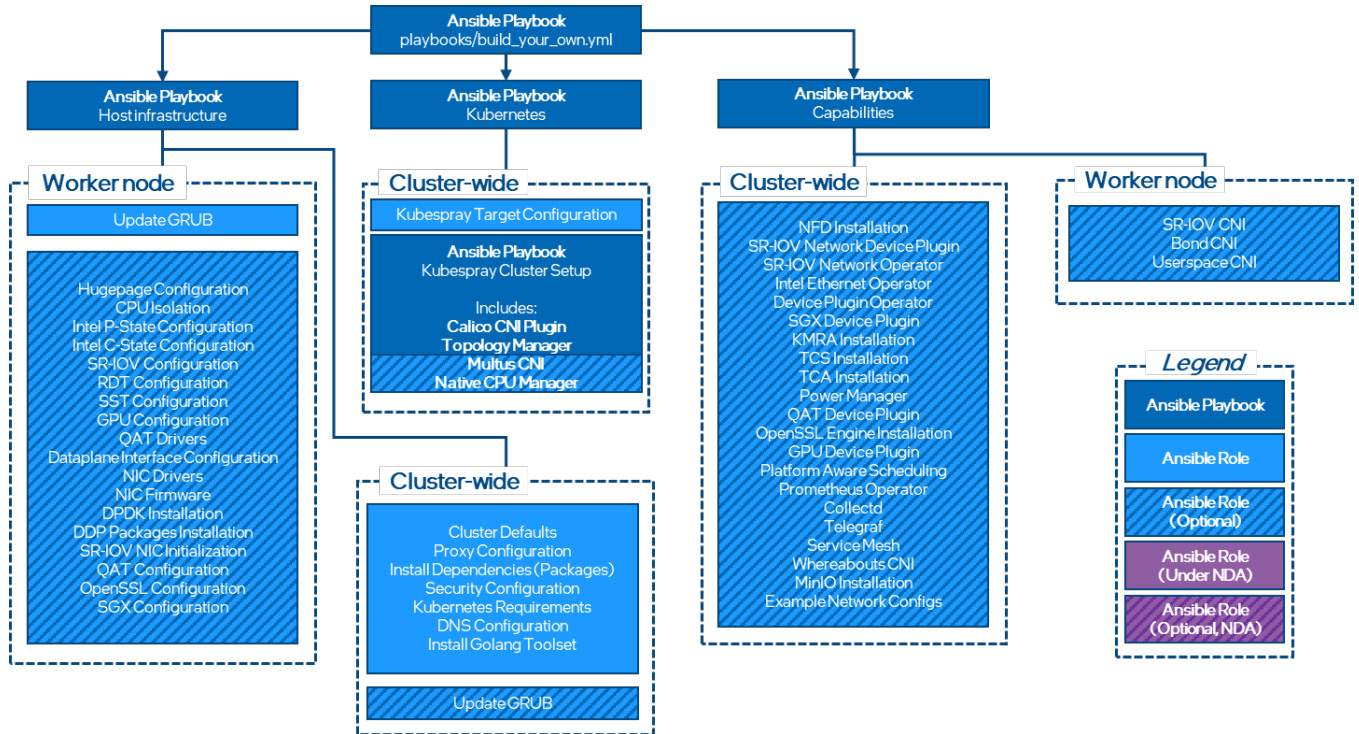


Figure 11. Build-Your-Own Configuration Profile Ansible Playbook

14.3 Step 3 - Set Up Build-Your-Own Configuration Profile

Review the optional Ansible group and host variables in this section and select options that match your desired configuration.

1. Update the `inventory.ini` file with your environment details as described in [Section 2.5.3](#).
2. Create `host_vars` files for all worker nodes as specified in [Section 2.5.4](#).
3. Update group and host variables to match your desired configuration as specified in [Section 2.3.4](#). Refer to the tables in [Section 13.3.1](#) and [Section 13.3.2](#).

Variables are grouped into two main categories:

1. Group variables – apply to both control and worker nodes and have cluster-wide impact.
2. Host variables – scope is limited to a single worker node.

The tables below are a summary of group and host variables. For lists showing all configurable properties, see [Section 6.3](#) and [Section 6.4](#). All of the variables are important but pay special attention to variables in **bold** as they almost always need to be updated to match the target environment.

14.3.1 Build-Your-Own Configuration Profile Group Variables

Table 66. Build-Your-Own Configuration Profile – Group Variables

COMPONENT	VALUE	
Kubernetes	false	For the list of all configurable properties, see Section 6.3
nfd_enabled	false	
topology_manager_enabled	false	
sriov_network_operator_enabled	false	
sriov_net_dp_enabled	false	
example_net_attach_defs	false	
collectd_enabled	false	
telegraf_enabled	false	

14.3.2 Build-Your-Own Configuration Profile Host Variables¹²

Table 67. Build-Your-Own Configuration Profile – Host Variables

COMPONENT	VALUE	
iommu_enabled	false	For the list of all configurable properties, see Section 6.4
sriov_cni_enabled	false	
install_dpdk	false	
isolcpus_enabled	false	
dataplane_interfaces	[]	

14.4 Step 4 - Deploy and Validate Build-Your-Own Configuration Profile Platform

Deploy the Build-Your-Own Configuration Profile Ansible playbook using the steps described in [Section 2.5.5](#).

Validate the setup of your Kubernetes cluster. Refer to the tasks in [Section 5](#) and run the validation processes according to the hardware and software components that you have installed.

¹² See backup for workloads and configurations or visit [Performance Index](#). Results may vary.

Part 3:

BMRA Applications

15 Workloads and Application Examples

This section provides examples of how to provision and deploy example applications or workloads.

15.1 Enabling Key Management NGINX Applications

KMRA source code and Dockerfiles: [Key Management Reference Application](#)

KMRA docker images on Docker Hub:

- AppHSM: <https://hub.docker.com/r/intel/apphsm>
- ctk_loadkey: https://hub.docker.com/r/intel/ctk_loadkey
- PCCS: <https://hub.docker.com/r/intel/pccs>

KMRA Helm charts are in `/roles/kmra_install/charts`.

Steps to deploy the full KMRA NGINX demo:

1. Generate a new PCCS primary API key and update the `kmra_pccs_api_key` variable in `group_vars/all.yml` (go to [Intel® Provisioning Certification Service for ECDSA Attestation](#) and subscribe).
2. Ensure that the `kmra_deploy_demo_workload` variable in the `group_vars/all.yml` is set to `true`.
3. Deploy the `full_nfv`, `on_prem`, or `remote_fp` profile to set up KMRA demo with NGINX. The `kmra` variable must be set to `on` in `profiles/profiles.yml`.

15.2 Enabling Trusted Certificate Service

Trusted Certificate Service (TCS) is a Kubernetes certificate signing solution that uses the security capabilities provided by Intel® SGX. The signing key is stored and used inside the SGX enclaves and is never stored in clear anywhere in the system. TCS is implemented as a [cert-manager external issuer](#) by supporting both cert-manager and Kubernetes certificate signing APIs.

To enable TCS on BMRA, follow the guide available at [Trusted Certificate Issuer](#).

15.2.1 Istio Custom CA Integration Using Kubernetes CSR

Istio supports [integrating custom certificate authority \(CA\) using Kubernetes CSR](#) as an experimental feature.

Detailed example steps described in the [Istio Custom CA with CSR](#) document show how to provision Istio workload certificates using an Issuer provided by the Trusted Certificate Service (TCS).

Note: Due to misconfiguration of the Istio Demo application, you might need to disable hugepages temporarily to avoid demo app becoming stuck in the `CrashLoopBackOff` state. To disable hugepages, execute following command on the worker node:

```
echo 0 > /proc/sys/vm/nr_hugepages
```

15.2.2 Remote Attestation and Manual Key Management

TCS supports SGX remote attestation and the sample key management reference application.

All required steps are described in the [Integrate Key Server](#) document.

15.3 Service Mesh Automated Remote Attestation and Key Management with KMRA, TCS, and TCA

Remote attestation is an advanced feature that allows an entity to gain the relying party's trust. Remote attestation gives the relying party increased confidence that the software is running inside an SGX enclave. The attestation results include the identity of the software being attested and an assessment of possible software tampering.

Key management enables external key management systems to deliver the certificates and keys via secure mechanisms into the SGX enclave. To enable the automated key management feature, KMRA AppHSM and KMRA PCCS applications must be enabled and configured as well as Trusted Certificate Service (TCS) and Trusted Certificate Attestation (TCA). BMRA tries to install all dependencies and configure the host with reasonable defaults.

KMRA application settings are collected under the `kmra` variable in the `group_vars/all.yml` file and all default values are available for reference in the `roles/kmra_install/defaults/main.yml` file. If you need to overwrite any default value, redefine it in the `group_vars/all.yml` file while keeping the variable structure.

In general, TCS does not require specific configuration. Default values used for TCS deployment are collected in the `roles/tcs_install/vars/main.yml` file and can be redefined in the `group_vars/all.yml` file.

TCA depends on settings of KMRA AppHSM, which should match. Refer to the default values, which can be found in the `roles/tca_install/vars/main.yml` file. Default value can be redefined in the `group_vars/all.yml` file.

Service Mesh default settings can be found in the `roles/service_mesh_install/vars/main.yml` file.

For detailed documentation on components involved in this feature, refer to:

- KMRA: [Key Management Reference Application](#)
- TCS: [Trusted Certificate Issuer](#)
- TCA: [Trusted Attestation Controller](#)

15.4 Istio TLS Splicing

To configure Istio with TLS splicing, first enable it in the `group_vars/all.yml` file.

```
service_mesh:
  enabled: true
  tls_splicing:
    enabled: true
```

The config creates an ingress gateway to act as a forward proxy, registers virtual service rule and external service entry to implement TLS passthrough for external service.

A client outside the mesh can use the cluster ingress gateway to access external services with TLS splicing.

```
export INGRESS_PORT=$(kubectl -n istio-system get service istio-ingressgateway -o
jsonpath='{.spec.ports[?(@.name=="http2")].nodePort}')
export SECURE_INGRESS_PORT=$(kubectl -n istio-system get service istio-ingressgateway -o
jsonpath='{.spec.ports[?(@.name=="https")].nodePort}')
export TCP_INGRESS_PORT=$(kubectl -n istio-system get service istio-ingressgateway -o
jsonpath='{.spec.ports[?(@.name=="tcp")].nodePort}')
export INGRESS_HOST=$(kubectl get po -l istio=ingressgateway -n istio-system -o
jsonpath='{.items[0].status.hostIP}')

curl -s -v --resolve www.example.com:$SECURE_INGRESS_PORT:$INGRESS_HOST
https://www.example.com:$SECURE_INGRESS_PORT
```

15.5 Web Application Firewall Using Traffic Analytics Development Kit

The functionality of the Web Application Firewall (WAF) running in the cluster can be tested from the command line. Start by getting the IP and port of the firewall:

```
# export NODE_PORT=$(kubectl get --namespace modsec-tadk -o
jsonpath="{.spec.ports[0].nodePort}" services tadk-intel-tadkchart)

# export NODE_IP=$(kubectl get nodes --namespace modsec-tadk -o
jsonpath="{.items[0].status.addresses[0].address}")

# curl http://$NODE_IP:$NODE_PORT
```

Now try sending a message with sample credentials to the firewall:

```
# curl -d "username=admin&password=unknown" "$NODE_IP:$NODE_PORT"
```

The resulting error code should be "403" (Forbidden), showing the firewall has blocked the request.

Part 4:

BMRA Release Notes

Appendix A BMRA Release Notes

This section lists the notable changes from the previous releases, including new features, bug fixes, and known issues.¹³

A.1 BMRA 22.05 Notable Facts

New Components:

- Support for cpusets, C-state, and uncore frequency scaling configuration
- Support for Trusted Certificate Attestation Controller and Trusted Certificate Issuer
- Support for Power Management Unit (PMU) telemetry for power C-state for 4th Generation Intel® Xeon® Scalable processor
- Support for Cloud Native Data Plane (CNDP) for enabling a cloud-based framework for networking and accelerated packet processing, provisioning, orchestrating, and managing microservices to scale efficiently with Kubernetes deployment.
- Added the Access Edge Configuration Profile to support vRAN software applications
- Added the Build-Your-Own Configuration Profile
- Support for Ethernet operator for firmware, driver, and flow control management
- Support for FEC Operator
- Support for Rocky Linux
- Support for Ubuntu 22.04 LTS

New Platforms:

- QCT's beta (2S-SPR XCC)
- DSG's Fox Creek Pass (2S-SPR XCC, MCC)
- Ruby Pass (1S-SPR XCC, MCC)

Updates/Changes:

- Support for Ansible 4.10
- Support for Kubernetes 1.23
- MinIO tenant load balancer option implementation
- Updated Key Management Reference Application (KMRA) version to 2.1
- Updated Data Plane Development Kit (DPDK) version to 22.03
- Updated Vector Packet Processing version to 21.10
- Updated Platform Aware Scheduling version to 0.8
- Updated Telegraf version to 1.1
- Updated Grafana and Prometheus versions
- Updated Intel Ethernet firmware and drivers
- Updated QAT drivers
- Update Power Manager version to 1.0.2

Removed Support:

- Ubuntu 21.04 as base operating system

Known Limitations (Errata):

- MinIO (Storage Configuration Profile) support is limited due to SR-IOV connectivity. See [Section A.9](#) for details.
- Intel SG1 Server Graphics cards have limited operating system support. See [Section A.9](#) for details.

The following table lists key features of the 4th Generation Intel Xeon Scalable processor and the support for those features in BMRA 22.05.

Table 68. Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 22.05

CATEGORY	FEATURE	KERNEL / OS	DPDK	KUBERNETES	BMRA 22.05 STATUS
CPU / Accelerator	IAX	5.11			BMRA OS includes the kernel support since BMRA 21.09 release.
	QAT	5.11			Supported and tested. Also validated as part of the NGINX workload since BMRA 21.09 release.
	DLB	N/A			Available as userspace library in DPDK since BMRA 21.09 release. DLB is not up-streamed in a Linux kernel yet, drivers available from 01.org.

¹³ See backup for workloads and configurations or visit www.intel.com/PerformanceIndex. Results may vary.

CATEGORY	FEATURE	KERNEL / OS	DPDK	KUBERNETES	BMRA 22.05 STATUS
	DSA	5.14			DSA supported and tested, including support for the DSA operator since BMRA 21.09 release.
Power Management	SST-PP, SST-TF SST-BF, SST-CP	5.3			SST-BF and SST-PP were available in previous generation. New SST-CP and SST-TF are supported and tested since BMRA 21.09 release.
Security	SGX	5.11			Supported and tested, including the SGX device plugin since BMRA 21.09 release.
	Cryptodev and CryptoNI	N/A			Supported and tested through DPDK 21.11 since BMRA 22.01 release. Not supported in BMRA 22.05.
RAS	RAS	5.11			collectd and Telegraf include RAS plugins since BMRA 21.09 release.
ISA	FP-16 (5G ISA)	5.11			BMRA OS includes the kernel support since BMRA 21.09 release.
	AMX (TMUL)	5.16			Not yet supported.
	VP2INTERSECT	5.4			BMRA OS includes the kernel support since BMRA 21.09 release.
	AIA (MOVDIRI, Power Instrs.)	5.10			Supported and tested as part of the DPDK 21.08 release since BMRA 21.09 release.
I/O	CXL 1.1	5.11			Supported but not tested as part of the DPDK 21.08 release and since RA 21.09 release.
	PCI Gen5	5.3			BMRA OS includes the kernel support since RA 21.09 release.
Virtualization	SIOV	N/A			BMRA OS includes the kernel support since RA 21.09 release.
	SVM	N/A			Not yet supported.

Refer to [Table 69](#) and [Table 70](#) for other features of 4th Generation Intel Xeon Scalable processor enabled in prior BMRA releases.

A.2 BMRA 22.05 Bug Fixes

The following bug fixes were completed in the BMRA 22.05 release:

- Settings related to Intel P-state and Intel Turbo Boost Technology were not enforced until the target was rebooted, leading certain tasks to not execute properly (see [Issue 74](#))
- Collectd pods show in CrashLoopBack. Fixed by restoring the `enable_pkgpower_plugin` var to allow disabling collectd-specific plug-ins with the flag that existed in 21.09 (see [Issue 82](#))

A.3 BMRA 22.01 Notable Facts

New Components:

- Intel Service Mesh
- CNDP (Cloud Native Data Plane)
- MinIO Object Storage
- Intel Power Manager (Power Operator)
- Platform Aware Scheduling (Telemetry Aware Scheduling + GPU Aware Scheduling)
- Intel DLB (Dynamic Load Balancer)

New Platforms:

- Taylors Falls Reference Design (Intel Xeon-D)

Updates/Changes:

- Playbooks and profile config files are generated automatically
- Profiles list expanded with 'Storage'
- RHEL bumped to version 8.5 as base operating system
- Version upgrades for key components:
 - DPDK = 21.11
 - Kubernetes = 1.22
 - Kubespray = 2.17
 - KMRA = 1.4

Removed Support:

- Intel CPU Manager for Kubernetes (CMK)
- CentOS (all distro versions) as base operating system

Known Limitations (Errata):

- Settings related to Intel P-state and Intel Turbo Boost Technology are not enforced until the target is rebooted, leading certain tasks to not execute properly (see: [Issue 74](#))

The following table lists key features of the 4th Generation Intel Xeon Scalable processor and the support for those features in BMRA 22.01.

Table 69. Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 22.01

CATEGORY	FEATURE	KERNEL / OS	DPDK	KUBERNETES	BMRA 22.01 STATUS
Security	CryptoDev and CryptoNI		21.11		Supported and tested through DPDK21.11
Virtualization	SVM in QAT				Supported and tested through out-of-tree driver version - QAT.L.4.16.0-00017

Refer to [Table 70](#) for other features of 4th Generation Intel Xeon Scalable processor enabled in BMRA 21.09.

A.4 BMRA 22.01 Bug Fixes

The following bug fixes were completed in the BMRA 22.01 release:

- Missing makefile: [Issue 72](#)
- Installing dependencies on localhost: [Issue 73](#)
- Duplicate code within the Profile specific playbooks: [Issue 78](#)
- Using role/vars/main.yml to set defaults makes changing almost impossible: [Issue 79](#)
- SRIOV-CNI build fails on Ubuntu20: [Issue 80](#)

A.5 BMRA 21.09 New Features

The following new features were updated or added in the BMRA 21.09 release:

- Support for Istio service mesh operator, Envoy, and control plane
- Support for Telegraf telemetry collection
- Support Intel Telemetry Insight Reports
- Support for additional container runtime: CRI-O
- Updated default network plugin: Calico
- Support Intel® device plugins operator (Intel® Data Streaming Accelerator (Intel® DSA)) for 4th Generation Intel® Xeon® Scalable processor
- Support Intel® Speed Select Technology - Performance Profile (Intel® SST-PP)
- Support for rendering profile config files from template
- Updated Intel® Ethernet 700 and 800 Network Adapter drivers
- Updated Intel® Software Guard Extensions (Intel® SGX) Software Development Kit (SDK)
- Updated Data Plane Development Kit (DPDK) and Open vSwitch (OVS) DPDK for use of AVX-512 instruction sets
- Updated Prometheus, Grafana, and Node Exporter telemetry packages
- Updated Node Feature Discovery (NFD)
- Updated Multus container network interface (CNI)
- Updated OpenSSL toolkit
- Updated Intel® QuickAssist Technology Engine for OpenSSL (Intel® QAT Engine for OpenSSL)
- Updated Intel® Multi-Buffer Crypto for IPSec (`intel-ipsec-mb`)

The following table lists key features of the 4th Generation Intel Xeon Scalable processor and the support for those features in BMRA 21.09. The versions of kernel/OS, DPDK, and Kubernetes in the table indicate the first ever versions to be enabled with the corresponding features that are part of this release.

Table 70. Status of Support for Key Features of 4th Generation Intel Xeon Scalable Processor in BMRA 21.09

CATEGORY	FEATURE	KERNEL / OS	DPDK	KUBERNETES	BMRA 21.09 STATUS
CPU/Accelerator	IAX	5.11 / Ubuntu 21.04			BMRA OS includes the kernel support

CATEGORY	FEATURE	KERNEL / OS	DPDK	KUBERNETES	BMRA 21.09 STATUS
	QAT	5.11 / Ubuntu 21.04, RHEL 8.4			Supported and tested. Also validated as part of the NGINX workload.
	DLB		20.11		Available as userspace library in DPDK
	DSA	5.11 / Ubuntu 21.04, RHEL 8.4		Version 1.19 – 1.21	DSA supported and tested, including support for the DSA operator
	Intel SST-PP / Intel SST-TF	5.3 / Ubuntu 20.04, RHEL 8.3, RHEL 8.4	21.05		SST-PP and SST-TF are supported and tested
Security	SGX	5.9 / Ubuntu 21.04, RHEL 8.4		Version 1.19 – 1.21	Supported and tested, including the SGX device plugin
RAS	RAS			Version 1.19 – 1.21	collectd and Telegraf include RAS plugins
ISA	FP-16 (5G ISA)	5.11 / Ubuntu 21.04, RHEL 8.4			BMRA OS includes the kernel support
	AMX (TMUL)	-	-	-	<i>Not enabled in this release</i>
	VP2INTERSECT	5.4 / Ubuntu 20.04, RHEL 8.4			BMRA OS includes the kernel support
	AIA (MOVDIRI)	5.9 / Ubuntu 21.04, RHEL 8.4	20.11		Supported and tested as part of the DPDK 21.08 release
	AIA (Power Instructions)	5.10 / Ubuntu 20.04, RHEL 8.4	21.02		Supported and tested as part of the DPDK 21.08 release
I/O	CXL 1.1	5.11 / Ubuntu 21.04			BMRA OS includes the kernel support
	PCIe Gen5	5.3 / Ubuntu 20.04, Ubuntu 21.04, RHEL 8.4			BMRA OS includes the kernel support
Virtualization	S-IOV	-	-	-	<i>Not enabled in this release</i>
	SVM	-	-	-	<i>Not enabled in this release</i>

A.6 BMRA 21.09 Bug Fixes

The following bug fixes were completed in the BMRA 21.09 release:

- Fixed inventory groups for inclusive terminology
- Fixed kubelet –cpu-cfs-quota to eliminate performance throttling
- Fixed QAT driver VF binding issue on RHEL 8.4
- Fixed inadvertent Intel SST-CP frequency throttling with proportional settings

A.7 BMRA 21.08 New Features

The following new features were updated or added in the BMRA 21.08 release:

- Support 4th Generation Intel® Xeon® Scalable processor (known as Sapphire Rapids - NDA)
- Support Intel® Xeon® D processor LCC and HCC (NDA)
- Updated Intel® Ethernet 700 and 800 Network Adapter drivers
- Updated Intel® Ethernet 800 Dynamic Device Personalization (DDP) profiles
- Support Intel® QuickAssist Technology (Intel® QAT) drivers for 4th generation processors (NDA)
- Updated Intel device plugins (Intel QAT, Intel® Software Guard Extensions (Intel® SGX), Intel® Server GPU)
- Support additional operating system versions: RHEL 8.4 and Ubuntu 21.04
- Support additional container runtime: containerd
- Support Kubernetes version 1.21
- Updated Kubernetes features: Node Feature Discovery (NFD), Telemetry Aware Scheduling (TAS), and SR-IOV device plugin (DP)
- Support Kubernetes Operators for: Intel® device plugin operator (Intel SGX, Intel Server GPU) and SR-IOV network
- Support Intel® QuickAssist Technology Engine for OpenSSL (Intel® QAT Engine for OpenSSL)
- Support Intel® QAT Engine for OpenSSL on 4th generation processors (NDA)
- Support containerized Intel® SGX Key Management Service (KMS) including integration of Key Management Reference Application (KMRA) version 1.2.1
- Updated collectd, Prometheus, and Grafana components
- Support for DPDK 21.05 and OVS 2.15
- Support Ansible Cluster Removal Playbook for cluster teardown and redeployment

A.8 BMRA 21.08 Bug Fixes

The following bug fixes were completed in the BMRA 21.08 release:

- Fixed cluster recovery errors on reboot
- Fixed TAS demo policy failure
- Fixed deployment failure with Intel® Turbo Boost Technology disabled
- Fixed SGX DP pod crashes
- Fixed mismatch in available Intel QAT resources
- Fixed deployment failure with missing Intel QAT configuration
- Fixed kernel header mismatch for compiled kernel modules
- Fixed package installation dependencies
- Fixed isolcpu generation for configurations with HT disabled
- Fixed DPDK installation failures
- Fixed DNS file permission issues
- Fixed IP dependency in inventory file

A.9 Known Issues

Issue:

MinIO tenant pods lose additional network connections after about 24 hours. BMRA 22.05 uses SR-IOV VFs for the second network interface.

Detail:

With `minio_vf: true` in the `host_vars/dataplane_interfaces` variable to attach additional SR-IOV VF connections on MinIO tenants pods, current Kubernetes cannot create the service that utilizes Multus second network connection as noted in [Issue 466](#). The second interface is an additional interface to Kubernetes and Kubernetes does not recognize the second interface that Multus creates.

Workaround:

After losing the second interface, BMRA 22.05 leaves the endpoints patch file in the controller node at `/opt/cek/charts/tenant/temp/minio-tenant-endpoints-patch.yml`. Use `kubectl patch endpoints minio-tenant-hl -n minio-tenant --patch-file /opt/cek/charts/tenant/temp/minio-tenant-endpoints-patch.yml` to bring back the lost second network connection. This assumes that BMRA 22.05 is deployed with default namespace settings.

Issue:

Support for deploying the `regional_dc` Configuration Profile on worker nodes with SG1 cards is limited.

Detail:

SG1 cards require the GPU kernel, which is not maintained for the newer Linux releases (Ubuntu 21.10, 22.04, and RHEL 8.5).

Workaround:

Select Ubuntu 20.04.4 when using SG1 cards.

Part 5: Abbreviations

Appendix B Abbreviations

The following abbreviations are used in this document.

ABBREVIATION	DESCRIPTION
AGF	Access Gateway Function
AI	Artificial Intelligence
AIC	Add In Card
AIA	Accelerator Interfacing Architecture
AMX	Advance Matrix Multiply
API	Application Programming Interface
BIOS	Basic Input/Output System
BKC	Best Known Configuration
BMRA	Bare Metal Reference Architecture
BOM	Bill of Material
CA	Certificate Authority
CDN	Content Delivery Network
CLOS	Class of Service
CMK	CPU Manager for Kubernetes
CMTS	Cable Modem Termination System
CNCF	Cloud Native Computing Foundation
CNDP	Cloud Native Data Plane (CNDP)
CNI	Container Network Interface
CO	Central Office
CTK	Crypto-API Toolkit
CU	Central Unit
CXL	Compute Express Link
DDP	Dynamic Device Personalization
DHCP	Dynamic Host Configuration Protocol
DLB	Intel® Dynamic Load Balancer (Intel® DLB)
DNS	Domain Name Service
DPDK	Data Plane Development Kit
DRAM	Dynamic Random Access Memory
DSA	Intel® Data Streaming Accelerator (Intel® DSA)
DU	Distribution Unit
EIST	Enhanced Intel SpeedStep® Technology
FPGA	Field-Programmable Gate Array
FW	Firmware
GAS	GPU Aware Scheduling
GPU	Graphics Processor Unit
HA	High Availability
HCC	High Core Count
HSM	Hardware Security Model
HT	Hyper Threading
IAX	In-Memory Analytics
IMC	Integrated Memory Controller
Intel® AVX	Intel® Advanced Vector Extensions (Intel® AVX)

ABBREVIATION	DESCRIPTION
Intel® AVX-512	Intel® Advanced Vector Extension 512 (Intel® AVX-512)
Intel® DCAP	Intel® Software Guard Extensions Data Center Attestation Primitives (Intel® SGX DCAP)
Intel® DLB	Intel® Dynamic Load Balancer (Intel® DLB)
Intel® DSA	Intel® Data Streaming Accelerator (Intel® DSA)
Intel® HT Technology	Intel® Hyper-Threading Technology (Intel® HT Technology)
Intel® QAT	Intel® QuickAssist Technology (Intel® QAT)
Intel® RDT	Intel® Resource Director Technology (Intel® RDT)
Intel® SecL – DC	Intel® Security Libraries for Data Center (Intel® SecL – DC)
Intel® SGX	Intel® Software Guard Extensions (Intel® SGX)
Intel® SST-BF	Intel® Speed Select Technology – Base Frequency (Intel® SST-BF)
Intel® SST-CP	Intel® Speed Select Technology – Core Power (Intel® SST-CP)
Intel® SST-PP	Intel® Speed Select Technology – Performance Profile (Intel® SST-PP)
Intel® SST-TF	Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF)
Intel® VT-d	Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-d)
Intel® VT-x	Intel® Virtualization Technology (Intel® VT) for IA-32, Intel® 64 and Intel® Architecture (Intel® VT-x)
IOMMU	Input/Output Memory Management Unit
IoT	Internet of Things
ISA	Instruction Set Architecture
I/O	Input/Output
K8s	Kubernetes
KMRA	Key Management Reference Application (KMRA)
KMS	Key Management Service (KMS)
LCC	Low Core Count
LLC	Last Level Cache
LOM	LAN on Motherboard
MEC	Multi-Access Edge Compute
NFD	Node Feature Discovery
NFV	Network Function Virtualization
NIC	Network Interface Card
NTP	Network Time Protocol
NUMA	Non-Uniform Memory Access
NVM/NVMe	Non-Volatile Memory
OAM	Operation, Administration, and Management
OCI	Open Container Initiative
OS	Operating System
OVS	Open vSwitch
OVS DPDK	Open vSwitch with DPDK
PBF	Priority Based Frequency
PCCS	Provisioning Certification Caching Service
PCI	Physical Network Interface
PCIe	Peripheral Component Interconnect Express
PF	Port Forwarding
PMD	Poll Mode Driver
PMU	Power Management Unit
PXE	Preboot Execution Environment
QAT	Intel® QuickAssist Technology

ABBREVIATION	DESCRIPTION
QoS	Quality of Service
RA	Reference Architecture
RAS	Reliability, Availability, and Serviceability
RDT	Intel® Resource Director Technology
RHEL	Red Hat Enterprise Linux
S3	Amazon Web Services Simple Storage Service
S-IOV	Intel® Scalable I/O Virtualization (Intel® Scalable IOV)
SA	Service Assurance
SGX	Intel® Software Guard Extensions (Intel® SGX)
SR-IOV	Single Root Input/Output Virtualization
SSD	Solid State Drive
SSH	Secure Shell Protocol
SVM	Shared Virtual Memory
TADK	Traffic Analytics Development Kit
TAS	Telemetry Aware Scheduling
TCA	Trusted Certificate Attestation
TCS	Intel® Trusted Certificate Service
TDP	Thermal Design Power
TLS	Transport Layer Security
TME	Total Memory Encryption
TMUL	Tile Multiply
UEFI	Unified Extensible Firmware Interface
UPF	User Plane Function
vBNG	Virtual Broadband Network Gateway
vCDN	Virtualized Content Delivery Network
vCMTS	Virtual Cable Modem Termination System
VF	Virtual Function
VMRA	Virtual Machine Reference Architecture
VPP	Vector Packet Processing
vRAN	Virtual Radio Access Network
WAF	Web Application Firewall



Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](https://www.intel.com/PerformanceIndex).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.