Solution Brief

Automated Onboarding
IoT and Edge Computing

Red Hat and Intel Streamline Edge Deployment and Provisioning with FDO

As IoT devices proliferate and edge servers are increasingly needed to deliver low latency connectivity, the FIDO Alliance has introduced a specification for automated, secure IoT provisioning technology: FIDO Device Onboard (FDO).

> “Edge computing doesn't compete with cloud computing but will complement and complete it. Edge computing is part of a distributed computing topology in which information processing is located close to the physical location where things and people connect with the networked digital world”[1]
>
> – Bob Gill,
> Gartner Analyst
> 2021 Strategic Roadmap for Edge Computing

Automating Edge Provisioning

Red Hat and Intel have co-engineered an automated deployment and provisioning solution, based on the FIDO Device Onboard (FDO) protocol, to meet enterprise requirements and simplify the installation and maintenance of open hybrid clouds. The vision of the open hybrid cloud builds upon a common infrastructure—from core to edge—that can handle diverse workloads distributed across an extended architecture, encompassing operations technology (OT), communication technology, and information technology (IT).

At the network edge, the machine interface meets the real world. In this environment, sensors and actuators connected to IoT gateways and other endpoints communicate using Internet protocols to link a computing cluster with a local network. Larger processing nodes run algorithms that require low-latency or high-bandwidth communications; these nodes can also interact with cloud-based computing and storage resources. Linux is the predominant operating system for enabling edge computing and Red Hat® Enterprise Linux® delivers the latest Linux features and capabilities backed by an outstanding support team.

Deploying, provisioning, configuring, and maintaining the servers that underlie edge computing installations has largely been accomplished using manual techniques. Typically, a skilled technician is needed at the edge site to carry out the time-consuming and costly procedures, and these technicians are often in short supply at edge site facilities. Even the best technician, however, can make mistakes. Automating processes lessens the chances of adding vulnerabilities or errors during critical installations.

As adoption of edge computing increases and supporting technologies gain momentum across multiple industries, a more efficient approach to provisioning can be implemented with the FDO framework.

Intel Management and Deployment Mechanisms for the Edge

Intel’s management technologies enhance the edge experience by providing platform-specific features to support edge computing, workload consolidation, and automated onboarding. Automated onboarding with the FDO framework from Intel interoperates effectively with Red Hat platforms to streamline installation and onboarding of edge servers.

What is FIDO Device Onboard?

FIDO is an open, standards-based authentication protocol developed by the FIDO Alliance, an industry association forged to strengthen authentication methods and reduce reliance on passwords. In collaboration with industry leaders, including Intel and Red Hat, the FIDO Alliance has introduced technical specifications defining open, scalable, interoperable mechanisms that simplify authentication using public key cryptography and protect user privacy.

These mechanisms include the creation of a digital proof of ownership known as the Ownership Voucher (represented by the white and red keys in Figure 1, which allows a digital transfer of ownership within the supply chain and securely identifies the device ownership during the onboarding process within the targeted device management system with FDO.

In the case of the FDO solution from Red Hat and Intel, this authentication process is used to protect the security and integrity of software images distributed to edge devices. Learn more about the FIDO Device Onboard specification by visiting this FIDO Alliance site.

Automated Provisioning Process

To streamline and simplify the deployment for high-volume scalability in edge computing environments, a device provisioning process is needed that supports:

- Automation – Delivering faster, parallel provisioning of devices with minimal exposure of passwords
- Security – Ensuring each device comes up reliably on the correct network

With the FDO solution from Red Hat and Intel, the device manufacturer creates a single SKU for edge deployments and users employ Red Hat tools to create a customized Red Hat Enterprise Linux image or other image for the edge deployment. FDO lets the user's image be installed automatically and securely in each device as the edge deployment proceeds.

Specifically, the manufacturer-initialized device boots into the Red Hat Enterprise Linux for Edge operating system and automatically enables the FDO agent to use embedded credentials. Then the manufacturer enables all device hardware and software security measures. The device is shipped to the customer along with an FDO credential, known as the Ownership Voucher. The Ownership Voucher informs the user's Device Management System about the new device.

When the device is booted, the device's OS invokes the FDO agent, and it automatically uses the embedded credentials and the Ownership Voucher to discover the user's Device Management System and authenticate securely. Then the Device Management System can instruct the device to download the user's selected target image securely, augmenting or replacing the factory-installed image. In this way, the user adds application software and can augment and repurpose the OS from the manufacturer or can download a new OS image. Red Hat provides tools to allow users to create custom OS images for such a purpose.

This solution from Red Hat and Intel automates the provision operations for large deployments, without requiring users and manufacturers to commit to custom device images for small manufacturing runs. Figure 1 shows the sequence followed for this provisioning solution.

An added benefit is that the user's software and OS are downloaded into the device on a "just in time" (JIT) basis as the device is onboarded. This gives the user fine-grained control over the versioning of the packages that are downloaded, and avoids a situation where a system is shelved "cold-spare" and installed with vulnerabilities present. FDO also ensures that security keys or tokens are allocated per device, and just in time to ensure that security is as tight as possible as the device comes up.

> "The open hybrid cloud isn't limited to an enterprise data center or public cloud environments; it includes the remote servers, advanced machinery, and other devices that exist on the far reaches of the enterprise network. The disparate nature of these footprints means that consistency is critical to success – Red Hat Enterprise Linux, as the backbone for the Red Hat Edge initiative, provides the consistent edge-native and intelligent platform to meet the dynamic demands of the hybrid cloud, from bare-metal servers to the cloud to the edge."[2]
>
> — Stefanie Chiras, Senior Vice President and General Manager, Red Hat Enterprise Linux Business Unit, Red Hat

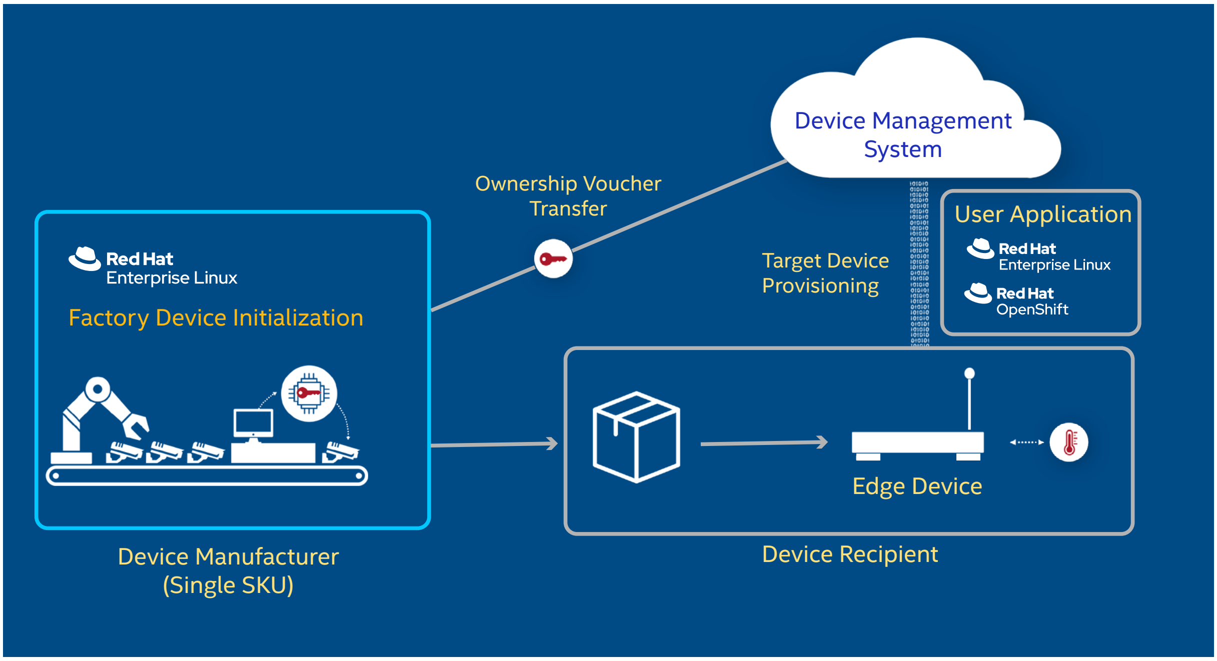

Figure 1. Establishing digital proof of ownership

Factory Device Preparation with FDO

Intel's FDO toolkit includes tools that allow manufacturers to automate creating FDO-based SKUs. The manufacturer starts by choosing a lightweight OS, for example, using tools provided by Red Hat to create a small ISO. Then the Intel Edge System Provisioning (ESP) and FDO tools are used to create manufacturing stations that initialize individual devices. ESP includes tools that use the ISO image as input and create a bootstrap server for a PXE boot. Existing ESP profiles include FDO components so that manufacturers can create a station that PXE boots, then installs a lightweight OS with FDO onto a storage partition, and stores the FDO Ownership Voucher. Manufacturers must also secure the BIOS and enable secure boot on target systems.

In a typical manufacturing situation, multiple target devices are connected to a manufacturing station over a low-cost, high-speed Ethernet link (1000BASE-T or higher). A single manufacturing station can image multiple target devices simultaneously, and multiple manufacturing stations can be replicated in parallel, uploading FDO credentials to a common database. See Figure 2.

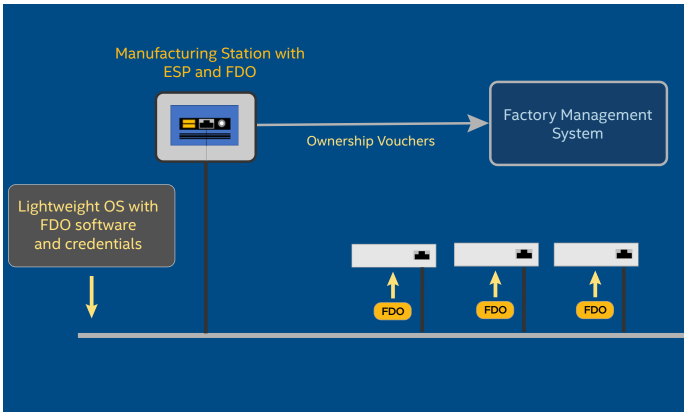
Manufacturing Station with ESP and FDO
Ownership Vouchers
Factory Management System
Lightweight OS with FDO software and credentials
FDO
FDO
FDO

Figure 2. ESP Manufacturing environment to initialize FDO with a lightweight OS

In typical factory installations, the ESP interfaces with a factory management system (FMS) at the target site to track inventory. This system can also automate storing and transmitting the Ownership Voucher along with the target system.

Target Device Provisioning with FDO

FDO permits secure onboarding to place an edge device under control of a Device Management System. Just-in-time OS or package installation under FDO can also ensure that the latest and most secure support software is running on the device.

The reference implementation from Intel supports OS downloads using Red Hat® OpenShift® discovery images and Red Hat Enterprise Linux images. FDO interfaces with the DMS using a REST interface, so it is straightforward to connect a virtual machine running FDO into an existing device manager.

As each device's Ownership Voucher is received, the user places it into the FDO software and chooses the image to download on the target device. Then the device is automatically initialized when it is first powered on.

FDO works with DMS systems deployed in public clouds, private clouds, and in closed network environments.

Deploying an OpenShift Cluster Using FDO

Deploying nodes into a Kubernetes cluster can be accomplished automatically as well. Red Hat supports an assisted installer mechanism to automate installation of nodes as Red Hat OpenShift worker nodes. The service provides validation and discovery of targeted hardware that improves success rates of installations and is accessible by means of the Red Hat Hybrid Cloud Console.

The installation process is based on an agent mechanism that establishes the communication with a cloud-based, assisted installer service, providing information about the device where the agent is being executed. In existing bare-metal deployments, this agent is loaded into the device by booting a discovery image, which is generated by the assisted installer service during the OpenShift cluster installation.

This process can be automated using FDO. In this case, the DMS chooses the installer service's discovery image as the ISO to download during FDO. When the device first boots, the FDO process downloads and installs the discovery image along with cluster credentials. Then the device reboots and automatically joins the Kubernetes cluster (as shown in Figure 3). This simplifies the OpenShift installation process by removing the need to manually deploy the discovery image on each device part of the cluster. Once devices have joined the cluster, the OpenShift assisted installer service can start the deployment of a new OpenShift Kubernetes cluster or add the device to an existing OpenShift Kubernetes cluster.

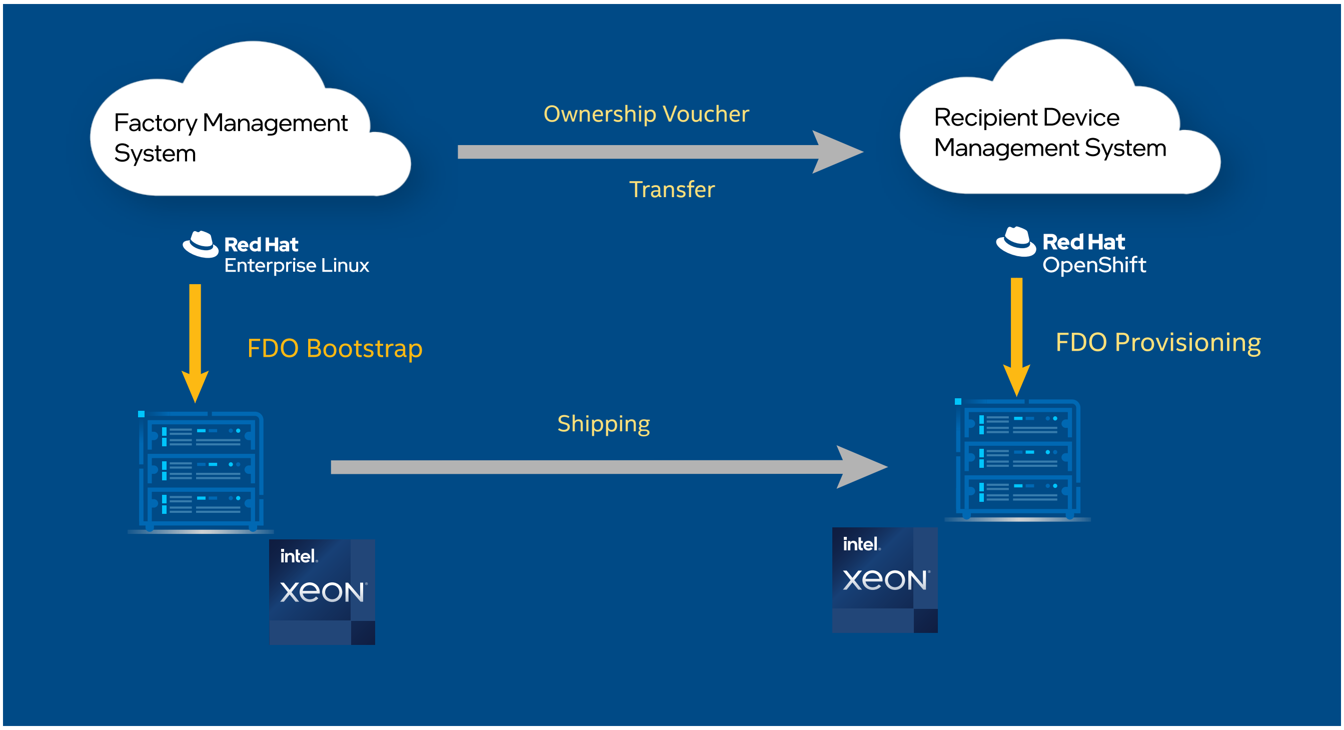

Figure 3. Installation process through an agent mechanism

> "Industry analysts speak to the expansion of edge computing, noting that by 2023, 75% of the data created will be outside the data center. It will be in factories, hospitals, retail stores and cities and driven by many forms of video. Additionally, more than 50% of that data will be processed, stored and analyzed closer to the creation of the data—at the edge—to deliver the right latency, bandwidth, reliability, security and privacy for a wide variety of uses across many markets."[3]
>
> — Tom Lantzsh, Senior Vice President, General Manager, Internet of Things Group, Intel
> — Dan Rodriguez, Corporate Vice President, General Manager, Network Platforms Group, Intel

Using Image Builder

The v8.4 release of Red Hat Enterprise Linux includes an edge-ready technology stack that serves as the foundation for the Red Hat Edge initiative. Image Builder, a tool that creates custom, deployable rpm-ostree images, lets developers, system architects, and network administrators generate packages tuned to the requirements of individual edge applications.

RHEL for Edge images can be deployed for bare metal, appliances, and edge servers, in any of three image types:

- **RHEL for Edge commit** – .tar archive for network deployment, but not directly bootable
- **RHEL for Edge container** – .tar archive for non-network deployments consisting of an OSTree commit embedded in an OCI-compliant container
- **RHEL for Edge installer** – .iso file for performing non-network based deployments where an OSTree commit is extracted based on timestamps from the running container to create an installable boot ISO.

Composing and deploying a RHEL for Edge image takes two steps:

- Creating a RHEL rpm-ostree image using Image Builder, either through the command-line interface in the composer-cli tool or the GUI available in the RHEL 8 Web Console.
- Deploying the image using the RHEL installer.

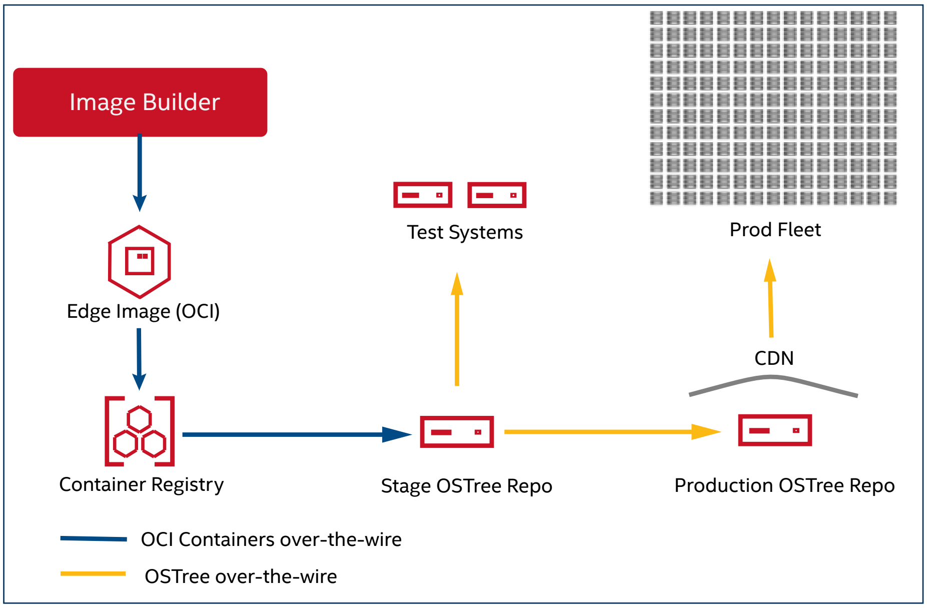

Figure 4. Image Builder operations from image to target

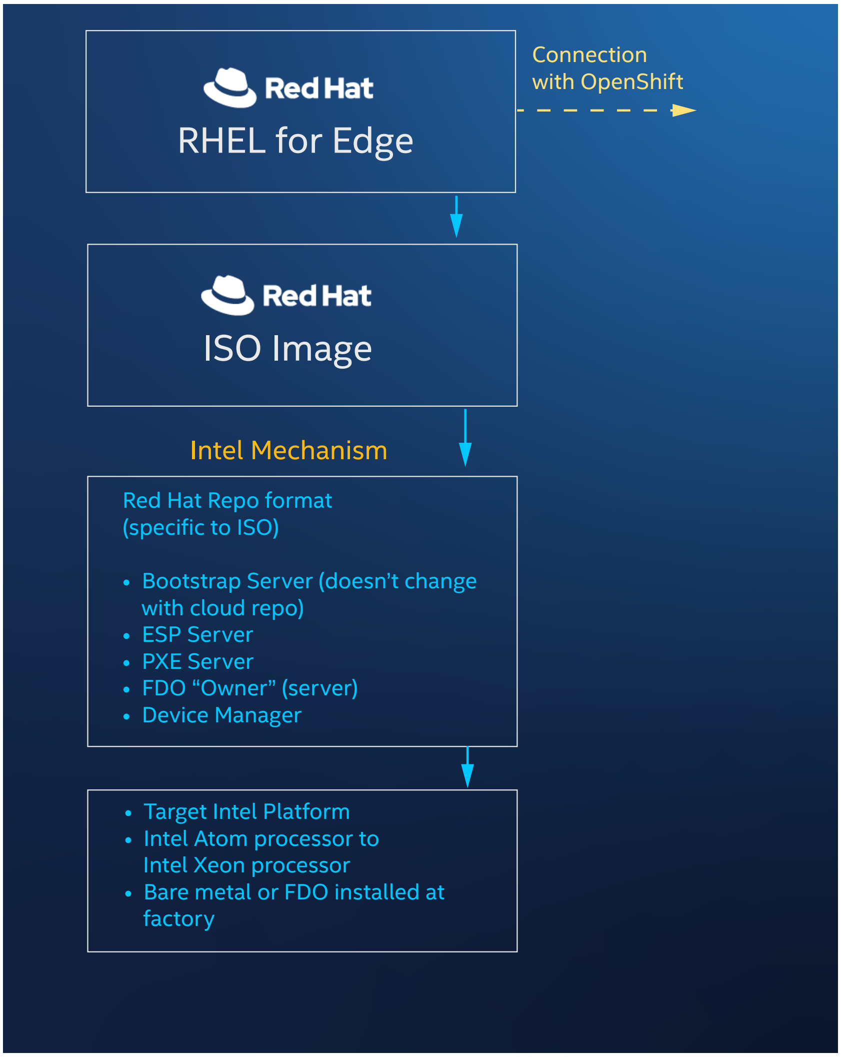
Red Hat
RHEL for Edge
Connection with OpenShift
Red Hat
ISO Image
Intel Mechanism
Red Hat Repo format (specific to ISO)
• Bootstrap Server (doesn't change with cloud repo)
• ESP Server
• PXE Server
• FDO "Owner" (server)
• Device Manager
• Target Intel Platform
• Intel Atom processor to Intel Xeon processor
• Bare metal or FDO installed at factory

Figure 5. Installing the ISO image to the target Intel platform

Quick, Reliable Edge Server Deployment

The Intel ESP and FDO mechanisms work well in combination with the RHEL for Edge Image Builder to generate ISO images for supporting slim and efficient systems. The Intel software helps create a manufacturing station to enable PXE installations to take place quickly. Using ESP, the new installations are initialized with FDO and secured with available mechanisms. For example, BIOS secure boot and SELinux/IMA mechanisms can be enabled.

Once the system is deployed to the target location, FDO installs the selected operating system and enables remote control of the target system. When FDO has completed installations of the operating system and application, full remote control can be used through a DMS platform.

The co-engineered solution from Intel and Red Hat delivers a powerful operating system and well-provisioned environment for deploying and managing edge devices. Legacy devices may also be supported, using ESP in a secure network within the target environment. Non-PXE versions of ESP are also available.

The solution supports adding nodes to legacy installations, targeted IoT platform configuration using RHEL for Edge, and Kubernetes capabilities at the edge through OpenShift. The open-source Edge Software Provisioner (ESP) and FIDO Device Onboard (FDO) solutions from Intel dovetail with Red Hat's installation process, allowing new or reconditioned computers to be installed rapidly with the latest software, automatically secured and connected to a remote Device Management System. The resulting solution provides reliable, quick, and secure installation of Red Hat components at the network edge.

Learn more

Intel Development Resources for IoT Professionals

Create complete, scalable, and optimized IoT solutions and speed your time to market.

Learn more ›

Red Hat's Vision for Edge Computing

Where your edge is depends on your organization, architecture, or use case. Red Hat's approach to edge computing focuses on three categories: Enterprise edge, Operations edge, and Provider Edge.

Learn more ›

About Red Hat

Red Hat is the world's leading provider of enterprise open source solutions, including high-performing Linux, cloud, container, and Kubernetes technologies. Visit redhat.com for more information.

intel. + Red Hat

1. Gill, Bob. *2021 Strategic Roadmap for Edge Computing*. Gartner. November 2020. https://www.gartner.com/en/documents/3992656-2021-strategic-roadmap-for-edge-computing
2. Red Hat Powers the Next Wave of Edge Computing with Latest Version of the World's Leading Enterprise Linux Platform. Red Hat Press Releases. April 2021. https://www.redhat.com/en/about/press-releases/red-hat-powers-next-wave-edge-computing-latest-version-worlds-leading-enterprise-linux-platform
3. Lantzsch, Tom, Dan Rodriguez. *Intel Fuels the Edge Today with Expanded Tech, Customer Deployments.* Intel Newsoom. September 2020. https://www.intel.com/content/www/us/en/newsroom/opinion/intel-fuels-edge.html

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's Global Human Rights Principles. Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

Intel does not control or audit third-party data. You should review this content, consult other sources, and confirm whether referenced data is accurate.

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

No product or component can be absolutely secure.

Your costs and results may vary.

Red Hat, the Red Hat logo, and other Red Hat marks are trademarks of Red Hat, Inc. or its affiliates.

© 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.

1121/RM/MESH/PDF