

QCT* Tests Autoscaling Capability of Intel® Rack Scale Design

In tests conducted by Quanta Cloud Technology* and Intel, the QCT Rackgo R* demonstrates 7-minute instantiation of virtual evolved packet core (vEPC) based on workload demand.¹ The demonstration shows the impact the prototype Intel® Rack Scale Design autoscaling feature can have in building agile networks.



Taking network functions virtualization (NFV) to the next level of deployment means fully utilizing the ability to scale composable resources (compute, storage, network) to match changing traffic levels or workload demands. Intel has developed and is trialing a prototype autoscaling framework for Intel® Rack Scale Design (Intel® RSD) that utilizes Redfish*-compliant application programming interfaces (APIs) to automate service scale up and scale down, including the identification of the right hardware resources required for the application.

Intel and its Intel® Network Builders ecosystem partner Quanta Cloud Technology Inc. (QCT)* have developed and conducted a demonstration utilizing the QCT Rackgo R* cloud system that shows how mobile network operators (MNOs) can deploy virtual evolved packet core (vEPC) instances and then activate them in response to an increase or decrease in workload in a best case of four minutes, but average between 5.5 and 7 minutes.¹

Overview of Intel® Rack Scale Design

The Intel RSD framework is designed to provide efficient, cloud-scale computing resources by enabling resource pooling and composability as well as pod telemetry and maintenance. The Intel RSD framework utilizes Redfish application programming interfaces (APIs) to allow data centers and cloud service providers to create unique configurations of compute, power, storage, and networking into a pod that is designed to meet the requirement of a particular workload. An Intel RSD pod is a rack or multiple racks of compute, storage, and accelerator resources. A pod's resources can be scaled up or down to meet evolving needs of the workload. If the workload goes away, the resources in the pod can be returned to the pool for reassignment or powered down to conserve power.

The Intel RSD pod manager is software and firmware that exposes the hardware in a pod to orchestration software and commands for management and policy enforcement. Pod manager uses a Redfish API for dynamic management and pooling of all hardware components: compute, memory, PCIe,* and storage resources.

Intel RSD is designed with several layers of abstraction, each managed by additional Redfish APIs, including the rack management module (RMM) at the rack level, the pooled system management engine (PSME) at the drawer level, the module management controller (MMC) at the module layer, and the baseboard management controller (BMC) on every blade.

Table of Contents

Overview of Intel® Rack Scale Design 1
Autoscaling Prototype 2
Overview of QCT Rackgo R 3
Intel® RSD Workload Autoscaling Prototype 3
vEPC Demonstration 4
Conclusion 5
About Quanta Cloud Technology (QCT) 6
About Intel® Network Builders... 6

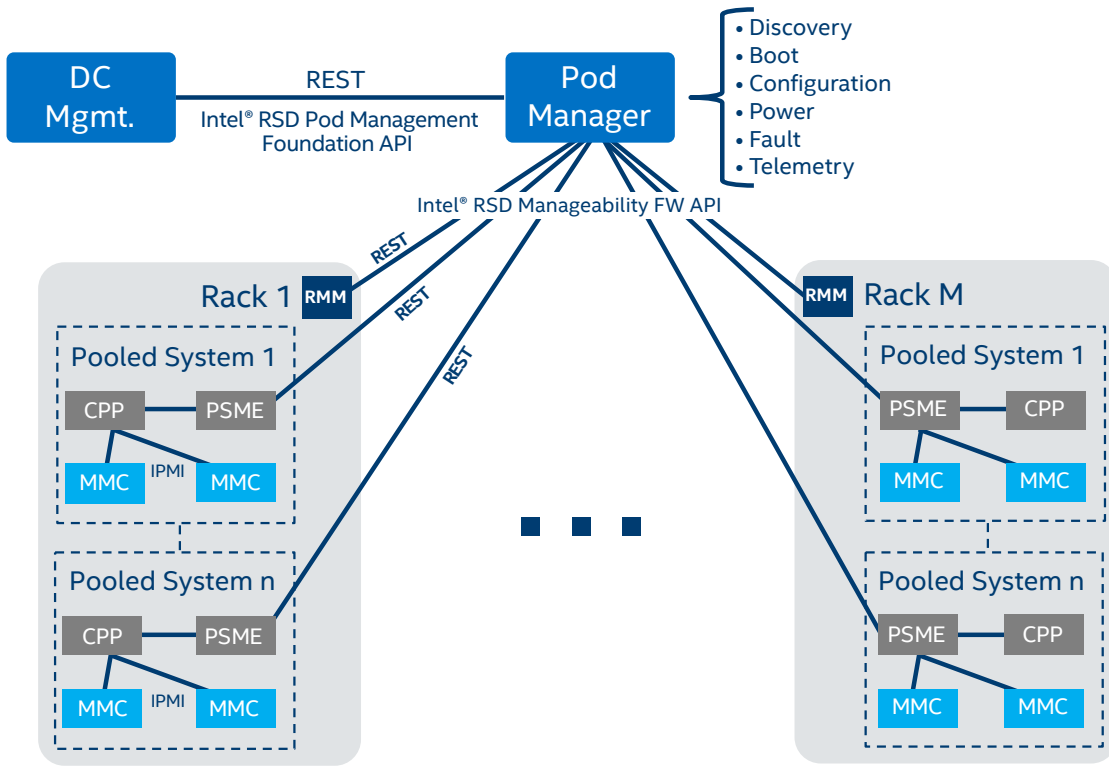


Figure 1. Overview of Intel RSD APIs connecting the pod manager and pooled equipment.

Autoscaling Prototype

The prototype autoscaling function that is described in this document is a response to requests by MNOs and communications service providers (CommSPs) for a faster way to adjust their network resources in response to workload demands. Figure 2 shows how complicated resource scaling can be, with a traditional approach taking 10 hours or more.² Only with the performance delivered by the autoscaling functionality can MNOs and CommSPs build networks that respond dynamically to workload needs.

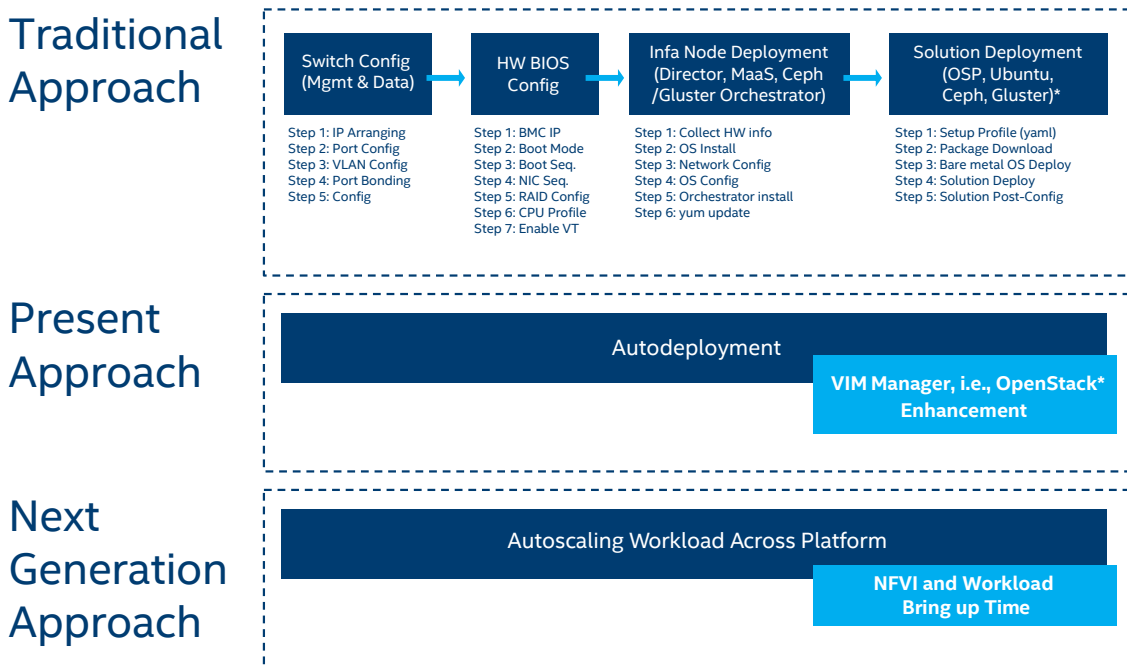


Figure 2. Actions required to deploy a new service required using traditional (manual) approach, autodeployment approach, and the prototype Intel RSD autoscaling approach.

Overview of QCT Rackgo R

QCT's Rackgo R cloud solution, based on Intel RSD, was launched in October 2017. The system leverages Intel RSD with QCT's own innovations and works with pods created using the company's line of compliant hardware products, including:

- PSME Compute Resources: QuantaGrid and QuantaPlex
- Pooled FPGA/GPU Blades: QuantaVault
- Pooled Storage: QuantaVault
- PSME Network: QuantaMesh

QCT Rackgo R features QCT System Manager (QSM) Premium, an optimized rack-level manager that lets IT teams create network mappings on the go, to visualize their infrastructure and redirect resources where they are needed. The network schematic also exposes resource capabilities, highlighting what equipment needs to be upgraded or replaced.

Management of the Rackgo R systems provides all network, compute, and storage resource configurations through a single management screen, simplifying administration. As virtual environments diversify and become more complex, Rackgo R provides a solution to manage heterogeneous virtual environments with its own complete resource composition and control solution.




Figure 3. QCT Rackgo R.³

Intel® RSD Workload Autoscaling Prototype

Intel has developed a prototype autoscaling functionality for the Intel RSD framework that can instantiate new VNFs in response to workload demands. The autoscaling framework is ideal for services that have wide swings in traffic levels and where the cloud, mobile, or communications service provider needs to be able to ramp up capacity or to scale it down to save on power costs or to free up resources for other applications.

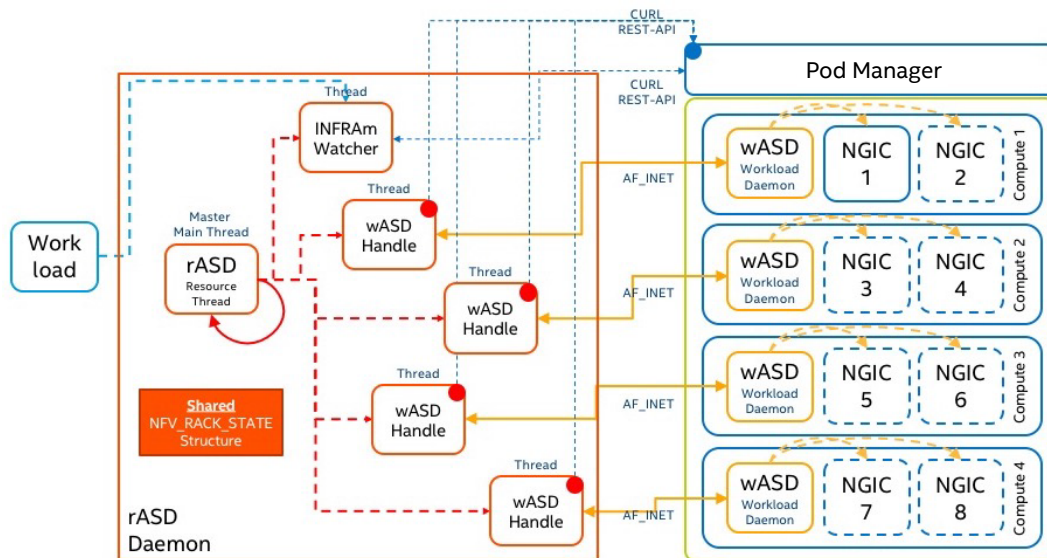


Figure 4. Intel RSD Autoscaling framework.

Figure 4 shows the Intel RSD autoscaling framework. Control functionality resides in the resource autoscale daemon (rASD), a process that runs on a server that has visibility to an entire network of Intel RSD compute nodes and networks pods.

vEPC Demonstration

To demonstrate the autoscaling capability, Intel and QCT set up a demonstration¹ utilizing virtual evolved packet core (vEPC) software. In 4G/LTE networks, the vEPC sits between the radio frequency (RF) network and the Internet protocol (IP) packet network providing:

- Device mobility management
- Serving gateway (SG) to route packets through the access network
- Packet gateway (PGW), an interface between LTE and Internet, managing the IP anchor, enforcing quality of service and charging policies via the policy and charging enforcement function (PCEF)
- Policy and charging rules function (PCRF), which is designed to provide rules for detecting service data flows, policy enforcement and flow-based charging for the PGW/PCEF

The vEPC provides a great workload for this demonstration because it aggregates traffic over a large number of widely distributed cell towers. Each EPC can handle traffic from about 10 million user devices connecting over 5,000 cell towers.⁴ The vEPC infrastructure is subject to high-impact loads as a large number of users connect to the system during typical busy hour periods that vary by the day of the week or with local or natural events that occur in the deployed area.

Figure 5 shows the communication events and triggers between the framework threads to instantiate the compute node and workload. For the purposes of this demonstration, the first node is already up and running. T_c represents the time taken to receive the threshold trigger and start a new compute node through PODm, and T_w is the time taken to trigger the workload daemon to instantiate the VNF.

NG4T* NG40-EPC*

A key component of the vEPC demonstration was the ng40 test system. The ng4T generated data and emulated mobile user equipment (UE) that was used in the demonstration to help trigger the autoscaling functionality. The ng40 test system provided for these tests features fully 3GPP compliant S1, S1-U, SGI, and S11 interfaces, and it simulates UEs, eNBs (LTE macro cells) and LTE small cells, mobile management entity (MME), home subscriber service (HSS), policy and charging rules function (PCRF), and application servers. The ng40 test system provides the following test features and services:

- Functional tests for eNB/UE and another LTE protocol stack and message flow conformance
- Load and capacity to test eRAN and EPC performance, analysis, and benchmarking
- Packet core emulation that offers a practical demonstration of eNBs and UEs for product trainings and exhibitions
- Lightweight network that can be used for application development, IT security, and deep packet inspection.

Workload Autoscaling Across RSD Compute (Automatic Compute Compose/Boot and Workload Instance Deployment)

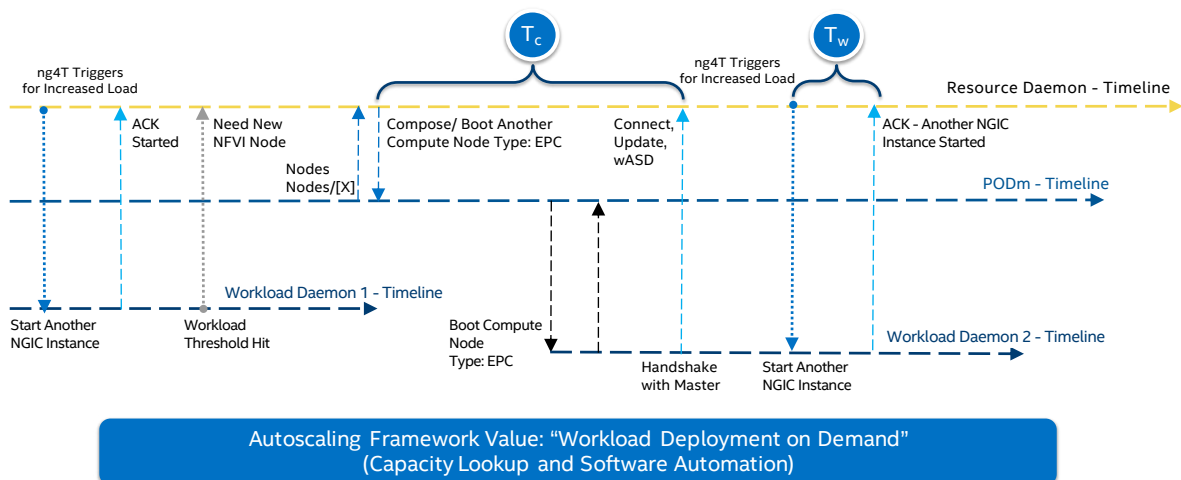


Figure 5. Workload autoscaling for the vEPC demonstration.

The INFRAM Watcher is a thread that polls for composed and assembled nodes and communicates with the pod manager to maintain an inventory of these compute resources. The INFRAM Watcher thread also watches for key workload performance metrics.

When there is an increase in the load and when it crosses the pre-set thresholds, the Watcher instructs the rASD resource thread to notify the current workload autoscale daemon (wASD) handle to instantiate next-generation integrated core (NGIC) on a particular compute node based on its inventory. The rASD resource thread shares the updated events with the wASD handle thread to start new compute node and respective wASD, to launch the OS and then provision the VNF (NGIC) to start the service.

The pod manager will then update the inventory in the INFRAM Watcher for the next change in traffic.

For the demonstration, Intel and QCT utilized a vEPC that meets the Control and User Plane Separation of EPC nodes (CUPS) standard. In this way, the data plane VNF and user plane VNF could be run in their own NGIC VMs.

The test also made use of QCT technology built into its QCT Rackgo R that builds on the functionality in the Intel RSD specification. QCT has enhanced the pod manager and developed its own management user interface that makes it easier to compose, assemble, and deploy resources. The QCT Rackgo R also features optimized BMC, BIOS, and intelligent platform management interface (IPMI) manager that all contribute to very fast deployment times.

In the demonstration, EPC traffic simulated by the ng40 test system is fed into the vEPC. Each NGIC in the system can serve 50,000 users.¹ The traffic levels were varied up and down in order to simulate increase and decrease of attached users. The autoscaling load profile of the system, when fully deployed, handles about 10 Gbps throughput with about 0.8 MPPS packet arrival rate. Load is applied to the system with users attaching at 500 transactions per second (TPS).¹ Figure 6 is a visual depiction of the workload autoscaling demonstration.

Figure 6 shows the systems response and timing to the increased load to give a detailed breakdown of the latency T_c and T_w .

Latency Breakdown

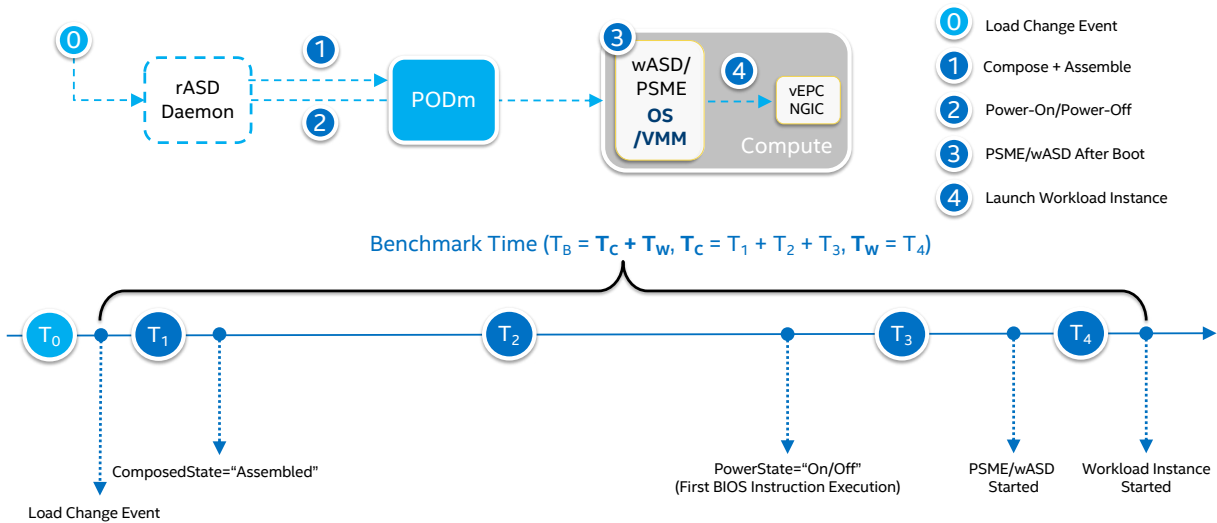


Figure 6. Systems response and timing of vEPC autoscaling demonstration.

The time for each event is as follows:¹

- Time from T_0 to T_1 = 30 seconds
- Time from T_1 to T_2 = ~ 60-90 seconds
- Time from T_2 to T_3 = ~ 90-180 seconds
- Time from T_3 - T_4 = ~ 120 seconds

Total time from load change event to launch workload instance: 330 seconds (5.5 minutes) to 420 seconds (7 minutes).

Conclusion


Autoscaling helps to fulfill the promise of NFV for more agile networks as MNOs expand their use of virtualized environments. This demonstration of the Intel RSD autoscaling framework prototype on a QCT Rackgo R is an example of the automation and flexibility benefits that can apply in a real use case. The use of the autoscaling capability holds promise for emerging mobile network use cases including 5G network slicing, edge cloud, and cloud native services. Future enhancements of the autoscaling functionality will include extending the autoscaling framework functionality to control multiple workloads based on the Intel RSD telemetry, and also utilizing rASD as an extension plugin for OpenStack.*

About Quanta Cloud Technology (QCT)

QCT is a global data center solution provider. It combines the efficiency of hyperscale hardware with infrastructure software from a diversity of industry leaders to solve next-generation data center design and operation challenges. QCT serves cloud service providers, telecoms, and enterprises running public, hybrid, and private clouds. Product lines include hyperconverged and software-defined data center solutions as well as servers, storage, switches, and integrated racks with a diverse ecosystem of hardware component and software partners. QCT designs, manufactures, integrates, and services cutting-edge offerings via its own global network. The parent of QCT is Quanta Computer, Inc., a Fortune Global 500 corporation. <http://www.QCT.io>

About Intel® Network Builders

Intel® Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at <http://networkbuilders.intel.com>.



¹ Testing conducted by Intel. Configurations of device under test (DUT): QCT Rackgo R, Hyper-Threading, Intel® Virtualization Technology (Intel® VT) for Directed I/O (Intel® VT-D) SR-IOV enabled on Intel® Xeon® Gold 6152 processor 22C @2.10 GHz with Intel® Ethernet Converged Network Adapter X710-DA4 NICs (4x10 GbE ports); Host OS Ubuntu® 16.04.3 Linux® version: 4.4.0-87; Guest OS Ubuntu 16.04.3 Linux version: 4.4.0-104. Configurations of ng4T traffic platform: Intel® Xeon® CPU E5-2695 v4 @ 2.10GHz; 16GB Micron® DDR4-2400 RDIMM, 4 channels; 2 DIMMs per channel, 128 GB total.

² Traditional Approach based on testing conducted by QCT. Configurations for Traditional Approach: QxStack Auto-Deployment tool used to facilitate the deployment of QxStack with Red Hat® OpenStack® Platform Director using a QuantaPlex T41S-2U with 2 x Intel Xeon processors E5-2660 v3. See full report, titled "QCT QxStack RedHat OpenStack Platform-RA," version 1.1, at <https://www.qct.io/Download/index/Reference-Architecture>.

³ Image courtesy of QCT.

⁴ See "Understanding the bottlenecks in virtualizing cellular core network functions," IEEE LANAMAN15 at [http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7114735&punumber%3D7113431%26filter%3DAND\(p_IS_Number%3A7114713\)](http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7114735&punumber%3D7113431%26filter%3DAND(p_IS_Number%3A7114713)); "High-performance evolved packet core signaling and bearer processing on general-purpose processors," at <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7113219>; and "Supporting Evolved Packet Core for One Million Mobile Subscribers with Four Intel® Xeon® Processor-Based Servers" at https://networkbuilders.intel.com/docs/MESH_Group_Intel_EPC_TB_FINAL.pdf

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

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

*QCT, the QCT logo, Rackgo, Quanta, and the Quanta logo are trademarks or registered trademarks of Quanta Computer Inc. ng4T, ng40, and the ng4T logo are trademarks of ng4T GmbH in Germany and/or other countries. Ubuntu is a trademark of Canonical Ltd.

* Other names and brands may be claimed as the property of others.

0218/DO/H09/PDF

Please Recycle

337074-001US