



NVIDIA RTX™ SERVER FOR BARE METAL
COMPUTE WITH ALTAIR ULTRAFLUIDX™
2019 DESIGN GUIDE

VERSION: 1.0



TABLE OF CONTENTS

- Chapter 1. SOLUTION OVERVIEW 1**
- 1.1 NVIDIA RTX Server Overview..... 1
- Chapter 2. SOLUTION DETAILS 2**
- 2.1 VALIDATED SERVER Configurations 3

Chapter 1.

SOLUTION OVERVIEW

Designed and tested through multi-vendor cooperation between NVIDIA and its system and ISV partners, NVIDIA RTX™ Server provides a trusted environment for Simulation Engineers/Aerodynamicists/Analysts to run overnight simulations for the Manufacturing & Design, Architecture, Engineering & Construction industries.

1.1 NVIDIA RTX SERVER OVERVIEW

Introduction:

Product Designs are increasingly complex making design validation and verification difficult with engineering simulations. Simulation models for such designs tend to be big as they need to be accurate. This leads to cost-prohibitive, time-consuming simulations. To perform more simulations in less time and to try out many product design ideas, one needs a powerful computer hardware to increase the engineering productivity by accelerating the engineering simulation throughput. NVIDIA RTX Server is a validated solution that brings GPU-accelerated power and performance to deliver the most efficient end-to-end design simulation solution, from interactive sessions in the desktop to batch simulation in the data center.

Audience:

The audience for this document include, but not limited to: Sales Engineers, Field Consultants, Professional Services, Partner Engineers, IT Managers and Customers who wish to take advantage of an appliance that is built and optimized to deliver on simulation workflows.

Chapter 2.

SOLUTION DETAILS

NVIDIA RTX Server is a highly flexible reference design that combines NVIDIA Quadro RTX 6000 and 8000 GPUs with NVIDIA virtual GPU software and leading third-party applications to deliver exceptional compute power in a validated system.

NVIDIA RTX Server for Bare Metal Compute with Altair ultraFluidX™ 2019 is a reference design comprised of (a) NVIDIA Quadro RTX 8000 graphics cards; (b) Altair ultraFluidX 2019 CFD solver software; and a (c) Qualified OEM server system. Combined, this validated solution provides unprecedented rendering and compute performance at a fraction of the cost, space, and power consumption of traditional CPU-based simulation nodes.

NVIDIA Quadro RTX 8000, powered by the NVIDIA Turing™ architecture and the NVIDIA RTX platform, brings the most significant advancement in computer graphics in over a decade to professional workflows. Not only can Designers and artists now wield the power of hardware-accelerated ray tracing, deep learning, and advanced shading to dramatically boost productivity and create amazing content faster than ever before, Simulation Engineers, and Aerodynamicists can benefit from fast performance single precision GPU computing.

Altair ultraFluidX is a simulation tool for ultra-fast prediction of the aerodynamic properties of passenger and heavy-duty vehicles as well as for the evaluation of building and environmental aerodynamics. Its cutting-edge technology is optimized for GPUs to deliver unbeatable performance and to allow for overnight simulations of complex cases on a single server. With ultraFluidX, the Lattice Boltzmann Method is a perfect fit for massively parallel architectures like GPUs and sets the stage for unprecedented turnaround times. Overnight runs on single servers become possible by utilizing state-of-the-art GPU optimized algorithms, while delivering the fidelity of a transient LES aerodynamics simulation. Conventional simulation approaches need thousands of CPU cores to achieve the turnaround times of ultraFluidX. GPU-based solution increases throughput while reducing hardware and energy cost.

RTX Servers, built by our OEM Partners, undergo NVIDIA’s Qualification test suite. Among systems that qualify as an RTX Server there is a subset that has gone through additional testing and validation for Altair ultraFluidX workload. These RTX Server Validated systems capture best practices from NVIDIA and its ecosystem partners.

Configurations for the Validated RTX Servers are listed in the below segment.

2.1 VALIDATED SERVER CONFIGURATIONS

Table 1 outlines the servers utilized to complete the NVIDIA RTX Server validation process.

Table 1: Validated Server Configurations

Server Model	Graphics	Configuration
<u>TYAN Thunder HX FT77D-B7109</u>	<p>8x Quadro RTX 8000</p> <p>4x Quadro RTX NVLink HB bridge 2-slot</p> <p>Quadro Driver Release 430 U2 (430.14)</p>	<ul style="list-style-type: none"> • Dual Intel® Xeon® Gold 6126 processor: 2.6-3.7GHz; 12 Cores, 24 Threads • 384 GB ECC DDR4 memory • 1TB M.2 PCIe NVMe SSD • RHEL 7.6

Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, and DGX are trademarks or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2019 NVIDIA Corporation. All rights reserved.