



NVIDIA EGX™ SERVER FOR BARE METAL
COMPUTE WITH ALTAIR NANOFLUIDX™
2019 DESIGN GUIDE

VERSION: 1.1



TABLE OF CONTENTS

- Chapter 1. SOLUTION OVERVIEW 1**
- 1.1 NVIDIA EGX 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 EGX™ Server provides a trusted environment for Simulation Engineers/Analysts to run overnight simulations for the Manufacturing & Design, Architecture, Engineering & Construction industries.

1.1 NVIDIA EGX 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 EGX 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 EGX Server is a highly configurable, partner-validated reference design that combines NVIDIA Quadro RTX™ A6000, A40, RTX™ 8000, RTX™ 6000 GPUs with NVIDIA virtual GPU software and leading third-party applications to deliver exceptional compute power in a validated system.

NVIDIA EGX Server for Bare Metal Compute with Altair nanoFluidX™ 2019 is a reference design comprised of (a) NVIDIA Quadro RTX A6000, A40, RTX 8000 or RTX 6000 graphics cards; (b) Altair nanoFluidX 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.

Built on the NVIDIA Ampere™ architecture, the NVIDIA® Quadro RTX™ A6000 and NVIDIA® A40, combine 48GB of graphics memory with the latest generation RT Cores, Tensor Cores, and NVIDIA Ampere architecture CUDA® cores for unprecedented graphics, rendering, and AI performance. Additional support for a range of commercially available remote access software means you can access the power of your Quadro desktop workstation from anywhere. Achieve breakthrough innovations with the world's most powerful graphics solution

NVIDIA Quadro RTX 8000 and RTX 6000, 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 Fluid Dynamicists can also benefit from fast performance single precision GPU computing.

Altair nanoFluidX is a particle-based hydrodynamics (SPH) simulation tool to predict fluid flow around complex geometries under complicated motion. It can be used, for example, to simulate

oiling within conventional high-performance or electric vehicle powertrain systems comprising rotating shafts or gears and analyze forces and torques on individual components of the system. In addition, sloshing can be predicted in tanks experiencing transient motions. Conceived and optimized for use on clusters of Graphical Processing Units (GPUs), the software is extremely fast. For such typical gear-train applications, without the need to simplify geometries, the code can run an order of magnitude faster than a Finite-Volume code. GPU computing provides a significant performance advantage and power savings with respect to their more cumbersome CPU counterparts. nanoFluidX is one of the pioneering commercial software packages written with NVIDIA CUDA bringing significant speed to the overall product development process with NVIDIA GPUs.

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

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

2.1 VALIDATED SERVER CONFIGURATIONS

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

Table 1: Validated Server Configurations

Server Model	Graphics	Configuration
ASUS ESC8000 G4	8x Quadro RTX 8000 or RTX 6000 4x Quadro RTX NVLink HB bridge 2-slot Quadro Driver Release 430 U2 (440.36)	<ul style="list-style-type: none"> • Dual Intel® Xeon® Gold 6248 processor: 2.5-3.9GHz; 20 Cores, 40 Threads • 384 GB ECC DDR4 memory • 1TB M.2 PCIe NVMe SSD • Ubuntu 18.04 LTS

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.