ALTAIR

ONLY FORWARD

Altair Simulation 2022.3

Hardware Recommendations and Certifications

Updated: 03/30/2023

Contents

Intellectual Property Rights Notice iii
Technical Support vii

Hardware Recommendations and Certifications 9

Recommended Graphics Boards 10
Recommended Workstation Desktop and Laptop/Notebook Hardware 13
Altair Simulation 2022.3 Solver Hardware Configuration Recommendations 28
Recommended GPU Computing Processor List 34
Additional Information on Driver Installations 35

Intellectual Property Rights Notice

Copyright © 1986-2023 Altair Engineering Inc. All Rights Reserved.

This Intellectual Property Rights Notice is exemplary, and therefore not exhaustive, of intellectual property rights held by Altair Engineering Inc. or its affiliates. Software, other products, and materials of Altair Engineering Inc. or its affiliates are protected under laws of the United States and laws of other jurisdictions. In addition to intellectual property rights indicated herein, such software, other products, and materials of Altair Engineering Inc. or its affiliates may be further protected by patents, additional copyrights, additional trademarks, trade secrets, and additional other intellectual property rights. For avoidance of doubt, copyright notice does not imply publication. Copyrights in the below are held by Altair Engineering Inc. or its affiliates. Additionally, all non-Altair marks are the property of their respective owners.

This Intellectual Property Rights Notice does not give you any right to any product, such as software, or underlying intellectual property rights of Altair Engineering Inc. or its affiliates. Usage, for example, of software of Altair Engineering Inc. or its affiliates is governed by and dependent on a valid license agreement.

Altair Simulation Products

Altair® AcuSolve® ©1997-2023

Altair Activate® ©1989-2023

Altair® Battery Designer™ ©2019-2023

Altair Compose® ©2007-2023

Altair® ConnectMe™ ©2014-2023

Altair® EDEM™ ©2005-2023

Altair® ElectroFlo™ ©1992-2023

Altair Embed® ©1989-2023

Altair Embed® SE ©1989-2023

Altair Embed®/Digital Power Designer ©2012-2023

Altair Embed® Viewer ©1996-2023

Altair® ESAComp® ©1992-2023

Altair® Feko® ©1999-2023

Altair® Flow Simulator™ ©2016-2023

Altair® Flux® ©1983-2023

Altair® FluxMotor® ©2017-2023

Altair® HyperCrash® ©2001-2023

Altair® HyperGraph® ©1995-2023

Altair® HyperLife® ©1990-2023

Altair® HyperMesh® ©1990-2023

Altair® HyperSpice™ ©2017-2023

Altair® HyperStudy® ©1999-2023

Altair® HyperView® ©1999-2023

Altair® HyperViewPlayer® ©2022-2023

Altair® HyperWorks® ©1990-2023

Altair® HyperXtrude® ©1999-2023

Altair® Inspire™ ©2009-2023

Altair® Inspire™ Cast ©2011-2023

Altair® Inspire™ Extrude Metal ©1996-2023

Altair® Inspire™ Extrude Polymer ©1996-2023

Altair® Inspire™ Form ©1998-2023

Altair® Inspire™ Mold ©2009-2023

Altair® Inspire™ PolyFoam ©2009-2023

Altair® Inspire™ Print3D ©2021-2023

Altair® Inspire™ Render ©1993-2023

Altair® Inspire™ Studio ©1993-2023

Altair® Material Data Center™ ©2019-2023

Altair® MotionSolve® ©2002-2023

Altair® MotionView® ©1993-2023

Altair® Multiscale Designer® ©2011-2023

Altair® nanoFluidX® ©2013-2023

Altair® OptiStruct® ©1996-2023

Altair® PollEx™ ©2003-2023

Altair® PSIM™ ©2022-2023

Altair® Pulse™ ©2020-2023

Altair® Radioss® ©1986-2023

Altair® romAI™ ©2022-2023

Altair® S-FRAME® ©1995-2023

Altair® S-STEEL™ ©1995-2023

Altair® S-PAD™ ©1995-2023

Altair® S-CONCRETE™ ©1995-2023

Altair® S-LINE™ ©1995-2023

Altair® S-TIMBER™ ©1995-2023

Altair® S-FOUNDATION™ ©1995-2023

Altair® S-CALC™ ©1995-2023

Altair® S-VIEW™ ©1995-2023

Altair® Structural Office™ ©2022-2023

Altair® SEAM® ©1985-2023

Altair® SimLab® ©2004-2023

Altair® SimLab® ST ©2019-2023

Altair SimSolid® ©2015-2023

Altair® ultraFluidX® ©2010-2023

Altair® Virtual Wind Tunnel™ ©2012-2023

Altair® WinProp™ ©2000-2023

Altair® WRAP™ ©1998-2023

Altair® GateVision PRO™ ©2002-2023

Altair® RTLvision PRO™ ©2002-2023

Altair® SpiceVision PRO™ ©2002-2023

Altair® StarVision PRO™ ©2002-2023

Altair® EEvision™ ©2018-2023

Altair Packaged Solution Offerings (PSOs)

Altair® Automated Reporting Director™ ©2008-2022

Altair® e-Motor Director™ ©2019-2023

Altair® Geomechanics Director™ ©2011-2022

Altair® Impact Simulation Director™ ©2010-2022

Altair® Model Mesher Director™ ©2010-2023

Altair® NVH Director™ ©2010-2023

Altair® NVH Full Vehicle™ ©2022-2023

Altair® NVH Standard™ ©2022-2023

Altair® Squeak and Rattle Director™ ©2012-2023

Altair® Virtual Gauge Director™ ©2012-2023

Altair® Weld Certification Director™ ©2014-2023

Altair® Multi-Disciplinary Optimization Director™ ©2012-2023

Altair HPC & Cloud Products

Altair® PBS Professional® ©1994-2023

Altair® PBS Works™ ©2022-2023

Altair® Control™ ©2008-2023

Altair® Access™ ©2008-2023

Altair® Accelerator™ ©1995-2023

Altair® Accelerator™ Plus ©1995-2023

Altair® FlowTracer™ ©1995-2023

Altair® Allocator™ ©1995-2023

Altair® Monitor™ ©1995-2023

Altair® Hero™ ©1995-2023

Altair® Software Asset Optimization (SAO) ©2007-2023

Altair Mistral™ ©2022-2023

Altair® Grid Engine® ©2001, 2011-2023

Altair® DesignAI™ ©2022-2023

Altair Breeze™ ©2022-2023

Altair® NavOps® ©2022-2023

Altair® Unlimited™ ©2022-2023

Altair Data Analytics Products

Altair Analytics Workbench™ © 2002-2023

Altair® Knowledge Studio® ©1994-2023

Altair® Knowledge Studio® for Apache Spark ©1994-2023

Altair® Knowledge Seeker™ ©1994-2023

Altair® Knowledge Hub™ ©2017-2023

Altair® Monarch® ©1996-2023

Altair® Panopticon™ ©2004-2023

Altair® SmartWorks™ ©2021-2023

Altair SLC™ ©2002-2023

Altair SmartWorks Hub™ ©2002-2023

Altair® RapidMiner® ©2001-2023

Altair One™ ©1994-2023

2022.3

March 17, 2023

Technical Support

Altair provides comprehensive software support via web FAQs, tutorials, training classes, telephone, and e-mail.

Altair One Customer Portal

Altair One (https://altairone.com/) is Altair's customer portal giving you access to product downloads, a Knowledge Base, and customer support. We recommend that all users create an Altair One account and use it as their primary portal for everything Altair.

When your Altair One account is set up, you can access the Altair support page via this link: www.altair.com/customer-support/

Altair Community

Participate in an online community where you can share insights, collaborate with colleagues and peers, and find more ways to take full advantage of Altair's products.

Visit the Altair Community (https://community.altair.com/community) where you can access online discussions, a knowledge base of product information, and an online form to contact Support. After you login to the Altair Community, subscribe to the forums and user groups to get up-to-date information about release updates, upcoming events, and questions asked by your fellow members.

These valuable resources help you discover, learn and grow, all while having the opportunity to network with fellow explorers like yourself.

Altair Training Classes

Altair's in-person, online, and self-paced trainings provide hands-on introduction to our products, focusing on overall functionality. Trainings are conducted at our corporate and regional offices or at your facility.

For more information visit: https://learn.altair.com/

If you are interested in training at your facility, contact your account manager for more details. If you do not know who your account manager is, contact your local support office and they will connect you with your account manager.

Telephone and E-mail

If you are unable to contact Altair support via the customer portal, you may reach out to technical support via phone or e-mail. Use the following table as a reference to locate the support office for your region.

Altair support portals are available 24x7 and our global support engineers are available during normal Altair business hours in your region.

When contacting Altair support, specify the product and version number you are using along with a detailed description of the problem. It is beneficial for the support engineer to know what type of workstation, operating system, RAM, and graphics board you have, so include that in your communication.

Location	Telephone	E-mail
Australia	+61 3 9866 5557	anzsupport@altair.com
Brazil	+55 113 884 0414	br_support@altair.com
Canada	+1 416 447 6463	support@altairengineering.ca
China	+86 400 619 6186	support@altair.com.cn
France	+33 141 33 0992	francesupport@altair.com
Germany	+49 703 162 0822	hwsupport@altair.de
Greece	+30 231 047 3311	eesupport@altair.com
India	+91 806 629 4500 +1 800 425 0234 (toll free)	support@india.altair.com
Israel		israelsupport@altair.com
Italy	+39 800 905 595	support@altairengineering.it
Japan	+81 3 6225 5830	jp-support@altair.com
Malaysia	+60 32 742 7890	aseansupport@altair.com
Mexico	+52 55 5658 6808	mx-support@altair.com
New Zealand	+64 9 413 7981	anzsupport@altair.com
South Africa	+27 21 831 1500	support@altair.co.za
South Korea	+82 704 050 9200	support@altair.co.kr
Spain	+34 910 810 080	support-spain@altair.com
Sweden	+46 46 460 2828	support@altair.se
United Kingdom	+44 192 646 8600	support@uk.altair.com
United States	+1 248 614 2425	hwsupport@altair.com

If your company is being serviced by an Altair partner, you can find that information on our web site at https://www.altair.com/PartnerSearch/.

See www.altair.com for complete information on Altair, our team, and our products.

Hardware Recommendations and Certifications

1

View the most recent recommended graphic boards, laptops and desktop hardware configurations.

This chapter covers the following:

- Recommended Graphics Boards (p. 10)
- Recommended Workstation Desktop and Laptop/Notebook Hardware (p. 13)
- Altair Simulation 2022.3 Solver Hardware Configuration Recommendations (p. 28)
- Recommended GPU Computing Processor List (p. 34)
- Additional Information on Driver Installations (p. 35)

Recommended Graphics Boards

Recommended CAE/CAD graphic boards to use with Altair Simulation applications.

The most recent vendor/manufacturer drivers should be used and all driver support for these cards should be addressed to the appropriate manufacturer of the graphic board.

Note: AMD graphics cards will no longer be supported on Linux x86_64 operating systems in Altair Simulation 2022 and higher products.

AMD Graphics Cards

Products	GPU Model	Driver Version
Radeon™ Pro	W6800 W6600 VII WX 9100 W5700 WX 8200 W5500 WX 7100 WX 5100 WX 4100 WX 3200 WX 3100 WX 2100	*Windows 10 (64-bit)* 22.Q4 *Linux (64-bit)* Not Supported
Radeon™ Pro Mobility	WX 7130 WX 7100 WX 4170 WX 4150 WX 4130 WX 3100 WX 2100	*Windows 10 (64-bit)* 22.Q4 *Linux (64-bit)* Not Supported

NVIDIA Graphics Boards

Products	GPU Model						Driver Version
	M (Maxwell)	**P (Pascal)**	**V (Volta)**	**RTX (Turing)**	**A* (Ampere)**	**ADA Generation**	
Quadro Series	M2000 M4000 M5000 M6000	P400 P420 P600 P620 P1000 P2000 P2200 P4000 P5000 P5200 P6000 GP100	GV100	RTX 3000 RTX 4000 RTX 5000 RTX 6000 RTX 8000 T400 T600 T1000	RTX A2000 RTX A3000 RTX A4000 RTX A5000 RTX A6000	RTX 6000 ADA	*Windows 10 (64-bit)* R525 U5 (528.49) *Linux (64-bit) ODE Long Live* 525.89.02
Quadro Mobility	M500M M520M M600M M620M M1000M M2000M M2200M M3000M M4000M M5000M	P500 P520 P600 P620 P1000 P2000 P3000 P3200 P4000 P4200 P5000 P5200	N/A	T400 T600 T1000 T1200 T2000 RTX 3000 RTX 4000 RTX 5000	RTX A2000 RTX A3000 RTX A4000 RTX A5000		*Windows 10 (64-bit)* R525 U5 (528.49) *Linux (64-bit) ODE Long Live* 525.89.02

Note:

Minimum OpenGL 3.2 and OpenCL 2.1 Requirement
Virtual server/clients and VirtualGL setups may work, but are not officially tested or supported.

NVIDIA Optimus or AMD Switchable Graphics
In order to ensure best performance, these options should be set to use discrete NVIDIA or AMD GPU and not the Intel GPU.

Power Options and Mobility Center
In order to ensure best performance, these options should be maximum performance for both GPU and CPU.

Graphics Driver Corruption or Installation Issues
In order to ensure best driver compatibility, it is recommended to use "Custom" and "Clean" install options instead of the general "Express" driver installer options.

Recommended Workstation Desktop and Laptop/ Notebook Hardware

DELL Workstations - Desktops

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
3260C	T400 T600 T1000 RTX A2000 RTX 3000	
3450 SFF	P400 P620 P1000 T400 T600 T1000	WX3200
3460 (mini)	RTX 3000	N/A
3460 SFF	T400 T600 T1000 RTX A2000	
3650	P400 P620 P1000 P2200 T400 T600 T1000 RTX 4000	W5500 W5700 W6600 W6800 WX3200

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
	RTX 5000 RTX A2000 RTX A4000 RTX A5000 RTX A6000	
3660T	T400 T600 T1000 RTX A2000 RTX A4000 RTX A4500 RTX A5000 RTX A5500 RTX A6000	
5860T	T400 T1000 RTX A2000 RTX A4000 RTX A4500 RTX A5000 RTX A5500 RTX A6000 GV100	
7865T	T400 T1000 RTX A2000 RTX A4000 RTX A4500 RTX A5000	

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
	RTX A5500 RTX A6000 GV100	
7960T	T400 T1000 RTX A2000 RTX A4000 RTX A4500 RTX A5000 RTX A5500 RTX A6000 GV100	
7960 Rack	T400 T1000 RTX A2000 RTX A4000 RTX A4500 RTX A5000 RTX A5500 RTX A6000 GV100	

DELL Workstations - Laptops

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
3470	T550	N/A
3550	P520	N/A

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
3551	P620	N/A
3560	T500	N/A
3561	T600 T1200	N/A
3570	T550 T600 RTX A500	N/A
3571	T600 RTX A1000 RTX A2000	N/A
5470	RTX A1000	N/A
5540 / 5740	T1000 T2000 RTX 3000*	N/A
5550 / 5750*	T1000 T2000 RTX 3000*	N/A
5560	T1200 RTX A2000	N/A
5570	RTX A1000 RTX A2000	N/A
5760	RTX A2000 RTX A3000	N/A
5770	RTX A2000 RTX A3000	

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
7540 / 7740	T1000 T2000 RTX 3000 RTX 4000 RTX 5000	WX 3200 WX 7130
7550 / 7750	T1000 T2000 RTX 3000 RTX 4000 RTX 5000	N/A
7560	T1200 RTX A2000 RTX A3000 RTX A4000 RTX A5000	N/A
7670	RTX A1000 RTX A2000 RTX A3000 RTX A4500 RTX A5500	N/A
7760	T1200 RTX A3000 RTX A4000 RTX A5000	N/A
7770	RTX A1000 RTX A3000 RTX A4500	N/A

Product	**Precision Workstation**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
	RTX A5500	

Lenovo Workstations - Desktops

Product	**Lenovo ThinkStation**	
Workstation Model	**NVIDIA Quadro GPU (442.92 or higher)**	**AMD Radeon™ Pro GPU**
P320 SFF / **P320 TWR***	P400 P600 P1000 P2000* P4000*	N/A
P330 SFF / **P330 TWR***	P400 P620 P1000 P2000* P2200* P4000*	N/A
P340 SFF / **P340 TWR***	RTX A2000* RTX 4000* RTX 5000* RTX A4000* RTX A5000* T400 T600 T1000	N/A
P340 Tiny	T1000	
P348	T400 T600	

Product	**Lenovo ThinkStation**	
Workstation Model	**NVIDIA Quadro GPU (442.92 or higher)**	**AMD Radeon™ Pro GPU**
	T1000	
P350 SFF / **P350 TWR***	P2200 RTX A2000 RTX A4000* RTX A5000* T400 T600 T1000	
P350 Tiny	T600 T1000	N/A
P358 TWR	RTX A2000 T600 T1000	N/A
P360 Ultra/ **P360 TWR***	RTX A2000 RTX A4000* RTX A4500* RTX A5000 RTX A5000m T400 T600* T1000	
P360 Tiny	T400 T1000	N/A
P520* / P520c	RTX 4000 RTX 5000 RTX 6000* RTX A2000	N/A

Product	**Lenovo ThinkStation**	
Workstation Model	**NVIDIA Quadro GPU (442.92 or higher)**	**AMD Radeon™ Pro GPU**
	RTX A4000 RTX A4500 RTX A5000 RTX A6000* T400 T600 T1000	
P620 (AMD Ryzen PRO 3975X) / **(AMD Ryzen Threadripper PRO 5955WX)**	GP100 RTX 4000 RTX 5000 RTX 6000 RTX 8000 RTX A2000 RTX A4000 RTX A4500 RTX A5000 RTX A5500 RTX A6000 T400 T600 T1000	W5500 W5700
P720 / P920	RTX 4000 RTX 5000 RTX 6000 RTX 8000 RTX A2000 RTX A4000 RTX A4500 RTX A5000	N/A

Product	**Lenovo ThinkStation**	
Workstation Model	**NVIDIA Quadro GPU (442.92 or higher)**	**AMD Radeon™ Pro GPU**
	RTX A5500 RTX A6000 T400 T600 T1000	
NEC (mini)	P1000	N/A

Lenovo Workstations - Laptops (*Windows 10 support only)

Product	**Workstation Model**	
Lenovo ThinkPad	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
P1 **Gen1**	P1000M P2000M	N/A
P1 **Gen2***	T1000 T2000	N/A
P1 **Gen3**	T1000 T2000	N/A
P1 **Gen4**	T1200 RTX A2000 RTX A3000 RTX A4000 RTX A5000	N/A
P1 **Gen5**	RTX A1000 RTX A2000 RTX A3000 RTX A4500 RTX A5500	N/A

Product	**Workstation Model**	
Lenovo ThinkPad	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
P14s	P520	N/A
P14s Gen1	P520	AMD Ryzen 7 PRO 4750U with Radeon Graphics (2 GB) AMD APU 4400
P14s Gen2	T500	AMD APU 5500
P14s Gen3a	N/A	Radeon PRO Graphics
P14s Gen3i	T550	N/A
P15 Gen1	T1000 T2000 RTX 3000 RTX 4000	N/A
P15 Gen2	T1200 RTX A2000 RTX A3000 RTX A4000 RTX A5000	N/A
P15s Gen1	P520	N/A
P15s Gen2	T500	
P15v Gen1	P620	N/A
P15v	T600 T1200	N/A

Product	**Workstation Model**	
Lenovo ThinkPad	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
Gen2	RTX A2000	
P15v **Gen3a/Gen3i**	T600 T1200 RTX A2000	N/A
P16 **Gen1**	RTX A1000 RTX A2000 RTX A3000 RTX A4500 RTX A5500	N/A
P16s **Gen1a**	N/A	Radeon PRO Graphics
P16s **Gen1i**	T550	N/A
P17 **Gen1**	T1000 T2000 RTX 3000 RTX 4000 RTX 5000	N/A
P17 **Gen2**	T1200 RTX A2000 RTX A3000 RTX A4000 RTX A5000	N/A
P40 **Yoga**	M500M	N/A
P43s*	P520	N/A
P50	M1000M	N/A

Product	**Workstation Model**	
Lenovo ThinkPad	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
	M2000M	
P50s	M500M	N/A
P51	M1200M M2200M	N/A
P51s	M520M	N/A
P52	P1000 P2000 P3200	N/A
P52s	P500	N/A
P53*	T1000 T2000 RTX 3000 RTX 4000	N/A
P53s*	P520	N/A
P71	M620M P3000 P4000 P5000	N/A
P72	P600 P2000 P3200 P4200 P5200	N/A
P73*	P620 T2000 RTX 3000 RTX 4000	N/A

Product	**Workstation Model**	
Lenovo ThinkPad	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
	RTX 5000	

Acer Workstations and Laptops (*Windows 10 support only)

Product	**Acer**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
ConceptD 500*	RTX 4000	N/A
ConceptD 700*	RTX 4000	N/A
Veriton K8	RTX 4000	N/A
Veriton K8690G	RTX A4000	N/A

Product	**Acer**	
Mobile Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
ConceptD 3 Pro*	T1000	N/A
ConceptD 3 Ezel Pro*	T1000	N/A
ConceptD 5 Pro*	RTX 3000	N/A
ConceptD 7 Pro*	RTX 3000	N/A

Product	**Acer**	
Mobile Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
ConceptD 7 **Ezel Pro***	RTX 5000	N/A

Altos Workstations and Laptops (*Windows 10 support only)

Product	**Altos** **Computing**	
Workstation Model	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
BrainSphere™ **P130 F5**	RTX 2000	N/A
BrainSphere™ **P530 F4**	K420 P400 K620 K1200 P1000 P2000 P4000 P5000 P6000 GP100 RTX 2000 RTX 4000 RTX 5000	N/A

SuperMicro Workstations

Product	**SuperMicro**	
System Model **AS- 5014A-TT**	**NVIDIA Quadro GPU**	**AMD Radeon™ Pro GPU**
Mainstream **Ryzen Threadripper Pro 3955WX**	RTX A4500	W6800
Performance **Ryzen Threadripper Pro 3975WX**	RTX A4500	W6800
Expert **Ryzen Threadripper Pro 3995WX**	RTX A4500	W6800
	(visit SuperMicro for a complete GPU support list)	

Comments

For NVIDIA GPU based laptops/notebooks the Optimus power saving option in the BIOS should be disabled and the NVIDIA drivers properly installed for optimal performance in Altair Simulation products.

For AMD GPU based laptops/notebooks; the Enduro/Switchable Graphics power saving option should be disabled and the AMD drivers properly installed for optimal performance in Altair Simulation products.

Optimus (Intel/NVIDIA) enabled drivers may create performance issues with notebooks/laptops compared to a dedicated non-shared GPU driver. Disabling the Optimus feature in BIOS, if available, will help give the best overall graphics performance.

Disable nView Window manager under NVIDIA drivers if you experience random crashes and/or issues.

All power saving modes, settings and governors for CPU frequencies and GPU performance should be set to maximum settings in order to get the optimal performance out of Altair Simulation products. This includes smooth graphics and high frame rates (FPS) on Windows and Linux platforms.

Altair Simulation 2022.3 Solver Hardware Configuration Recommendations

Recommended hardware configurations for Altair Solvers.

AcuSolve Solver

Table 1:

Problem Size	Small	Medium	Large
Typical Workload Steady State or Transient	Steady state: Up to 1M nodes Transient: Up to 100K nodes	Steady state: Up to 10M nodes Transient: Up to 1M nodes	Steady state: Greater than 10M nodes Transient: Greater than 1M nodes
Throughput [1]	Single job	Single job	Single job
CPU [2]	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or AMD EPYC 7002 or 7003 series	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or AMD EPYC 7002 or 7003 series	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or AMD EPYC 7002 or 7003 series
Number of CPU / node	1-4	1-4	1-4
Number of cores / node	32 – 128	32 – 128	32 – 128
Number of nodes	1-8	8 – 48	>48
Minimum Memory Configuration / node [3]	300MB to 3GB	3GB to 30GB	More than 30GB (3KB per CFD node)
Storage (minimum)	500 GB SATA or SSD	1.5 TB local storage	1.5 TB local storage
Network Interconnect	Gigabit Ethernet Or Infiniband	Infiniband	Infiniband
Operating System	Linux kernel 2.6.32 or higher Windows 7 or 10	Linux kernel 2.6.32 or higher	Linux kernel 2.6.32 or higher
GPU	Yes	Yes	Yes
MPI	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher

Problem Size	Small	Medium	Large
Setup (2000-3000 computational nodes per core)	Pure OpenMP or Hybrid OpenMP/MPI	Hybrid OpenMP/MPI	Hybrid OpenMP/MPI
Hyper Threading	Not recommended	Not recommended	Not recommended

Feko Solver

Table 2:

Problem Size	Small	Medium	Large
General recommendations given for MoM and MLFMM dependent on problem size in terms of number of unknowns / mesh elements. For other solution methods (FEM, FDTD, RL-GO, PO, UTD) many factors to be considered.	Pure MoM: less than 50k unknowns. MLFMM: between 100k and 500k unknowns	Pure MoM: between 50k and 100k unknowns. MLFMM: between 500k and 5M unknowns	Pure MoM: >100k unknowns MLFMM: >5M unknowns
Throughput [1]	Single job	Single large job or few jobs in parallel	Single very large job or multiple jobs
CPU [2]	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or later	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or later	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or later
Number of CPU / node	2	2	2
Number of cores / node	32 – 56	32 – 56	32 – 56
Number of nodes	1	8 – 16	> 16
Minimum Memory Configuration / node [3]	64 GB	128 GB	256 GB
Storage (minimum)	500 GB SATA or SSD	500 GB SATA or SSD	500 GB SATA or SSD

Problem Size	Small	Medium	Large
Network Interconnect	Gigabit Ethernet	Infiniband	Infiniband
Operating System	RHEL or CentOS 8.3 Windows 10	Linux kernel 2.6.32 or higher	Linux kernel 2.6.32 or higher
GPU	Yes	No	No
MPI	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher
Setup	Pure MPI	Pure MPI	Pure MPI
Hyper Threading	Not recommended	Not recommended	Not recommended

Flux Solver

Table 3:

Problem type	Small	Medium	Large
Typical Workload (depending on number of DOF, element type, and other factors)	< 300 000 DOF	Around 500 000 DOF	Around 5M DOF
Throughput [1]	Single	Single	Single
CPU [2]	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake"	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake"	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake"
Number of CPU / node	1	2	2
Number of cores / node	8	16	16+
Number of nodes	1	1	1-4
Minimum Memory Configuration / node [3]	8 GB	16-32 GB	300GB
Storage (minimum)	500 GB SATA or SSD	1 TB local storage SSD	1.5 TB local storage SSD
Network Interconnect			Infiniband

Problem type	Small	Medium	Large
Operating System	Linux kernel 3.10.0-693 or higher Windows 7 or 10 with SSD	Linux kernel 3.10.0-693 or higher Windows 7 or 10 with SSD	Linux kernel 3.10.0-693 or higher Windows 7 or 10 with SSD
GPU	No	No	No
MPI	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher
Setup	SMP	SMP or Hybrid 2MPI/node	SMP or Hybrid 2MPI/node
Hyper Threading	Not recommended	Not recommended	Not recommended

Radioss Solver

Table 4:

Problem Size	Small	Medium	Large
Typical Workload Crash & Impact	Component tests, sled test, drop test, … Less than 500K elements	Medium crash model, between 1 and 6 millions of elements model	Accurate car crash model (rupture), very large model with size > 6 million elements
Throughput [1]	Single job	Single large job or few jobs in parallel	Single very large job or multiple jobs
CPU [2]	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or AMD EPYC 7002 or 7003 series	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lake" or AMD EPYC 7002 or 7003 series	Dual CPU socket For example, Intel Xeon Gold "Cascade Lake" or "Ice Lke" or AMD EPYC 7002 or 7003 series
Number of CPU / node	2	2	2
Number of cores / node	32 – 64	32 – 128	32 – 128
Number of nodes	1	8 – 16	> 16
Minimum Memory Configuration / node [3]	64-128GB	64-256GB	64-256GB

Problem Size	Small	Medium	Large
Storage (minimum)	500 GB SATA or SSD	1,5 TB local storage	1,5 TB local storage
Network Interconnect	Gigabit Ethernet	Infiniband	Infiniband
Operating System	Linux kernel 2.6.32 or higher Windows 7 or 10	Linux kernel 2.6.32 or higher	Linux kernel 2.6.32 or higher
GPU	No	No	No
MPI	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher	Intel MPI 2018.4 or higher
Setup	Pure MPI	Pure MPI or Hybrid with 2 or 4 OpenMP threads per MPI	Hybrid with 2 or 4 OpenMP threads per MPI
Hyper Threading [5]	Yes, Hybrid with 2 OpenMP per MPI	Not recommended	Not recommended

OptiStruct Solver

Table 5:

Problem type	Small or medium	Large static	Large dynamic
Typical Workload (depending on number of DOF, element type, and other factors)	Nonlinear - less than 2M DOF; linear static - less than 5M DOF; NVH - less than 5M DOF	Nonlinear - more than 2M DOF; linear static - more than 5M DOF	NVH - more than 5M DOF
Throughput [1]	Single	Single	Single or few jobs in parallel
CPU [2]	Dual CPU socket For example, Intel Xeon Gold “Cascade Lake” or “Ice Lake”	Dual CPU socket For example, Intel Xeon Gold “Cascade Lake” or “Ice Lake	Dual CPU socket For example, Intel Xeon Gold “Cascade Lake” or “Ice Lake"
Number of CPU / node	2	2	2
Number of cores / node	8-24	24+	24+
Number of nodes	1	1-8	1-8

Problem type	**Small or medium**	**Large static**	**Large dynamic**
Minimum Memory Configuration / node [3]	16-64GB	128GB	256GB
Storage (minimum)	512GB local storage	1TB local storage	3 TB local storage, SSD and RAID0 recommended
Network Interconnect		Infiniband	Infiniband
Operating System	Linux kernel 2.6.32 or higher Windows 7 or 10 with SSD	Linux kernel 2.6.32 or higher Windows 7 or 10 with SSD	Linux kernel 2.6.32 or higher Windows 7 or 10 with SSD
GPU	Yes	Yes	Yes (Eigenvalue Extraction – AMSES or Lanczos)
MPI	Linux: Intel MPI 2018.4 (recommended) Windows: Intel MPI 5.1.0.078 (recommended)	Linux: Intel MPI 2018.4 (recommended) Windows: Intel MPI 5.1.0.078 (recommended)	Linux: Intel MPI 2018.4 (recommended) Windows: Intel MPI 5.1.0.078 (recommended)
Setup	SMP or DDM hybrid or SMP+GPU	DDM hybrid or SMP +GPU	SMP or DDM hybrid or SMP+GPU
Hyper Threading	Not recommended	Not recommended	Not recommended

1. Number of simultaneous jobs. Use of a workload management middleware like Altair PBS is highly recommended to insure optimal and dedicated usage of the CPU resource
2. Typical node configuration is based on dual CPU socket processors
3. It is extremely important to populate all the memory banks on the mother board.
4. In Hybrid mode, it is recommended to set a number of MPIs that is a multiple of the number of sockets and then set the number of OpenMP in a way that number of MPIs x number of OpenMP equal number of physical cores.
5. Hyper Threading (HT) may increase performance by around 10% on single node. In this case, recommended setup is to run 2 OpenMP per MPI, with a number of MPIs that matches the total number of physical cores on the node. On multi-node, it is better not using HT

Recommended GPU Computing Processor List

Recommended graphic boards for use with the Altair Solver applications for high-powered GPU computing.

The following table lists the recommended graphic boards for use with the Altair Solver applications for high-powered GPU computing.

	Manufacturer and Model	Graphics Card	Driver Version (Minimum or Higher)
AcuSolve	NVIDIA (Tesla)	P100 V100	Linux (64-bit) 418.39 Windows (64-bit) 396.26
	NVIDIA (Quadro)	GP100 GV100	Linux (64-bit) 418.39 Windows (64-bit) 396.26
OptiStruct	NVIDIA (Tesla)	P100 V100	Linux (64-bit) 440.33 Windows (64-bit) 441.22
	NVIDIA (Quadro)	GP100 GV100	Linux (64-bit) 440.33 Windows (64-bit) 441.22

Note:

- The most recent vendor/manufacturer drivers should be used and all driver support for these cards should be addressed to the appropriate manufacturer of the graphics board.
- Double Precision GPU cards should be used to run OptiStruct. Single precision GPU cards (such as RTX 600, etc) are not recommended for OptiStruct runs.

Additional Information on Driver Installations

The NVIDIA Driver Update recommendation is to use the **Custom installation** option and select the **Perform clean installation** option to validate that there are no conflicts in DLL/drivers.

The same should be done with AMD hardware and drivers as well using AMD's custom uninstall tools.

Figure 1: