



# GPU-ACCELERATED APPLICATIONS



## Test Drive the World's Fastest Accelerator – Free!

Take the GPU Test Drive, a free and easy way to experience accelerated computing on GPUs. You can run your own application or try one of the preloaded ones, all running on a remote cluster. Try it today.

[www.nvidia.com/gputestdrive](http://www.nvidia.com/gputestdrive)



# GPU-ACCELERATED APPLICATIONS

Accelerated computing has revolutionized a broad range of industries with over five hundred applications optimized for GPUs to help you accelerate your work.

## CONTENTS

- 1 Computational Finance
- 2 Climate, Weather and Ocean Modeling
- 2 Data Science and Analytics
- 5 Artificial Intelligence
  - DEEP LEARNING AND MACHINE LEARNING
- 9 Federal, Defense and Intelligence
- 10 Design for Manufacturing/Construction: CAD/CAE/CAM
  - COMPUTATIONAL FLUID DYNAMICS
  - COMPUTATIONAL STRUCTURAL MECHANICS
  - DESIGN AND VISUALIZATION
  - ELECTRONIC DESIGN AUTOMATION
  - INDUSTRIAL INSPECTION
- 19 Media & Entertainment
  - ANIMATION, MODELING AND RENDERING
  - COLOR CORRECTION AND GRAIN MANAGEMENT
  - COMPOSITING, FINISHING AND EFFECTS
  - EDITING
  - ENCODING AND DIGITAL DISTRIBUTION
  - ON-AIR GRAPHICS
  - ON-SET, REVIEW AND STEREO TOOLS
  - WEATHER GRAPHICS
- 26 Medical Imaging
- 29 Oil and Gas
- 30 Research: Higher Education and Supercomputing
  - COMPUTATIONAL CHEMISTRY AND BIOLOGY
  - NUMERICAL ANALYTICS
  - PHYSICS
  - SCIENTIFIC VISUALIZATION
- 41 Safety and Security
- 44 Tools and Management

# Computational Finance

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Accelerated Computing Engine</b>	Elsen	Secure, accessible, and accelerated back-testing, scenario analysis, risk analytics and real-time trading designed for easy integration and rapid development.	<ul style="list-style-type: none"> <li>• Web-like API with Native bindings for Python, R, Scala, C</li> <li>• Custom models and data streams are easy to add</li> </ul>	Multi-GPU Single Node
<b>Adaptiv Analytics</b>	SunGard	A flexible and extensible engine for fast calculations of a wide variety of pricing and risk measures on a broad range of asset classes and derivatives.	<ul style="list-style-type: none"> <li>• Existing models code in C# supported transparently, with minimal code changes</li> <li>• Supports multiple backends including CUDA and OpenCL</li> <li>• Switches transparently between multiple GPUs and CPUs depending on the deal support and load factors.</li> </ul>	Multi-GPU Single Node
<b>Alea.cuBase F#</b>	QuantAlea's	F# package enabling a growing set of F# capability to run on a GPU	<ul style="list-style-type: none"> <li>• F# for GPU accelerators</li> </ul>	Multi-GPU Single Node
<b>Esther</b>	Global Valuation	In-memory risk analytics system for OTC portfolios with a particular focus on XVA metrics and balance sheet simulations.	<ul style="list-style-type: none"> <li>• High quality models not admitting closed form solutions</li> <li>• Efficient solvers based on full matrix linear algebra powered by GPUs and Monte Carlo algorithms</li> </ul>	Multi-GPU Single Node
<b>Global Risk</b>	MISYS	Regulatory compliance and enterprise wide risk transparency package	<ul style="list-style-type: none"> <li>• Risk analytics</li> </ul>	Multi-GPU Single Node
<b>Hybridizer C#</b>	Altimesh	Multi-target C# framework for data parallel computing.	<ul style="list-style-type: none"> <li>• C# with translation to GPU or Multi-Core Xeon</li> </ul>	Multi-GPU Single Node
<b>MACS Analytics Library</b>	Murex	Analytics library for modeling valuation and risk for derivatives across multiple asset classes	<ul style="list-style-type: none"> <li>• Market standard models for all asset classes paired with the most efficient resolution methods (Monte Carlo simulations and Partial Differential Equations)</li> </ul>	Multi-GPU Single Node
<b>MiAccLib 2.0.1</b>	Hanweck Associates	Accelerated libraries which encompasses high speed multi-algorithm search engines, data security engine and also video analytics engines for text processing, encryption/decryption and video surveillance respectively.	<ul style="list-style-type: none"> <li>• Text Processing: Exact Match, Approximate\Similarity Text, Wild Card, MultiKeyword and MultiColumnMultiKeyword, etc</li> <li>• Data Security: Accelerated Encryption/Description for AES-128</li> <li>• Video Analytics: Accelerated Intrusion Detection Algorithm</li> </ul>	Multi-GPU Single Node
<b>NAG</b>	Numerical Algorithms Group	Random number generators, Brownian bridges, and PDE solvers	<ul style="list-style-type: none"> <li>• Monte Carlo and PDE solvers</li> </ul>	Single GPU Single Node
<b>O-Quant options pricing</b>	O-Quant	Offering for risk management and complex options / derivatives pricing using GPU	<ul style="list-style-type: none"> <li>• Cloud-based interface to price complex derivatives representing large baskets of equities</li> </ul>	Multi-GPU Multi-Node
<b>Oneview</b>	Numerix	Numerix introduced GPU support for Forward Monte Carlo simulation for Capital Markets and Insurance	<ul style="list-style-type: none"> <li>• Equity/FX basket models with BlackScholes/Local Vol models for individual equities and FX</li> <li>• Algorithms: AAD (Automatic Algebraic Differential)</li> <li>• New approaches to AAD to reduce time to market for fast Price Greeks and XVA Greeks</li> </ul>	Multi-GPU Multi-Node
<b>Pathwise</b>	Aon Benfield	Specialized platform for real-time hedging, valuation, pricing and risk management	<ul style="list-style-type: none"> <li>• Spreadsheet-like modeling interfaces, Python-based scripting environment, and Grid middleware</li> </ul>	Multi-GPU Single Node
<b>SciFinance</b>	SciComp, Inc	Derivative pricing (SciFinance)	<ul style="list-style-type: none"> <li>• Monte Carlo and PDE pricing models</li> </ul>	Single GPU Single Node
<b>Synerscope Data Visualization</b>	Synerscope	Visual big data exploration and insight tools	<ul style="list-style-type: none"> <li>• Graphical exploration of large network datasets including geo-spatial and temporal components</li> </ul>	Single GPU Single Node

<b>Volera</b>	Hanweck Associates	Real-time options analytical engine (Volera)	<ul style="list-style-type: none"> <li>Real-time options analytics engine</li> </ul>	Multi-GPU Single Node
<b>Xcelerit SDK</b>	Xcelerit	Software Development Kit (SDK) to boost the performance of Financial applications (e.g. Monte-Carlo, Finite-difference) with minimum changes to existing code.	<ul style="list-style-type: none"> <li>C++ programming language, cross-platform (back-end generates CUDA and optimized CPU code), supports Windows and Linux operating systems</li> </ul>	Multi-GPU Single Node

## Climate, Weather and Ocean Modeling

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ACME-Atm</b>	US DOE	Global atmospheric model used as component to ACME global coupled climate model	<ul style="list-style-type: none"> <li>Dynamics only</li> </ul>	Multi-GPU Multi-Node
<b>COSMO</b>	COSMO Consortium	Regional numerical weather prediction and climate research model	<ul style="list-style-type: none"> <li>Radiation only in the trunk release, all features in the MCH branch used for operational weather forecasting</li> </ul>	Multi-GPU Multi-Node
<b>Gales</b>	KNMI, TU Delft	Regional numerical weather prediction model	<ul style="list-style-type: none"> <li>Full Model</li> </ul>	Multi-GPU Multi-Node
> <b>WRF AceCAST-WRF</b>	TempoQuest Inc.	WRF model from NCAR, but now commercialized by TQI. Used for numerical weather prediction and regional climate studies. All popular aspects of WRF model are GPU developed.	<ul style="list-style-type: none"> <li>All popular aspects of WRF model are GPU developed: - ARW dynamics - 19 physics options including enough to run the full WRF model on GPUs</li> </ul>	Multi-GPU Multi-Node

## Data Science and Analytics

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ANACONDA</b>	Anaconda	Anaconda is the leading python package manager, that is the lead contributor to several open source data science libraries. Anaconda includes Numba, a Python-to-GPU compiler that compiles easy-to-read Python code to many-core and GPU architectures. Also includes single-line install of key deep learning packages for GPUs such as pytorch. Anaconda has been downloaded over 15M times and is used for AI & ML data science workloads using TensorFlow, Theano, Keras, Caffe, Neon, Lasagne, NLTK, spaCY.	<ul style="list-style-type: none"> <li>Includes Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND, and sorting algorithms from the CUB and Modern GPU libraries Includes Numba (JIT python compiler) and Dask (python scheduler) Includes single-line install of numerous DL frameworks such as pytorch</li> </ul>	Multi-GPU Single Node
<b>ArgusSearch</b>	Planet AI	Deep Learning driven document search	<ul style="list-style-type: none"> <li>fast full text search engine: searches hand-written and text documents, including PDF</li> <li>allows almost any arbitrary requests (Regular Expressions are supported)</li> <li>provides a list of matches sorted by confidence</li> </ul>	Multi-GPU Single Node
<b>Automatic Speech Recognition</b>	Capio	In-house and Cloud-based speech recognition technologies	<ul style="list-style-type: none"> <li>Real-time and offline (batch) speech recognition</li> <li>Exceptional accuracy for transcription of conversational speech</li> <li>Continuous Learning (System becomes more accurate as more data is pushed to the platform)</li> </ul>	Multi-GPU Single Node
<b>Blazegraph GPU</b>	Blazegraph	First and fastest GPU accelerated platform for graph query. It provides drop-in acceleration for existing RDF/ Sparql and Tinkerpop/ Blueprints graph applications. It provides high-level graph database APIs with transparent GPU acceleration for graph query.	<ul style="list-style-type: none"> <li>GPU-accelerated SPARQL graph query</li> <li>Data Management using the RDF interchange model</li> <li>Tinkerpop/Blueprints Graph Support</li> <li>Billions of edges on a single multi-GPU node</li> <li>SaaS and Appliance models available.</li> </ul>	Multi-GPU Single Node

<b>BlazingDB</b>	BlazingDB	GPU-accelerated relational database for data warehousing scenarios available for AWS and on-premise deployment.	<ul style="list-style-type: none"> <li>• Modern data warehousing application supporting petabyte scale applications</li> </ul>	Multi-GPU Single Node
<b>BrytlytDB</b>	Brytlyt	In-GPU-memory database built on top of PostgreSQL	<ul style="list-style-type: none"> <li>• GPU-Accelerated joins, aggregations, scans, etc. on PostgreSQL. Visualization platform bundled with database is called SpotLyt.</li> </ul>	Multi-GPU Multi-Node
<b>CuPy</b>	Preferred Networks	CuPy ( <a href="https://github.com/cupy/cupy">https://github.com/cupy/cupy</a> ) is a GPU-accelerated scientific computing library for Python with a NumPy compatible interface.	<ul style="list-style-type: none"> <li>• CUDA</li> <li>• multi-GPU support</li> </ul>	Multi-GPU Single Node
<b>Datalogue</b>	Datalogue	AI powered pipelines that automatically prepare any data from any source for immediate & compliant use	<ul style="list-style-type: none"> <li>• Data transformation</li> <li>• Ontology mapping</li> <li>• Data standardization</li> <li>• Data augmentation</li> </ul>	Multi-GPU Single Node
<b>DeepGram</b>	DeepGram	Voice processing solution for call centers, financials and other scenarios	<ul style="list-style-type: none"> <li>• Speech to text and phonetic search using GPU deep learning</li> </ul>	Multi-GPU Single Node
<b>Driverless AI</b>	H2O.ai	*Automated Machine Learning with Feature Extraction. Essentially BI for Machine Learning and AI, with accuracy very similar to Kaggle Experts.	<ul style="list-style-type: none"> <li>• Automated machine learning and feature extraction</li> <li>• Automated statistical visualization</li> <li>• Interpretability toolkit for machine learning models</li> </ul>	Multi-GPU Single Node
<b>Gunrock</b>	UC Davis	Gunrock is a library for graph processing on the GPU. Gunrock achieves a balance between performance and expressiveness by coupling high performance GPU implementations with a high-level programming model, that requires minimal GPU programming knowledge.	<ul style="list-style-type: none"> <li>• Direction-optimizing BFS, SSSP, PageRank, Connected Components, Betweenness-centrality, triangle counting</li> <li>• Multi-GPU support for frontier-based methods</li> </ul>	Multi-GPU Single Node
<b>H2O4GPU</b>	H2O.ai	H2O is a popular machine learning platform which offers GPU-accelerated machine learning. In addition, they offer deep learning by integrating popular deep learning frameworks.	<ul style="list-style-type: none"> <li>• Currently supporting tree based methods (GBM &amp; Random Forest), GLM, Kmeans and are working on a bunch of other algorithms that are coming soon</li> <li>• Supports TensorFlow, Caffe and MXNet</li> </ul>	Multi-GPU Single Node
<b>IntelligentVoice</b>	Intelligent Voice	Far more than a transcription tool, this speech recognition software learns what is important in a telephone call, extracts information and stores a visual representation of phone calls to be combined with text/instant messaging and E-mail. Intelligent Voice's search and alert makes it possible to tackle issues before they arise, address data security concerns and monitor physical access to data.	<ul style="list-style-type: none"> <li>• Advanced Speech Recognition across large data sets, JumpTo Technology, for data visualisation, E-Discovery, extraction from phone calls, IM &amp; Email defining key phrases and emotional analysis</li> <li>• Compliance, defining key conversations and interactions</li> </ul>	Multi-GPU Single Node
<b>Jedox</b>	Jedox	Helps with portfolio analysis, management consolidation, liquidity controlling, cash flow statements, profit center accounting, treasury management, customer value analysis and many more applications, all accessible in a powerful web and mobile application or Excel environment.	<ul style="list-style-type: none"> <li>• This database holds all relevant data in GPU memory and is thus an ideal application to utilize the Tesla K40 &amp; 12 GB on-board RAM</li> <li>• Scale that up with multiple GPUs and keep close to 100 GB of compressed data in GPU memory on a single server system for fast analysis, reporting, and planning.</li> </ul>	Multi-GPU Single Node

<b>Kinetica</b>	Kinetica	A GPU accelerated engine with an in-memory database for advanced analytics, location-based visualization, and streamlined machine learning.	<ul style="list-style-type: none"> <li>• Run SQL queries on static, streaming, geospatial and temporal data</li> <li>• Simultaneously ingest &amp; analyze streaming with historical data</li> <li>• In-memory, distributed image processing &amp; rendering</li> <li>• Consolidate database workloads including traditional BI workloads and newer AI workloads</li> <li>• Distribute and linearly scale out data and workloads to meet required SLAs</li> <li>• OLAP optimized and SQL 92 compliant</li> <li>• Native geospatial functionality including native geospatial data types, functions and joins</li> <li>• Build dashboards with Reveal - a built-in visualization tool that allows adhoc analysis of data including geospatial</li> <li>• Use the User Defined Functions capability to execute open source or proprietary algorithms and ML models in-database</li> <li>• Native RESTful database with a full stack of APIs from C++ to NodeJS to Java</li> <li>• Connect to a range of streaming, ETL and data technologies using certified connectors</li> <li>• Leverage full text search functionality for text-based queries out of the box</li> <li>• Highly secure with a range of security capabilities including fine grained security control &amp; auditing</li> <li>• Deploy mission critical installations using high availability option</li> </ul>	Multi-GPU Multi-Node
<b>Labellio</b>	KYOCERA Communication Systems Co	The world's easiest deep learning web service for computer vision, which allows everyone to build own image classifier with only web browser.	<ul style="list-style-type: none"> <li>• Neural net fine-tuning for image data, data crawling, data browsing as well as drag-and-drop style data cleansing backed by AI support</li> </ul>	Multi-GPU Single Node
<b>Numba</b>	Anaconda	JIT compilation of Python functions for execution on various targets (including CUDA)	<ul style="list-style-type: none"> <li>• JIT compilation of Python functions for execution on various targets (including CUDA)</li> </ul>	Multi-GPU Single Node
<b>OmniSci</b>	OmniSci	OmniSci is GPU-powered big data analytics and visualization platform that is hundreds of times faster than CPU in-memory systems. OmniSci uses GPUs to execute SQL queries on multi-billion row datasets and optionally render the results, all in milliseconds.	<ul style="list-style-type: none"> <li>• Uses LLVM's nvptx backend to generate CUDA kernels. OpenGL- (EGL) based rendering is not open source.</li> <li>• Can run in a docker container using NVIDIA-docker.</li> </ul>	Multi-GPU Single Node
<b>Polymatica</b>	Polymatica	Analytical OLAP and Data Mining Platform	<ul style="list-style-type: none"> <li>• Visualization, Reporting, OLAP in-memory with GPU acceleration, Data Mining, Machine Learning, Predictive Analytics</li> </ul>	Multi-GPU Multi-Node
<b>Sqream DB</b>	Sqream	GPU accelerated SQL database engine for big data analytics. Sqream speeds SQL analytics by 100X by translating SQL queries into highly parallel algorithms run on the GPU.	<ul style="list-style-type: none"> <li>• Up to 100TB of raw data can be stored and queried in a standard 2U server</li> <li>• Inserts and analyzes hundreds of billions of records in seconds</li> <li>• No indexes required</li> <li>• No changes to SQL code or data science paradigms required</li> </ul>	Multi-GPU Single Node

<b>SynerScope</b>	SynerScope	Big data visualization and data discovery, for combining Analytics on Analytics with IoT compute-at-the-edge smart sensors.	<ul style="list-style-type: none"> <li>• Real-time Interaction with data</li> </ul>	Single GPU Single Node
<b>ZX Lib (Fuzzy Logic)</b>	Tanay	Financial analytics and data mining library	<ul style="list-style-type: none"> <li>• Monte Carlo simulations, pricing of vanilla and exotic options, fixed income analytics, data mining</li> </ul>	Multi-GPU Single Node

## Artificial Intelligence

### DEEP LEARNING AND MACHINE LEARNING

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AlphaSense</b>	AlphaSense	*PaaS for Financial analysis based on public corporate information *Geared at financial analysts within financial services. *Allows very fast searches of public corporate information, and allows questring answering format ("the Google for Analyst research")	<ul style="list-style-type: none"> <li>• PaaS for Financial analysis based on public corporate information</li> <li>• Geared at financial analysts within financial services.</li> <li>• Allows very fast searches of public corporate information, and allows questring answering format ("the Google for Analyst research")</li> </ul>	Multi-GPU Single Node
<b>ANACONDA Enterprise</b>	Anaconda	Anaconda Enterprise takes Anaconda to the next level and makes it easy, secure, and manageable to scale powerful analytics workflows from the laptop to the server and then scaled out to your cluster, while also incorporating collaboration, publishing, security, and Hadoop-optimized deployment.	<ul style="list-style-type: none"> <li>• Anaconda Enterprise opens up the full capabilities of your GPU or multi-core processor to the Python programming language. Common operations like linear algebra, random number generation, FFT and Monte Carlo simulation run faster, and take advantage of multiple cores</li> <li>• Identify and remedy performance bottlenecks easily with data, code and in-notebook profilers</li> <li>• Includes Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND, and sorting algorithms from the CUB and Modern GPU libraries</li> </ul>	Multi-GPU Single Node
<b>Apache Mahout</b>	Apache Mahout	Mahout is building an environment for quickly creating scalable performant machine learning applications.	<ul style="list-style-type: none"> <li>• Designed to make it extremely easy to add new algorithms. Designed to be distributed instead of single machine.</li> </ul>	Multi-GPU Multi-Node
<b>ARYA.ai</b>	ARYA.ai	Deep learning platform with end-to-end workflows for Enterprise. Incorporates TensorFlow. Focusing on consumer banking and insurance industries.	<ul style="list-style-type: none"> <li>• Deep learning platform with end-to-end workflows for Enterprise. Incorporates TensorFlow.</li> </ul>	Multi-GPU Multi-Node
<b>Avitas Systems - Inspection as a Service</b>	Avitas Systems	Inspection solution offering enhanced, robotic-based autonomous inspection, advanced predictive analytics, digital inspection data warehousing, and intelligent inspection planning recommendations available in a web-based interface	<ul style="list-style-type: none"> <li>• DL for visual inspection</li> </ul>	
<b>BIDMach -</b>	UC Berkeley	The fastest machine learning library available. Holds the record for many common machine learning algorithms. Both BIDMach and its sister library BIDmat were originated at UC Berkley.	<ul style="list-style-type: none"> <li>• Written in Scala and supports Scala and Java interfaces.</li> <li>• Supports linear regression, logistic regression, SVM, LDA, K-Means and other operations.</li> </ul>	Multi-GPU Single Node
<b>Bons.ai</b>	Bons.ai	Bons.ai is an artificial intelligence platform which abstracts away the low-level, inner workings of machine learning systems to empower more developers to integrate richer intelligence models into their work.	<ul style="list-style-type: none"> <li>• Easy to use programming interface. Bons.ai has its own programming language called Inkling. Primary focus on reinforcement learning.</li> </ul>	Multi-GPU Single Node



<b>C3 Fraud Detection</b>	C3 IoT	C3 IoT is a Platform-as-a-Service for industrial customers including utilities, manufacturing, retail, finance, and healthcare. GPUs consumed exclusively through Amazon AWS. Their first DL use cases targets fraud detection from time series power consumption data.	<ul style="list-style-type: none"> <li>• Deep learning models, including RNNs</li> </ul>	Multi-GPU Single Node
<b>Caffe2</b>	Facebook	This is a faster framework for deep learning, it's forked from BVLC/caffe (master branch). This allows data-parallel via MPI.	<ul style="list-style-type: none"> <li>• Using the GPU cluster processing mass image data</li> </ul>	Multi-GPU Single Node
> <b>CatBoost</b>	Yandex	CatBoost is an open-source gradient boosting library with categorical features support	<ul style="list-style-type: none"> <li>• Extremely fast learning on GPU Multi-GPU Multi-Node</li> </ul>	Multi-GPU Multi-Node
<b>Chainer</b>	Preferred Networks, Inc.	DL framework that makes the construction of neural networks (NN) flexible and intuitive.	<ul style="list-style-type: none"> <li>• Dynamic NN construction, which makes debugging easier</li> <li>• CPU/GPU-agnostic coding, which is promoted by CuPy, partially NumPy-compatible multidimensional array library for CUDA</li> <li>• Data-dependent NN construction, which fully exploits the control flows of Python without magic</li> </ul>	Multi-GPU Multi-Node
<b>Clarifai</b>	Clarifai	Clarifai brings a new level of understanding to visual content through deep learning technologies. Clarifai uses GPUs to train large neural networks to solve practical problems in advertising, media, and search across a wide variety of industries.	<ul style="list-style-type: none"> <li>• GPU-based training and inference</li> <li>• Recognizes and indexes images with predefined classifiers, or with custom classifiers</li> </ul>	Multi-GPU Single Node
<b>CNTK</b>	Microsoft	Microsoft's Computational Network Toolkit (CNTK) is a unified computational network framework that describes deep neural networks as a series of computational steps via a directed graph.	<ul style="list-style-type: none"> <li>• Supports many applications, including Speech Recognition, Machine Translation, Image Recognition, Image Captioning, Text Processing and Relevance, Language Understanding, Language Modeling</li> </ul>	Multi-GPU Single Node
<b>Databricks cloud</b>	Databricks	Databricks cloud is GPU-accelerated PaaS offering build on top of AWS.	<ul style="list-style-type: none"> <li>• GPU instances available with CUDA drivers included,</li> <li>• GPU support provided by Spark scheduler,</li> <li>• integration of TensorFlow, Keras</li> <li>• TensorFrames data connector</li> <li>• Deep learning pipelines/workflows</li> <li>• transfer learning and image loading.</li> </ul>	Multi-GPU Multi-Node
<b>DeepBench</b>	BaiDu Research	The primary purpose of DeepBench is to benchmark operations that are important to deep learning on different hardware platforms.	<ul style="list-style-type: none"> <li>• DeepBench consists of a set of basic operations (dense matrix multiplies, convolutions and communication) as well as some recurrent layer types</li> <li>• Both forward and backward operations are tested</li> <li>• This first version of the benchmark will focus on training performance in 32-bit floating-point arithmetic</li> </ul>	Multi-GPU Single Node
<b>DeepInstinct</b>	DeepInstinct	Zero day end point malware detection solution offered to enterprise markets.	<ul style="list-style-type: none"> <li>• Zero-day threats &amp; APT attack detection on endpoints, servers and mobile devices</li> </ul>	Multi-GPU Single Node
<b>Deeplearni.ng</b>	Deeplearni.ng	Deep Learning Platform based on TensorFlow. Allows end-to-end workflows. Targeting consumer banking and insurance industries.	<ul style="list-style-type: none"> <li>• Deep learning workflows can be built</li> <li>• Based on TensorFlow</li> <li>• Use cases in consumer banking and Insurance</li> </ul>	Multi-GPU Multi-Node

<b>Deeplearning4j</b>	SkyMind	Deeplearning4j is the most popular deep learning framework for the JVM, and includes all major neural nets such as convolutional, recurrent (LSTMs) and feedforward.	<ul style="list-style-type: none"> <li>Integrates with Hadoop and Spark to run distributed</li> <li>Java and Scala APIs</li> <li>Composable framework that facilitates building your own nets</li> <li>Includes ND4J, the Numpy for Java.</li> </ul>	Multi-GPU Single Node
<b>Dextro</b>	Axon	Dextro's API uses deep learning systems to analyze and categorize videos in real-time.	<ul style="list-style-type: none"> <li>Object and scene detection</li> <li>Machine transcription for audio</li> <li>Motion and movement detection.</li> </ul>	Multi-GPU Single Node
<b>Gridspace</b>	Gridspace	Voice analytics to turn your streaming speech audio into useful data and service metrics. Instrument your contact / call center and work communications today with powerful deep learning-driven voice analytics	<ul style="list-style-type: none"> <li>Speech-to-text transcription</li> <li>Compliance</li> <li>Call grading</li> <li>Call topic modeling</li> <li>Customer service enhancement</li> <li>Customer churn prediction</li> </ul>	N/A
<b>Keras</b>	Open Source	Keras is a minimalist, highly modular neural networks library, written in Python, and capable of running on top of either TensorFlow or Theano. Keras was developed with a focus on enabling fast experimentation.	<ul style="list-style-type: none"> <li>cuDNN version depends on the version of TensorFlow and Theano installed with Keras</li> <li>Supported Interfaces: Python</li> </ul>	Multi-GPU Single Node
<b>MatConvNet</b>		CNNs for MathWorks MATLAB, allows you to use MATLAB GPU support natively rather than writing your own CUDA code	<ul style="list-style-type: none"> <li>Building Blocks, Simple CNN wrapper, DagNN wrapper, cuDNN implemented</li> </ul>	Multi-GPU Single Node
<b>Matroid</b>	Matroid	Matroid offers video classification service in the cloud. Matroid allows training video detections on a set of images and then applying those video detection. For example, video detectors can be trained to identify Steve Jobs and then search for Steve Jobs appearance anywhere in the video.	<ul style="list-style-type: none"> <li>Matroid is multi-cloud and allows it customers to easily switch between AWS, Azure and Google Cloud.</li> </ul>	Multi-GPU Multi-Node
<b>MetaMind</b>	Einstein Platform Services	Provides a deep learning API for image recognition and text sentiment analysis. Uses either prebuilt, public, or custom classifiers.	<ul style="list-style-type: none"> <li>GPU-based training and inference</li> <li>Recognizes image and analyzes text, creates and trains classifiers with tooling for uploading and managing datasets</li> </ul>	Multi-GPU Single Node
<b>MXNet</b>	Amazon	MXnet is a deep learning framework designed for both efficiency and flexibility that allows you to mix the flavors of symbolic programming and imperative programming to maximize efficiency and productivity.	<ul style="list-style-type: none"> <li>MXnet supports cuDNN v5 for GPU acceleration</li> </ul>	Multi-GPU Multi-Node
<b>Neon</b>	Nervana Systems	Neon is a fast, scalable, easy-to-use Python based deep learning framework that has been optimized down to the assembler level. Neon features a rich set of example and pre-trained models for image, video, text, deep reinforcement learning and speech applications.	<ul style="list-style-type: none"> <li>Training, inference and deployment of deep learning models. Process over 442M images per day on a Titan X</li> </ul>	Multi-GPU Single Node
<b>NVCaffe</b>	Berkeley AI Research	The Caffe deep learning framework makes implementing state-of-the-art deep learning easy.	<ul style="list-style-type: none"> <li>Process over 40M images per day with a single NVIDIA K40 or Titan GPU</li> </ul>	Single GPU Single Node

<b>PaddlePaddle</b>	PaddlePaddle	PaddlePaddle (PARallel Distributed Deep LEarning) is an easy-to-use, efficient, flexible and scalable deep learning platform, which is originally developed by Baidu scientists and engineers for the purpose of applying deep learning to many products at Baidu.	<ul style="list-style-type: none"> <li>Optimized math operations through SSE/AVX intrinsics, BLAS libraries (e.g. MKL, ATLAS, cuBLAS) or customized CPU/GPU kernels</li> <li>Highly optimized recurrent networks which can handle variable-length sequence without padding</li> <li>Optimized local and distributed training for models with high dimensional sparse data</li> </ul>	Multi-GPU Single Node
<b>Sentient</b>	Sentient	Sentient is an AI platform company with special focus on digital marketing, ecommerce and finance trading applications.	<ul style="list-style-type: none"> <li>Sentient is using GPU deep learning in its commercially available ecommerce, digital marketing and financial trading applications. Studio.ml is a new project designed to make AI development easier by hiding most of the complexity. Studio.ml runs on-premise and in the cloud.</li> </ul>	
<b>SpaceKnow PaaS</b>	SpaceKnow	PaaS for deep learning extraction of satellite data information, targeted at Financial Services and Defense / Intelligence Can track macro/micro-economic activity by applying deep learning to satellite images.	<ul style="list-style-type: none"> <li>Extracts economic activity from satellite images using deep learning</li> <li>Can provide batch mode extraction</li> </ul>	Multi-GPU Multi-Node
<b>Tensorflow</b>	Google	Google's TensorFlow is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them.	<ul style="list-style-type: none"> <li>TensorFlow is flexible, portable and performant creating an open standard for exchanging research ideas and putting machine learning in products</li> </ul>	Multi-GPU Single Node
<b>Theano</b>	LISA Lab	Theano is a symbolic expression compiler that powers large-scale computationally intensive scientific investigations.	<ul style="list-style-type: none"> <li>Abstract expression graphs for transparent GPU acceleration.</li> </ul>	Multi-GPU Single Node
<b>Torch7</b>	Open Source	Torch7 is an interactive development environment for machine learning and computer vision.	<ul style="list-style-type: none"> <li>Computational back-ends for multicore GPUs.</li> </ul>	Multi-GPU Single Node
<b>UETorch</b>	Facebook	It provides an embedded Torch environment within the powerful Unreal Engine 4. This allows one to have deep learning models directly interact with the game world, and paves way for powerful research. An example of doing AI Research using UETorch is for a neural network to learn physics and intuition about the real world.	<ul style="list-style-type: none"> <li>Game interaction and physics, CUDA-optimized deep learning and neural networks</li> <li>CuDNN supported</li> </ul>	Multi-GPU Single Node
<b>Unify.ID</b>	Unify.ID	Behavioral user authentication service	<ul style="list-style-type: none"> <li>Identifies individuals based on unique factors, such as the way they walk, type and sit.</li> </ul>	Multi-GPU Single Node
<b>Visual Intelligence API</b>	Deep Vision	Deep Vision specializes in understanding visual content and getting the most value of data by applying visual recognition for enterprises.	<ul style="list-style-type: none"> <li>Visual Intelligence API allows leader enterprises in verticals like e-commerce and online auctions, media and entertainment and retailers, to analyze content related with faces, brands and context tags to perform actions like: Curate and organize visual content Search and recommend visually Get insights and analytics visually</li> </ul>	

# Federal, Defense and Intelligence

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Advanced Ortho Series</b>	DigitalGlobe	Geospatial visualization	<ul style="list-style-type: none"> <li>Image orthorectification</li> </ul>	Multi-GPU Single Node
<b>ArcGIS Pro</b>	ESRI	Viewshed2 - Determines the raster surface locations visible to a set of observer features, using geodesic methods. Aspect - Determines the compass direction that the downhill slope faces for each location Slope - Determines the slope (gradient or steepness) from each cell of a raster	<ul style="list-style-type: none"> <li>Viewshed2 - transforms the elevation surface into a geocentric 3D coordinate system and runs 3D sightlines to each transformed cell center.</li> <li>Aspect - The values of each cell in the output raster indicate the compass direction the surface faces at that location. It is measured clockwise in degrees from 0 (due north) to 360 (again due north), coming full circle.</li> <li>Slope - The output slope raster can be calculated in two types of units, degrees or percent (percent rise).</li> </ul>	Multi-GPU Single Node
<b>Blaze Terra</b>	Eternix	Geospatial visualization	<ul style="list-style-type: none"> <li>3D visualization of geospatial data</li> </ul>	Multi-GPU Single Node
<b>Elcomsoft</b>	Elcomsoft	High-performance distributed password recovery software with NVIDIA GPU acceleration and scalability to over 10,000 workstations.	<ul style="list-style-type: none"> <li>GPU acceleration for password recovery</li> <li>10-100x speedup for password recovery</li> </ul>	Multi-GPU Single Node
<b>ENVI</b>	Harris	Image Processing and Analytics	<ul style="list-style-type: none"> <li>Image orthorectification</li> <li>Image transformation</li> <li>Atmospheric correction</li> <li>Panchromatic co-occurrence texture filter</li> </ul>	Multi-GPU Single Node
<b>Geomatics GXL</b>	PCI	Image processing	<ul style="list-style-type: none"> <li>Image orthorectification and additional image processing</li> </ul>	Multi-GPU Single Node
<b>GeoWeb3d Desktop</b>	Geoweb3d	Geospatial visualization of 3D and 2D data; mensuration; mission planning	<ul style="list-style-type: none"> <li>3D visualization and analysis of geospatial data</li> </ul>	Multi-GPU Single Node
<b>Graphistry</b>	Graphistry	Graphistry is the first visual investigation platform to handle increasing enterprise-scale workloads.	<ul style="list-style-type: none"> <li>Graph reasoning</li> <li>GPU-accelerated visual analytics, visual pivoting, and rich investigation templating</li> </ul>	Multi-GPU Single Node
<b>Ikena ISR</b>	MotionDSP	Real-time full motion video (FMV) and wide-area motion imagery (WAMI) enhancement and computer-vision-based analytics software for intelligence analysts	<ul style="list-style-type: none"> <li>Real-time super-resolution-based video enhancement on live streams, geospatial visualization, target detection and tracking, and fast 2-D mapping</li> </ul>	Multi-GPU Single Node
<b>LuciadLightspeed</b>	Luciad	Geospatial visualization and analysis	<ul style="list-style-type: none"> <li>Geospatial situational awareness</li> </ul>	Single GPU Single Node
<b>Manifold Systems</b>	Manifold Systems	Full-featured GIS, vector/raster processing & analysis	<ul style="list-style-type: none"> <li>Manifold surface tools</li> </ul>	Multi-GPU Single Node
> <b>OmniSIG</b>	deepsig.io	The OmniSig sensor provides a new class of RF sensing and awareness using DeepSig's pioneering application of Artificial Intelligence (AI) to radio systems. Going beyond the capabilities of existing spectrum monitoring solutions, OmniSIG is able to not only detect and classify signals but understand the spectrum environment to inform contextual analysis and decision making. Compared to traditional approaches, OmniSIG provides higher sensitivity and accuracy, is more robust to harsh impairments and dynamic spectrum environments, and requires less computational resources and dynamic range.	<ul style="list-style-type: none"> <li>The OmniSIG sensor typically operates in a real-time streaming fashion, ingests radio samples from many common radio interfaces, and can make use of packet formats like VITA49 or SDDS. The web-based user interface can be used from any device with a browser, including mobile handsets, and the OmniSIG software also provides its metadata output stream in JSON form for use by other applications.</li> </ul>	Multi-GPU Single Node

<b>SNEAK</b>	OpCoast	Electromagnetic signals propagation modeling for complex urban and terrain environments.	<ul style="list-style-type: none"> <li>• Ray tracing, DTED and remote sensing inputs.</li> </ul>	Multi-GPU Single Node
<b>SocetGXP</b>	BAE Systems	The Automatic Spatial Modeler (ASM) is designed to generate 3-D point clouds with accuracy similar to LiDAR, which can extract 3-D objects from stereo images. ASM can extract dense 3-D point clouds from stereo images, and extract accurate building edges and corners from stereo images with high resolution, large overlaps, and high dynamic range.	<ul style="list-style-type: none"> <li>• Automated 3D feature extraction</li> </ul>	Multi-GPU Single Node
<b>Terrabuilder PhotoMesh</b>	Skyline Software	PhotoMesh integrates a GPU-based, fast algorithm, able to automatically build 3D models from simple photographs. PhotoMesh revolutionizes the use of geospatial data by fully automating the generation of high-resolution, textured, 3D mesh models from standard 2D images.	<ul style="list-style-type: none"> <li>• 3D model building from imagery; building texture generation.</li> </ul>	Multi-GPU Single Node

## Design for Manufacturing/Construction: CAD/CAE/CAM

### COMPUTATIONAL FLUID DYNAMICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AcuSolve</b>	Altair	General purpose CFD software	<ul style="list-style-type: none"> <li>• Linear equation solver</li> </ul>	Multi-GPU Single Node
> <b>ADS Flow Solver - Code LEO</b>	ADSCFD Inc.	CFD simulation for turbochargers, turbines, and compressors	<ul style="list-style-type: none"> <li>• URANS, structured/unstructured solver, explicit density based</li> </ul>	Multi-GPU Multi-Node
<b>ANSYS Fluent</b>	ANSYS	General purpose CFD software	<ul style="list-style-type: none"> <li>• Linear equation solver</li> <li>• Radiation heat transfer model</li> <li>• Discrete Ordinate Radiation model</li> </ul>	Multi-GPU Multi-Node
<b>ANSYS Icepak</b>	ANSYS	CFD software for electronics thermal management	<ul style="list-style-type: none"> <li>• Linear Equation Solver</li> </ul>	Multi-GPU Multi-Node
<b>ANSYS Polyflow</b>	ANSYS	CFD software for the analysis of polymer and glass processing	<ul style="list-style-type: none"> <li>• Direct Solvers</li> </ul>	Multi-GPU Single Node
<b>CPFD Barracuda-VR and Barracuda</b>	CPFD	Fluidized bed modeling software	<ul style="list-style-type: none"> <li>• Linear equation solver, particle calculations</li> </ul>	Single GPU Single Node
<b>Culises for OpenFOAM</b>	FluiDyna	Solver library for general purpose CFD software	<ul style="list-style-type: none"> <li>• Linear equation solvers</li> </ul>	Multi-GPU Single Node
<b>EXN/Aero</b>	Envenio	On-demand HPC-cloud CFD solver	<ul style="list-style-type: none"> <li>• Multiphase, heat transfer, buoyancy, multi-grid, concurrent single/double precision, ideal gas, incompressible &amp; compressible flows, RANS/LES/DES, conjugate heat transfer</li> </ul>	Multi-GPU Multi-Node
<b>FFT Actran</b>	FFT	Simulation of acoustics propagation at high frequency or in huge domains such as exhaust of turbomachines, full truck cabin exterior acoustics, and ultrasonic parking sensors.	<ul style="list-style-type: none"> <li>• Discontinuous Galerkin Method (DGM) solver</li> </ul>	Multi-GPU Single Node
<b>FINE/Turbo</b>	Numeca International	Structured, multi-block, multi-grid CFD solver targeting the turbo machinery industry	<ul style="list-style-type: none"> <li>• Multi-grid solver</li> </ul>	Multi-GPU Single Node

> <b>Pacefish</b>	Numeric Systems	CFD application for ground transportation and building aerodynamics	<ul style="list-style-type: none"> <li>• Lattice-Boltzmann Method for single-phase flows</li> <li>• transient simulation</li> <li>• isothermal modeling</li> <li>• integrated fast and robust pre-processor for complex geometries</li> <li>• local grid refinement</li> <li>• uRANS (K-Omega-SST), hybrid uRANS-LES (SST-DDES &amp; SST-IDDES)</li> <li>• LES (Smagorinsky) turbulence modeling</li> </ul>	Multi-GPU Single Node
<b>GeoPlat-RS</b>	GridPoint Dynamics (GPD)	Geoplat Pro-RS is a parallel hydrodynamic simulator with a flexible architecture. This enables to reduce the time for writing the entire simulator by 2/3, and, as consequence, to quickly bring new physical processes into the algorithm. Current stage of development: Implementation of BlackOil model; Creation of pre- and post-processing modules (BlackOil model). Implementation of compositional model; Creation of pre- and post-processing modules (BlackOil model and compositional model).	<ul style="list-style-type: none"> <li>• CUDA, Spectral Decomposition with CUFFT library</li> </ul>	Multi-GPU Single Node
<b>HiFUN</b>	SANDI	High Resolution Flow Solver on Unstructured Meshes. State-of-the-art Euler/RANS solver. Super scalability on massively parallel HPC platforms. The code is ported using OpenACC directives for NVIDIA GPU	<ul style="list-style-type: none"> <li>• HiFUN imbibes most recent CFD technologies; many of them home grown</li> <li>• HiFUN exhibits highly scalable parallel performance with its ability to scale upto several thousand processors on massively parallel computing platforms</li> <li>• Capable of handling complex geometries and flow physics arising in high lift flows</li> </ul>	Multi-GPU Single Node
<b>midas NFX(CFD)</b>	Midas	General purpose CFD software based on FEM	<ul style="list-style-type: none"> <li>• Linear equation solver (Iterative Solver and AMG Preconditioner)</li> </ul>	Single GPU Single Node
<b>MIKE 21</b>	DHI	2D hydrological modelling of coast and sea	<ul style="list-style-type: none"> <li>• Hydrodynamics</li> <li>• Advection-dispersion</li> <li>• Sand and mud transport</li> <li>• Coupled modelling</li> <li>• Particle tracking</li> <li>• Oil spill</li> <li>• Ecological modelling</li> <li>• Agent based modelling</li> <li>• Various wave models</li> </ul>	Multi-GPU Single Node
<b>MIKE FLOOD</b>	DHI	1D & 2D urban, coastal, and riverine flood modelling	<ul style="list-style-type: none"> <li>• Hydrodynamics</li> </ul>	Multi-GPU Single Node
<b>Moldflow</b>	Autodesk	Plastic mold injection software	<ul style="list-style-type: none"> <li>• Linear equation solver</li> </ul>	Single GPU Single Node
<b>nanoFluidX</b>	Altair	Meshless CFD solver (Smoothed Particle Hydrodynamics, SPH)	<ul style="list-style-type: none"> <li>• Single/multi-phase flows, thermal, moving/rotating geometries, inlet/outlet boundary conditions</li> </ul>	Multi-GPU Multi-Node
<b>Numerix</b>	Zeus	Simulation of flow around buildings	<ul style="list-style-type: none"> <li>• Discrete computational technique</li> </ul>	Multi-GPU Single Node

<b>Particleworks</b>	Prometech	CFD software using MPS (Moving Particle Simulation) method for automotive, energy, material, chemical processing, medical, food, and civil engineering industries where free surface fluid flow and fluid mixing phenomena occur.	<ul style="list-style-type: none"> <li>viscosity model, viscosity/pressure term solution, turbulence model, airflow, surface tension model, rigid body, thermal properties, external force and aeration.</li> <li>boundary conditions: particle wall, polygon wall, inflow and outflow boundary, simulation domain, and pump.</li> </ul>	Multi-GPU Multi-Node
<b>PowerViz</b>	Exa	Industry proven, modern post-processing app for EXA POWERFLOW CFD	<ul style="list-style-type: none"> <li>Rendering</li> <li>Ray tracing</li> </ul>	Multi-GPU Single Node
<b>Simcenter 3D</b>	Siemens PLM Software	Industry proven, modern pre- & post-processing app for multidiscipline CAE	<ul style="list-style-type: none"> <li>Rendering, Raytracing</li> </ul>	Multi-GPU Single Node
<b>Speed IT FLOW</b>	Vratis	Incompressible single-phase CFD software	<ul style="list-style-type: none"> <li>Finite-volume solver: Simple and piso, incompressible single-phase flows with k-OmegaSST turbulence</li> </ul>	Single GPU Single Node
> <b>TACOMA</b>	GE GRC Brian Mitchell			
<b>Turbostream</b>	Turbostream Ltd.	CFD software for turbomachinery flows	<ul style="list-style-type: none"> <li>Explicit solver</li> </ul>	Multi-GPU Single Node
<b>ultraFluidX</b>	Altair	Lattice-Boltzmann-based CFD solver for ground transportation aerodynamics	<ul style="list-style-type: none"> <li>Single-phase flows, isothermal, integrated volume mesh generation, local refinement, LES turbulence modeling</li> </ul>	Multi-GPU Multi-Node
> <b>zCFD</b>	Zenotech Simulation Unlimited	General purpose CFD solver		

## CFD (Research Developments)

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
> <b>ALYA</b>	Barcelona Supercomputing Center (BSC)	Alya is a high performance computational mechanics code to solve complex coupled multi-physics / multi-scale problems, which are mostly coming from the engineering realm.	<ul style="list-style-type: none"> <li>Incompressible Flows Compressible Flows Non-linear Solid Mechanics Species transport equations Excitable Media Thermal Flows N-body collisions</li> </ul>	Multi-GPU Multi-Node
<b>DualSPHysics</b>	University of Manchester	SPH-based CFD software	<ul style="list-style-type: none"> <li>SPH model</li> </ul>	Multi-GPU Single Node
<b>GIN3D</b>	Boise St - Senocak	General purpose CFD software for incompressible flows	<ul style="list-style-type: none"> <li>Implicit solver</li> </ul>	Multi-GPU Single Node
<b>HiFILES</b>	Stanford - Jameson	General purpose CFD software for compressible flows.	<ul style="list-style-type: none"> <li>Explicit solver</li> </ul>	Multi-GPU Multi-Node
<b>HiPSTAR</b>	University of Southampton and University of Melbourne - Sandberg	CFD software for compressible reacting flows	<ul style="list-style-type: none"> <li>Explicit solver</li> </ul>	Multi-GPU Single Node
<b>JENRE, Propel (NRL)</b>	US Naval Research Lab	CFD software for compressible flows	<ul style="list-style-type: none"> <li>Explicit solver</li> </ul>	Multi-GPU Single Node
<b>MUST</b>	ITP Aero Engines	An edge-based CFD solver for RANS eqns on unstructured grid. It is proprietary to IPT, a division of Rolls Royce, and used in-house. GPU-based version of the MUST solver is used in production for the design of turbines and compressors	<ul style="list-style-type: none"> <li>Flux computations over the edges using a multigrid solver</li> </ul>	Multi-GPU Single Node
<b>PyFR</b>	Imperial College - Vincent	General purpose CFD software for compressible flows.	<ul style="list-style-type: none"> <li>High-order explicit solver based on flux reconstruction method</li> </ul>	Multi-GPU Multi-Node
<b>RAPTOR</b>	US DOE	CFD formulation of turbulent combustion for fuel injector and other engine applications	<ul style="list-style-type: none"> <li>Flow solver</li> </ul>	Multi-GPU Multi-Node
<b>S3D</b>	Sandia and Oak Ridge NL	Direct numerical solver (DNS) for turbulent combustion	<ul style="list-style-type: none"> <li>Chemistry model</li> </ul>	Multi-GPU Multi-Node

> Indicates new application

## COMPUTATIONAL STRUCTURAL MECHANICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Adams</b>	MSC Software	Multi-Body Dynamics simulation software	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>ANSYS Mechanical</b>	ANSYS	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Direct and iterative solvers</li> </ul>	Multi-GPU Multi-Node
<b>EDEM</b>	DEM Solutions, Ltd.	Market-leading Discrete Element Method (DEM) software for bulk material simulation	<ul style="list-style-type: none"> <li>• DEM solver</li> </ul>	Single GPU Single Node
> <b>GranuleWorks</b>	Prometech	DEM-based advanced simulator for granular materials in pharma and powder metallurgy: granular material segregation, screening, grinding, screw conveying, mixing, compaction, filling, dustproof, toner transport, electrode materials filling, cliff collapses/debris flow, etc.	<ul style="list-style-type: none"> <li>• size distribution, contact force model, rolling resistance model, liquid bridge force model, van der Waals force model, heat transfer and external force.</li> <li>• boundary conditions: polygon wall, inflow and outflow boundary, and simulation domain.</li> <li>• coupling with Particleworks MPS solver: support for aeration and pumps</li> </ul>	Multi-GPU Multi-Node
<b>Helyx PEM</b>	Engys	Specialised add-on solver for HELYX to simulate large numbers of solid objects in motion using the Polyhedral Element Method (PEM)	<ul style="list-style-type: none"> <li>• Polyhedral Elements Method solver</li> </ul>	Single GPU Single Node
<b>HyperWorks</b>	Altair	Industry proven, modern pre- & post-processing app for CAE	<ul style="list-style-type: none"> <li>• Rendering</li> <li>• Anti-aliasing</li> <li>• Ambient Occlusion</li> </ul>	Single GPU Single Node
<b>Impetus Afea</b>	Impetus Afea	Predicts large deformations of structures and components exposed to extreme loading conditions	<ul style="list-style-type: none"> <li>• Non-linear Explicit Finite-Element Solver</li> </ul>	Multi-GPU Single Node
> <b>Irazu</b>	Geomechanica Inc.	Simulation and analysis tool for rock mechanics, involving large deformations, fracturing and multi-physics phenomena	<ul style="list-style-type: none"> <li>• Explicit 2D and 3D FEM and FDEM solvers; Coupled hydraulic, mechanical, transport, thermal and fracture processes</li> </ul>	Single GPU Single Node
<b>LS-Dyna Implicit</b>	LSTC	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Linear equation solver</li> </ul>	Multi-GPU Single Node
<b>Marc</b>	MSC Software	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Direct sparse solver</li> </ul>	Multi-GPU Single Node
<b>midas GTS NX</b>	Midas	Simulation tool for geo-technical analysis	<ul style="list-style-type: none"> <li>• Linear equation solver(Multi Frontal Solver)</li> </ul>	Single GPU Single Node
<b>midas NFX(Structural)</b>	Midas	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Linear equation solver(Multi Frontal Solver)</li> </ul>	Single GPU Single Node
<b>MSC Nastran</b>	MSC Software	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Direct sparse solver</li> </ul>	Multi-GPU Single Node
<b>NX Nastran</b>	Siemens PLM Software	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Linear and nonlinear equation solver</li> </ul>	Multi-GPU Multi-Node
<b>OptiStruct</b>	Altair	Industry proven, modern structural analysis solver and solution for structural design and optimization.	<ul style="list-style-type: none"> <li>• Direct and iterative solvers</li> </ul>	Single GPU Single Node
<b>RADIOSS Implicit</b>	Altair	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Direct and iterative solvers</li> </ul>	Multi-GPU Single Node
> <b>RecurDyn</b>	FunctionBay, Inc.	Multi-Flexible Body Dynamics simulation software	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>Rocky DEM</b>	Rocky DEM	Discrete Element Modeling (DEM)-based particle simulation software	<ul style="list-style-type: none"> <li>• Explicit DEM solver (dry/sticky contact rheologies)</li> <li>• 1-way &amp; 2-way coupling with ANSYS Fluent and ANSYS Mechanical</li> </ul>	Multi-GPU Single Node
<b>SIMULIA 3DEXPERIENCE</b>	Dassault Systèmes SIMULIA Corp.	Realistic simulation solution (Uses Abaqus Standard for GPU computing)	<ul style="list-style-type: none"> <li>• Direct sparse solver</li> </ul>	Single GPU Single Node



<b>SIMULIA Abaqus/Standard</b>	Dassault Systèmes SIMULIA Corp.	Simulation and analysis tool for structural mechanics	<ul style="list-style-type: none"> <li>• Direct sparse solver</li> <li>• AMS Solver</li> <li>• Steady State Dynamics</li> </ul>	Multi-GPU Multi-Node
--------------------------------	------------------------------------	-------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------	-------------------------

## DESIGN AND VISUALIZATION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3DEXCITE DeltaGen</b>	Dassault Systèmes	High-end 3D visualization and realtime interaction to help increase visual quality, speed, and flexibility.	<ul style="list-style-type: none"> <li>• Interactive ray tracing and global illumination.</li> <li>• Integration with Siemens TeamCenter.</li> <li>• Cluster support Realtime &amp; Offline Production Process Integration and scene building.</li> <li>• Scene Analysis, Xplore DeltaGen, SDK for DeltaGen.</li> </ul>	Multi-GPU Single Node
<b>Abaqus/CAE</b>	Dassault Systèmes SIMULIA Corp.	Complete solution for Abaqus finite element modeling, visualization, and process automation	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Multi-GPU Single Node
<b>Accelerad</b>	MIT Sustainable Design Lab	Accelerad is a free suite of programs for fast and accurate lighting and daylighting analysis and visualization. It was developed by Nathaniel Jones at the MIT Sustainable Design Lab and modeled after the popular Radiance software suite developed by Greg Ward at Lawrence Berkeley National Laboratory.	<ul style="list-style-type: none"> <li>• Up to forty times faster using OptiX</li> <li>• Renderings with large numbers of ambient bounces</li> <li>• Calculations over many thousands of sensor points</li> <li>• Fast simulation of annual climate-based daylighting metrics</li> <li>• AcceleradRT - Interactive interface for real-time daylighting, glare, and visual comfort analysis with validated accuracy. includes AcceleradVR, an immersive visualization interface compatible with most virtual reality headsets.</li> </ul>	N/A
<b>Allplan</b>	Nemetschek	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction	<ul style="list-style-type: none"> <li>• OpenGL and DirectX based GPU rendering</li> </ul>	Single GPU Single Node
<b>ANSA</b>	BETA CAE Systems	Industry proven, modern pre-processing app for CAE	<ul style="list-style-type: none"> <li>• OpenGL</li> </ul>	Single GPU Single Node
<b>ANSYS Discovery Live</b>	ANSYS	Interactive and CAD-agnostic Windows-based app that gives engineers instantaneous simulation results to help them explore and refine product designs	<ul style="list-style-type: none"> <li>• OpenGL-based visualization &amp; CUDA-based Structural, Fluid Dynamics and Thermal simulations</li> </ul>	Single GPU Single Node
<b>ANSYS Workbench</b>	ANSYS	Industry proven, modern pre- & post-processing app for CAE	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Multi-GPU Single Node
<b>Apex</b>	MSC Software	Industry proven, modern pre- & post-processing app for CAE	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>ArchiCAD</b>	Nemetschek	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction	<ul style="list-style-type: none"> <li>• OpenGL and DirectX based GPU rendering</li> </ul>	Single GPU Single Node
<b>AutoCAD</b>	Autodesk	2D and 3D CAD design, drafting, modeling, architectural drawing, and engineering software. Supports Open GL. Native DWG support.	<ul style="list-style-type: none"> <li>• Surface, mesh, and solid modeling tools, model documentation tools, parametric drawing capabilities</li> <li>• Native DWG support</li> <li>• GRID Support.</li> </ul>	Single GPU Single Node
<b>CATIA</b>	Dassault Systèmes	The reference CAD application for advanced engineering with batching capability and extreme reliability used by 80 of the automotive industry and the entire aerospace industry.	<ul style="list-style-type: none"> <li>• GPU performance scaling</li> <li>• VR native integration with HTC Vive</li> </ul>	Single GPU Single Node

<b>CATIA Live Rendering</b>	Dassault Systèmes	Realistic 3D Rendering on full CATIA 3D CAD model	<ul style="list-style-type: none"> <li>Physically Based Rendering with no data preparation thanks to native NVIDIA Iray Photoreal integration and interactive realistic rendering using NVIDIA Iray IRT</li> </ul>	Multi-GPU Single Node
<b>COMSOL</b>	COMSOL	General Purpose CAE simulation software	<ul style="list-style-type: none"> <li>OpenGL</li> </ul>	Multi-GPU Single Node
<b>Creo Parametric</b>	PTC	Parametric design solution suite.	<ul style="list-style-type: none"> <li>Anti-aliasing, better lighting and enhanced shaded-with-edges mode</li> <li>Immersive design environment with realistic materials. GRID Support.</li> </ul>	Single GPU Single Node
<b>FEMAP</b>	Siemens PLM	Industry proven, modern pre- & post-processing app for CAE	<ul style="list-style-type: none"> <li>Rendering</li> </ul>	Single GPU Single Node
<b>HIM</b>	Optis	Human Ergonomics Evaluation software for MRO (Maintenance Repair Overall) with large model dynamic review in VR with accurate physics behavior	<ul style="list-style-type: none"> <li>PhysX</li> </ul>	
<b>IC.IDO</b>	ESI Group	Immersive VR solution for engineering and virtual prototyping. The Helios rendering engine is highly optimized for NVIDIA GPUs.	<ul style="list-style-type: none"> <li>NV Pro Pipeline (RiX) for OpenGL rendering, VRWorks SPS and VR SLI (NVLink support), and DesignWorks, including VR Occlusion Culling open source sample and OptiX</li> </ul>	Multi-GPU Single Node
<b>Inventor</b>	Autodesk	3D mechanical design, documentation, and product simulation.	<ul style="list-style-type: none"> <li>Uses BIM for intelligent building components to improve design accuracy.</li> </ul>	Single GPU Single Node
<b>Iray</b>	NVIDIA	A ready-to-integrate, physically-based, photorealistic rendering solution.	<ul style="list-style-type: none"> <li>Iray Interactive</li> <li>Iray Photoreal</li> <li>Iray Cluster</li> <li>Fast interactive ray tracing</li> <li>Physically-based, global-illumination rendering</li> <li>Distributed cluster rendering.</li> </ul>	Multi-GPU Single Node
<b>Iray for 3ds Max</b>	Lightworks Design	Iray physically-based renderer plugin for Autodesk 3ds Max	<ul style="list-style-type: none"> <li>Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising</li> </ul>	Multi-GPU Multi-Node
<b>Iray for Maya</b>	0x1 Software & Consulting GmbH	Iray physically-based renderer plugin for Autodesk Maya.	<ul style="list-style-type: none"> <li>Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising</li> </ul>	Multi-GPU Multi-Node
<b>Iray for Rhino</b>	migenius	Iray plugin for Rhino.	<ul style="list-style-type: none"> <li>Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support.</li> </ul>	
<b>Iray Server</b>	NVIDIA	The scaling solution for any Iray based application	<ul style="list-style-type: none"> <li>Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising</li> </ul>	Multi-GPU Multi-Node
<b>NX</b>	Siemens PLM	Siemens PLM Software premium design app, has full Iray integration and support multi-gpu rendering. Still CPU bound for most tasks otherwise	<ul style="list-style-type: none"> <li>Iray, MDL</li> </ul>	Multi-GPU Multi-Node
<b>Patran</b>	MSC Software	Industry proven, modern pre- & post-processing app for CAE	<ul style="list-style-type: none"> <li>Rendering</li> </ul>	Single GPU Single Node
<b>PLM Software NX and Teamcenter</b>	Siemens	Product lifecycle management solutions from design to simulation to production to service.	<ul style="list-style-type: none"> <li>Design software, NX, and PLM viewer applications, TcVis and Active Workspace</li> <li>GRID support</li> </ul>	Single GPU Single Node

<b>Recap PRO</b>	Autodesk	ReMake is a solution for converting reality captured with photos or scans into high-definition 3D meshes. These meshes that can be cleaned up, fixed, edited, scaled, measured, re-topologized, decimated, aligned, compared and optimized for downstream workflows entirely in ReMake.	<ul style="list-style-type: none"> <li>• Generation of 3D meshed models from laser scans or photos of an object</li> <li>• GPU accelerated photogrammetry process from 2D to 3D. 3D model display accelerated by GPU's for smooth navigation of converted models in all display modes</li> </ul>	Multi-GPU Single Node
<b>Review</b>	PiXYZ	Import any CAD data to prepare and experience your content with VR	<ul style="list-style-type: none"> <li>• Large CAD file support with NVIDIA Pascal Single Pass Stereo extension integration</li> </ul>	Single GPU Single Node
<b>Revit</b>	Autodesk	Building Information Modeling (BIM) for architecture, engineering, and construction.	<ul style="list-style-type: none"> <li>• Modeling (BIM) to design, build, and maintain higher-quality, more energy-efficient buildings</li> <li>• GRID support</li> </ul>	Single GPU Single Node
> <b>Rhino</b>	McNeel & Assoc.	General purpose conceptual/ industrial design software for AEC and Manufacturing industries, including a real-time ray-traced display mode that is CUDA-based.		Single GPU Single Node
> <b>Siemens STAR-CCM+</b>	Siemens PLM	Post-processing for CFD-focused multiphysics simulation	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
> <b>Siemens STAR-CCM+ VR</b>	Siemens PLM	Immersive VR for CFD results visualization	<ul style="list-style-type: none"> <li>• HTC Vive virtual reality headset</li> </ul>	Single GPU Single Node
<b>Simpleware</b>	Synopsys	3D image data visualization, analysis and model generation software	<ul style="list-style-type: none"> <li>• OpenGL</li> </ul>	Single GPU Single Node
> <b>SolidEdge</b>	Siemens PLM	Mid range CAD option from Siemens	<ul style="list-style-type: none"> <li>• n/a</li> </ul>	Single GPU Single Node
<b>SOLIDWORKS</b>	Dassault Systèmes	Covers all aspects of product development process with a seamless, integrated workflow; design, verification, sustainable design, communication and data management.	<ul style="list-style-type: none"> <li>• High performance in Shaded, Shaded w/ Edges, and RealView modes, FSAA for sharp edges, Order Independent Transparency</li> <li>• Real time photorealistic renderings with SOLIDWORKS Visualize, an Iray-based application.</li> </ul>	Single GPU Single Node
<b>SOLIDWORKS Visualize</b>	Dassault Systèmes	Easy to use photorealistic rendering software	<ul style="list-style-type: none"> <li>• Iray-based ray-tracing, animation support, network rendering.</li> </ul>	Single GPU Single Node
> <b>Speos</b>	Optis	Light simulation and ray tracing software		Multi-GPU Single Node
<b>Studio</b>	PiXYZ	Interactively prepare & optimize any CAD data before using your favorite staging tool	<ul style="list-style-type: none"> <li>• Large scale CAD format, support for multi-CAD file standard, prepare, optimize and heal your geometry before experiencing it in VR</li> </ul>	Single GPU Single Node
<b>Substance Designer</b>	Allegorithmic	Material shader edition, market reference for procedural texture creation.	<ul style="list-style-type: none"> <li>• Iray rendering including textures/ substances and bitmap texture export to render in any Iray powered compatible with MDL</li> </ul>	Multi-GPU Single Node
<b>Substance Painter</b>	Allegorithmic	Intuitive interactive 3D painting software with physics and particle support.	<ul style="list-style-type: none"> <li>• Iray rendering to enhance all artwork released with the software</li> </ul>	Multi-GPU Single Node
<b>T-FLEX CAD</b>	Top Systems	3D and 2D parametric design, simulation, photorealistic rendering	<ul style="list-style-type: none"> <li>• High performance visualization, real time photorealistic rendering, CUDA</li> </ul>	Multi-GPU Single Node
<b>Theia-RT</b>	Optis	Light simulation stand-alone validation tool with OpenGL PBR engine, deterministic ray tracer, Monte-Carlo progressive GPU renderer (beta) with a strong physics focus for engineering decision making.	<ul style="list-style-type: none"> <li>• Fast real-time engine integrating complex precomputed light effects</li> </ul>	Multi-GPU Single Node
<b>UE4</b>	Epic Games	Unreal Engine 4 is a suite of integrated tools for game developers to design and build games, simulations, and visualizations	<ul style="list-style-type: none"> <li>• - GPU Accelerated Rendering on OpenGL, DirectX and Vulkan. - Phys-X implemented.</li> </ul>	Multi-GPU Single Node

<b>VRED</b>	Autodesk	VRED 3D visualization software helps automotive designers and engineers create product presentations, design reviews, and virtual prototypes. Use Digital Prototyping to quickly visualize ideas and evaluate designs.	<ul style="list-style-type: none"> <li>• Enhanced geometry behavior</li> <li>• Automotive product interoperability</li> <li>• Navigation in a scene</li> <li>• Import Alias layer structure</li> <li>• Asset Manager improvements</li> <li>• Integrated file converter</li> <li>• Analytic rendering modes</li> <li>• Gap Analysis tool</li> <li>• Oculus Rift support</li> <li>• Animation module</li> <li>• Multiple rendering modes</li> <li>• Subsurface scattering</li> <li>• Displacement mapping</li> </ul>	Multi-GPU Single Node
<b>VRX</b>	Optis	Driving, headlamp, ADAS/AD simulators leveraging Theia-RT render engine and dynamic environment details (traffic, pedestrians...). VRX connects to Matlab/Simulink and ingests camera, IR, Lidar, Ultrasound and Radar sensor simulation.	<ul style="list-style-type: none"> <li>• VR simulation with HMD and CAVE support</li> </ul>	Multi-GPU Single Node
<b>Vectorworks</b>	Nemetschek	Building Information Modeling (BIM) enabled design software for the Architecture, Landscape, and Entertainment industries	<ul style="list-style-type: none"> <li>• OpenGL based GPU rendering</li> </ul>	Multi-GPU Single Node
<b>WYSIWYG</b>	Cast Software	The WYSIWYG software products, designed specifically for lighting professionals, offers a range of solutions to meet the needs of designers, assistants, electricians, console operators, teachers, and students.	<ul style="list-style-type: none"> <li>• The speed of wysiwyg's Shaded Views depends entirely on GPU, the GPU will have an easier time rendering ten risers consolidated into one Mesh, than rendering them as individual risers, Wysiwyg also support NVIDIA SLI technologies</li> </ul>	Multi-GPU Single Node

## ELECTRONIC DESIGN AUTOMATION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ANSYS HFSS</b>	ANSYS	Simulation tool for modeling 3-D full-wave electromagnetic fields in high-frequency and high-speed electronic components.	<ul style="list-style-type: none"> <li>• Transient solver</li> <li>• FEM solver</li> </ul>	Multi-GPU Single Node
<b>&gt; ANSYS HFSS SBR+</b>	ANSYS	Simulation tool for installed antenna performance and antenna-to-antenna coupling.	<ul style="list-style-type: none"> <li>• High-frequency solver</li> </ul>	Multi-GPU Single Node
<b>ANSYS Maxwell</b>	ANSYS	Industry-leading electromagnetic field simulation software for the design and analysis of electric motors, actuators, sensors, transformers and other electromagnetic and electromechanical devices.	<ul style="list-style-type: none"> <li>• Eddy Current Solver</li> </ul>	Multi-GPU Single Node
<b>ANSYS Nexxim</b>	ANSYS	Circuit simulation engine for RF/analog/mixed-signal IC design; IBIS-AMI analysis speedup with GPU computing.	<ul style="list-style-type: none"> <li>• AMI analysis</li> </ul>	Single GPU Single Node
<b>CDP</b>	D2S	GPU acceleration of real-time in-line enhancement of semiconductor manufacturing equipment such as the NuFlare EBM-9500 and MBM-1000 mask writers	<ul style="list-style-type: none"> <li>• Simulation-based processing</li> </ul>	Multi-GPU Multi-Node
<b>CST MPhysics STUDIO</b>	Dassault Systèmes SIMULIA Corp.	Multiphysics simulation including thermal, CFD, and mechanical capabilities. Tightly integrated with CST's electromagnetic solvers.	<ul style="list-style-type: none"> <li>• Conjugated Heat Transfer Solver</li> </ul>	Multi-GPU Multi-Node

<b>CST STUDIO SUITE</b>	Dassault Systèmes SIMULIA Corp.	Accurate and efficient computational solution for 3D simulation of electromagnetic devices in a wide range of frequencies	<ul style="list-style-type: none"> <li>• Transient Solver</li> <li>• Integral Equation Solver</li> <li>• Asymptotic Solver</li> <li>• Multilayer Solver</li> </ul>	Multi-GPU Multi-Node
<b>EMPro</b>	KeySight	Modeling and simulation environment for analyzing 3D EM effects of high speed and RF/Microwave components.	<ul style="list-style-type: none"> <li>• FDTD solver</li> </ul>	Multi-GPU Single Node
<b>FEKO</b>	Altair	3D EM modeling and simulation	<ul style="list-style-type: none"> <li>• FDTD solver</li> <li>• MoM solver</li> <li>• RL-GO solver</li> <li>• CMA Solver</li> </ul>	Multi-GPU Single Node
<b>JMAG</b>	JMAG	FEA software for electromechanical design. Fast solver / High quality mesh / Advanced modeling technologies.	<ul style="list-style-type: none"> <li>• EM transient solver</li> <li>• EM time harmonic solver</li> <li>• EM static solver</li> </ul>	Multi-GPU Single Node
<b>KeySight ADS</b>	KeySight	Simulation tool for design of RF, microwave and high speed digital circuits	<ul style="list-style-type: none"> <li>• Transient Convolution simulation with BSIM4 models</li> </ul>	Single GPU Single Node
<b>REMCOM XFDTD</b>	REMCOM	3D EM Simulation	<ul style="list-style-type: none"> <li>• FDTD Solver</li> </ul>	Multi-GPU Multi-Node
<b>SEMCAD-X</b>	SPEAG	3D EM modeling and simulation	<ul style="list-style-type: none"> <li>• FDTD solver</li> </ul>	Single GPU Single Node
<b>Serenity</b>	Lucernhammer	EM Simulation (RCS) tool	<ul style="list-style-type: none"> <li>• MoM solver</li> </ul>	Multi-GPU Single Node
<b>Sim4Life</b>	ZMT Zurich MedTech AG	3D Electromagnetics & Acoustic modeling and simulation	<ul style="list-style-type: none"> <li>• FDTD and Acoustics Solvers</li> </ul>	Multi-GPU Single Node
<b>TrueMask MDP</b>	D2S	GPU-accelerated simulation and data preparation for mask writing	<ul style="list-style-type: none"> <li>• Simulation-based processing</li> </ul>	Multi-GPU Multi-Node
<b>TrueModel</b>	D2S	GPU-accelerated simulation and geometric checking of curvilinear shapes	<ul style="list-style-type: none"> <li>• Simulation-based processing</li> </ul>	Multi-GPU Multi-Node
<b>Virtuoso - Cadence</b>	Cadence Design Systems	EDA design simulation. Primary app = Allegro	<ul style="list-style-type: none"> <li>• Visualization and acceleration for EDA and CAD design software.</li> </ul>	Multi-GPU Multi-Node
<b>VSim for Electromagnetics</b>	Tech-X Corporation	Physics Simulation and modeling software for EM	<ul style="list-style-type: none"> <li>• FDTD solver</li> </ul>	Single GPU Single Node
<b>WIPL-D 2D Solver</b>	WIPL-D	2D EM modeling and simulation	<ul style="list-style-type: none"> <li>• MoM Solver</li> </ul>	Multi-GPU Single Node
<b>WIPL-D Pro</b>	WIPL-D	3D EM modeling and simulation	<ul style="list-style-type: none"> <li>• MoM Solver</li> </ul>	Multi-GPU Multi-Node
<b>Wireless InSite</b>	REMCOM	Uses OptiX 3.8 for Ray-tracing and Propagation prediction	<ul style="list-style-type: none"> <li>• X3D Ray Tracer</li> </ul>	Multi-GPU Single Node

## INDUSTRIAL INSPECTION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
> <b>Cognex VisionPro ViDi</b>	Cognex	Deep learning-based software dedicated to industrial image analysis. Cognex ViDi Suite is a field-tested, optimized and reliable software solution based on a state-of-the-art set of algorithms in machine learning.	<ul style="list-style-type: none"> <li>• Feature localization &amp; identification</li> <li>• Segmentation &amp; defect detection</li> <li>• Object &amp; scene classification</li> <li>• Text &amp; character recognition</li> </ul>	Single GPU Single Node
> <b>HALCON</b>	MVTec Software	MVTec HALCON is the comprehensive standard software for machine vision with an integrated development environment. HALCON allows models to be trained on GPUs, and outputs trained models for inference on CPU, GPU, or Jetson.	<ul style="list-style-type: none"> <li>• Deep learning - pre-trained networks optimized for latency or precision. HALCON also provides an IDE for training neural networks. Features include sub-pixel detection, edge detection, counting, OCR, barcode reading, 3D reconstruction from stereo</li> </ul>	Single GPU Single Node

> <b>IBM Visual Insights</b>	IBM	IBM Visual Insights uses cognitive capabilities to review and analyze parts, components, and products and identify defects by matching patterns to images of defects that it has previously analyzed and classified. Deploy models to edge computing on production lines to facilitate rapid image capture by camera and cognitive identification of defects. Quickly assess quality inspection metrics across manufacturing processes.	<ul style="list-style-type: none"> <li>• Cloud-based DL training, deployment on (spec'ed) edge server</li> </ul>	Multi-GPU Single Node
------------------------------	-----	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------	--------------------------

## Media & Entertainment

### ANIMATION, MODELING AND RENDERING

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3ds Max</b>	Autodesk	3D modeling, animation, and rendering	<ul style="list-style-type: none"> <li>• Iray interactive</li> <li>• Photorealistic, physically correct rendering</li> </ul>	Multi-GPU Single Node
<b>Agisoft PhotoScan</b>	Agisoft	Agisoft PhotoScan is a stand-alone software product that performs photogrammetric processing of digital images and generates 3D spatial data to be used in GIS applications, cultural heritage documentation, and visual effects production as well as for indirect measurements of objects of various scales.	<ul style="list-style-type: none"> <li>• CUDA,</li> </ul>	Multi-GPU Single Node
<b>Blender</b>	Blender Inst	3D modeling, rendering and animation	<ul style="list-style-type: none"> <li>• GPU-accelerated viewport</li> </ul>	Single GPU Single Node
<b>Blender Cycles</b>	Blender Inst.	GPU renderer	<ul style="list-style-type: none"> <li>• CUDA-accelerated rendering</li> </ul>	Multi-GPU Single Node
<b>Cinema 4D</b>	Maxon	3D modeling, animation, and rendering	<ul style="list-style-type: none"> <li>• Increased model complexity at interactive rates</li> </ul>	Single GPU Single Node
<b>Clarisse</b>	Isotropix	Integrated environment for creation and assembly of immense environments, look development, lighting and rendering.	<ul style="list-style-type: none"> <li>• OptiX AI-accelerated de-noising accelerates interactivity over 5X.</li> </ul>	
<b>DAZ Studio</b>	Daz3D	Powerful and free 3D creation software tool that is not only easy to use but rich in features and functionality. Whether you are a novice or proficient 3D artist or 3D animator, Daz Studio enables you to create amazing 3D art.	<ul style="list-style-type: none"> <li>• Iray Interactive</li> <li>• Iray Photoreal</li> <li>• MDL support</li> </ul>	Multi-GPU Single Node
> <b>finalRender</b>	Cebas	GPU renderer	<ul style="list-style-type: none"> <li>• CUDA-based GPU rendering</li> </ul>	Multi-GPU Single Node
<b>FurryBall</b>	AAA Studio	GPU renderer	<ul style="list-style-type: none"> <li>• NVIDIA OptiX and DirectX GPU rendering</li> </ul>	Single GPU Single Node
<b>HIERO Player</b>	Foundry	Shot management, conform and review timeline	<ul style="list-style-type: none"> <li>• Fluid, interactive playback</li> </ul>	Single GPU Single Node
<b>Houdini</b>	SideFX	Procedural 3D modeling, animation and rendering	<ul style="list-style-type: none"> <li>• Faster simulations</li> </ul>	Multi-GPU Single Node
> <b>Indigo</b>	Glare Technology	Unbiased, physically-based renderer.	<ul style="list-style-type: none"> <li>• GPU-accelerated rendering.</li> </ul>	Multi-GPU Single Node
<b>KATANA</b>	The Foundry	Powerful look development and lighting tool	<ul style="list-style-type: none"> <li>• Faster interactive graphics</li> </ul>	Single GPU Single Node
<b>Kilton/Megaton</b>	Blastcode	Physics-based simulation plug in	<ul style="list-style-type: none"> <li>• Faster simulation</li> </ul>	Single GPU Single Node
> <b>Lightwave</b>	NewTek	3D modeling, animation, and rendering	<ul style="list-style-type: none"> <li>• Increased model complexity at interactive rates</li> </ul>	Single GPU Single Node
> <b>LuxRender</b>	LuxRender	GPU 3D Renderer	<ul style="list-style-type: none"> <li>• GPU-accelerated ray tracing</li> </ul>	Single GPU Single Node

<b>MARI</b>	The Foundry	3D paint tool allows painting directly onto 3D models	<ul style="list-style-type: none"> <li>• Faster interactive painting</li> </ul>	Single GPU Single Node
> <b>Mantra</b>	SideFX	Houdini Mantra renderer using OptiX AI de-noising for acceleration.	<ul style="list-style-type: none"> <li>• Much faster rendering</li> </ul>	
<b>Maxwell</b>	Next Limit	CUDA-accelerated interactive and final-frame renderer	<ul style="list-style-type: none"> <li>• unrestricted image resolution</li> <li>• network rendering</li> <li>• de-noising</li> </ul>	Multi-GPU Single Node
<b>Maya</b>	Autodesk	3D modeling, animation, and rendering	<ul style="list-style-type: none"> <li>• Increased model complexity, larger scenes</li> </ul>	Single GPU Single Node
<b>MODO</b>	Foundry	3D modeling, animation and rendering	<ul style="list-style-type: none"> <li>• Increased model complexity, larger scenes</li> </ul>	Single GPU Single Node
<b>Motion Builder</b>	Autodesk	Character animation and motion capture	<ul style="list-style-type: none"> <li>• Increased model complexity at interactive rates</li> </ul>	Single GPU Single Node
<b>Mudbox</b>	Autodesk	3D sculpting	<ul style="list-style-type: none"> <li>• Increased model complexity at interactive rates</li> </ul>	Single GPU Single Node
<b>Octane Render</b>	Otoy	CUDA-accelerated GPU renderer	<ul style="list-style-type: none"> <li>• GPU rendering</li> </ul>	Multi-GPU Single Node
<b>Realflow</b>	Next Limit	Fluid simulation system	<ul style="list-style-type: none"> <li>• GPU-accelerated simulation</li> </ul>	Single GPU Single Node
<b>RealityCapture</b>	Capturing Reality	Photogrammetry	<ul style="list-style-type: none"> <li>• CUDA-accelerated, fast photogrammetry</li> </ul>	Multi-GPU Single Node
<b>Redshift Renderer</b>	Redshift	GPU-accelerated, biased renderer	<ul style="list-style-type: none"> <li>• CUDA-based GPU final-frame rendering</li> <li>• Mac and Windows supported</li> </ul>	Multi-GPU Single Node
> <b>RenderMan</b>	Pixar	Leading film renderer.	<ul style="list-style-type: none"> <li>• OptiX-accelerated de-noising</li> </ul>	
<b>Sculptris</b>	Pixologic	3D sculpting	<ul style="list-style-type: none"> <li>• Increased model complexity at interactive rates</li> </ul>	Single GPU Single Node
<b>Thea Render</b>	Altair	Physically-based progressive spectral CPU/GPU Renderer supporting fast interactive changes and bucket rendering for high resolution images	<ul style="list-style-type: none"> <li>• advanced material layering system, subsurface scattering, displacement mapping, physical sun-sky, IES support and many other features</li> </ul>	Multi-GPU Single Node
<b>TurbulenceFD</b>	Jawset	Voxel-based gaseous fluid dynamics plug-in	<ul style="list-style-type: none"> <li>• 12X performance boost using NVIDIA GPUs</li> </ul>	Single GPU Single Node
<b>Twinmotion</b>	Abvent	AEC project review/phasing/marketing showcase including VR viewer with multi CAD/BIM/AEC import and a simple to use interface but powered by UE4	<ul style="list-style-type: none"> <li>• UE4 PBR material VR experience</li> </ul>	Single GPU Single Node
<b>V-Ray GPU</b>	Chaos Group	GPU renderer with CPU Hybrid rendering	<ul style="list-style-type: none"> <li>• CUDA interactive GPU rendering</li> </ul>	Multi-GPU Single Node

## COLOR CORRECTION AND GRAIN MANAGEMENT

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ARRI de-bayering SDK</b>	ARRI	RAW de-bayering SDK	<ul style="list-style-type: none"> <li>• De-bayering of ARRI RAW and primary color grading.</li> </ul>	Single GPU Single Node
<b>Baselight</b>	FilmLight	Color grading	<ul style="list-style-type: none"> <li>• Real-time color correction</li> </ul>	Multi-GPU Single Node
<b>Cinema RAW SDK</b>	Canon	RAW de-bayering	<ul style="list-style-type: none"> <li>• GPU-accelerated de-bayering</li> </ul>	Single GPU Single Node
<b>DaVinci Resolve</b>	Blackmagic Design	Color grading, editing and 3D effects	<ul style="list-style-type: none"> <li>• Real-time color correction and de-noising</li> </ul>	Multi-GPU Single Node
<b>Dark Energy</b>	Cinnafilm	Application and plug-in for image enhancement	<ul style="list-style-type: none"> <li>• Image de-noising and restoration</li> </ul>	Multi-GPU Single Node
<b>Diamant-Film Restoration</b>	HS-Art	Film cleanup and restoration	<ul style="list-style-type: none"> <li>• CUDA accelerated optical flow, de-flicker, in-painting and over 30 filters</li> </ul>	Multi-GPU Single Node
<b>Effects Suite</b>	Red Giant	Visual effects plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Grain and Noise Reducer</b>	Wavelet Beam	Video noise reduction	<ul style="list-style-type: none"> <li>• CUDA-accelerated grain and noise reduction</li> </ul>	Multi-GPU Single Node

<b>Magic Bullet Looks</b>	Red Giant	Color and finishing tools	• Faster effects	Single GPU Single Node
<b>Mist</b>	Marquise Technologies	Mastering tool for cinema, broadcast and over-the-top content	• CUDA-accelerated de-bayering, color grading, transcoding and image enhancement	Multi-GPU Single Node
<b>Nucoda</b>	Digital Vision	Color grading	• GPU-accelerated color grading	Single GPU Single Node
<b>Pablo family</b>	Grass Valley	Color grading and finishing	• Real time color correction	Multi-GPU Single Node
<b>PFClean</b>	The Pixel Farm	Image restoration and remastering	• CUDA-based image processing acceleration	Multi-GPU Single Node
<b>RAW Converter</b>	ARRI	RAW de-Bayering and primary color grading	• CUDA-accelerated de-bayering and primary grading	Single GPU Single Node
<b>REDCINE-X PRO</b>	Red Digital Cinema	Primary color grading	• CUDA-accelerated de-bayering and grading	Single GPU Single Node
> <b>Red Digital Cinema R3D SDK</b>	Red Digital Cinema	Red Digital Cinema camera SDK decodes and de-bayers Red RAW camera data and allows primary color grading. It is used by many color grading and video editing applications.	• CUDA accelerated wavelet decoding, de-bayering and primary color grading	Single GPU Single Node
<b>Scratch</b>	Assimilate	Color grading and finishing	• Accelerated debayering for real-time digital finishing	Single GPU Single Node
<b>SpeedGrade CC</b>	Adobe	Color grading	• Real-time grading and finishing with Lumetri Deep Color Engine	Single GPU Single Node

## COMPOSITING, FINISHING AND EFFECTS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>After Effects CC</b>	Adobe	Motion graphics and effects	<ul style="list-style-type: none"> <li>• Faster effects plus 3D ray tracing based on NVIDIA OptiX and up to 10x faster perf on about 10% effects</li> <li>• More effects dev &amp; optimization in progress</li> <li>• Moving the effects pipeline on top of CUDA for the modular pipeline that will also be used as an effect in Premiere Pro</li> </ul>	Multi-GPU Single Node
<b>Clipster</b>	Rohde & Schwarz	Video and film player and DCI Packager	<ul style="list-style-type: none"> <li>• Video scaling</li> <li>• Color space conversion</li> <li>• Data format conversion</li> </ul>	Multi-GPU Single Node
<b>Complete</b>	CoreMelt	Visual effects plug-in	• Faster effects	Single GPU Single Node
<b>Continuum Complete</b>	Boris FX	Visual effects plug-in	• Faster effects	Single GPU Single Node
<b>Element 3D</b>	Video Copilot	3D object based particle system	• Faster effects	Single GPU Single Node
<b>FilmTouch</b>	Pixelan	Video effects plug-in	• Faster effects	Single GPU Single Node
<b>Flame Premium</b>	Autodesk	Finishing and color grading	• Integrated toolset for 3D VFX, editorial, and color grading	Multi-GPU Single Node
<b>Fusion</b>	Blackmagic Design	Effects and compositing	• Faster effects	Single GPU Single Node
<b>HIERO</b>	Foundry	Multi-shot management tool that supports collaborative working, review and approval, quick production turnaround and delivery	• Fluid, interactive playback	Single GPU Single Node
<b>Mamba FX</b>	SGO	High-end compositing	• Faster keying, tracking, painting and restoration	Single GPU Single Node
<b>Mistika Ultima</b>	SGO	Color grading and finishing	• Faster keying, tracking, painting and restoration, de-bayering	Single GPU Single Node



<b>Mistika VR</b>	SGO	Near real-time optical flow stitching	<ul style="list-style-type: none"> <li>• GPU-accelerated video stitching with manual controls</li> <li>• compatible with most camera rigs</li> <li>• stitch, review and improve results in seconds</li> <li>• export clips in many formats, including DPX and ProRes</li> </ul>	Single GPU Single Node
<b>Monsters GT</b>	Boris FX	Visual effects plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>NUKE</b>	Foundry	Compositing tool with 3D tracker	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Open FX</b>	Neat Video	Video noise reduction plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>PFTTrack</b>	The Pixel Farm	3D scene creation and tracking	<ul style="list-style-type: none"> <li>• CUDA-accelerated tracking</li> </ul>	Multi-GPU Single Node
<b>ROBUSKEY</b>	Robuskey	Chroma keyer plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Sapphire</b>	Boris FX	Visual effects plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Twitch</b>	Video Copilot	Video effects plug-in for After Effects	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Twixtor</b>	Re:Vision Effects	Visual effects plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Video Essentials</b>	NewBlueFX	Video effects plug-in	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node

## EDITING

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Catalyst Browse/Prepare/Edit</b>	Sony	4K, Sony RAW, and HD video editing	<ul style="list-style-type: none"> <li>• Faster effects, transitions and encoding</li> </ul>	Single GPU Single Node
<b>Edius Pro</b>	Grass Valley	Video editing	<ul style="list-style-type: none"> <li>• Faster effects</li> <li>• RAW de-bayering</li> </ul>	Single GPU Single Node
<b>Final Cut Pro</b>	Apple	Video editing	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Kolor Autopano</b>	GoPro	360 video stitching for GoPro Cameras	<ul style="list-style-type: none"> <li>• GPU accelerated 360 stitching 2-4X faster than CPU</li> </ul>	Single GPU Single Node
<b>Lightroom CC</b>	Adobe	Photo editing	<ul style="list-style-type: none"> <li>• Develop module is GPU accelerated   speed up is only modest   Minimal AI features like auto-tagging and intelligent image adjustments added - this are DL training only, NOT inference</li> </ul>	Single GPU Single Node
<b>Lightworks</b>	EditShare	Video editing	<ul style="list-style-type: none"> <li>• Faster effects</li> <li>• CUDA-accelerated de-bayering</li> </ul>	Single GPU Single Node
<b>Media Composer</b>	Avid	Video editing	<ul style="list-style-type: none"> <li>• Faster video effects, unique stereo 3D capabilities</li> </ul>	Single GPU Single Node
<b>MXF</b>	Film Partners	Collaborative editing system supporting Avid Media Composer, Adobe Premiere Pro, Grass Valley Edius and Blackmagic Resolve	<ul style="list-style-type: none"> <li>• NVIDIA Video Codec allows remote GPU-accelerated production workflows</li> </ul>	Single GPU Single Node
<b>Photoshop CC</b>	Adobe	Image editing	<ul style="list-style-type: none"> <li>• Natural canvas OpenGL accelerated, Blur galleries OpenCL accelerated</li> </ul>	Single GPU Single Node
<b>Premiere Pro CC</b>	Adobe	Video editing	<ul style="list-style-type: none"> <li>• Real-time video editing &amp; accelerated output rendering via Mercury Playback Engine - CUDA OpenCL Summer 2018 - Video Codec SDK - both encode &amp; decode (including HEVC for 8K)</li> </ul>	Multi-GPU Single Node

<b>Qube</b>	Grass Valley	Broadcast video editing	<ul style="list-style-type: none"> <li>Faster video effects, unique stereo 3D capabilities</li> </ul>	Single GPU Single Node
<b>Skybox 360/VR Tools</b>	Mettle	VR design plugin for Premiere Pro	<ul style="list-style-type: none"> <li>Fast VR processing via OpenGL and CUDA acceleration</li> </ul>	Single GPU Single Node
<b>Skybox Studio</b>	Mettle	VR design plugin for After Effects	<ul style="list-style-type: none"> <li>Fast VR processing via OpenGL and CUDA acceleration</li> </ul>	Single GPU Single Node
<b>Smoke</b>	Autodesk	Finishing and editing	<ul style="list-style-type: none"> <li>Faster effects</li> </ul>	Single GPU Single Node
<b>Vegas Pro</b>	Magix	Video editing	<ul style="list-style-type: none"> <li>Faster video effects and encoding</li> </ul>	Single GPU Single Node
<b>Velocity</b>	Imagine Communications	Video editing	<ul style="list-style-type: none"> <li>Faster effects</li> </ul>	Single GPU Single Node

## ENCODING AND DIGITAL DISTRIBUTION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Alchemist on Demand</b>	Grass Valley	Video standards conversion	<ul style="list-style-type: none"> <li>GPU-accelerated video procession and encoding</li> </ul>	Multi-GPU Single Node
<b>Amberfin</b>	Dalet	Transcoding and video quality analysis	<ul style="list-style-type: none"> <li>GPU-accelerated video procession and encoding</li> </ul>	Single GPU Single Node
<b>Aurora</b>	Tektronix	Automated video quality measurement	<ul style="list-style-type: none"> <li>GPU-accelerated video quality assessment</li> </ul>	Single GPU Single Node
> <b>AW-360C10</b>	Panasonic	360-degree Live Camera designed for live sporting events, concerts and stadium events	<ul style="list-style-type: none"> <li>low-latency, real-time 4K 360 degree stitching from four camera inputs using Jetson TX-1</li> <li>control from tablet or over wi-fi from PC</li> <li>automatic exposure and white balance adjustment</li> </ul>	Single GPU Single Node
<b>Content Agent</b>	Root6	Automated transcoding and workflow management	<ul style="list-style-type: none"> <li>GPU-accelerated video procession and encoding</li> </ul>	Multi-GPU Single Node
<b>Core</b>	ArcVideo	Video processing and transcoding Live	<ul style="list-style-type: none"> <li>Accelerated transcoding and encoding</li> </ul>	Multi-GPU Single Node
<b>Elemental Live</b>	Elemental	Live streaming video processing and encoding	<ul style="list-style-type: none"> <li>Video encoding and video processing</li> </ul>	Multi-GPU Single Node
<b>Elemental Server</b>	Elemental	File-based video processing and encoding	<ul style="list-style-type: none"> <li>Video encoding and video processing</li> </ul>	Multi-GPU Single Node
<b>Fast CinemaDNG Processor</b>	Fastvideo	RAW video debayering, denoising and color correction completely on GPU side	<ul style="list-style-type: none"> <li>High-quality GPU-based RAW video processing up to 160 fps</li> <li>Wavelet, realtime de-noising (pre and post bayer)</li> <li>Standard color correction features and monitoring options</li> <li>Export to 16-bit TIF or 10-bit ProResFull-sized video processing</li> <li>Realtime 4K, 6K, and 8K playback supported</li> </ul>	Multi-GPU Single Node
<b>JPEG2000 Codec</b>	Comprimato	JPEG2000 encoding and decoding for DCP, IMF, video editing, broadcast contribution, and archiving.	<ul style="list-style-type: none"> <li>Faster-than-real-time UltraHD / 4K</li> <li>Lossy and mathematically lossless</li> <li>High-bit-depth (HDR)</li> </ul>	Multi-GPU Single Node
<b>Lightspeed Live</b>	Telestream	Enterprise-class live streaming system that can ingest, encode, package and deploy multiple sources to multiple destinations. System utilizes the latest technologies to deliver pristine quality and exceptional processing speed. Video processing and transcoding can be accelerated with GPU for up to 9x speed improvements	<ul style="list-style-type: none"> <li>Video processing and transcoding</li> </ul>	Multi-GPU Single Node

<b>Live</b>	ArcVideo	High-density, real-time video processing and encoding.	<ul style="list-style-type: none"> <li>Accelerated broadcast encoding with NVIDIA CUDA and NVENC</li> </ul>	Multi-GPU Single Node
<b>Media Encoder CC</b>	Adobe	Output aggregator and encoder for Premiere Pro & After Effects	<ul style="list-style-type: none"> <li>Faster output rendering based on Mercury Playback Engine</li> </ul>	Multi-GPU Single Node
<b>Medialooks SDK</b>	Medialooks	MFormats SDK provides complete control over the video pipeline	<ul style="list-style-type: none"> <li>NVIDIA Video Codec used for accelerated encoding and decoding</li> </ul>	Single GPU Single Node
<b>Multiplatform Transcoder</b>	ERLAB	Video processing and encoding software	<ul style="list-style-type: none"> <li>Pre-processing encoding, decoding, post-processing and delivery</li> </ul>	Single GPU Single Node
<b>Piko TV</b>	Kizil Elektronik	Linear broadcast encoder	<ul style="list-style-type: none"> <li>H.264 and HEVC 4K encoding for broadcast channels</li> </ul>	Single GPU Single Node
<b>PixelStrings</b>	Cinnafilm	Cloud-based image processing Platform-as-a-Service (PaaS) delivers the high-quality, automated video conversion and frame optimization	<ul style="list-style-type: none"> <li>motion-compensated frame rate conversion</li> <li>high-quality de-interlacing</li> <li>texture-aware scaling</li> <li>degrain/regrain to any film look,</li> <li>denoise/retexture to limit banding</li> <li>reverse telecine/pulldown pattern correction</li> <li>interlace artifact and dust removal</li> <li>runtime retiming</li> </ul>	Multi-GPU Single Node
<b>Server 2</b>	Sorenson Media	Video transcoding for server app	<ul style="list-style-type: none"> <li>H.264 video encoding and video processing</li> </ul>	Multi-GPU Single Node
<b>Squeeze Desktop 7</b>	Sorenson Media	Video transcoding application and plug-In	<ul style="list-style-type: none"> <li>H.264 video encoding and video processing</li> </ul>	Multi-GPU Single Node
<b>Tachyon</b>	Cinnafilm	Standards conversion	<ul style="list-style-type: none"> <li>Video processing and frame rate conversion</li> </ul>	Multi-GPU Single Node
<b>Tornado</b>	Marquise Technologies	Transcoding engine for IMF and DCP facilities	<ul style="list-style-type: none"> <li>Image re-sizing up to 8K</li> <li>Color space conversion: 601/709, REC 2020, DCI XYZ, ACES 1.0</li> <li>De-bayering: ARRIRAW, DNG, RED R3D, SONY F65, F55 RAW, Phantom flex 4K, Canon C500</li> <li>Mezzanine: ProRes 444, Avid DNxHD 444, XDCAM, AVC Intra, AS-11 DPP, IMF</li> <li>Uncompressed: DPX, TIFF, OpenEXR</li> </ul>	Single GPU Single Node
<b>Transkoder</b>	Colorfront	Encoding and transcoding for DCP and IMF mastering	<ul style="list-style-type: none"> <li>JPEG2000 encoding and decoding</li> <li>32-bit floating point processing on multiple GPUs</li> <li>MXF wrapping, accelerated checksums and AES encryption and decryption,</li> <li>IMF/IMP and DCI/DCP package authoring, editing, transwrapping</li> </ul>	Multi-GPU Single Node
<b>Vantage LightSpeed</b>	Telestream	Enterprise-class live streaming system that can ingest, encode, package and deploy multiple sources to multiple destinations. System utilizes the latest technologies to deliver pristine quality and exceptional processing speed. Video processing and transcoding can be accelerated with GPU for up to 9x speed improvements	<ul style="list-style-type: none"> <li>Video transcoding and processing</li> </ul>	Multi-GPU Single Node
<b>Viarte</b>	Isovideo	Video standards conversion	<ul style="list-style-type: none"> <li>CUDA-accelerated video processing and encoding</li> </ul>	Multi-GPU Single Node
<b>VidiCert</b>	Joanneum Research	Video and film quality assurance	<ul style="list-style-type: none"> <li>CUDA accelerated video quality analysis</li> </ul>	Multi-GPU Single Node
<b>Wowza Streaming Engine Transcoder</b>	Wowza	H.264 video encoding	<ul style="list-style-type: none"> <li>NVENC accelerated video encoding</li> </ul>	Single GPU Single Node

## ON-AIR GRAPHICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Air</b>	Cinegy	Broadcast play-out server	<ul style="list-style-type: none"> <li>Real-time on-air graphics</li> </ul>	Single GPU Single Node
<b>Broadcast Dscript 3D</b>	Monarch	3D on-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Capture</b>	Cinegy	Video ingest	<ul style="list-style-type: none"> <li>Uses NVENC to encode/decode multiple H.264 and HEVC streams</li> </ul>	Single GPU Single Node
<b>Clarity</b>	Pixel Power	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Cube</b>	Dalet	On-air Graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>eStudio</b>	Brainstorm	Virtual sets and motion graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
> <b>GS2 Graphics Engine</b>	ChyronHego	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
> <b>InfinitySet</b>	Brainstorm			
<b>Livebook GFX</b>	AJT Systems	The LiveBook is designed to fit every production environment and facilitate evolving work flows. Whether you are broadcasting over IP, or using SDI for internal or downstream keying, the LiveBook will be able to adapt to your environment.	<ul style="list-style-type: none"> <li>GRAPHICS SOLUTION FOR COMPACT LIVE SPORTS PRODUCTIONS</li> </ul>	Multi-GPU Single Node
<b>Mosaic</b>	ChyronHego	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Multiviewer</b>	Evertz	Broadcast multiviewer	<ul style="list-style-type: none"> <li>Uses NVENC H.264 and HEVC encoding and decoding</li> </ul>	Single GPU Single Node
<b>Nexio Channelbrand</b>	Imagine Communications	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Multi-GPU Single Node
<b>Nexio G8</b>	Imagine Communications	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Nexio TitleOne</b>	Imagine Communications	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Reality Virtual Studio</b>	Zero Density	Photorealistic virtual studio solution in broadcast industry, powered by Epic Unreal Engine 4	<ul style="list-style-type: none"> <li>node-based compositing system designed for real-time production</li> <li>same content can be used for broadcast and VR</li> <li>image quality is achieved by on NVIDIA GPUs through deferred rendering methods, unique anti-aliasing technology and advanced features such as depth of field, motion blur, light maps, screen space reflections and refraction</li> </ul>	Single GPU Single Node
<b>Titler Pro</b>	NewBlueFX	video titling	<ul style="list-style-type: none"> <li>GPU-accelerated graphics</li> </ul>	Single GPU Single Node
<b>tOG</b>	RT Software	On-air graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Type</b>	Cinegy	On-air Graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Vertigo</b>	Grass Valley	On-air Graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Virtuoso</b>	Monarch	Virtual sets and motion graphics	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Viz Engine</b>	Vizrt	On-air graphics and virtual sets	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node
<b>Wasp3D - CG</b>	Wasp3D	On-air graphics and virtual sets	<ul style="list-style-type: none"> <li>Real-time rendering</li> </ul>	Single GPU Single Node

## ON-SET, REVIEW AND STEREO TOOLS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Cortex Dailies</b>	MTI Film	Review, color grading and transcoding on set	<ul style="list-style-type: none"> <li>• CUDA accelerated grading and transcoding</li> </ul>	Multi-GPU Single Node
<b>Dimension</b>	Blackmagic Design	3D stereoscopic workflow	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node
<b>Fluid 4K Review</b>	BlueFish444	Review and approval of 4K content	<ul style="list-style-type: none"> <li>• Real-time video review</li> </ul>	Single GPU Single Node
<b>ICE</b>	Marquise Technologies	IMF reference video player	<ul style="list-style-type: none"> <li>• RAW data support for ARRIRAW, DNG, RED R3D, SONY F65, F55 RAW, Phantom flex 4K and Canon C500</li> <li>• HDR content encoded in Dolby Vision, HDR10, HDR10+ or HLG</li> <li>• Uncompressed formats support: DPX, TIFF and OpenEXR</li> </ul>	Single GPU Single Node
<b>On-Set Dailies</b>	Colorfront	Review, color grading and transcoding on set	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Multi-GPU Single Node
<b>Previzion</b>	Lightcraft	On-set virtual production	<ul style="list-style-type: none"> <li>• Real-time, virtual set production</li> </ul>	Single GPU Single Node
<b>RV</b>	Autodesk	Review and approval of 4K content	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node

## WEATHER GRAPHICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Cinemative HD</b>	Accuweather	Weather graphics	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node
<b>Max Weather</b>	WSI	Weather graphics	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node
<b>Metacast</b>	ChyronHego	Weather graphics	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node
<b>MeteoEarth</b>	MeteoGraphics	Weather graphics	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node
<b>Storyteller</b>	Accuweather	Weather graphics	<ul style="list-style-type: none"> <li>• Real-time</li> </ul>	Single GPU Single Node

## Medical Imaging

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Carestream Vue PACS</b>	Carestream Health	Carestream's Vue Clinical Collaboration Platform gives all those who provide, manage, receive and reimburse care the ability to access the clinical images they need, using the preferred device and application for each workflow and setting.		
<b>Centricity Universal Viewer</b>	General Electric	PACS Featuring a single image repository across 2D and 3D studies, Centricity Universal Viewer intuitively brings together the tools needed by radiologists, cardiologists and other clinicians to provide enterprise-wide access on a single desktop.	<ul style="list-style-type: none"> <li>• Intelligent productivity tools, Advanced visualization applications, An advanced mammography workflow, Cross Enterprise Display, Advanced Cardiology tools, Access from anywhere1 Image enable your EMR.</li> </ul>	
<b>IntelliSpace</b>	Phillips	PACS		
<b>IntelliSpace Portal 8.0</b>	Philips	IntelliSpace Portal 8.0 is an advanced visualization platform that offers a single integrated solution to help you work quickly and efficiently with increased diagnostic confidence, especially during reading and follow up of complex cases.		
<b>iNtuition iEMV</b>	Terarecon, Inc.	Advanced Image Post-Processing		

<b>JiveX</b>	VISUS Health IT	free DICOM Viewer provides comparable functionality as the JiveX Review Client such as measuring tools, zoom, etc. The hanging protocols, which are less relevant outside the clinical setting, are not available.	
<b>McKesson Radiology</b>	McKesson Corporation	McKesson Radiology Analytics is a new cloud-based analytics solution that helps you interactively visualize and explore operational metrics, while discovering business improvement opportunities. With the introduction of McKesson Radiology Analytics, we're helping customers achieve their operational goals by unlocking the real value of their data, with a focus on data accuracy, visualization and interaction.	<ul style="list-style-type: none"> <li>• Volumetric CT and MR Solutions, Teaching File Solutions, McKesson Radiology Collaboration, Improved Enhancements Program</li> </ul>
<b>Merge PACS</b>	IBM Watson Health	Merge PACS is a powerful reading workflow platform that simplifies physician's reading activities and centralizes the management of studies. Beyond consolidating multiple reading specialty systems, Merge PACS is uniquely designed to handle high-enterprise imaging volumes, perform in diverse reading environments (local, remote, dispersed and teleradiology), and enable providers to seamlessly scale their care delivery.	<ul style="list-style-type: none"> <li>• Unleash the true power of the worklist (Harness a highly modular and flexible worklist that allows studies to have numerous statuses for different users, enabling multiple workflows to coexist natively and simultaneously.) Focus on patient studies without interruptions (Avoid reading disruptions by benefiting from one platform for all specialties, seamless access to data and an automated pre-load of your next study while reviewing the current one.) Drive critical communications and compliance (Enable active consultations among experts through the embedded instant messaging application and auditable worklist exchanges.) Measure and improve outcomes (Go above and beyond monitoring your systems proactively by understanding how your organization performs with enterprise-wide reports on reading activities.)</li> </ul>
<b>MicroDicom</b>	MicroDicom	MicroDicom DICOM viewer is equipped with most common tools for manipulation of DICOM images and it has an intuitive user interface. It also has the advantage of being free for use and accessible to everyone.	<ul style="list-style-type: none"> <li>• Supported DICOM images - without compression and RLE, JPEG Lossy, JPEG Lossless, JPEG 2000 Lossy, JPEG2000 Lossless compressions Structured Reports, MPEG-2 and MPEG-4 transfer syntaxes, Encapsulated PDF, Open DICOM directory files, Display patient list from DICOMDIR, Load via drag&amp;drop or double-click Open images in common graphic formats(jpeg, bmp, png, gif, tiff), Convert DICOM images to JPEG, BMP, PNG, GIF, TIFF, Convert DICOM images to movie file format(AVI), Copy DICOM image to clipboard, Mouse-driven level-window, user-defined window presets Zooming and panning, Medical image processing operation, Measurements and annotations, Brightness/contrast control, Image resize, Rotate, Flip, Invert, Displaying DICOM attributes of selected image, Suited for patient CD/DVD to show DICOM images without installation Run on Windows XP, Windows Vista, Windows 7, Windows 8, Windows 8.1 and Windows 10, Available for x86 and x64 platforms, Portable version</li> </ul>

<b>NovaPACS</b>	Novarad	With NovaPACS, patient studies can be viewed from any computer at any of our facilities or from a referring physician's office. NovaPACS also allows the radiologists to read studies performed at any of our facilities, from any of our facilities, making radiology workflow much more efficient and greatly reducing turnaround time for report dictation.	<ul style="list-style-type: none"> <li>• Customizable, Disaster Recovery, Nova RIS Integration, Product Improvement, Evergreen Product Program, Support</li> </ul>	
<b>PowerGrid</b>	University of Illinois Urbana-Champaign	Advanced MRI reconstruction modeling	<ul style="list-style-type: none"> <li>• Discrete Fourier Transform</li> </ul>	Multi-GPU Single Node
<b>RadiAnt</b>	Medixant	RadiAnt DICOM Viewer provides basic tools for the manipulation and measurement of images	<ul style="list-style-type: none"> <li>• Fluid zooming and panning, Brightness and contrast adjustments, negative mode, Preset window settings for Computed Tomography (lung, bone, etc.), Ability to rotate (90, 180 degrees) or flip (horizontal and vertical) images, Segment length, Mean, minimum and maximum parameter values (e.g. density in Hounsfield Units in Computed Tomography) within circle/ellipse and its area, Angle value (normal and Cobb angle), Pen tool for freehand drawing</li> </ul>	
<b>RemotEye</b>	NeoLogica	RemotEye Suite represents a very complete solution for all your medical image viewing needs. It offers full DICOM compliance, web-based architecture, cross-platform compatibility, support for a wide variety of devices, including desktop computers, tablets smartphones.		
<b>Sectra PACS</b>	Sectra AB	DICOM Viewer		
<b>SYNAPSE</b>	Fujifilm Corporation			
<b>syngo.via</b>	Siemens Healthcare	With our PACS, you can benefit from software solutions for routine reading that increase quality and productivity in your radiology department.	<ul style="list-style-type: none"> <li>• Robust performance, great speed, intuitive operation, and intelligent working aids.</li> </ul>	
<b>Visage 7</b>	Pro Medicus, Ltd.	DICOM viewer, also referred to as: enterprise viewer, a universal viewer (UniViewer), or an archive neutral viewer (ANV).		
<b>VistARad</b>	U.S. Department of Veterans Affairs	The VistA Imaging system integrates clinical images, scanned documents, and other non-textual data into the patient's electronic medical record. VistA Imaging can capture and manage many different kinds of images	<ul style="list-style-type: none"> <li>• The VistA Imaging System's primary functions are: Clinical Image Display, Image Capture, Filmless Radiology, Image Management</li> </ul>	
<b>Vitre Vision</b>	Vital Images	Advanced Image Post-Processing		
<b>XERO VIEWER</b>	Agfa-Gevaert Group	XERO Viewer provides secure access to imaging data from different departments and multiple sources, in one view, to anyone who needs it.	<ul style="list-style-type: none"> <li>• Zero-footprint, Web 2.0 technology, Image-enable your EHR, One consolidated view, Collaborate using the Web, Images on your mobile device, Mobile upload of multimedia objects, View, edit and compare ECGs</li> </ul>	

# Oil and Gas

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>6X</b>	Ridgeway Kite	Reservoir Simulation on Tesla	<ul style="list-style-type: none"> <li>CUDA Simulation Parallelization</li> </ul>	
<b>AISight for SCADA</b>	BRS Labs	Proactive integrity management and real-time precursor alerts for enhanced SCADA operations in oil and gas.	<ul style="list-style-type: none"> <li>24/7 real-time analysis and alerting scaling to thousands of sensors across remote and geographically dispersed locations including historical analysis and trend reports.</li> </ul>	Multi-GPU Single Node
<b>AxRTM</b>	Acceleware	Reverse Time Migration Software	<ul style="list-style-type: none"> <li>CUDA accelerated libraries for building RTM software</li> </ul>	Multi-GPU Multi-Node
<b>DecisionSpace</b>	Halliburton (Landmark)	E&P platform for geoscience, well planning, drilling, earth modeling	<ul style="list-style-type: none"> <li>CUDA acceleration of fault extraction</li> </ul>	Multi-GPU Single Node
<b>Echelon</b>	Stoneridge Technology	Reservoir simulator	<ul style="list-style-type: none"> <li>Fully GPU-accelerated reservoir model, including dual-perm, dual porosity, pressure varying perm and porosity</li> <li>Eclipse compatible input deck</li> </ul>	Multi-GPU Multi-Node
> <b>GeoCluster</b>	CGG	Seismic Processing Suite	<ul style="list-style-type: none"> <li>CUDA accelerated processing</li> </ul>	Multi-GPU Multi-Node
<b>GeoDepth</b>	Paradigm Geophysical	Seismic Interpretation Suite	<ul style="list-style-type: none"> <li>CUDA-accelerated RTM</li> </ul>	Multi-GPU Multi-Node
<b>Geoteric</b>	Geoteric	Seismic interpretation	<ul style="list-style-type: none"> <li>Attributes calculations</li> <li>Geobodies extraction</li> </ul>	Multi-GPU Single Node
<b>Graydient S (SCADA)</b>	Giant Grey	Machine learning anomaly detection for large scale industrial data.	<ul style="list-style-type: none"> <li>Proactive integrity management and real-time precursor alerts for enhanced SCADA operations in oil and gas</li> <li>24/7 real-time analysis and alerting scaling to thousands of sensors across remote and geographically dispersed location</li> </ul>	Multi-GPU Single Node
<b>HUESpace</b>	Bluware	Library SDK toolkit for creating applications for seismic compression and seismic/geospatial imaging and interpretation	<ul style="list-style-type: none"> <li>CUDA acceleration for compression and large-scale visualization</li> </ul>	Multi-GPU Single Node
<b>InsightEarth</b>	CGG	Seismic Interpretation Suite	<ul style="list-style-type: none"> <li>OpenCL acceleration for AFE and 3D Curvature attributes</li> </ul>	Multi-GPU Single Node
<b>Omega2 RTM</b>	Schlumberger	Seismic processing	<ul style="list-style-type: none"> <li>Multiple algorithms (RTM, etc)</li> </ul>	Multi-GPU Multi-Node
<b>PumaFlow IFP</b>	Beicip-Franlab	Reservoir simulation	<ul style="list-style-type: none"> <li>GPU-accelerated linear solver</li> </ul>	Multi-GPU Single Node
<b>Roxar RMS</b>	Emerson	Reservoir modeling	<ul style="list-style-type: none"> <li>Multi GPU capabilities via HUESpace</li> </ul>	Multi-GPU Single Node
<b>RTM</b>	Tsunami	Seismic processing	<ul style="list-style-type: none"> <li>RTM algorithm</li> </ul>	Multi-GPU Multi-Node
<b>Seismic City RTM</b>	Seismic City	RTM Seismic Processing	<ul style="list-style-type: none"> <li>CUDA acceleration</li> </ul>	Multi-GPU Multi-Node
<b>SKUA</b>	Paradigm	Reservoir modeling	<ul style="list-style-type: none"> <li>Faults, Horizons and Flow Simulation Grid</li> </ul>	Multi-GPU Single Node
<b>tNavigator []</b>	Rock Flow Dynamics (RFD)	tNavigator Solver is a software package, offered as a single executable, which allows to build static and dynamic reservoir models, run dynamic simulations, calculate PVT properties of fluids, build surface network model, calculate lifting tables, and perform extended uncertainty analysis as a part of one integrated workflow.	<ul style="list-style-type: none"> <li>CUDA, Pascal/Volta architecture, Multi-Node GPU</li> </ul>	Multi-GPU Multi-Node
<b>VoxelGeo</b>	Paradigm Geophysical	Seismic Interpretation Package	<ul style="list-style-type: none"> <li>Multi-GPU volume rendering</li> <li>Horizon-flattening</li> <li>Attribute calculations</li> </ul>	Multi-GPU Single Node

> Indicates new application



# Research: Higher Education and Supercomputing

## COMPUTATIONAL CHEMISTRY AND BIOLOGY

### Bioinformatics

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Arioc</b>	Johns Hopkins University	High-throughput read alignment with GPU-accelerated exploration of the seed-and-extend search space	<ul style="list-style-type: none"> <li>• Single-end alignment, paired-end alignment</li> <li>• Output in SAM or database-ready binary formats</li> <li>• Multiple GPU implementation</li> </ul>	Multi-GPU Single Node
<b>BarraCUDA</b>	University of Cambridge Metabolic Research Labs	Sequence mapping software	<ul style="list-style-type: none"> <li>• Alignment of short sequencing reads, alignment of indels with gap openings and extensions.</li> </ul>	Multi-GPU Multi-Node
<b>BEAGLE-lib</b>	Open Source	BEAGLE is a high-performance library that can perform the core calculations at the heart of most Bayesian and Maximum Likelihood phylogenetics packages. It can make use of highly-parallel processors such as those in graphics cards (GPUs) found in many PCs.	<ul style="list-style-type: none"> <li>• Evaluation of likelihood for sequence evolution on trees and Arbitrary models (e.g. nucleotide, amino acid, codon)</li> <li>• Speed-ups (over CPU only version): nucleotide model = up to 25x, codon model = up to 50x.</li> </ul>	Multi-GPU Single Node
<b>Campaign</b>	SimTK	An open-source library of GPU-accelerated data clustering algorithms and tools.	<ul style="list-style-type: none"> <li>• K-means (and Kps-means, a K-means variant for GPUs with parallel sorting for improved performance)</li> <li>• K-medoids</li> <li>• K-centers (a K-medoids variant in which medoids are placed only once according to a heuristic)</li> <li>• Hierarchical clustering and Self-organizing map</li> </ul>	Multi-GPU Multi-Node
<b>CUDASW++</b>	Open Source	Open source software for Smith-Waterman protein database searches on GPUs.	<ul style="list-style-type: none"> <li>• Parallel search of Smith-Waterman database.</li> </ul>	Multi-GPU Single Node
<b>CUSHAW</b>	Open Source	Parallelized short read aligner	<ul style="list-style-type: none"> <li>• Parallel, accurate long read aligner for large genomes</li> </ul>	Multi-GPU Single Node
<b>G-BLASTN</b>	Hong Kong Baptist University	GPU-accelerated nucleotide alignment tool based on the widely used NCBI-BLAST.	<ul style="list-style-type: none"> <li>• Blastn and megablast modes of NCBI-BLAST</li> </ul>	Single GPU Single Node
<b>GHOST-Z GPU</b>	Akiyama_Laboratory, Tokyo Institute of Technology	Sequence homology search tool.	<ul style="list-style-type: none"> <li>• Good for Shotgun Metagenome Analysis.</li> </ul>	Multi-GPU Multi-Node
<b>GPU-Blast</b>	Carnegie Mellon University	Local search with fast k-tuple heuristic	<ul style="list-style-type: none"> <li>• Protein alignment according to BLASTP</li> </ul>	Single GPU Single Node
<b>mCUDA-MEME</b>	Open Source	Ultrafast scalable motif discovery algorithm based on MEME .	<ul style="list-style-type: none"> <li>• Scalable motif discovery algorithm based on MEME</li> </ul>	Multi-GPU Single Node
<b>MUMmer GPU</b>	Open Source	High-throughput local sequence alignment program	<ul style="list-style-type: none"> <li>• Aligns multiple query sequences against reference sequence in parallel</li> </ul>	Single GPU Single Node
<b>NVBIO</b>	Open Source	NVBIO is an open source C++ library of reusable components designed to accelerate bioinformatics applications using CUDA.	<ul style="list-style-type: none"> <li>• Data structures, algorithms, and utility routines useful for building complex computational genomics applications on CPU-GPU systems</li> </ul>	Multi-GPU Single Node
<b>NVBowtie</b>	Open Source	A largely complete implementation of the Bowtie2 aligner on top of NVBIO.	<ul style="list-style-type: none"> <li>• Good coverage of Bowtie2 features and comparable quality results</li> </ul>	Multi-GPU Single Node
<b>PEANUT</b>	Open Source	Read mapper for DNA or RNA sequence reads to a known reference genome.	<ul style="list-style-type: none"> <li>• Achieves supreme sensitivity and speed compared to current state of the art read mappers like BWA MEM, Bowtie2 and RazerS3</li> <li>• PEANUT reports both only the best hits or all hits</li> </ul>	Single GPU Single Node

> Indicates new application

<b>REACTA</b>	Open Source	A modified version of GCTA with improved computational performance, support for Graphics Processing Units (GPUs), and additional features. The purpose of REACTA is to quantify the contribution of genetic variation to phenotypic variation for complex traits.	<ul style="list-style-type: none"> <li>GRM creation, REML analysis, Regional Heritability (including multi-GPU)</li> </ul>	Multi-GPU Single Node
<b>SOAP3</b>	Genomics	GPU-based software for aligning short reads with a reference sequence. It can find all alignments with k mismatches, where k is chosen from 0 to 3.	<ul style="list-style-type: none"> <li>Short read alignment tool that is not heuristic based; reports all answers.</li> </ul>	Multi-GPU Multi-Node
<b>SOAP3-dp</b>	The University of Hong Kong	SOAP3-dp: Ultra-fast GPU-based tool for short read alignment via index-assisted dynamic programming.	<ul style="list-style-type: none"> <li>Borrows-Wheeler Transformation, Dynamic Programming</li> </ul>	Multi-GPU Single Node
<b>SeqNFind</b>	Accelerated Technology Laboratories	SeqNFind; is a powerful tool suite that addresses the need for complete and accurate alignments of many small sequences against entire genomes utilizing a unique hardware/software cluster system for facilitating bioinformatics research in Next Generation sequencing and genomic comparisons.	<ul style="list-style-type: none"> <li>Hardware and software for reference assembly, blast, SW, HMM, de novo assembly</li> </ul>	Multi-GPU Single Node
<b>Synomics Studio</b>	Row Analytics	Multi-Omics Biomarker Network Discovery and Validation Synomics Studio is a new, highly scalable analysis platform that enables researchers and clinicians to discover novel associations between multiple genotypic, phenotypic and clinical attributes of their patients and their disease risk /therapy responses.	<ul style="list-style-type: none"> <li>Multi-SNP association studies (GWAS studies with up to 30 SNPs/SNVs in combination)</li> <li>Configurable number of cycles of fully random permutation for validation of SNP networks Speed-up on GPU = 170x vs multi-core CPU alone (further speed-up available on multi-GPU and NVLink devices)</li> <li>Representative performance for 15,000 case:controls, 200,000 SNPs</li> <li>2 SNP associations found and validated in 12 mins on single 20 core IBM POWER8NVL with 4x Tesla P100 GPU</li> <li>17 SNP associations found and validated in 6 days on single 20 core IBM POWER8NVL with 4x Tesla P100 GPU</li> </ul>	Multi-GPU Single Node
<b>UGene</b>	Unipro	Open source Smith-Waterman for SSE/CUDA, Suffix array based repeats finder and dotplot.	<ul style="list-style-type: none"> <li>Fast short read alignment</li> </ul>	Multi-GPU Single Node
<b>WideLM</b>	Open Source	Fits numerous linear models to a fixed design and response.	<ul style="list-style-type: none"> <li>Parallel linear regression on multiple similarly-shaped models</li> </ul>	Multi-GPU Single Node

## Microscopy

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>BioEM</b>	Max Planck Institute	GPU-accelerated computing of Bayesian inference of electron microscopy images	<ul style="list-style-type: none"> <li>BioEM can use CUDA for the cross-correlation step, which essentially consists of an image multiplication in Fourier space and a Fourier back-transformation</li> </ul>	Multi-GPU Single Node
<b>cryoSPARC</b>	cryoSPARC	CryoSPARC is an easy to use software tool that enables rapid, unbiased structure discovery of proteins and molecular complexes from cryo-EM data.	<ul style="list-style-type: none"> <li>Ab-initio reconstruction, heterogeneous reconstruction, and high-speed high-resolution refinement of 3D protein structures implemented on GPUs</li> <li>Lean memory usage: 768 X 768 X 768 box size on a 12GB GPU for refinement</li> <li>Multiple simultaneous jobs on multiple GPUs</li> </ul>	Multi-GPU Single Node

<b>Huygens</b>	Scientific Volume Imaging	Huygens Products: Greatly improve your microscope images	<ul style="list-style-type: none"> <li>• Deconvolution of volumetric images and time series from widefield, confocal, light sheet, super-resolution STED microscopes and more.</li> <li>• Chromatic aberration and cross-talk correction, image stabilization and stitching</li> <li>• Visualization, tracking, colocalization and object analysis</li> <li>• Multi-GPU and cluster support</li> </ul>	Multi-GPU Single Node
<b>Microvolution</b>	Microvolution	Nearly instantaneous 3D deconvolution & up to 200 times faster.	<ul style="list-style-type: none"> <li>• 3D deconvolution for fluorescence microscopy</li> <li>• Written for use only on GPUs</li> <li>• Multi-GPU support</li> </ul>	Single GPU Single Node
<b>Phasefocus II Box</b>	phasefocus	The Phasefocus box product implements data processing aspects of Phasefocus's imaging methods that are known collectively as the Phasefocus Virtual Lens & ptychography.	<ul style="list-style-type: none"> <li>• Computational diffractive imaging engine (ptychography)</li> <li>• Multi-GPU support</li> </ul>	Single GPU Single Node
<b>RELION-2</b>	MRC Laboratory of Molecular Biology	RELION (for REGularised Likelihood OptimisatioN, pronounce rely-on) is a stand-alone computer program that employs an empirical Bayesian approach to refinement of (multiple) 3D reconstructions or 2D class averages in electron cryo-microscopy (cryo-EM).	<ul style="list-style-type: none"> <li>• Both image classification and high-resolution refinement accelerated up to 40-fold</li> <li>• Template-based particle selection accelerated almost 1000-fold</li> <li>• Reduced memory requirements</li> <li>• High-resolution cryo-EM structure determination in a matter of day on a single workstation</li> </ul>	Multi-GPU Single Node

## Molecular Dynamics

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ACEMD</b>	Accellera Ltd	GPU simulation of molecular mechanics force fields, implicit and explicit solvent	<ul style="list-style-type: none"> <li>• Written for use only on GPUs.</li> </ul>	Multi-GPU Multi-Node
<b>AMBER</b>	University of California at San Francisco	Suite of programs to simulate molecular dynamics on biomolecule.	<ul style="list-style-type: none"> <li>• PMEMD Explicit Solvent and GB Implicit Solvent</li> </ul>	Multi-GPU Single Node
<b>CHARMM</b>	Harvard University	MD package to simulate molecular dynamics on biomolecule.	<ul style="list-style-type: none"> <li>• Implicit (5x)</li> <li>• Explicit (2x)</li> <li>• Solvent via OpenMM, now ported natively to GPUs</li> </ul>	Multi-GPU Single Node
<b>DESMOND</b>	David E. Shaw Research	High-speed molecular dynamics simulations of biological systems.	<ul style="list-style-type: none"> <li>• The code uses novel parallel algorithms and numerical techniques to achieve high performance and accuracy</li> </ul>	Multi-GPU Single Node
<b>ESPReso</b>	ESPReso	Highly versatile software package for performing and analyzing scientific Molecular Dynamics many-particle simulations of coarse-grained atomistic or bead-spring models as they are used in soft-matter research in physics, chemistry and molecular biology.	<ul style="list-style-type: none"> <li>• Hydrodynamic / Electrokinetic forces</li> <li>• P3M electrostatics.</li> </ul>	Multi-GPU Single Node
<b>Folding@Home</b>	Stanford University	A distributed computing project that studies protein folding, misfolding, aggregation, and related diseases.	<ul style="list-style-type: none"> <li>• Powerful distributed computing molecular dynamics system</li> <li>• Implicit solvent and folding.</li> </ul>	Multi-GPU Single Node
<b>GENESIS</b>	RIKEN	GENESIS (GENERALized-Ensemble Simulation System) is a software package for molecular dynamics simulations and trajectory analyses.	<ul style="list-style-type: none"> <li>• Powerful parallelization for hybrid (CPU+GPU) systems</li> <li>• Full electrostatics with PME</li> <li>• Large (1-100 million atoms) biological systems</li> </ul>	Multi-GPU Single Node

<b>Genesis</b>	Diamond Visionics	GenesisRTX, is an advanced high-fidelity runtime rendering engine which eliminates the need for traditional off-line database compiling or formatting.	<ul style="list-style-type: none"> <li>• Powerful parallelization for hybrid (CPU+GPU) systems</li> <li>• Full electrostatics with PME</li> <li>• Large (1-100 million atoms) biological systems</li> </ul>	Multi-GPU Single Node
<b>GPUgrid.net</b>	Acellera Ltd	A distributed computing project that uses GPUs for molecular simulations.	<ul style="list-style-type: none"> <li>• High-performance all-atom biomolecular simulations</li> <li>• Explicit solvent and binding</li> </ul>	Multi-GPU Single Node
<b>GROMACS</b>	GROMACS	Simulation of biochemical molecules with complicated bond interactions.	<ul style="list-style-type: none"> <li>• Implicit (5x)</li> <li>• Explicit (2x) Solvent</li> </ul>	Multi-GPU Single Node
<b>HALMD</b>	HALMD	Large-scale simulations of simple and complex liquids.	<ul style="list-style-type: none"> <li>• Simple fluids and binary mixtures (pair potentials, high-precision NVE and NVT, dynamic correlations)</li> </ul>	Single GPU Single Node
<b>HOOMD-Blue</b>	University of Michigan	Particle dynamics package written grounds up for GPUs.	<ul style="list-style-type: none"> <li>• Written for use only on GPUs</li> </ul>	Multi-GPU Single Node
<b>HTMD</b>	Acellera Ltd	High throughput molecular dynamics simulations	<ul style="list-style-type: none"> <li>• Available via Conda and github</li> <li>• Support ACEMD, PMEMD, NAMD, GROMACS</li> <li>• AMBER and CHARMM force fields</li> <li>• Adaptive sampling, Markov State Models, visualization, protein preparation and ligand parameterization</li> </ul>	Multi-GPU Single Node
<b>LAMMPS</b>	Sandia National Lab	Classical molecular dynamics package	<ul style="list-style-type: none"> <li>• Lennard-Jones, Gay-Berne, Tersoff, and dozens more potentials</li> </ul>	Multi-GPU Multi-Node
<b>MELD</b>	University of Calgary	OpenMM plugin written for GPUs	<ul style="list-style-type: none"> <li>• Integrative approach to combine physics and information</li> <li>• Orders of magnitude faster protein folding than brute force MD</li> </ul>	Multi-GPU Single Node
<b>myPresto</b>	N2PC/AIST/JBIC, Japan	Open Source Computational Drug Discovery Suite.	<ul style="list-style-type: none"> <li>• High performance virtual screening by MD binding free energy calculation.</li> </ul>	Multi-GPU Multi-Node
<b>NAMD</b>	University of Illinois at Champaign Urbana	Designed for high-performance simulation of large molecular systems.	<ul style="list-style-type: none"> <li>• Full electrostatics with PME and most simulation features; 100M atom capable.</li> </ul>	Multi-GPU Single Node
<b>OpenMM</b>	Stanford University	Library and application for molecular dynamics for HPC with GPUs.	<ul style="list-style-type: none"> <li>• Molecular Dynamics toolkit that is the heart of Folding@Home and is used by several/a growing number of other applications.</li> <li>• Extensible and growing</li> <li>• Implicit and explicit solvent, custom forces</li> </ul>	Multi-GPU Single Node
<b>PolyFTS</b>	University of California at Santa Barbara	Classical molecular simulation code for studying polymer self-assembly and thermodynamics.	<ul style="list-style-type: none"> <li>• Uses auxiliary fields as the fundamental simulation degrees of freedom</li> <li>• Uses cuFFT extensively (~ 80%)</li> <li>• CUDA code is ~20%</li> <li>• Multi CPU or single GPU per job</li> <li>• 1x = Ivy Bridge E5-2690 CPU all 10 cores</li> <li>• 3-8X on K40 or K80 (utilizing 1/2 of the K80)</li> </ul>	Single GPU Single Node

<b>SOP-GPU</b>	SOP-GPU	SOP-GPU package, where SOP stands for the Self Organized Polymer Model fully implemented on a GPU, is a scientific software package designed to perform Langevin Dynamics Simulations of the mechanical or thermal unfolding, and mechanical indentation of large biomolecular systems in the experimental subsecond (millisecond-to-second) timescale.	<ul style="list-style-type: none"> <li>Langevin dynamics simulations using the coarse-grained Self Organized Polymer (SOP) model, Multiple simulation trajectories can be performed simultaneously on a single GPU, Calpha and Calpha-Cbeta models are supported, Simulations of protein forced unfolding, Novel simulations of nanoindentation in silico, Support for hydrodynamic interactions, Up to ~100 ms of simulation time per day, Systems of up to 1,000,000 amino-acids (on GPUs with 6GB or great memory)</li> </ul>	Single GPU Single Node
----------------	---------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------

## Quantum Chemistry

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Abinit</b>	ABINIT	Allows to find total energy, charge density and electronic structure of systems made of electrons and nuclei within DFT.	<ul style="list-style-type: none"> <li>Local Hamiltonian</li> <li>Non-local Hamiltonian</li> <li>LOBPCG algorithm</li> <li>Diagonalization/ orthogonalization.</li> </ul>	Multi-GPU Single Node
<b>ACES III</b>	University of Florida	Takes best features of parallel implementations of quantum chemistry methods for electronic structure.	<ul style="list-style-type: none"> <li>Integrating scheduling GPU into SIAL programming language and SIP runtime environment.</li> </ul>	Multi-GPU Multi-Node
<b>ACES 4</b>	University of Florida	New SIA/aces4 development A new super instruction architecture with interface applications for quantum chemistry (aces4) is now under development. Details and downloads can be found on the project's git hub site.	<ul style="list-style-type: none"> <li>Integrating scheduling GPU into SIAL programming language and SIP runtime environment.</li> </ul>	Multi-GPU Single Node
<b>ADF</b>	Software for Chemistry & Materials	Density Functional Theory (DFT) software package that enables first-principles electronic structure calculations.	<ul style="list-style-type: none"> <li>Geometry optimizations and frequency calculations with GGA functionals.</li> <li><a href="https://www.scm.com/doc/ADF/Input/Technical_Settings.html?highlight=GPU?">https://www.scm.com/doc/ADF/Input/Technical_Settings.html?highlight=GPU?</a></li> </ul>	Multi-GPU Single Node
<b>BigDFT</b>	BigDFT	Implements density functional theory by solving the Kohn-Sham equations describing the electrons in a material.	<ul style="list-style-type: none"> <li>DFT; Daubechies wavelets, part of Abinit</li> </ul>	Multi-GPU Multi-Node
<b>CP2K</b>	CP2K	Program to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems.	<ul style="list-style-type: none"> <li>DBCSR (space matrix multiply library)</li> </ul>	Multi-GPU Multi-Node
<b>GAMESS-UK</b>	Open Source	The general purpose ab initio molecular electronic structure program for performing SCF-, DFT- and MCSCF-gradient calculations.	<ul style="list-style-type: none"> <li>(ss ss) type integrals within calculations using Hartree-Fock ab initio methods and density functional theory</li> <li>Supports organics and inorganics.</li> </ul>	Multi-GPU Multi-Node
<b>GAMESS-US</b>	Ames Laboratory/Iowa State University	Computational chemistry suite used to simulate atomic and molecular electronic structure.	<ul style="list-style-type: none"> <li>Libqc with Rys Quadrature Algorithm, Hartree-Fock, MP2 and CCSD</li> </ul>	Multi-GPU Multi-Node
<b>Gaussian</b>	Gaussian, Inc.	Predicts energies, molecular structures, and vibrational frequencies of molecular systems.	<ul style="list-style-type: none"> <li>Joint NVIDIA, PGI and Gaussian collaboration</li> </ul>	Multi-GPU Single Node
<b>GPAW</b>	GPAW	Real-space grid DFT code written in C and Python	<ul style="list-style-type: none"> <li>Electrostatic poisson equation</li> <li>Orthonormalizing of vectors</li> <li>Residual minimization method (rmm-diis)</li> </ul>	Multi-GPU Multi-Node
<b>gWL-LSMS</b>	ORNL	Materials code for investigating the effects of temperature on magnetism.	<ul style="list-style-type: none"> <li>Generalized Wang-Landau method</li> </ul>	Multi-GPU Multi-Node
<b>LATTE</b>	Open Sourcee	Density matrix computations	<ul style="list-style-type: none"> <li>CU_BLAS, SP2 Algorithm</li> </ul>	Multi-GPU Single Node

<b>LSDalton</b>	LSDalton	Linear-scaling HF and DFT code suitable for large molecular systems, now also with some CCSD capabilities Tensor Algebra Library Routines for Shared Memory Systems which is being used to GPU accelerate three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	<ul style="list-style-type: none"> <li>• (T) correction to the CCSD energy</li> <li>• RI-MP2 energy/gradient (in development)</li> <li>• CCSD energy (in development)</li> <li>• GPU-based ERI generator (in development)</li> </ul>	Multi-GPU Single Node
<b>MOLCAS</b>	MOLCAS	Methods for calculating general electronic structures in molecular systems in both ground and excited states.	<ul style="list-style-type: none"> <li>• CU_BLAS</li> </ul>	Multi-GPU Single Node
<b>MOPAC2012</b>	MOPAC	Semiempirical Quantum Chemistry	<ul style="list-style-type: none"> <li>• Pseudodiagonalization, full diagonalization, and density matrix assembling via Magma libraries</li> </ul>	Single GPU Single Node
<b>NWChem</b>	NWChem	NWChem aims to provide its users with computational chemistry tools that are scalable both in their ability to treat large scientific computational chemistry problems efficiently, and in their use of available parallel computing resources from high-performance parallel supercomputers to conventional workstation clusters. Tensor Algebra Library Routines for Shared Memory Systems which is being used to GPU accelerate three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	<ul style="list-style-type: none"> <li>• Triples part of Reg-CCSD(T), CCSD and EOMCCSD task schedulers</li> </ul>	Multi-GPU Single Node
<b>Octopus</b>	Harvard University	Used for ab initio virtual experimentation and quantum chemistry calculations.	<ul style="list-style-type: none"> <li>• Full GPU support for ground-state, real-time calculations</li> <li>• Kohn-Sham Hamiltonian, orthogonalization, subspace diagonalization, poisson solver, time propagation DFT application</li> </ul>	Single GPU Single Node
<b>PEtot</b>	Lawrence Berkeley Laboratories	First principles materials code that computes the behavior of the electron structures of materials.	<ul style="list-style-type: none"> <li>• Density functional theory (DFT) plane wave pseudopotential calculations</li> </ul>	Multi-GPU Single Node
<b>Q-CHEM</b>	Q-Chem Inc.	Computational chemistry package designed for HPC clusters.	<ul style="list-style-type: none"> <li>• Various features including RI-MP2</li> </ul>	Single GPU Single Node
<b>QBox</b>	Open Source	Qbox is a C++/MPI scalable parallel implementation of first-principles molecular dynamics (FPMD) based on the plane-wave, pseudopotential formalism. Qbox is designed for operation on large parallel computers.		Single GPU Single Node
<b>QMCPACK</b>	QMCPACK	QMCPACK, an open-source production level many-body ab initio Quantum Monte Carlo code for computing the electronic structure of atoms, molecules, and solids.	<ul style="list-style-type: none"> <li>• Main features</li> </ul>	Multi-GPU Multi-Node
<b>Quantum Espresso/ PWscf</b>	Quantum Espresso Foundation	An integrated suite of computer codes for electronic structure calculations and materials modeling at the nanoscale.	<ul style="list-style-type: none"> <li>• PWscf package: linear algebra (matrix multiply), explicit computational kernels, 3D FFTs</li> </ul>	Multi-GPU Multi-Node
<b>QUICK</b>	Michigan State University	QUICK is a GPU-enabled ab initio quantum chemistry software package.	<ul style="list-style-type: none"> <li>• Running Hartree-Fock and DFT energy on GPU, Supports s, p, d, f orbitals on energy calculation, HF gradient with s,p,d orbital support, GPU-based ERI generator</li> </ul>	Multi-GPU Single Node
<b>RMG</b>	North Carolina State University	RMG is a density functional theory (DFT) based electronics structure code that uses real space grids to represent wavefunctions, charge densities, and ionic potentials. Designed for scalability, it has been run successfully on systems with thousands of nodes (including GPU nodes) and hundreds of thousands of CPU cores.	<ul style="list-style-type: none"> <li>• Supports 10k+ GPU nodes, multipetaflops capable</li> <li>• Handles thousands of atoms with full DFT precision</li> <li>• Supports multiple GPUs per node</li> <li>• Fully open source, with installation support, Downloads, documentation, forums <a href="http://www.rmgdft.org">www.rmgdft.org</a></li> <li>• Cray XE6/XK7</li> </ul>	Multi-GPU Single Node

> Indicates new application

<b>TAL-SH</b>	Oak Ridge National Lab	Tensor Algebra Library Routines for Shared Memory Systems which is being used to accelerated three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	<ul style="list-style-type: none"> <li>TAL-SH: Tensor Algebra Library for Shared Memory Computers: Nodes equipped with multicore CPU, Nvidia GPU, and Intel Xeon Phi (in progress). Author: Dmitry I. Lyakh (Liakh): quant4me@gmail.com, liakhdi@ornl.gov Copyright (C) 2014-2016 Dmitry I. Lyakh (Liakh) Copyright (C) 2014-2016 Oak Ridge National Laboratory (UT-Battelle) LICENSE: GNU Lesser General Public License v.3 API reference manual: DOC/TALSH_manual.txt BUILD: Modify the header of the Makefile accordingly and run make.</li> </ul>	Multi-GPU Multi-Node
<b>TeraChem</b>	PetaChem LLC	Quantum chemistry software designed to run on NVIDIA GPU.	<ul style="list-style-type: none"> <li>Full GPU-based solution; Performance compared to GAMESS CPU version</li> </ul>	Multi-GPU Single Node
<b>VASP</b>	University of Vienna	Complex package for performing ab-initio quantum-mechanical molecular dynamics (MD) simulations using pseudopotentials or the projector-augmented wave method and a plane wave basis set.	<ul style="list-style-type: none"> <li>Blocked Davidson (ALGO = NORMAL &amp; FAST), RMM-DIIS (ALGO = VERYFAST &amp; FAST), K-Points and optimization for critical step in exact exchange calculations</li> </ul>	Multi-GPU Single Node

## Visualization and Docking

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Amira</b>	Thermo fisher Scientific	A multifaceted software platform for visualizing, manipulating, and understanding Life Science and bio-medical data.	<ul style="list-style-type: none"> <li>3D visualization of volumetric data and surfaces</li> </ul>	Single GPU Single Node
<b>BINDSURF</b>	Universidad Catolica de Murcia	A virtual screening methodology that uses GPUs to determine protein binding sites.	<ul style="list-style-type: none"> <li>Allows fast processing of large ligand databases</li> </ul>	Single GPU Single Node
<b>BUDE</b>	Bristol University Docking Station	Molecular docking program	<ul style="list-style-type: none"> <li>Empirical Free Energy Force field</li> </ul>	Single GPU Single Node
<b>FastROCS</b>	Open Eye Scientific Software, Inc.	Molecule shape comparison application	<ul style="list-style-type: none"> <li>Real-time shape similarity searching/ comparison</li> </ul>	Multi-GPU Multi-Node
<b>Interactive Molecule Visualizer</b>	University of Illinois	Experimental interactive molecule visualizer based on a ray-tracing engine.	<ul style="list-style-type: none"> <li>Targeting high quality images and ease of interaction, IMV uses the latest GPU computing acceleration techniques, combined with natural user interfaces such as Kinect and Wiimotes</li> </ul>	Single GPU Single Node
<b>Molegro Virtual Docker 6</b>	QIAGEN	Method for performing high accuracy flexible molecular docking.	<ul style="list-style-type: none"> <li>Energy grid computation, pose evaluation and guided differential evolution</li> </ul>	Single GPU Single Node
<b>PIPER Protein Docking</b>	Boston University	Protein-protein docking program	<ul style="list-style-type: none"> <li>Molecule docking</li> </ul>	Single GPU Single Node
<b>PyMol</b>	Schrodinger, Inc.	User-sponsored molecular visualization system on an open-source foundation	<ul style="list-style-type: none"> <li>Lines: 460% increase</li> <li>Cartoons: 1246% increase</li> <li>Surface: 1746% increase</li> <li>Spheres: 753% increase</li> <li>Ribbon: 426% increase</li> </ul>	Single GPU Single Node
<b>VEGA ZZ</b>	University of California, San Francisco	Molecular Modeling Toolkit	<ul style="list-style-type: none"> <li>Virtual logP, molecular surface values</li> </ul>	Single GPU Single Node
<b>VMD</b>	University of Illinois	Visualization and analyzing large bio-molecular systems in 3-D graphics	<ul style="list-style-type: none"> <li>High quality rendering, large structures (100M atoms), analysis and visualization tasks, multiple GPU support for display of molecular orbitals</li> </ul>	Multi-GPU Single Node

## NUMERICAL ANALYTICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Accelereyes-ArrayFire</b>	ArrayFire	Comprehensive GPU function library	<ul style="list-style-type: none"> <li>Hundreds of functions for math, signal/image processing, statistics, and more.</li> </ul>	Multi-GPU Single Node
<b>HiPLAR</b>		3High Performance Linear Algebra in R	<ul style="list-style-type: none"> <li>Supports GPU and multi-core platforms, compatible with legacy R code, no new data types or operators, auto-tuning, support for R Matrix package</li> </ul>	Multi-GPU Single Node
<b>MATLAB</b>	Mathworks	GPU acceleration for MATLAB (high-level technical computing language).	<ul style="list-style-type: none"> <li>Acceleration for 200+ of most used -- and more than 500 -- MATLAB functions (incl. Parallel Computing Toolkit, Signal Processing, Image Processing, Communications Systems, etc)</li> </ul>	Multi-GPU Single Node
<b>Mathematica</b>	Wolfram	A symbolic technical computing language and development environment.	<ul style="list-style-type: none"> <li>Development environment for CUDA and OpenCL</li> <li>GPU acceleration for Wolfram Finance Platform.</li> </ul>	Multi-GPU Single Node
<b>NMath Premium</b>	NMath	GPU-accelerated math and statistics for .NET, automatically detects the presence of a CUDA-enabled GPU at runtime and seamlessly redirects appropriate computations to it.	<ul style="list-style-type: none"> <li>Automatically offloads computations to the GPU.</li> </ul>	Single GPU Single Node

## PHYSICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AWP</b>	AWP	The Anelastic Wave Propagation, AWP-ODC, independently simulates the dynamic rupture and wave propagation that occurs during an earthquake. Dynamic rupture produces friction, traction, slip, and slip rate information on the fault. The moment function is constructed from this fault data and used to currentize wave propagation.	<ul style="list-style-type: none"> <li>3D Finite Difference Computation</li> </ul>	Single GPU Single Node
<b>BQCD</b>	USQCD	Lattice quantum chromodynamics application, used for nuclear and high energy physics calculations. Most usage is concentrated in EMEA	<ul style="list-style-type: none"> <li>Wilson-clover fermion linear solver</li> </ul>	Multi-GPU Single Node
<b>CASTRO</b>	CASTRO	A multicomponent compressible hydrodynamic code for astrophysical flows including self-gravity, nuclear reactions and radiation. CASTRO uses an Eulerian grid and incorporates adaptive mesh refinement (AMR). The approach uses a nested hierarchy of logically-rectangular grids with simultaneous refinement in both space and time.	<ul style="list-style-type: none"> <li>Gravitational Field Solver</li> </ul>	Multi-GPU Single Node
<b>Changa</b>	CHANGA	Astrophysics code performs collisionless N-body simulations. It can perform cosmological simulations with periodic boundary conditions in comoving coordinates or simulations of isolated stellar systems.	<ul style="list-style-type: none"> <li>Gravitational Model has been accelerated using CUDA</li> </ul>	Single GPU Single Node
<b>Chemora</b>	CHEMORA	Chemora is a system for performing simulations of systems described by differential equations running on accelerated computational clusters.	<ul style="list-style-type: none"> <li>Chemora embeds the equations' computational kernels into dynamically compiled loop nests shaped for input size and GPU structure</li> </ul>	Multi-GPU Single Node



<b>Cholla</b>	Cholla	Computational Hydrodynamics On ParaLLeL Architectures for Astrophysics	<ul style="list-style-type: none"> <li>Models the Euler equations on a static mesh and evolves the fluid properties of thousands of cells simultaneously using GPUs</li> <li>It can update over ten million cells per GPU-second while using an exact Riemann solver and PPM reconstruction, allowing computation of astrophysical simulations with physically interesting grid resolutions (<math>&gt;256^3</math>) on a single device; calculations can be extended onto multiple devices with nearly ideal scaling beyond 64 GPUs</li> </ul>	Multi-GPU Single Node
<b>Chroma</b>	USQCD	Lattice Quantum Chromodynamics (LQCD)	<ul style="list-style-type: none"> <li>Wilson-clover fermions, Krylov solvers, Domain-decomposition</li> </ul>	Multi-GPU Multi-Node
<b>CPS</b>	USQCD	Lattice quantum chromodynamics application, used for nuclear and high energy physics calculations.	<ul style="list-style-type: none"> <li>Wilson, domain-wall and Mbius fermion linear solvers</li> </ul>	Multi-GPU Single Node
<b>CPS (GRID)</b>	USQCD	CPS is developed for lattice QCD and written by C++, with some machine-specific assembly routines. It is being developed by members of Columbia University, Brookhaven National Laboratory. The CPS consists of code to build a library which can be statically linked to your code to create an executable. CPS has optimized codes for QCDOC, IBM Blue Gene machines, and builds for scalar machines or parallel machines with QMP.	<ul style="list-style-type: none"> <li>QUDA is supported. The GRID code from Edinburgh is currently being optimized.</li> </ul>	Multi-GPU Multi-Node
<b>CST PARTICLE STUDIO</b>	Dassault Systèmes SIMULIA Corp.	Self-consistent simulation of charged particles in electromagnetic fields	<ul style="list-style-type: none"> <li>Particle-in-Cell Solver</li> </ul>	Multi-GPU Multi-Node
<b>GADGET</b>	Max Planck Institute	A code for cosmological simulations of structure formation	<ul style="list-style-type: none"> <li>MPI</li> </ul>	Multi-GPU Multi-Node
<b>GAMER</b>	OpenSource	A GPU-accelerated Adaptive Mesh Refinement Code for astrophysical applications. Currently the code solves the hydrodynamics with self-gravity.	<ul style="list-style-type: none"> <li>Adaptive mesh refinement (AMR). Hydrodynamics with self-gravity</li> <li>A variety of GPU-accelerated hydrodynamic and Poisson solvers. Hybrid OpenMP/MPI/GPU parallelization</li> <li>Concurrent CPU/GPU execution for performance optimization. Hilbert space-filling curve for load balance</li> </ul>	Multi-GPU Single Node
> <b>GPU-AH</b>	Universidade do Porto	Developed at Centro de Astrofisica e Astronomia da Universidade do Porto, GPU-AH simulates the evolution of a network of line-like topological defects - Abelian-Higgs cosmic strings - in a cosmic context.	<ul style="list-style-type: none"> <li>Besides evolving the network timestep-by-timestep, it also calculates the average network density and velocity. These quantities confirm that the simulation reproduces the behaviour of local string networks when compared to literature results.</li> </ul>	Single GPU Single Node
> <b>GPUwalls</b>	Universidade do Porto	Developed at Centro de Astrofisica e Astronomia da Universidade do Porto, GPUwalls simulates the evolution of a network of the simplest topological defect - domain wall - in a cosmic context.	<ul style="list-style-type: none"> <li>Besides evolving the network timestep-by-timestep, it also calculates the average network density and velocity. These quantities confirm that the simulation reproduces the behaviour of wall networks in the literature up to machine precision (compared to our earlier CPU version of the code).</li> </ul>	Single GPU Single Node

<p>&gt; <b>GTC Irvine</b></p>	<p>GTC</p>	<p>The gyrokinetic toroidal code (GTC) is a massively parallel, particle-in-cell code for turbulence simulation in support of the burning plasma experiment ITER, the crucial next step in the quest for fusion energy. GTC is the production code for the multi-institutional US Department Of Energy (DOE) Scientific Discovery through Advanced Computing (SciDAC) project, GSEP Center (Gyrokinetic Simulation of Energetic Particle Turbulence and Transport), and DOE INCITE project that was awarded 35M hours of CPU time for 2011. Currently maintained at UC Irvine, GTC was the first fusion code to reach in production simulations the teraflop in 2001 on the seaborg computer at NERSC and the petaflop in 2008 on the jaguar computer at ORNL. GTC simulation of the turbulence self-regulation by zonal flows was published in a 1998 Science paper, which has received the most citations for any magnetic fusion research paper published since 1996.</p>	<ul style="list-style-type: none"> <li>• Key accelerated features are the PUSHe, Collision and Poisson Solver</li> </ul>	<p>Multi-GPU Multi-Node</p>
<p><b>GTC-P</b></p>	<p>Princeton Plasma Physics Lab</p>	<p>A development code for optimization of plasma physics. Full science and data sets are included, but in a simplified form to allow performance testing and tuning.</p>	<ul style="list-style-type: none"> <li>• Optimized with CUDA. OpenACC development underway</li> </ul>	<p>Multi-GPU Single Node</p>
<p><b>HACC</b></p>	<p>HACC</p>	<p>Simulates N-Body Astrophysics. The HACC (Hardware/Hybrid Accelerated Cosmology Code) framework exploits this diverse landscape at the largest scales of problem size, obtaining high scalability and sustained performance. Developed to satisfy the science requirements of cosmological surveys, HACC melds particle and grid methods using a novel algorithmic structure that flexibly maps across architectures, including CPU/GPU, multi/many-core, and Blue Gene systems. We demonstrate the success of HACC on two very different machines, the CPU/GPU system Titan and the BG/Q systems Sequoia and Mira, attaining unprecedented levels of scalable performance. We demonstrate strong and weak scaling on Titan, obtaining up to 99.2% parallel efficiency, evolving 1.1 trillion particles.</p>	<ul style="list-style-type: none"> <li>• This code has been optimized with CUDA runs in full production mode</li> </ul>	<p>Multi-GPU Single Node</p>
<p><b>HAMR GPU</b></p>	<p>HAMR</p>	<p>GPU accelerated General Relativistic Magneto Hydrodynamic application</p>	<ul style="list-style-type: none"> <li>• Active galactic nuclei which assumes a radiatively inefficient sub-eddington rate torus</li> <li>• Axisymmetric ideal MHD. Viscosity and resistivity through use of Riemann solver (HLL)</li> <li>• Density floors to mass load the jet. Uses grids that can resolve the substructure of the jet over 5 orders of magnitude</li> </ul>	<p>Multi-GPU Single Node</p>
<p><b>MAESTRO</b></p>	<p>MAESTRO</p>	<p>A low Mach number stellar hydrodynamics code that can be used to simulate long-time, low-speed flows that would be prohibitively expensive to model using traditional compressible code.</p>	<ul style="list-style-type: none"> <li>• Gravitational Field Solver</li> </ul>	<p>Multi-GPU Single Node</p>
<p><b>MILC</b></p>	<p>USCQD</p>	<p>Lattice Quantum Chromodynamics (LQCD) codes simulate how elemental particles are formed and bound by the strong force to create larger particles like protons and neutrons.</p>	<ul style="list-style-type: none"> <li>• Staggered fermions</li> <li>• Krylov solvers</li> <li>• Gauge-link fattening</li> </ul>	<p>Multi-GPU Multi-Node</p>

<b>NekCEM</b>	ANL	A high-fidelity, open-source electromagnetics solver based on spectral element and spectral element discontinuous Galerkin methods, written in Fortran and C. The code is actively developed at Mathematics and Computer Science Division of Argonne National Laboratory.	<ul style="list-style-type: none"> <li>The OpenACC implementation covers all solution routines for the Maxwell equation solver in NekCEM, including a highly tuned element-by-element operator evaluation and a GPUDirect gather-scatter kernel to effect nearest-neighbor flux exchanges</li> </ul>	Multi-GPU Multi-Node
<b>OSIRIS</b>	UCLA Plasma Physics Group	Simulates Plasma Physics including Laser interaction	<ul style="list-style-type: none"> <li>2 dimensions of the particle push have been optimized with CUDA. Additional optimization is being planned with OpenACC</li> </ul>	Multi-GPU Single Node
<b>PIConGPU</b>	HZDR	A relativistic Particle-in-Cell code that describes the dynamics of a plasma by computing the motion of electrons and ions subject to the Maxwell-Vlasov equation.	<ul style="list-style-type: none"> <li>Simulation of laser-wakefield acceleration of electrons.</li> </ul>	Multi-GPU Single Node
<b>PPM</b>	PPM	Piecewise parabolic method, a higher-order extension of Godunov's method which uses spatial interpolation and allows for a steeper representation of discontinuities, particularly contact discontinuities.	<ul style="list-style-type: none"> <li>Turbulent, compressible mixing of gases in the context of stars near the ends of their lives and also in inertial confinement fusion</li> </ul>	Single GPU Single Node
<b>QUDA</b>	USQCD	Library for Lattice QCD calculations using GPUs.	<ul style="list-style-type: none"> <li>QUDA supports the following fermion formulations: Wilson, Wilson-clover, Twisted mass, Improved staggered (asqtad or HISQ) and Domain wall</li> </ul>	Multi-GPU Single Node
<b>RAMSES</b>	CEA	Simulates astrophysical problems on different scales (e.g. star formation, galaxy dynamics, cosmological structure formation).	<ul style="list-style-type: none"> <li>GPU acceleration is applied for radiative transfer for reionization, and the hydrodynamic solver using AMR</li> </ul>	Multi-GPU Multi-Node
<b>XGC</b>	PPPL	Simulates edge effects for MHD plasma physics	<ul style="list-style-type: none"> <li>The particle push portion has been optimized with CUDA and is being fully optimized with OpenACC and CUDA</li> </ul>	Multi-GPU Multi-Node

## SCIENTIFIC VISUALIZATION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3D Slicer</b>	Open Source	Medical visualization & segmentation	<ul style="list-style-type: none"> <li>Rendering, image processing</li> </ul>	Single GPU Single Node
<b>Animator</b>	GNS	Industry proven, modern post-processing app for CAE	<ul style="list-style-type: none"> <li>Rendering</li> </ul>	Multi-GPU Single Node
<b>ANSYS EnSight</b>	ANSYS	Industry proven post-processing app for CAE	<ul style="list-style-type: none"> <li>Rendering</li> <li>Ray tracing</li> </ul>	Multi-GPU Single Node
<b>FieldView</b>	IntelligentLight	Visualization application for CFD	<ul style="list-style-type: none"> <li>Rendering</li> </ul>	Single GPU Single Node
<b>FluoRender (SCI, U of Utah)</b>	University of Utah	Interactive rendering tool for confocal microscopy data visualization.	<ul style="list-style-type: none"> <li>Multi-channel volume rendering</li> </ul>	Single GPU Single Node
<b>HVR (LCSE, U of Minnesota)</b>	University of Minnesota	Interactive volume rendering application	<ul style="list-style-type: none"> <li>Volume rendering</li> </ul>	Multi-GPU Single Node
<b>ImageVis3D</b>	SCI, University of Utah	Simple, scalable, and interactive volume rendering application.	<ul style="list-style-type: none"> <li>Out-of-core volume rendering</li> </ul>	Single GPU Single Node
<b>Index</b>	NVIDIA	Interactive distributed volumetric compute and visualization framework.	<ul style="list-style-type: none"> <li>Parallel distributed 3D rendering of dense or sparse volumes</li> <li>Accurate ray casting or ray tracing at high resolution of full size datasets</li> <li>Plug-in to ParaView also available.</li> </ul>	Multi-GPU Multi-Node
<b>ParaView</b>	Kitware	Scalable data analysis and visualization application. One of the main vis tools at HPC sites.	<ul style="list-style-type: none"> <li>Rendering and analysis tasks</li> <li>Plugin for NVIDIA Index</li> <li>OptiX rendering backend</li> <li>CUDA accelerated filters (data transformation routines)</li> </ul>	Multi-GPU Multi-Node

<b>Seg3D (SCI, U of Utah)</b>	SCI, University of Utah	Segmentation application for medical data	• Rendering, image processing	Single GPU Single Node
<b>SPECFEM3D</b>	CIG	There are two modules/apss in the SPECFEM family: GLOBE and CARTESIAN. The global model is the former Gordon Bell Awardee code. Used for global inversion. Also part of the CAAR effort (although, that one is mostly focused on workflow, rather than the actual model). The regional model is CARTESIAN and it is the app used for seismic simulations, earthquake models, submarine acoustics etc In addition to being used as a community app, Specfem3D is also use as a proxy app for proprietary codes	• OpenCL and CUDA hardware accelerators, based on an automatic source-to-source transformation library. Simulates acoustic (fluid), elastic (solid), coupled acoustic/elastic, poroelastic or seismic wave propagation in any type of conforming mesh of hexahedra (structured or not).	Multi-GPU Single Node
> <b>Tecplot</b>	Tecplot	General purpose scientific visualization software for Aerodynamics, O&G, Internal Combustion and Geoscience applications	• Rendering	
<b>VisIt</b>	LLNL	Scalable data anlysis and visualization application	• Rendering and analysis tasks	Multi-GPU Single Node
<b>Visulalization Toolkit (VTK)</b>	Open Source	Data anlysis and visualization toolkit	• Rendering	Single GPU Single Node
<b>v13 (Argonne National Lab)</b>		Large dataset visualization in cosmology, astrophysics, and biosciences fields.	• Volume rendering of particles	Multi-GPU Single Node

## Safety and Security

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AI-NVR</b>	IronYun	Search in Video	• Search in Video	
> <b>Alert</b>	Irvine Sensors	Alert provides people counting and intrusion detection	• People counting, Intrusion detection	
> <b>Arvas</b>	VI Dimensions			
> <b>Better Tomorrow</b>	Anyvision	Face recognition for multiple industries	• Face recognition	Multi-GPU Single Node
<b>BioSurveillance NEXT, BioFinder</b>	Herta Security	Real time facial recognition and forensic alerts against multiple watchlists.	• Supports crowded scenes, difficult lighting, faster than real-time analysis, partial face concealment	Multi-GPU Single Node
<b>Cezurity EVO</b>	Cezurity	Event Observer (EvO): engine for detecting malicious activity on user computers. Centralized detection engine; Event chains; Context; Real-time analysis - Cezurity Cloud: Cloud-based technology for detecting malware. Cezurity Cloud has the flexibility to fit into diverse solutions. Different information can be sent and processed by the server, depending on the needs of each product or solution. For example, Cezurity Cloud is currently used as a subsystem to supply data for the Cezurity EvO detection engine. Cezurity Cloud helps the Anti-Virus Scanner to detect malware. In addition, the technology is used for monitoring and analyzing changes in our APT-D solution designed to detect persistent threats against corporate networks.	• CUDA	Multi-GPU Single Node
<b>Cylance</b>	Cylance	Advanced AI-based end point malware detection	• End point malware detection solution build using GPU deep learning technology	Multi-GPU Single Node
> <b>Discover Experencer</b>	Viisights			
> <b>Face++</b>	Megvii	Face Recognition for many verticals	• Face Recognition	

<b>FaceControl</b>	VOCORD	Detects and recognizes the faces of people, freely passing-by cameras, providing an instant alert to people on a watchlist, recognizes age and gender, counts people by faces, tags newcomers and regular visitors. The system uses deep neural network algorithms and performs recognition with extremely high accuracy in field applications.	<ul style="list-style-type: none"> <li>• Non-cooperative biometrical facial recognition system, ALPR, video analytics and pattern recognition, video processing and video enhancement</li> </ul>	Multi-GPU Single Node
> <b>FindFace Enterprise Server SDK</b>	NTechLab	Powered by Ntechlab face recognition algorithm, FindFace Enterprise Server SDK effectively processes face recognition and works on the client, no biometric data is transferred or stored by NtechLab. It detects and identifies people faces in live video streams and video footage addressing a wide range of business tasks, such as precise people count, demographic information, people flow and client behavior. FindFace Enterprise Server SDK allows for integration into any web, mobile, or desktop application using the cross-platform REST API. The FindFace Enterprise Server SDK 2.0 can be widely applied in a variety cases, including customer analytics, client verification, fraud prevention, hospitality, and access control.	<ul style="list-style-type: none"> <li>• CUDA, NVDec, DeepStream (testing)</li> </ul>	Multi-GPU Single Node
<b>Glueck Media; Glueck Analytics</b>	Glueck	Deep Learning/Machine Learning based Computer Vision technology enabling understanding of how human feels and perceives the environment around them, focusing on face and people analytics.	<ul style="list-style-type: none"> <li>• Specific features supported includes; Facial Expression, Age Estimation, Gender, Ethnicity, Multi Face Tracking, Attention Time</li> </ul>	Multi-GPU Single Node
<b>Graydient V (Video)</b>	Giant Grey	Machine learning anomaly detection for enhanced video analytics.	<ul style="list-style-type: none"> <li>• Proactive event detection and real-time alerts for safety, unauthorized access prevention, and loss prevention</li> <li>• 24/7 real-time analysis and alerting scaling to thousands of video streams across remote and geographically dispersed locations</li> </ul>	Multi-GPU Single Node
<b>Ikena Forensic, Ikena Spotlight</b>	MotionDSP	Real-time (render-less) super-resolution-based video enhancement and redaction software for forensic analysts and law enforcement professionals	<ul style="list-style-type: none"> <li>• Multi-filter, render-less video reconstruction (super-resolution, stabilization, light/color correction), and automatic tracking for redaction video from body cameras, CCTV and other sources</li> </ul>	Multi-GPU Single Node
<b>iMotionFocus</b>	iCetana	Intelligent analysis of video on 1,000+ camera streams to significantly filter and reduce the camera streams requiring an operator view.	<ul style="list-style-type: none"> <li>• GPU accelerated machine learning to identify abnormal activity within video streams</li> </ul>	Multi-GPU Single Node
<b>innovi</b>	Agent Video Intelligence (Agent Vi)	Agent Video Intelligence's (Agent Vi) solutions allow users to achieve optimal value from their video surveillance networks by automating video analysis to detect and alert for events of interest, expedite search in recorded video and extract statistical data from the footage captured by surveillance cameras	<ul style="list-style-type: none"> <li>• The solution provides real-time video analysis and alerts, video search and investigation, big data analysis, geospatial mapping and more</li> </ul>	
<b>LUNA</b>	VisionLabs	LUNA PLATFORM is a biometric data management system for facial verification and identification. The platform offers a great flexibility to create scenarios of varying complexity for integrated facial recognition on GPU. LUNA SDK, a facial recognition engine developed by VisionLabs, is the core technology of the LUNA PLATFORM.	<ul style="list-style-type: none"> <li>• Face detection, face alignment, facial descriptor extraction, face matching, facial attribute classification and face spoofing prevention. Optimized scalability using multithreading; Computationally efficient and compact face descriptors; Broad range of working conditions with domain-specific face descriptors</li> </ul>	Multi-GPU Single Node

<b>OpenALPR</b>	OpenALPR	Automatic license plate and vehicle make/model/year recognition software applied to video streams from IP cameras.	<ul style="list-style-type: none"> <li>• High accuracy license plate character recognition spanning North America, Europe, United Kingdom, Australia, Korea, Singapore and Brazil</li> <li>• APIs and source code available for embedded applications and web services</li> </ul>	Multi-GPU Single Node
<b>Recotraffic; Recosecure; Recohospital</b>	Recogine	Intelligent Transportation Systems covering complex multi-modal surface transportation solutions at a regional, sub-regional, corridor and small area level using deep computer vision technologies.	<ul style="list-style-type: none"> <li>• Traffic Data Collection, Incident Detection, Integrated Management, Vehicle Classification and supporting related application</li> </ul>	Multi-GPU Single Node
<b>SenDISA Platform</b>	Sensen Networks	SenSen provides Video-IoT data analytic software solutions targeted at increasing revenue and reducing the cost of operations of customers. SenSen software can process and fuse data from cameras and other sensors like GPS, Radar, and Lidar in real time for parking guidance, parking enforcement, speed enforcement, traffic data analytics and road safety applications. Casinos use SenSen solutions for table game analytic solutions and customer analytics. SenSen solutions are also used in retail, security and tolling applications.	<ul style="list-style-type: none"> <li>• Intelligent Transportation - parking enforcement Casino game table analytics</li> </ul>	
<b>SenseFace</b>	Sensetime	Intelligent engine of visual surveillance system with ex ante, concurrent & ex post technical support service. Applicable to security, missing people & suspects search.	<ul style="list-style-type: none"> <li>• Face Recognition</li> </ul>	
<b>Syndex Pro</b>	Briefcam	Improved security and operations by turning video data into useful information. Based on Video Synopsis technology, Syndex Pro allows users to review hours of video in minutes, while applying search filters for achieving accurate results and faster time-to-target. Data can be processed on-demand or in real time to support a wide range of use cases.	<ul style="list-style-type: none"> <li>• Review hours of video in minutes, Search in Video</li> </ul>	
<b>XIntelligence XHound XTransport</b>	Xjera Labs	AI-based image and video analytics solution. This solution is ideal for people counting and recognition and vehicle counting for various commercial applications, with proven accuracy, high-level customization, and robust security.	<ul style="list-style-type: none"> <li>• People counting, face recognition, license plate recognition</li> </ul>	
> <b>Xpose</b>	Lentix			

# Tools and Management

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Allinea Forge</b>	Allinea now owned by ARM Ltd	Allinea Forge Professional provides all you will need to debug, profile and optimize for high performance - from single threads through to complex parallel HPC and scientific codes with MPI, OpenACC, OpenMP, threads or NVIDIA CUDA	<ul style="list-style-type: none"> <li>• CUPTI, cudagdb</li> </ul>	Multi-GPU Multi-Node
<b>Bright Cluster Manager</b>	Bright Computing	Bright Cluster Manager lets you administer clusters as a single entity, provisioning the servers, GPUs, operating system, and workload manager from a unified interface. We make it easy to build an NVIDIA GPU cluster by packaging all the relevant software including CUDA, NVIDIA driver, DCGM, NCCL, and a full deep learning stack. With Bright, you can configure GPUs individually or in groups, which is a real time saver for those with a large cluster. You can even set properties on your NVIDIA GPUs using BrightView. Once up and running, we monitor GPU metrics and run GPU health checks to make sure everything is working as it should. Bright makes managing GPU clusters easy.	<ul style="list-style-type: none"> <li>• Intuitive web app provides comprehensive view of GPU and cluster metrics; Powerful Cluster Management Shell as alternative user interface; Fully Support for NVIDIA libraries, CUDA, OpenCL, OpenACC, CUDA-aware libraries, NCCL, and CUB; Comprehensive monitoring of GPU; Brings in GPU resources from public (AWS, Azure) and private (OpenStack) clouds within minutes &amp;; Automated scaling of the cluster based on pre-defined policies &amp; Supports several popular Linux distributions: RHEL and derivatives, SUSE SLES and Ubuntu LTS; GPU-enabled Docker containers; Offers a complete deep learning stack; Deployment for popular HPC filesystems and management of fast interconnects</li> </ul>	Multi-GPU Multi-Node
<b>cmake</b>	Kitware	CMake is an open-source, cross-platform family of tools designed to build, test and package software. CMake is used to control the software compilation process using simple platform and compiler independent configuration files, and generate native makefiles and workspaces that can be used in the compiler environment of your choice. The suite of CMake tools were created by Kitware in response to the need for a powerful, cross-platform build environment for open-source projects.		N/A
<b>HPCToolkit</b>	Rice University	HPCToolkit is an integrated suite of tools for measurement and analysis of program performance on computers ranging from multicore desktop systems to the nation's largest supercomputers.	<ul style="list-style-type: none"> <li>• CUPTI</li> </ul>	Multi-GPU Multi-Node
<b>IBM Spectrum LSF</b>	IBM Corporation	IBM Spectrum LSF is a complete workload management solution for demanding HPC environments. Featuring intelligent, policy-driven scheduling, it helps organizations to improve competitiveness by accelerating research and design while controlling costs through superior resource utilization and ease of use. Building on over 20 years of experience, IBM Spectrum LSF features a highly scalable and available architecture designed to address the challenge of aligning compute resources with business priorities. With the ability to detect, monitor and schedule GPU enabled workloads to the appropriate resources, IBM Spectrum LSF enables users to easily take advantage of the benefits provided by GPUs.	<ul style="list-style-type: none"> <li>• Enforcement of GPU allocations via cgroups; Exclusive allocation and round robin shared mode allocation; CPU-GPU affinity; Boost control; Power management/li; Multi-Process Server (MPS) support; NVIDIA Volta and DCGM support</li> </ul>	Multi-GPU Multi-Node

<b>Magma</b>	ICL - University of Tennessee Knoxville	The MAGMA project aims to develop a dense linear algebra library similar to LAPACK but for heterogeneous/hybrid architectures, starting with current "Multicore+GPU" systems. The MAGMA research is based on the idea that, to address the complex challenges of the emerging hybrid environments, optimal software solutions will themselves have to hybridize, combining the strengths of different algorithms within a single framework. Building on this idea, we aim to design linear algebra algorithms and frameworks for hybrid manycore and GPU systems that can enable applications to fully exploit the power that each of the hybrid components offers.		Multi-GPU Single Node
<b>open SpeedShop</b>	Krell Institute	Open SpeedShop (OISS) is an open source multi-platform performance tool enabling performance analysis of HPC applications running on both single node and large scale platforms. OISS gathers and displays several types of information to aid in solving performance problems, including: program counter sampling for a quick overview of the applications performance, call path profiling to add caller/callee context and locate critical time consuming paths, access to the machine hardware counter information, input/output tracing for finding I/O performance problems, MPI function call tracing for MPI load imbalance detection, memory analysis, POSIX thread tracing, NVIDIA CUDA analysis, and OpenMP analysis. OISS offers a command-line interface (CLI), a graphical user interface (GUI) and a python scripting API user interface.	• CUPTI	Multi-GPU Multi-Node
<b>PAPI</b>	ICL - University of Tennessee Knoxville	PAPI provides the tool designer and application engineer with a consistent interface and methodology for use of the performance counter hardware found in most major microprocessors. PAPI enables software engineers to see, in near real time, the relation between software performance and processor events. In addition, PAPI provides access to a collection of components that expose performance measurement opportunities across the hardware and software stack.	• CUPTI	Multi-GPU Multi-Node



<b>PBS Professional</b>	Altair	PBS Works is an HPC Resource management suite consisting of PBS Professional (Workload scheduler), PBS Access (Remote Job submission and Remote Display) and PBS Control (Manage the compute infrastructure).	<ul style="list-style-type: none"> <li>• Pre-job node risk identification and GPU resource allocation; Automated monitoring of node health; Reduced job terminations due to GPU failure; Increased system resiliency via intelligent routing decisions; Increased job throughput via topology optimization; Optimized job scheduling through GPU load and health monitoring; Real-time monitoring and troubleshooting; Report usage (hardware and software) by individual, group, project, or organization; Understand trends to plan for future needs; Simulate infrastructure changes to answer “what if” questions using real historical workloads with various hypothetical changes in quantities of nodes and licenses &amp; Define once, deploy anywhere: public cloud, private cloud, or bare metal; Leverage cloud bursting for peak workloads; Submit and manage jobs from anywhere with drag-and-drop simplicity; Manage remote files and data directly from the secure web portal; Watch progress of running jobs, both graphically and via tail -f.</li> </ul>	Multi-GPU Multi-Node
<b>Parallware Trainer</b>	Appentra Solutions	Parallware Trainer is the new interactive, real-time editor with GUI features to facilitate the learning, usage, and implementation of parallel programming. Users are actively involved in learning parallel programming through observation, comparison, and hands-on experimentation. Parallware Trainer provides support for widely used parallel programming strategies using OpenACC and OpenMP with execution on multicore processors and GPUs.		N/A
<b>SLURM</b>	SchedMD	SLURM is a highly configurable open source workload and resource manager. In its simplest configuration, Slurm can be installed and configured in a few minutes. Use of optional plugins provides the functionality needed to satisfy the needs of demanding HPC centers with diverse job types, policies and work flows. Advanced configurations use plug-ins to provide features like accounting, resource limit management, by user or bank account, and support for sophisticated scheduling algorithms.	<ul style="list-style-type: none"> <li>• Scales to millions of cores and tens of thousands of GPGPUs. Military grade security ; Heterogenous platform support allowing users to take advantage of GPGPUs. Flexible plugin framework enables Slurm to meet complex customization requirements ; Topology aware job scheduling for maximum system utilization; Extensive scheduling options including advanced reservations, suspend/resume, backfill, fair-share and preemptive scheduling for critical jobs; No single point of failure</li> </ul>	Multi-GPU Multi-Node
<b>TAU - Tuning and Analysis Utilities</b>	University of Oregon	TAU is a program and performance analysis tool framework. TAU provides a suite of static and dynamic tools to form an integrated analysis environment for parallel Fortran, C++, C, Java, and Python applications. Also, recent advancements in TAU’s code analysis capabilities have allowed new static tools to be developed, such as an automatic instrumentation tool	<ul style="list-style-type: none"> <li>• CUPTI</li> </ul>	Multi-GPU Multi-Node

<b>Torque / Moab</b>	Adaptive Computing	Moab HPC Suite is a workload and resource orchestration platform that automates the scheduling, managing, monitoring, and reporting of HPC workloads on massive scale. TORQUE provides control over batch jobs and distributed computing resources. It is an advanced open-source product based on the original PBS project* and incorporates the best of both community and professional development.	<ul style="list-style-type: none"> <li>Request/schedule gpus based on gpu location in NUMA systems ; Collect and report metrics and status information; Set gpu mode at job run time</li> </ul>	Multi-GPU Multi-Node
<b>Totalview for HPC</b>	Rogue Wave Software	TotalView for HPC allows simultaneous debug many processes and threads in a single window. Work backwards from failure through reverse debugging, isolating the root cause faster by eliminating repeated restarts of the application. Reproduce difficult problems that occur in concurrent programs that use threads, OpenACC, OpenMP, MPI and CUDA	<ul style="list-style-type: none"> <li>CUDA debug API</li> </ul>	Multi-GPU Multi-Node
<b>Univa Grid Engine</b>	Univa	The Univa Grid Engine suite is a leading workload management system. The solution maximizes the use of shared resources in a data center and applies advanced management policy enforcement to deliver results faster, more efficiently, and with lower overall costs. The product suite can be deployed in any technology environment, including containers: on-premise, hybrid or in the cloud.	<ul style="list-style-type: none"> <li>Manage Nvidia CUDA, OpenACC, OpenCL plus MPI hybrid apps; Optimize scheduling with resource-mapped GPUs; Manage GPU apps within or without Docker containers; Obtain visibility with CUDA-specific metrics for GPU monitors and reports; Extend on-premise deployments to incorporate cloud-based GPU instances</li> </ul>	Multi-GPU Multi-Node
<b>Vampir</b>	TU Dresden	Vampir provides an easy-to-use framework that enables developers to quickly display and analyze arbitrary program behavior at any level of detail. The tool suite implements optimized event analysis algorithms and customizable displays that enable fast and interactive rendering of very complex performance monitoring data.	<ul style="list-style-type: none"> <li>CUPTI</li> </ul>	Multi-GPU Multi-Node

