



# GPU-ACCELERATED APPLICATIONS



# GPU-ACCELERATED APPLICATIONS

Accelerated computing has revolutionized a broad range of industries with over six 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
- 13 Public Sector
- 14 Design for Manufacturing/Construction: CAD/CAE/CAM
  - CFD (MFG)
  - CFD (RESEARCH DEVELOPMENTS)
  - COMPUTATIONAL STRUCTURAL MECHANICS
  - DESIGN AND VISUALIZATION
  - ELECTRONIC DESIGN AUTOMATION
  - INDUSTRIAL INSPECTION
- 29 Media and Entertainment
  - ANIMATION, MODELING AND RENDERING
  - COLOR CORRECTION AND GRAIN MANAGEMENT
  - COMPOSITING, FINISHING AND EFFECTS
  - (VIDEO) EDITING
  - (IMAGE & PHOTO) EDITING
  - ENCODING AND DIGITAL DISTRIBUTION
  - ON-AIR GRAPHICS
  - ON-SET, REVIEW AND STEREO TOOLS
  - WEATHER GRAPHICS
- 44 Medical Imaging
- 47 Oil and Gas
- 48 Life Sciences
  - BIOINFORMATICS
  - MICROSCOPY
  - MOLECULAR DYNAMICS
  - QUANTUM CHEMISTRY
  - (MOLECULAR) VISUALIZATION AND DOCKING
- 62 Research: Higher Education and Supercomputing
  - NUMERICAL ANALYTICS
  - PHYSICS
  - SCIENTIFIC VISUALIZATION
- 68 Safety and Security
- 71 Tools and Management
- 79 Agriculture
- 79 Business Process Optimization

## 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)



# Computational Finance

APPLICATION NAME	COMPANYNAME	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</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>• Codes 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>	QuantAleas	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</li> <li>• 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.	<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>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>O-Quant options pricing</b>	O-Quant	Offering for risk management and complex options and derivatives pricing using GPUs.	<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>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</li> <li>• Python-based scripting environment</li> <li>• 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 analytics</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)</li> <li>Supports Windows and Linux operating systems</li> </ul>	Multi-GPU Single Node

## Climate, Weather and Ocean Modeling

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<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</li> <li>All features in the MCH branch used for operational weather forecasting</li> </ul>	Multi-GPU Multi-Node
<b>E3SM-EAM</b>	US DOE	Global atmospheric model used as component to E3SM global coupled climate model.	<ul style="list-style-type: none"> <li>Dynamics and most physics</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>GRAF</b>	IBM/TWC	New GPU-based global weather model based on MPAS from NCAR	<ul style="list-style-type: none"> <li>Full application</li> </ul>	Multi-GPU Multi-Node
<b>WRF AceCAST-WRF</b>	TempoQuest Inc.	WRF model from NCAR 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>ARW dynamics</li> <li>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	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Anaconda Distribution</b>	Anaconda	The open-source Anaconda Distribution is the easiest way to perform Python/R data science and machine learning on Linux, Windows, and Mac OS X. With over 11 million users worldwide, it is the industry standard for developing, testing, and training on a single machine.	<ul style="list-style-type: none"> <li>Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND</li> <li>Sorts algorithms from the CUB and Modern GPU libraries</li> <li>Includes Numba (JIT Python compiler) and Dask (Python scheduler)</li> <li>Includes single-line install of numerous DL frameworks such as Pytorch</li> </ul>	Multi-GPU Multi-Node
<b>AnswerRocket</b>	AnswerRocket	AnswerRocket leverages AI and machine learning techniques to automate the hard work of business analysis, empowering teams to generate business intelligence and advanced analysis in seconds.	<ul style="list-style-type: none"> <li>Pluggable machine learning models</li> <li>Ask Questions in Plain English</li> <li>Create Interactive Visualizations &amp; Dashboards</li> <li>Provides Augmented Analytics</li> <li>Supports a wide variety of data sources</li> </ul>	Multi-GPU Multi-Node
<b>ArgusSearch</b>	Planet AI	Deep Learning driven document search tool.	<ul style="list-style-type: none"> <li>Fast full text search engine</li> <li>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>BlazingSQL</b>	BlazingSQL	GPU-accelerated SQL Engine for analytics available on all major CSP and on-premise deployment.	<ul style="list-style-type: none"> <li>• Distributed SQL Query Engine</li> <li>• Supports petabyte scale applications</li> <li>• Supports traditional big data formats and data stores</li> </ul>	Multi-GPU Multi-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</li> <li>• 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	<p>Automated Machine Learning with Feature Extraction. Essentially BI for Machine Learning and AI, with accuracy very similar to Kaggle Experts.</p> <p>H2O Driverless AI is an artificial intelligence (AI) platform for automatic machine learning. Driverless AI automates some of the most difficult data science and machine learning workflows such as feature engineering, model validation, model tuning, model selection and model deployment. It aims to achieve highest predictive accuracy, comparable to expert data scientists, but in much shorter time thanks to end-to-end automation. Driverless AI also offers automatic visualizations and machine learning interpretability (MLI). Especially in regulated industries, model transparency and explanation are just as important as predictive performance. Modeling pipelines (feature engineering and models) are exported (in full fidelity, without approximations) both as Python modules and as pure Java standalone scoring artifacts.</p>	<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>GPUdb</b>	Kinetica	Multi-GPU, Multi-Machine distributed object store providing SQL style query capability, advanced geospatial query capability, heatmap generation, and distributed rasterization services.	<ul style="list-style-type: none"> <li>• Query against big data in real time</li> <li>• No pre-indexing allows for complex, ad-hoc query chains</li> <li>• Interactively explore large, streaming data sets</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>• Available algorithms include Gradient Boosting Machines (GBM's)</li> <li>• Generalized Linear Models (GLM's)</li> <li>• K-Means Clustering</li> <li>• SVD</li> <li>• PCA</li> <li>• K-means</li> <li>• XGBoost.</li> <li>• It can be used as a drop-in replacement for scikit-learn with support for GPUs on selected (and ever-growing) algorithms.</li> <li>• A new R API brings the benefits of GPU-accelerated machine learning to the R user community. The R package is a wrapper around the H2O4GPU Python package, and the interface follows standard R conventions for modeling.</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</li> <li>• JumpTo Technology, for data visualisation</li> <li>• E-Discovery</li> <li>• Extraction from phone calls</li> <li>• 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</li> <li>• Tesla K40 &amp; 12 GB on-board RAM</li> <li>• Scales up with multiple GPUs</li> <li>• Keeps close to 100 GB of compressed data in GPU memory on a single server system</li> <li>• Fast analysis, reporting, and planning</li> </ul>	Multi-GPU Single Node
<b>Labellio</b>	KYOCERA Communication Systems Co	The world's easiest deep learning web service for computer vision, allowing everyone to build own image classifier with only web browser.	<ul style="list-style-type: none"> <li>• Neural net fine-tuning for image data</li> <li>• Data crawling and data browsing</li> <li>• Drag-and-drop style data cleansing backed by AI support</li> </ul>	Multi-GPU Single Node
<b>Numba</b>	Anaconda	<p>Numba is a compiler for Python array and numerical functions that gives you the power to speed up your applications with high performance functions written directly in Python.</p> <p>Numba generates optimized machine code from pure Python code using the LLVM compiler infrastructure. With a few simple annotations, array-oriented and math-heavy Python code can be just-in-time optimized to performance similar as C, C++ and Fortran, without having to switch languages or Python interpreters.</p>	<ul style="list-style-type: none"> <li>• On-the-fly code generation (at import time or runtime, at the user's preference)</li> <li>• Native code generation for the CPU (default) and GPU hardware</li> <li>• Integration with the Python scientific software stack (enabled via Numpy)</li> <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</li> <li>• OpenGL- (EGL) based rendering</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</li> <li>• Data Mining</li> <li>• Machine Learning</li> <li>• Predictive Analytics</li> </ul>	Multi-GPU Multi-Node
<b>Sqream DB</b>	SQream Technologies	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</li> <li>• Pricing of vanilla and exotic options</li> <li>• Fixed income analytics</li> <li>• Data mining</li> </ul>	Multi-GPU Single Node

## Artificial Intelligence

### DEEP LEARNING AND MACHINE LEARNING

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AiFi Nano</b>	AiFi Inc.	Cashier-free (like Amazon grab and go solution) and stock out retail software	<ul style="list-style-type: none"> <li>• cuDNN</li> <li>• TensorRT</li> <li>• DeepStream</li> </ul>	Multi-GPU Single Node
<b>AI Image Labeling</b>	Frenzy	Builds robust self-labeling training datasets for classifying exact objects and products in visual scenes at a fraction of the time and cost	<ul style="list-style-type: none"> <li>• GPU in the cloud</li> </ul>	Multi-GPU Single Node
<b>AI Lifecycle</b>	Clarifai	Clarifai brings a new level of understanding to visual content through deep learning technologies. Uses GPUs to train large neural networks to solve practical problems in advertising, media, and search across a wide variety of industries such as automated tagging, visual search, and recommendation engine, predictive maintenance, demographic analysis and more.	<ul style="list-style-type: none"> <li>• GPU-based training and inference</li> <li>• Recognizes and indexes images with predefined classifiers or custom classifiers</li> </ul>	Multi-GPU Single Node
<b>Allganize NLU APIs for Enterprises</b>	Allganize, Inc.	Natural Language Understanding APIs for enterprise: Answer-bot based on documents with unstructured data (text + table), e.g., manuals, instructions, FAQ documents; Review analysis; sentiment analysis, summarizing etc. Provided as APIs.	<ul style="list-style-type: none"> <li>• Training and inferencing using V100</li> </ul>	Multi-GPU Multi-Node
<b>AlphaSense</b>	AlphaSense, Inc.	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 questing 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 questing answering format ("the Google for Analyst research")</li> </ul>	Multi-GPU Single Node
<b>AlwaysAI</b>	Always AI	Easy-to-use platform to build and deploy computer vision applications for embedded devices at the edge. Apply for an early access on the product link	<ul style="list-style-type: none"> <li>• Jetson Nano</li> </ul>	Single GPU Single Node

<b>Anaconda Enterprise</b>	Anaconda	Anaconda Enterprise combines core AI technologies, governance, and cloud-native architecture. Each piece - core AI, governance, and cloud native - are critical components to enabling organizations to automate AI at speed and scale.	<ul style="list-style-type: none"> <li>• Access 1,500+ secure Python and R data science packages and libraries from Anaconda</li> <li>• Curate a private package repository controlled by IT.</li> <li>• Craft package policies by blacklisting and whitelisting license types and versions.</li> <li>• Leverage code and GPU-specific Conda packages designed to accelerate computation and train models.</li> <li>• Share centralized GPU clusters across teams, using custom resource profiles to establish resource limits.</li> <li>• Create Anaconda installers with custom sets of packages for Windows, Mac, and Linux.</li> <li>• Easily distribute your own proprietary packages to share code, algorithms, and models.</li> <li>• AE5.3 v2.19</li> <li>• Use Open Source Software Securely</li> <li>• Quickly and easily share notebooks with others, using your preferred IDE.</li> <li>• Grant or restrict access to individual notebooks by user or group.</li> <li>• Benefit from automated version control in data science projects.</li> <li>• Connect to Hadoop/Spark clusters and other data sources for distributed workloads.</li> <li>• Create custom resource profiles by role to efficiently allocate resources across teams.</li> </ul>	Multi-GPU Single Node
<b>Antuit Demand Planning and Forecasting</b>	Antuit	Extracts maximum predictability from the available data. Proprietary “Dynamic Aggregation” logic with attribute-based disaggregation generates forecasts for all products, including new, slow-moving, and end-of-life. Spark and GPU clusters, along with optimized AI algorithms, provide scaling for the largest retailers. Incorporates all available demand drivers, such as price elasticities, promotional lifts, weather, and hyper-local event data.	<ul style="list-style-type: none"> <li>• CUDA 10.1</li> <li>• CuDNN 7.6</li> <li>• CuBLAS 10.2</li> </ul>	Multi-GPU Multi-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>• Extremely easy to add new algorithms</li> <li>• Distributed instead of single machine</li> </ul>	Multi-GPU Multi-Node
<b>Artificial Intelligence Radio Transceiver (AIR-T)</b>	Deepwave Digital	The Artificial Intelligence Radio Transceiver (AIR-T) is software defined radio designed and developed for RF deep learning applications. The app is equipped with three signal processors including a 256 core NVIDIA Jetson TX2, a field programmable gate array (FPGA), and dual embedded CPUs.	<ul style="list-style-type: none"> <li>• The AIR-T is designed to be an edge-compute inference engine for deep learning algorithms.</li> </ul>	N/A
<b>ARYA.ai</b>	ARYA.ai	Deep learning platform with end-to-end workflows for Enterprise, incorporating TensorFlow. Focuses on consumer banking and insurance industries.	<ul style="list-style-type: none"> <li>• Deep learning</li> <li>• TensorFlow.</li> </ul>	Multi-GPU Multi-Node



<b>Aura Vision</b>	Aura Vision	Capture unique insights from every visitor, using your existing cameras	<ul style="list-style-type: none"> <li>• Segmented footfall</li> <li>• Shopper motivation</li> <li>• Product engagement</li> <li>• Window display ROI</li> <li>• Store utilization</li> <li>• Service wait times</li> </ul>	Single GPU Single Node
<b>Avitas Systems - Inspection as a Service</b>	Avitas Systems	Avitas Systems configures various multi rotor and helicopter drones with multiple sensor kits including RGB cameras, laser sensors, infrared and others collecting inspection data to meet different customer use cases. Ingests inspection data where an AI back-end turns the raw data into inspection findings such as corrosion levels, damaged/missing parts, encroaching vegetation volumes.	<ul style="list-style-type: none"> <li>• Drone based data capture</li> <li>• RGB Camera, Laser and Infrared sensing</li> <li>• Deep learning driven Object detection for Inspection</li> <li>• Detect corrosion levels, damaged/missing parts, encroaching vegetation volumes.</li> <li>• AI workbench</li> <li>• Photogrammetry</li> </ul>	Multi-GPU Multi-Node
<b>AWM Smart Shelf</b>	Adroit Worldwide Media, Inc.	Application for Automated Inventory Intelligence (view and track virtually in a retail environment), Content Management System (manage inventory, prices and content), Led Display (prices, promotions and advertisements at the click of a button) and Product Mapper (automate creation of planograms and auditing process)	<ul style="list-style-type: none"> <li>• kubernetes</li> <li>• Docker</li> <li>• RTX 2080</li> </ul>	Multi-GPU Single Node
<b>Badger Insights</b>	Badger Technologies	Badger Technologies provides data and analytics for retail operations through automation solutions that include a fully autonomous robot to address out-of-stock, planogram compliance, and price integrity	<ul style="list-style-type: none"> <li>• GPU accelerated</li> </ul>	Single GPU Single Node
<b>BIDMach -</b>	UC Berkeley	The fastest machine learning library available. Holds the record for many common machine learning algorithms.	<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</li> <li>• Novel programming language called Inklings</li> <li>• Primary focus on reinforcement learning</li> </ul>	Multi-GPU Single Node
<b>Brain Frame</b>	Aotu	BrainFrame platform provides Out-Of-The-Box Smart Vision Applications for multiple verticals. The drag-and-drop VisionCapsules system allows you to pick from a wide selection of custom algorithms to extract exactly the information you want	<ul style="list-style-type: none"> <li>• Jetpack</li> <li>• Jetson</li> </ul>	Single GPU Single Node
<b>Caffe2</b>	Facebook	This is a faster framework for deep learning, it's forked from BVLC/caffe (master branch). Allows data-parallel via MPI.	<ul style="list-style-type: none"> <li>• GPU cluster processing</li> <li>• Mass image data</li> </ul>	Multi-GPU Single Node
<b>Cartwatch Checkout</b>	Signatrix	Protect the checkout area and reduce the workload of your checkout staff	<ul style="list-style-type: none"> <li>• Real-time alerts on theft (mis-scan) at the checkout lanes</li> <li>• Featuring Jetpack and TensorRT</li> </ul>	Single 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</li> <li>• Multi-GPU</li> <li>• 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>CNTK</b>	Microsoft Corp.	Microsoft 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>• Speech Recognition</li> <li>• Machine Translation</li> <li>• Image Recognition</li> <li>• Image Captioning</li> <li>• Text Processing and Relevance</li> <li>• Language Understanding</li> <li>• Language Modeling</li> </ul>	Multi-GPU Single Node
<b>ConundrumAI</b>	Conundrum Industrial Limited	Conundrum, a UK-based company, develops AI solutions for predictive maintenance and optimization of industrial processes.	<ul style="list-style-type: none"> <li>• Automated deep learning significantly speeds up a build of the applications based on DL models;</li> <li>• Transfer Learning enables to boost the performance of the applications by transferring knowledge between them;</li> <li>• Data based digital twins and reinforcement learning for optimization.</li> </ul>	Multi-GPU Single Node
<b>Darwin</b>	SparkCognition	Darwin is a machine learning product that accelerates data science at scale by automating the building and deployment of models. Based on a proprietary neuro-evolutionary algorithm, Darwin uses a combination of ML methods and genetic algorithms, to arrive at a new generation of designs.	<ul style="list-style-type: none"> <li>• Unique neuro-evolutionary algorithm on GPU</li> <li>• Automated ML for model building on GPUs</li> <li>• GPU accelerated PyTorch</li> </ul>	Multi-GPU Single Node
<b>Databricks Unified Analytics Platform</b>	Databricks	Databricks provides a cloud-based platform designed to make big data and machine learning simple.	<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>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>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>Dessa</b>	Dessa	Deep Learning Platform based on TensorFlow. Allows end-to-end workflows. Targets 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>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>Discourse.ai training automation for chatbot</b>	Discourse.ai	NLU model training/re-training/fine-tuning for contact center operation automation trained from raw transcripts to identify the intentions automatically, complemented by human annotation. Models are used for post-call analysis, chatbot design etc.	<ul style="list-style-type: none"> <li>• V100</li> <li>• P100</li> <li>• T4 GPUs</li> <li>• cuDNN</li> </ul>	Multi-GPU Multi-Node
<b>Dr. Retail</b>	SkyREC Inc.	Instore data analytics	<ul style="list-style-type: none"> <li>• TensorRT 5.1</li> <li>• nvJPEG</li> <li>• NVEnc</li> <li>• NVDec</li> </ul>	Single GPU Single Node
<b>Frenzy Enterprise Solutions</b>	Frenzy	Frenzy Enterprise Solutions provides retailers and brands with the tools to provide customer's the best experience and more purchasing opportunities including Similar Product Recommendations, Inventory Tagging, Camera Search, Complimentary Product Recommendations, How To Wear It, Influencer Matching	<ul style="list-style-type: none"> <li>• GPU on the cloud</li> </ul>	Multi-GPU Single Node
<b>G3C.AI</b>	Graymatics	Retail in store analytics solutions through Deep CCTV Streaming Analytics	<ul style="list-style-type: none"> <li>• In store analytics: heat-maps, shopper tracking, dwell time, people counting, mood detection, demographics</li> <li>• Featuring TensorRT and Deepstream</li> </ul>	Multi-GPU Single Node
<b>Gridspace</b>	Gridspace	Voice analytics to turn streaming speech audio into useful data and service metrics. Instrumental to contact call center and work communications 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>Insights</b>	AnyVision	Insight delivers in-store analytics with features such as: heavy shoppers, gaze estimation, heatmaps, customer journey, and offline to online	<ul style="list-style-type: none"> <li>• NVIDIA Tesla T4 and Jetson</li> </ul>	Multi-GPU Single Node
<b>Keras</b>	Open Source	Keras is a minimalist, highly modular neural networks library, written in Python. Capable of running on top of either TensorFlow or Theano and 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>	Mathworks	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</li> <li>• Simple CNN wrapper</li> <li>• DagNN wrapper</li> <li>• 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.	<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</li> <li>• Creates and trains classifiers with tooling for uploading and managing datasets</li> </ul>	Multi-GPU Single Node

<b>Mobiliya ThirdEye</b>	Mobiliya	Artificial Intelligence powered solution to automate security and surveillance for your building, parking premise, and retail. Complements and boosts your existing CCTV and/or IP Camera infrastructure.  Supported functionally: object identification, facial recognition, product inspection	<ul style="list-style-type: none"> <li>• CUDA/cuDNN</li> <li>• TensorRT</li> <li>• Deespstream</li> <li>• Jetpack</li> </ul>	Single 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>	Intel	Neon is a fast, scalable, easy-to-use Python based deep learning framework that has been optimized down to the assembler level. 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</li> <li>• Processes 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>out of stock detection</b>	Focal Systems	Deep Learning Computer Vision track your On-Shelf Availability throughout your entire store 100+ times a day	<ul style="list-style-type: none"> <li>• On-Shelf Availability Analytics per hour</li> <li>• Real-time Alerts on your “never be outs”</li> </ul>	Multi-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>Preciate</b>	Preciate	Face recognition capabilities providing a 360 Omni channel view of your customers to identify shoppers who are registered to a loyal program and pull info about their purchasing behavior to personalize service	<ul style="list-style-type: none"> <li>• Real time recognition and staff alert</li> </ul>	Single GPU Single Node
<b>Protects &amp; Insights</b>	Briefcam	Transform video into actionable intelligence. features: video synopsis and real time alerts, loss prevention, customer engagement and tying info to POS data, heatmaps, shopper tracking	<ul style="list-style-type: none"> <li>• NVIDIA Tesla and Jetson.</li> <li>• TesnorRT</li> </ul>	Multi-GPU Single Node
<b>Retail Analytics</b>	Pilot AI Labs	Retail in-store analytics for stock out (cameras in shelves), demographics (age/gender), shopper tracking/counting, anomaly detection, drive through solutions and more	<ul style="list-style-type: none"> <li>• Jetpack</li> <li>• Jetson Tx2</li> <li>• RTX 2080</li> </ul>	Single GPU Single Node

<b>samadii/dem</b>	Metariver Technology	Software for computing various behaviors of massive solid particles of various size particles from small particle with Brownian motion to large particle such as ore with DEM(Discrete Element Method).	<ul style="list-style-type: none"> <li>• Solid particle simulator, DEM solver</li> <li>• Multi-Physics module(Drag and Buoyancy force, Magnetic force, Coulomb force, adhesion force, Van der Waals force, Brownian motion and heat effect)</li> <li>• VPS(Virtual Particle System), Cluster model</li> <li>• Co-simulation with MBD(Multi Body Dynamics) solvers (ADAMS, DADS, RecurDyn, Daful)</li> <li>• Co-simulation with ANSYS Mechanical (Flexible body).</li> </ul>	Multi-GPU Multi-Node
<b>SAS</b>	SAS	SAS Machine Learning. SAS Viya Visual Data Mining and Visualization suites now leverage GPU deep learning	<ul style="list-style-type: none"> <li>• Volta V100 with tensor cores</li> <li>• TensorRT for inference on the NVIDIA Jetson TX2 box</li> <li>• RNN</li> <li>• Multiple GPUs on a single SMP node</li> <li>• Homogeneous and heterogeneous MPP with synchronized Stochastic Gradient Descent</li> </ul>	Multi-GPU Multi-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</li> <li>• Studio.ml is a new project designed to make AI development easier by hiding most of the complexity</li> <li>• Studio.ml runs on-premise and in the cloud</li> </ul>	Single GPU Single Node
<b>SENTINEL-IQ</b>	FaceFirst	Face recognition surveillance platform for real time face recognition, loss prevention, dwell time, conversation, unique visitor and others	<ul style="list-style-type: none"> <li>• NVIDIA Tesla &amp; Quadro</li> </ul>	Multi-GPU Single Node
<b>Shopic Frictionless Shopping</b>	Shopic	Frictionless Shopping - using smart cart	<ul style="list-style-type: none"> <li>• NVIDIA Xavier NX</li> </ul>	Single GPU Single Node
<b>SmartCart</b>	Imagr	SmartCart comprised of four tiny cameras and AI vision recognition system	<ul style="list-style-type: none"> <li>• NVIDIA Jetson, Xaiver</li> <li>• TensorRT</li> </ul>	Single GPU Single Node
<b>Smart Skin</b>	Human engine	AI-enhanced processing of 3D and 4D data. Used to create high quality 3D characters for interactive media (games, mobile apps, VFX, VR/AR and mixed reality experiences, etc) <ul style="list-style-type: none"> <li>- automatic retopology of 3D and 4D data using machine learning</li> <li>- photogrammetry : noise-reduction and hole-patching using machine learning</li> <li>- realistic lip-sync using 4D-trained neural network</li> </ul>	<ul style="list-style-type: none"> <li>• CUDA</li> <li>• Hairworks</li> <li>• PhysX</li> <li>• cuDNN</li> <li>• OptiX</li> </ul>	Multi-GPU Multi-Node
<b>SpaceKnow PaaS</b>	SPACEKNOW	PaaS for deep learning extraction of satellite data information targeted at Financial Services and Public Sector. Tracks 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>• Provides 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>The Deep North Video Analytics platform</b>	Deep North	The Deep North platform includes Occupancy Management, Gesture Analysis, Zone Management, Vehicle Analysis, Dashboard and reporting	<ul style="list-style-type: none"> <li>TensorRT</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>TrigoVison</b>	TrigoVision	Retail automation platform that provides seamless checkout, shoplifting prevention, and real-time inventory updates.	<ul style="list-style-type: none"> <li>TensorRT</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>Veesion</b>	Veesion	Shoplifting detection using deep learning algorithm that continuously analyses the content of security cameras. It automatically detects gestures associated with shoplifting in real-time. Sends a video alert to a human operator who confirms the theft and takes action.	<ul style="list-style-type: none"> <li>Real-time shoplifting prospects alerts</li> </ul>	Multi-GPU Single Node
<b>ViMo</b>	Motionloft	Video analytics using Tx1/Tx2, people counting, queue management, bounce rate, gender, age, heatmap and path tracing.	<ul style="list-style-type: none"> <li>Jetpack</li> <li>TensorRT</li> </ul>	Single 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: <ul style="list-style-type: none"> <li>&gt; Curate and organize visual content</li> <li>&gt; Search and recommend visually</li> <li>&gt; Get insights and analytics visually</li> </ul> </li> </ul>	Single GPU Single Node
<b>Voca's Virtual Agent</b>	voca.ai	Human like cell center conversation AI	<ul style="list-style-type: none"> <li>Jasper</li> <li>NeMo</li> </ul>	Multi-GPU Multi-Node
<b>vuForecast</b>	deepVu	ML/DL enabled vuForecast learns from historical inventory, point of sale, promotions and logistics data augmented with DeepVu's real-time data platform aggregating numerous external micro and macro economic signals to accurately forecast future demand	<ul style="list-style-type: none"> <li>ML (dmlc/XGBoost) + Dask for distributed training</li> <li>DL (RNN/LSTM networks) + PyTorch 1.1</li> <li>DL (RL) + TensorFlow 1.14 and 2.0</li> </ul>	Multi-GPU Single Node
<b>Walkout</b>	walkout	Autonomous check out - smart cart	<ul style="list-style-type: none"> <li>NVIDIA Jetson Tx2</li> </ul>	Single GPU Single Node
<b>Zippin</b>	Zippin	Checkout-free technology offering inventory tracking and insights to ensure the right products are in the right place, at the right time.	<ul style="list-style-type: none"> <li>Jetpack</li> </ul>	Multi-GPU Single Node



# Public Sector

APPLICATION NAME	COMPANYNAME	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. Transforms the elevation surface into a geocentric 3D coordinate system and runs 3D sightlines to each transformed cell center. Takes advantage of Tensor Cores for both training and inference .	<ul style="list-style-type: none"> <li>Viewshed2</li> <li>Deep Learning</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 Multi-Node
<b>Blaze Terra</b>	Eternix	Geospatial visualization tool	<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>Deep Learning training</li> <li>Deep learning inferencing</li> <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>ERDAS Imagine</b>	Hexagon Geospatial	Remote sensing, photogrammetry and GIS toolset for the interactive, semi-automated and automated extraction of information from remotely sensed imagery and point clouds.	<ul style="list-style-type: none"> <li>Gray Level co-occurrence matrix (CLCM) image processing operation</li> <li>NNDiffuse image pan sharpening operation</li> <li>Deep learning capabilities using the GPU accelerated versions of Tensorflow</li> </ul>	Single GPU Single Node
<b>Geomatics GXL</b>	PCI	Image processing	<ul style="list-style-type: none"> <li>Image orthorectification</li> <li>Additional image processing</li> </ul>	Multi-GPU Single Node
<b>GeoWeb3d Desktop</b>	Geoweb3d	Geospatial visualization of 3D and 2D data, mensuration and 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</li> <li>Visual pivoting</li> <li>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.	<ul style="list-style-type: none"> <li>Real-time super-resolution-based video enhancement on live streams</li> <li>Geospatial visualization</li> <li>Target detection and tracking</li> <li>Fast 2-D mapping</li> </ul>	Multi-GPU Single Node
<b>LuciadLightspeed</b>	Hexagon Geospatial	Geospatial visualization and analysis	<ul style="list-style-type: none"> <li>GPU accelerated line of sight and view shed calculations</li> <li>GPU accelerated hypsometry calculations, including terrain slope, ridge and valley detection, terrain orientation and azimuth calculations</li> <li>GPU accelerated imaging operator for geospatially referenced imagery</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>• Operates in a real-time streaming fashion</li> <li>• Ingests radio samples from many common radio interfaces</li> <li>• Make use of packet formats like VITA49 or SDDS.</li> <li>• Can be used from any device with a browser, including mobile handsets</li> <li>• 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. Extracts 3-D objects and 3_D dense point clouds from stereo images. Also extracts 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</li> <li>• Building texture generation</li> </ul>	Multi-GPU Single Node

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

### CFD (MFG)

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>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 Multi-Node
<b>ADS Flow Solver - Code LEO</b>	ADSCFD, Inc.	A Compressible, explicit time-marching CFD solver for aerospace applications. Capable of handling both internal and external flows with robustness and accuracy	<ul style="list-style-type: none"> <li>• Unstructured/Structured Meshes</li> <li>• Multigrid Accelerations</li> <li>• Multiple Turbulence Models</li> <li>• Rotor-stator Interfaces</li> </ul>	Multi-GPU Multi-Node
<b>Altair AcuSolve</b>	Altair	Computational Fluid Dynamics (CFD) tool, providing users with a full range of physical models. Simulations involving flow, heat transfer, turbulence, and non-Newtonian materials are handled with ease by AcuSolve's robust and scalable solver technology.	<ul style="list-style-type: none"> <li>• Linear solvers for flow, temperature, turbulence model, and mesh movement equations</li> </ul>	Single GPU Single Node

<b>Altair nanoFluidX</b>	Altair	State-of-the-art particle-based (SPH) fluid dynamics code for simulation of single and multiphase flows in complex geometries with complex motion.	<ul style="list-style-type: none"> <li>• Extremely fast</li> <li>• Single and Multiphase Flows</li> <li>• Arbitrary motion definition</li> <li>• Time-dependent acceleration</li> <li>• Inlets/outlets</li> <li>• Surface tension and adhesion</li> <li>• Steady-state thermal solutions through coupling</li> </ul>	Multi-GPU Multi-Node
<b>Altair ultraFluidX</b>	Altair	Simulation tool for ultra-fast prediction of the aerodynamic properties of passenger and heavy-duty vehicles as well as for the evaluation of building and environmental aerodynamics.	<ul style="list-style-type: none"> <li>• CUDA-accelerated high-fidelity flow field computations based on the Lattice Boltzmann method</li> <li>• CUDA-aware MPI support for multi-GPU and multi-node usage</li> <li>• Efficient implementation of tailor-made automotive features, including rotating wheels, belt systems, boundary layer suction and porous media support</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	Modeling software for simulating Fluidized Reactors	<ul style="list-style-type: none"> <li>• Linear equation solver for isothermal, non-reacting simulations and for thermal reacting cases</li> <li>• Discrete multi-component particle calculations</li> </ul>	Single GPU Single Node
<b>DYVERSO</b>	Next Limit	Multi-physics simulation engine for liquids and granular substances. Can be used to mimic behavior of rigid and soft bodies	<ul style="list-style-type: none"> <li>• Fluid solver in Real Flow 10.5 based on Smoothed particle hydrodynamics (SPH)</li> <li>• Fluid solver in Real Flow 10.5 based on Position based dynamics (PBD)</li> </ul>	Single GPU Single Node
<b>Fine/Open</b>	Numeca International	FINE/Open with OpenLabs is a powerful CFD Flow Integrated Environment dedicated to complex internal and external flows. It allows users to freely develop and exchange physical models in CFD, with a new open approach to CFD. Complex programming tasks are avoided through the usage of an easy meta-language.	<ul style="list-style-type: none"> <li>• Incompressible, low and high speed flows</li> <li>• Efficient preconditioned compressible solver with fast agglomerated multigrid acceleration and adaptation techniques to combine completely unstructured hexahedral grids</li> </ul>	Multi-GPU Multi-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 Multi-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.	<ul style="list-style-type: none"> <li>• CUDA</li> <li>• 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, with code 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 up to 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>JSCAST</b>	Qualica Inc.	Integrated CAE product for studying and predicting the casting process. Includes high precision mold filling and solidification solvers.	<ul style="list-style-type: none"> <li>• Solvers for mold filling and solidification</li> <li>• Rendering</li> </ul>	Single 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 for simulating physical, chemical, and biological processes	<ul style="list-style-type: none"> <li>• Flexible Mesh (FM) engines use GPUs.</li> <li>• Hydrodynamic and turbulence calculations</li> </ul>	Multi-GPU Single Node
<b>MIKE 3</b>	DHI	3D Modeling of Coast and Sea	<ul style="list-style-type: none"> <li>• Hydrodynamic part of the flexible mesh engines (MIKE 3 HD FM).</li> </ul>	Multi-GPU Multi-Node
<b>MIKE FLOOD</b>	DHI	1D & 2D urban, coastal, and riverine flood modelling	<ul style="list-style-type: none"> <li>• Hydrodynamics</li> <li>• 2D Overland flow</li> <li>• Coupling of 1D and 2D models for complex flooding issues</li> </ul>	Multi-GPU Single Node
<b>Numerix</b>	Zeus	Custom software development in the areas of CFD, FEA and Electromagnetics	<ul style="list-style-type: none"> <li>• Lattice Boltzmann Method (LBM) for flow around buildings</li> <li>• SPH based flow solver for simulating flow over urban environments</li> </ul>	Multi-GPU Single Node
<b>Pacefish</b>	Numeric Systems GmbH	CFD application for Automotive Aerodynamics, Pedestrian Comfort and Wind Loading	<ul style="list-style-type: none"> <li>• Transient Lattice-Boltzmann Method for single-phase flows</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> <li>• Scalable up to 16 GPUs</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>• Explicit and Implicit methods</li> </ul>	Multi-GPU Multi-Node
<b>PowerViz</b>	Dassault Systèmes SIMULIA Corp.	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 Digital Industries Software	A unified, scalable, open and extensible environment for 3D CAE with connections to design, 1D simulation, test, and data management.	<ul style="list-style-type: none"> <li>• Rendering</li> <li>• Raytracing</li> </ul>	Multi-GPU Single Node
<b>Simcenter STAR-CCM+</b>	Siemens Digital Industries Software	Integrated solution for CFD-focused Multiphysics simulation	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single 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>Turbostream</b>	Turbostream Ltd.	CFD software for turbomachinery flows	<ul style="list-style-type: none"> <li>• Finite Volume explicit solver for RANS/URANS calculations</li> <li>• Variable time-steps and multigrid for convergence acceleration</li> </ul>	Multi-GPU Multi-Node
<b>zCFD</b>	Zenotech Simulation Unlimited	General purpose CFD solver	<ul style="list-style-type: none"> <li>• Turbulent flow (RANS, URANS, DDES or LES) including automatic scalable wall functions</li> </ul>	Multi-GPU Single Node

## CFD (RESEARCH DEVELOPMENTS)

APPLICATION NAME	COMPANYNAME	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</li> <li>• Compressible Flows</li> <li>• Non-linear Solid Mechanics</li> <li>• Species transport equations</li> <li>• Excitable Media</li> <li>• Thermal Flows</li> <li>• 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>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>Project Chrono</b>	University of Wisconsin-Madison	Chrono is a physics-based modelling and simulation infrastructure based on a platform-independent open-source design implemented in C++. Systems can be made of rigid and flexible/compliant parts with constraints, motors and contacts; parts can have three-dimensional shapes for collision detection	<ul style="list-style-type: none"> <li>• Robotics</li> <li>• Wheeled vehicle dynamics</li> <li>• Tracked vehicle dynamics</li> <li>• Nonlinear finite element analysis</li> <li>• Mechatronics</li> <li>• Off-road vehicle mobility</li> <li>• Terramechanics</li> <li>• Virtual reality</li> <li>• Granular flows</li> <li>• Collision detection</li> <li>• Autonomous vehicles</li> <li>• Seismic engineering</li> <li>• Augmented reality</li> </ul>	Multi-GPU Multi-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

## COMPUTATIONAL STRUCTURAL MECHANICS

APPLICATION NAME	COMPANYNAME	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>Altair EDEM</b>	Altair	Software for bulk material simulation that uses the Discrete Element Modeling (DEM) technology to simulate and analyze behavior of bulk materials	<ul style="list-style-type: none"> <li>• EDEM Simulator, a DEM solver</li> <li>• Integration with Ansys and Abaqus for FEA for bulk material simulation</li> <li>• Integration with Adams, Siemens and RecurDyn for Multi-body Dynamics</li> <li>• Integration with Ansys Fluent for Particle-Fluid Systems</li> </ul>	Multi-GPU Single Node
<b>Altair HyperWorks</b>	Altair	Comprehensive, open architecture CAE simulation suite in the industry, offering the best technologies to design and optimize high performance, weight efficient and innovative products. It includes a full set of modeling and visualization tools.	<ul style="list-style-type: none"> <li>• OpenGL v3.2</li> <li>• OpenCL v2.0 support</li> <li>• Anti-aliasing</li> </ul>	Single GPU Single Node

<b>Altair OptiStruct</b>	Altair	Industry proven, modern structural analysis solver for linear and nonlinear problems under static and dynamic loadings. It is also the market-leading solution for structural design and optimization.	<ul style="list-style-type: none"> <li>• Direct solver (BCS)</li> <li>• Eigenvalue solvers (AMSES and Lanczos)</li> <li>• Iterative solver (PCG)</li> </ul>	Single GPU Single Node
<b>Amphyon</b>	AdditiveWorks	Simulation-based process software for powder bed based, laser beam melting additive manufacturing processes	<ul style="list-style-type: none"> <li>• Mechanical Process Simulation</li> <li>• Thermal Process Simulation</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>Autodesk Nastran</b>	Autodesk	Autodesk Nastran FEA software analyzes linear and nonlinear stress, dynamics, and heat transfer characteristics of structures and mechanical components.	<ul style="list-style-type: none"> <li>• Double Precision on GPU</li> </ul>	Multi-GPU Multi-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>Helixx 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>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</li> <li>• Coupled hydraulic, mechanical, transport, thermal and fracture processes</li> </ul>	Single 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>MatDEM</b>	Nanjing University	MatDEM is a software for Fast GPU Matrix computing of Discrete Element Method. The software implements automatic stacking modeling, layered material, joint surface and load settings, rich post-processing functions and secondary development.	<ul style="list-style-type: none"> <li>• Full product support on GPU</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	Multidisciplinary structural analysis application used to perform static, dynamic, and thermal analysis across linear and nonlinear domains	<ul style="list-style-type: none"> <li>• Direct sparse solver</li> </ul>	Multi-GPU Single Node
<b>PERMAS-XPU</b>	INTES GmbH	General purpose structural simulation software	<ul style="list-style-type: none"> <li>• Linear Equation Solver</li> </ul>	Single 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 for simulating behavior of bulk materials with complex particle shapes and size distributions	<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>Simcenter Nastran</b>	Siemens Digital Industries Software	Finite element method (FEM) solver for computational performance, accuracy, reliability and scalability	<ul style="list-style-type: none"> <li>Linear and nonlinear equation solver</li> <li>Frequency response module</li> <li>Matrix decomposition computations</li> </ul>	Multi-GPU Multi-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
<b>ThreeParticle/CAE</b>	BECKER 3D GmbH	Multiphysics Discrete Element Method (DEM) simulation platform for bulk materials with complex shapes and built-in multi-body dynamics (MBD), Finite Element Analysis (FEA) & Smoothed Particle Hydrodynamics (SPH)	<ul style="list-style-type: none"> <li>GPU accelerated Smoothed Particle Hydrodynamics</li> <li>Simulate complex and real particle shapes using DEM combined with SPH, FEA, MBD, Wear</li> </ul>	Single GPU Single Node

## DESIGN AND VISUALIZATION

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3D CAT.live</b>	Shenzhen Rayvision Technology Co Ltd	Real-time rendering cloud service for 3D applications. The massive GPU computing power in the cloud is used to process heavy image rendering calculations and stream output to the terminal device synchronously, thereby realizing light weight of the terminal device and making high-quality 3D graphics applications ubiquitous. Users can use any common networked device to access the 3D application hosted in the 3DCAT cloud without downloading and installing the application. Supports almost all rendering engines that can run on the Windows platform, and supports the opening of NVIDIA RTX real-time ray tracing function.	<ul style="list-style-type: none"> <li>Cloud XR SDK</li> <li>DLSS (potential)</li> </ul>	Multi-GPU Multi-Node
<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.	<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>Additive Mfg Toolkit</b>	Dyndrite	Dyndrite has developed a GPU-based geometry kernel with CUDA. The initial application for this kernel is an Additive Manufacturing Toolkit which speeds up the process of 3D printing, especially for complex parts.	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	N/A
<b>ALLPLAN</b>	Nemetschek ALLPLAN	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction.	<ul style="list-style-type: none"> <li>• OpenGL 4, and now moving to Vulkan</li> <li>• Vulkan for wireframe rendering already with plan to ship full integration with Version 2022 in September 2021</li> </ul>	Single GPU Single Node
<b>ANSA</b>	BETA CAE Systems	Multidisciplinary CAE pre-processing tool for full model build up, from CAD data to ready-to-run solver input file, in a single integrated environment	<ul style="list-style-type: none"> <li>• OpenGL</li> <li>• OpenCL</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</li> <li>• CUDA-based Structural Stress, Modal, Fluid Dynamics, Thermal, Electrical Conduction and Coupled Multi-Physics simulations</li> </ul>	Single GPU Single Node
<b>Ansys SPEOS</b>	ANSYS	Physically accurate optical simulation software dedicated to predictive illumination and optical performance of systems. High-fidelity visualization of the final result, based on unique human vision algorithm.	<ul style="list-style-type: none"> <li>• SPEOS Live Preview</li> <li>• 360 degrees for immersive or observer view</li> <li>• Optical part design</li> <li>• Optical sensors test</li> <li>• HUD design and analysis</li> <li>• Infrared modeling</li> </ul>	Single GPU Single Node
<b>Ansys VRXPERIENCE for HMI and Perceived Quality</b>	ANSYS	Predictive physics-based real time lighting simulation with VR capabilities to experience and validate the impact of your design proposition on appearance and perceived quality.	<ul style="list-style-type: none"> <li>• Physics-based real time lighting simulation with VR capabilities from HMD to CAVEs (multi-GPU, multi-node)</li> <li>• SPEOS Live Preview (raytracing) based on CUDA/OptiX benefiting from RTX architecture (single GPU)</li> <li>• Scalable rendering capabilities, ranging from rasterization to fully GPU ray-traced SPEOS Live Preview</li> </ul>	Single GPU Single Node
<b>Ansys VRXPERIENCE Lighting and Sensors</b>	ANSYS	Predictive validation of vehicle systems for the optimization of intelligent headlamp units and sensors dedicated to ADAS and AD. Rapid and simple virtual test of systems, relying on the unique combination of visually realistic driving simulator, and physics-based simulation. Real-time and interactive driving simulator to virtually create, test and experience future vehicle driving in real-world like conditions.	<ul style="list-style-type: none"> <li>• Multispectral Physics-based real time lighting simulation with multi-display capabilities (driving simulator).</li> </ul>	Multi-GPU Multi-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	Unified environment for virtual product development	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>ARCHICAD</b>	Nemetschek GRAPHISOFT	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction.	<ul style="list-style-type: none"> <li>• OpenGL based GPU rendering</li> <li>• Fast, efficient graphics in the viewport</li> <li>• RTX photorealistic rendering with Twinmotion</li> </ul>	Single GPU Single Node
<b>Arch-Log</b>	Luminova Japan	A web service based on NVIDIA Iray and RealityServer (from migenius) for rendering and configuring building materials.	<ul style="list-style-type: none"> <li>• Iray</li> <li>• RealityServer</li> <li>• Quadro</li> <li>• DGX</li> </ul>	Multi-GPU Multi-Node

<b>AutoCAD</b>	Autodesk	2D and 3D CAD designing, drafting, modeling, architectural drawing, and engineering software.	<ul style="list-style-type: none"> <li>• Surface, mesh and solid modeling tools, model documentation tools, parametric drawing capabilities</li> <li>• Open GL</li> <li>• Native DWG support</li> <li>• GRID Support.</li> </ul>	Single GPU Single Node
<b>Avatar VR</b>	NeuroDigital Technologies	Haptic VR gloves for training design or remote operation.	<ul style="list-style-type: none"> <li>• PhysX</li> </ul>	Single GPU Single Node
<b>CATIA 3DEXPERIENCE</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 OpenGL performance scaling in R2017x</li> <li>• VR native integration with HTC Vive in R2017x</li> <li>• VR SLI in R2018x</li> <li>• Stellar GPU in R2019x FD01</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>Clarisse</b>	Isotropix	Set dressing and layout tool with integrated renderer	<ul style="list-style-type: none"> <li>• GPU accelerated interactive rendering 50-100X faster than with CPU</li> <li>• OptiX AI-accelerated de-noising</li> </ul>	Single GPU Single Node
<b>Clo3D</b>	CLO Virtual Fashion Inc	3D garment simulation and design	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	Single GPU Single Node
<b>COMSOL</b>	COMSOL	Multiphysics general-purpose simulation software for modeling designs, devices and processes in all fields of engineering, manufacturing, and scientific research	<ul style="list-style-type: none"> <li>• OpenGL version 2.0</li> <li>• DirectX version 9</li> </ul>	Multi-GPU Single Node
<b>Creo Generative Topology Optimization Extension (GTO)</b>	PTC	Creo Generative Topology Optimization Extension (GTO) creates optimized product designs based on your constraints and requirements - including materials and manufacturing processes	<ul style="list-style-type: none"> <li>• CUDA accelerated Generative Design</li> </ul>	Multi-GPU Single Node
<b>Creo Parametric</b>	PTC	Professional 3D CAD software for product design and development, including parametric modeling, simulation/analysis, and product documentation for companies ranging from SMB to Enterprise.	<ul style="list-style-type: none"> <li>• GPU accelerated real-time engineering simulation with Creo Simulation Live</li> <li>• Full scene anti-aliasing</li> <li>• Order independent transparency</li> <li>• Better lighting and enhanced shaded-with-edges mode</li> <li>• Immersive design environment with realistic materials</li> </ul>	Single GPU Single Node
<b>Easy 3D Scan</b>	Cappasity	3D digitizing software that creates and embeds 3D product images into your website, mobile and AR/VR apps, and gives your customer a near real shopping experience.	<ul style="list-style-type: none"> <li>• OpenCL</li> </ul>	Single GPU Single Node
<b>Enscape</b>	Enscape GmbH	Plug-in for Revit, Rhino, SketchUp, ARCHICAD, and Vectorworks	<ul style="list-style-type: none"> <li>• One-click to VR experience</li> <li>• Design reviews for buildings</li> <li>• 3D and VR visualization of CAD data for AEC</li> </ul>	Single GPU Single Node
<b>Grasshopper</b>	McNeel & Assoc.	Grasshopper is a graphical algorithm editor tightly integrated with Rhino's 3-D modeling tools. Unlike RhinoScript, Grasshopper requires no knowledge of programming or scripting, but still allows designers to build form generators from the simple to the awe-inspiring.	<ul style="list-style-type: none"> <li>• Fast, scalable OpenGL 3.3 pipeline leverages latest NVIDIA GPUs</li> <li>• GPU computed shaders and memory optimizations</li> <li>• Rhino 6 leverages NVIDIA RT Cores for Real-time ray tracing viewport mode</li> <li>• Rendering engine is CYCLES, fully integrated inside Rhino 6 now</li> </ul>	Single GPU Single Node

<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</li> <li>• VRWorks SPS and VR SLI (NVLink support)</li> <li>• DesignWorks, including VR Occlusion Culling open source sample and OptiX</li> </ul>	Multi-GPU Single Node
<b>Inspire Studio/Render (formerly known as Evolve)</b>	Altair	Inspire Studio is a high quality 3D Hybrid Modeling and Rendering environment that enables industrial designers to evaluate, research and visualize various designs faster than ever before. Inspire Studio runs on both Mac OS X and Windows.	<ul style="list-style-type: none"> <li>• NURBS modeling</li> <li>• PolyNURBS modeling</li> <li>• OpenGL 4.5 Core</li> <li>• OpenGL-based real-time high-quality rendering</li> <li>• Interactive high-quality rendering using Thea Render</li> <li>• Production rendering using Thea Render</li> <li>• Integrated “dark room” environment to manage render queue and post-processing of rendered images</li> </ul>	Single 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 Server</li> <li>• Fast interactive ray tracing</li> <li>• Physically-based, global-illumination rendering</li> <li>• Distributed cluster rendering.</li> </ul>	Multi-GPU Multi-Node
<b>Iray for 3ds Max</b>	Siemens Digital Industries Software	A 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 and AI based denoising</li> </ul>	Multi-GPU Multi-Node
<b>Iray for Maya</b>	0x1 Software & Consulting GmbH	A 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 Pty Ltd	Iray plugin for Rhino	<ul style="list-style-type: none"> <li>• Iray Photoreal and Iray Interactive support</li> <li>• VCA clustering</li> <li>• Cloud rendering</li> <li>• MDL support.</li> </ul>	Single GPU Single Node
<b>Iray Server</b>	migenius Pty Ltd	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 and AI-based denoising</li> </ul>	Multi-GPU Multi-Node
<b>KeyShot</b>	Luxion	Physically correct real time and batch CPU GPU photorealistic renderer, popular in manufacturing, AEC, and M&E	<ul style="list-style-type: none"> <li>• GPU accelerated real time and batch rendering with NVIDIA OptiX</li> <li>• GPU accelerated AI Denoising with NVIDIA OptiX Denoiser</li> <li>• Network rendering on GPU accelerated nodes</li> <li>• Support for 30 different native file formats, many free plugins and live linked applications</li> </ul>	Multi-GPU Multi-Node

<b>LensMechanix</b>	Zemax	LensMechanix is the best application for mechanical engineers to package optical systems in CAD software. It is available for SOLIDWORKS users and for Creo Parametric users.	<ul style="list-style-type: none"> <li>• Optical product teams need an easier and faster way to get from design to manufacture</li> <li>• LensMechanix is the answer</li> <li>• LensMechanix is software for mechanical engineers who design housing for optical products in CAD</li> <li>• With LensMechanix, mechanical engineers can access the complete design data of optical systems designed in OpticStudio and start designing the mechanical envelope right away</li> <li>• They can then validate their mechanical design and fix issues before building a physical prototype</li> </ul>	Single GPU Single Node
<b>LumenRT</b>	Bentley Systems	Easily integrate life-like digital nature into your simulated infrastructure designs, and create high-impact visuals for stakeholders. Easy for any professional in the AEC industry to use and produce stunningly beautiful and easily understandable visualizations. Best for very large infrastructure, i.e. 100s of square kilometers rendering.	<ul style="list-style-type: none"> <li>• RT Cores for real time ray tracing</li> <li>• TensoRT for denoising</li> <li>• All using the DXR API</li> </ul>	Single GPU Single Node
<b>Medium</b>	Adobe	VR tool to sculpt, model, and paint. For beginners as well as pros. Adobe acquired from Oculus in December 2019	<ul style="list-style-type: none"> <li>• GLSL shaders</li> <li>• Vulkan</li> <li>• NVENC</li> </ul>	Single GPU Single Node
<b>Meshroom Studio</b>	Meshroom VR	Real-time rendering tool specially made for industrial design reviews, allowing to import, edit materials, set up your scene and showcase your model in real-time.	<ul style="list-style-type: none"> <li>• RTX real-time ray tracing</li> </ul>	Single GPU Single Node
<b>META</b>	BETA CAE Systems	High-performance multi-disciplinary CAE post-processor	<ul style="list-style-type: none"> <li>• OpenGL</li> <li>• OpenCL</li> </ul>	Single GPU Single Node
<b>META VR</b>	BETA CAE Systems	Powerful processing and visualization environment for interaction with full-scale simulation models with collaboration capabilities	<ul style="list-style-type: none"> <li>• OpenGL</li> <li>• OpenCL</li> </ul>	Single GPU Single Node
<b>MicroStation Connect</b>	Bentley Systems	MicroStation is the world's leading 3D computer-aided design and visualization software for the architecture, engineering, construction, and operation of all infrastructure types. Largest CAD in AEC for Civil Engineering users. <ul style="list-style-type: none"> <li>• Very tight collaboration with Autodesk Revit.</li> <li>• MicroStation has internal Rendering tool called Vue, shipping with the base CAD tool.</li> </ul>	<ul style="list-style-type: none"> <li>• Digital Nature modeling is Full Ray Tracing-enabled</li> <li>• Reality Modeling leveraging NVIDIA AI acceleration</li> <li>• GPU acceleration for Viz, Rendering, Simulation Bentley apps are optimized for NV Quadro RTX</li> </ul>	Single GPU Single Node
<b>Notch Builder</b>	10bit FX	A motion graphics and VFX tool designed by games artists and VJs. Compositing, grading and strong inter-operability with other packages.	<ul style="list-style-type: none"> <li>• GPU accelerated graphics and effects</li> </ul>	Single GPU Single Node
<b>NX</b>	Siemens Digital Industries Software	Siemens PLM Software premium design app with full Iray integration, supporting multi-gpu rendering. Still CPU bound for most tasks otherwise	<ul style="list-style-type: none"> <li>• GRID support</li> <li>• Iray, MDL (see NX Ray Traced Studio Database Entry)</li> </ul>	Multi-GPU Multi-Node

<b>OpticStudio</b>	Zemax	OpticStudio combines complex physics and interactive visuals so you can analyze, simulate, and optimize optics, lighting and illumination systems, and laser systems, all within tolerance specifications.	<ul style="list-style-type: none"> <li>• Share designs between OpticStudio and CAD packages as native files, giving mechanical engineers full access to the optical coordinate system and all critical dimension there is no need for file format conversions which can cause loss of design data</li> <li>• Simulate the impact of mechanical components on optical performance to uncover any issues and make informed design decisions</li> <li>• Check for, and resolve errors, before building costly physical prototypes</li> </ul>	N/A
<b>Painter</b>	Corel	Raster-based digital art application for drawing, sketching and painting.	<ul style="list-style-type: none"> <li>• GPU accelerated brushes</li> </ul>	Single GPU Single 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>Quark VR</b>	Quark VR	QuarkVR is an ultra-fast software solution which provides low-latency compression and wireless transmission. It offloads the heavy processing on the GPU, and is hardware-agnostic.	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	Single GPU Single Node
<b>QUINDOS</b>	Hexagon Manufacturing Intelligence	Coordinate metrology software	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>RealityServer</b>	migenius Pty Ltd	3D rendering and collaborative visualization and model manipulation platform based on NVIDIA Iray.	<ul style="list-style-type: none"> <li>• NVIDIA Iray.</li> </ul>	Multi-GPU Multi-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 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</li> <li>• 3D model display accelerated by GPU for smooth navigation of converted models in all display modes</li> </ul>	Multi-GPU Single Node
<b>RECOM WaveFarer</b>	RECOM	WaveFarer is a high-fidelity radar simulator for drive scenario modeling at frequencies up to and beyond 100GHz.	<ul style="list-style-type: none"> <li>• Near-field propagation method</li> <li>• Targeted ray casting, dynamic scenario, radiation patterns from antennas</li> </ul>	Multi-GPU Single Node
<b>RETOMO</b>	BETA CAE Systems	New software for the generation of 3D-tesellated models from CT-scan images	<ul style="list-style-type: none"> <li>• OpenGL</li> </ul>	Single GPU Single Node
<b>Review</b>	PiXYZ	Imports 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 CYCLES (their Renderer based on open source Blender) a real-time ray-traced display mode that is CUDA-based.	<ul style="list-style-type: none"> <li>• Fast, scalable OpenGL 3.3 pipeline leverages latest NVIDIA GPUs</li> <li>• GPU computed shaders and memory optimizations</li> <li>• Rhino 6 leverages NVIDIA RT Cores for Real-time ray tracing viewport mode</li> <li>• Rendering engine is CYCLES, fully integrated inside Rhino 6 now</li> </ul>	Single GPU Single Node
<b>Simcenter Femap</b>	Siemens Digital Industries Software	Engineering simulation application for creating, editing, and importing/re-using mesh-centric finite element analysis models of complex products or systems	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node



<b>Simcenter Prescan</b>	Siemens Digital Industries Software	virtually validate ADAS and automated vehicle functionalities by replicating real world scenarios, adding sensor models, and interface for control systems to design and verify algorithms for data processing, sensor fusion, decision making and control	<ul style="list-style-type: none"> <li>• Speed up the TIS sensor used for radar, lidar, PMD and ultrasonic sensors</li> <li>• Camera sensor and fisheye camera sensor</li> </ul>	Multi-GPU Multi-Node
<b>Simcenter STAR-CCM+ VR</b>	Siemens Digital Industries Software	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>SketchUp Pro</b>	Trimble SketchUp	SketchUp, formerly Google SketchUp, now part of Trimble in Sunnyvale, CA. SketchUp is a 3D modeling computer program for a wide range of drawing applications such as architectural, interior design, landscape architecture, civil and mechanical engineering, film and video game design.	<ul style="list-style-type: none"> <li>• OpenGL now but moving to DirectX 11 for SketchUp, and DirectX 12 for TEKLA</li> <li>• Fast, efficient graphics in the viewport</li> <li>• RTX photorealistic rendering</li> <li>• 3rd party plug-ins supported by SketchUp Pro</li> </ul>	Single GPU Single Node
<b>Solid Edge</b>	Siemens Digital Industries Software	SMB CAD option from Siemens	<ul style="list-style-type: none"> <li>• KeyShot rendering</li> </ul>	Single GPU Single Node
<b>SOLIDWORKS</b>	Dassault Systèmes	3D design and product development solution including design, simulation, cost estimation, manufacturability checks, CAM, sustainable design, 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 based on NVIDIA Iray	<ul style="list-style-type: none"> <li>• Iray-based ray-tracing</li> <li>• Animation support</li> <li>• Network rendering</li> <li>• OptiX-based Artificial Intelligence denoiser</li> </ul>	Single GPU Single Node
<b>Spotscale</b>	Spotscale	3D reconstruction algorithms are tailored for buildings and urban environments. using drones to captured data.	<ul style="list-style-type: none"> <li>• cuDNN</li> </ul>	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</li> <li>• 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 Alchemist</b>	Adobe	Allows to simply create material from picture or by blending pre-existing materials, create and manage your material libraries	<ul style="list-style-type: none"> <li>• DL powered material recognition</li> <li>• Material scan, edit and blend</li> </ul>	Single GPU Single Node
<b>Substance Designer</b>	Adobe	Material shader edition and market reference for procedural texture creation.	<ul style="list-style-type: none"> <li>• RTX bakers</li> <li>• Iray viewport/rendering</li> </ul>	Multi-GPU Single Node
<b>Substance Painter</b>	Adobe	Intuitive interactive 3D painting software with physics and particle support.	<ul style="list-style-type: none"> <li>• RTX bakers</li> <li>• Iray viewport</li> </ul>	Multi-GPU Single Node
<b>Sunata</b>	Siemens Digital Industries Software	Cloud-based thermal modeling for additive manufacturing. Recommends optimal parameters for the print, including print orientation and support structures.	<ul style="list-style-type: none"> <li>• Thermal simulation</li> </ul>	Multi-GPU Single Node
<b>Teamcenter Active Workspace</b>	Siemens Digital Industries Software	Active Workspace is an IT-friendly client for Teamcenter product lifecycle management, with zero-install footprint and web browser access that provides an identical and seamless experience on any computing or smart device.	<ul style="list-style-type: none"> <li>• GRID support</li> </ul>	Single 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</li> <li>• Real time photorealistic rendering</li> <li>• CUDA</li> </ul>	Multi-GPU Single Node
<b>UE4</b>	Epic Games	Unreal Engine 4 is a suite of integrated tools for developers to design and build games, simulations, and visualizations.	<ul style="list-style-type: none"> <li>• GPU Accelerated Rendering on OpenGL, DirectX and Vulkan</li> <li>• Phys-X implemented</li> </ul>	Single GPU Single Node
<b>Vectorworks</b>	Nemetschek VECTORWORKS	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>Volumetric Camera Systems</b>	Volumetric Camera Systems	4D capture service with high quality and realistic “holograms-in-motion” of people, animals, or any moving subject.  Secondly, we offer “photo-realistic 3D environment captures” using industrial grade Leica Laser Scanners and advanced high-resolution multi-camera systems.	<ul style="list-style-type: none"> <li>• CUDA</li> <li>• Quadro GPUs</li> </ul>	Single GPU Single Node
<b>VRED</b>	Autodesk	VRED 3D visualization software for automotive designers and engineers to create product presentations, design reviews, and virtual prototypes. Uses 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>WYSIWYG</b>	Cast Software	Wysiwyg is an all-in-one lighting design software with fully integrated CAD, plots, data, visualization and virtual show control. Features the largest CAD library with thousands of 3D objects you can choose from to design your entire show.	<ul style="list-style-type: none"> <li>• GPU accelerated Shaded Views and Virtual Views</li> </ul>	Multi-GPU Single Node
<b>ZLVE</b>	Zerolight	Immersive customer experience with VR or web GPU streaming	<ul style="list-style-type: none"> <li>• VRS and foveated rendering for VR and 3D experience through AWS GPU streaming</li> </ul>	Multi-GPU Single Node

## ELECTRONIC DESIGN AUTOMATION

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Advanced Design System (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>Altair Feko</b>	Altair	Comprehensive computational electromagnetics (CEM) code used widely in the telecommunications, automobile, space and public sector industries to solve high-frequency problems.	<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>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> <li>• OpenGL rendering</li> </ul>	Multi-GPU Single Node
<b>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> <li>• OpenGL rendering</li> </ul>	Multi-GPU Multi-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, and IBIS-AMI analysis speedup with GPU computing.	<ul style="list-style-type: none"> <li>AMI analysis</li> </ul>	Single GPU Single Node
<b>Cadence Allegro</b>	Cadence Design Systems	EDA/ECAD tool for PCB (Printed Circuit Board) Design	<ul style="list-style-type: none"> <li>OpenGL extensions</li> <li>Scalable Vector Graphics (SVG), Path Rendering SDK</li> </ul>	Multi-GPU Multi-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>Computational lithography simulations for mask synthesis on GPUs</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>	Single GPU Single 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>Finite Difference Time Domain (FDTD) 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>REMCOM XFDTD</b>	REMCOM	3D EM Simulation solver.	<ul style="list-style-type: none"> <li>FDTD Solver</li> </ul>	Multi-GPU Multi-Node
<b>samadii/em</b>	Metariver Technology	Software for computing the electromagnetic field in three dimensional space using the Maxwell equation, a governing equation that can comprehensively represent these electromagnetic phenomena	<ul style="list-style-type: none"> <li>Electromagnetics simulator, FEM solver(scalar FEM, vector FEM)</li> <li>Electrostatics solver, Electromagnetic wave solver</li> <li>Magnetostatics solver, Electric current solver, Electrodynamics solver</li> <li>Co-simulation with samadii/sciv, samadii/dem and fluid flow solvers.</li> </ul>	Multi-GPU Multi-Node
<b>samadii/plasma</b>	Metariver Technology	Software for computing plasma phenomenon with PIC(Particle-in-Cell) method. Two-way coupled simulation with samadii/em and samadii/sciv.	<ul style="list-style-type: none"> <li>Plasma simulator, Charged particle motion analysis</li> <li>Particle and surface reaction calculation, Field analysis, Sheath range prediction</li> <li>DSMC collision module, PIC module</li> <li>Co-simulation with samadii/em, Ansys Maxwell and COMSOL.</li> </ul>	Multi-GPU Multi-Node
<b>SEMCAD-X</b>	SPEAG	3D Full wave electromagnetic and computational life sciences simulation solver	<ul style="list-style-type: none"> <li>FDTD solver</li> </ul>	Multi-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>Transient, Broadband, and Harmonic simulations FDTD solver</li> <li>Linear and non-linear 3D full wave acoustics solvers</li> </ul>	Multi-GPU Single Node

<b>Synopsys LucidShape</b>	Synopsys	LucidShape is a computer aided lighting (CAL) design software for automotive lighting design tasks. Supports algorithms optimized for automotive applications, LucidShape facilitates the design of automotive forward, rear and signal lighting, and reflectors.	<ul style="list-style-type: none"> <li>• Ray Tracing</li> <li>• Monte Carlo simulations using OptiX 6.5 and CUDA 10.2</li> </ul>	Single 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>VSim for Electromagnetics</b>	Tech-X Corporation	Conformal FDTD for electromagnetics for a variety of material types, yielding engineering outputs that can be used for design of electromagnetic devices	<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 for long cylindrical structures	<ul style="list-style-type: none"> <li>• MoM Solver</li> <li>• Matrix fill-in and near-field calculations</li> </ul>	Multi-GPU Single Node
<b>WIPL-D Pro</b>	WIPL-D	Solver for fast and accurate electromagnetic analysis of arbitrary composite 3D metallic and dielectric structures	<ul style="list-style-type: none"> <li>• MoM (Method of Moments) Solver</li> <li>• DDS (Domain Decomposition Solver)</li> </ul>	Multi-GPU Multi-Node
<b>WIPL-D Pro CAD</b>	WIPL-D	Modeling and simulation environment uniting versatile, yet simple geometry modeling, with signature WIPL-D simulation accuracy	<ul style="list-style-type: none"> <li>• MoM (Method of Moments) Solver</li> </ul>	Multi-GPU Single Node
<b>Wireless InSite</b>	REMCOM	Uses Optix 4.1 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	COMPANYNAME	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 and identification</li> <li>• Segmentation and defect detection</li> <li>• Object and 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</li> <li>• HALCON also provides an IDE for training neural networks</li> <li>• 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 Corporation	IBM Visual Insights uses cognitive capabilities to review and analyze parts, components, and products. Identifies 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 and Entertainment

## ANIMATION, MODELING AND RENDERING

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3ds Max</b>	Autodesk	3D modeling, animation, and rendering	<ul style="list-style-type: none"> <li>Faster interactive graphics</li> <li>Availability of Arnold with AI denoising</li> <li>Availability of Chaos V-Ray, Otoy Octane, Redshift, cebas finalRender third-party GPU renderers</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. Generates 3D spatial data to be used in GIS applications, and cultural heritage documentation for visual effects production and indirect measurements of objects of various scales.	<ul style="list-style-type: none"> <li>CUDA-accelerated photogrammetry solution</li> <li>RTX opportunity</li> </ul>	Multi-GPU Single Node
<b>Altair 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>GPU-accelerated hybrid renderer</li> <li>Advanced material layering system with subsurface scattering, displacement mapping, physical sun-sky and IES support</li> </ul>	Multi-GPU Single Node
<b>ArmorPaint</b>	Armory	ArmorPaint is a software designed for physically-based texture painting. There is a standalone version, or you can use as an Armory3D project. Draw textures directly using node based materials and brushes.	<ul style="list-style-type: none"> <li>GPU accelerated painting processes</li> </ul>	Single GPU Single Node
<b>Arnold</b>	Autodesk	Solid Angle Arnold film and animation renderer	<ul style="list-style-type: none"> <li>RTX</li> </ul>	Multi-GPU Single Node
<b>Beauty Box</b>	Digital Anarchy	Automatic masking and skin retouching.	<ul style="list-style-type: none"> <li>GPU accelerated graphics and compute</li> </ul>	Single GPU Single Node
<b>Blender</b>	Blender Institute	3D modeling, rendering and animation	<ul style="list-style-type: none"> <li>GPU-accelerated interactive viewport</li> </ul>	Single GPU Single Node
<b>Blender Cycles</b>	Blender Institute	GPU renderer	<ul style="list-style-type: none"> <li>CUDA-accelerated rendering</li> <li>RTX-accelerated ray tracing</li> </ul>	Multi-GPU Single Node
<b>Character Animator</b>	Adobe	Character animator that uses your expressions & movements to animate characters in real-time	<ul style="list-style-type: none"> <li>Auto lip syncing</li> <li>Deep Learning</li> </ul>	Single 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> <li>Support for Redshift and Chaos V-Ray and Otoy Octane and third-party GPU renderers</li> </ul>	Single GPU Single Node
<b>Corona</b>	Chaos Group	High-performance photorealistic renderer	<ul style="list-style-type: none"> <li>OptiX AI de-noising</li> </ul>	Single GPU Single Node
<b>D5 Render</b>	D5 Innovation	D5 Render, based on NVIDIA RTX GPU's real-time ray tracing and rasterization technology, aims to bring unprecedented real-time rendering experience for architecture and interior design.	<ul style="list-style-type: none"> <li>Real-time GPU accelerated physically based global illumination and ray tracing.</li> </ul>	Single GPU Single Node
<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.	<ul style="list-style-type: none"> <li>GPU accelerated compute</li> <li>Rendering via NVIDIA IRAY and Optix</li> </ul>	Multi-GPU Single Node
<b>Dimension</b>	Adobe	3D design tool enabling graphic designers to compose, adjust, and render photorealistic images.	<ul style="list-style-type: none"> <li>Accelerated graphics &amp; MDL (Material Definition Language)</li> </ul>	Multi-GPU Single Node

<b>finalRender</b>	Cebas	PLUGIN for 3dsMAX Physically Based (Spectral) Wavelength Simulation Biased + Unbiased Hybrid Rendering Unlimited Network Rendering	<ul style="list-style-type: none"> <li>• CUDA-accelerated renderer for Autodesk 3DS Max</li> <li>• OptiX AI de-noising</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>	Foundry	Powerful look development and lighting tool	<ul style="list-style-type: none"> <li>• Faster interactive graphics</li> </ul>	Single GPU Single Node
<b>Lightwave 3D</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>	Foundry	3D paint tool that allows painting directly onto 3D models	<ul style="list-style-type: none"> <li>• Faster interactive painting</li> </ul>	Single GPU Single Node
<b>Mars</b>	sheencity	Real-time architectural visualization tool with advanced features such as real-time ray tracing, DLSS, and VR.	<ul style="list-style-type: none"> <li>• RTX Ray tracing</li> <li>• DLSS</li> </ul>	Single GPU Single Node
<b>Massive</b>	Massive	Simulation and visualization tools for autonomous agent driven animation for film, games, television, architecture and transportation.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>Maverick Renderer</b>	Maverick	CUDA-based GPU renderer	<ul style="list-style-type: none"> <li>• CUDA-accelerated ray-tracing</li> <li>• OptiX 7 de-noising</li> </ul>	Single GPU Single Node
<b>Maxwell</b>	Next Limit	CUDA-accelerated interactive and final-frame renderer	<ul style="list-style-type: none"> <li>• CUDA-accelerated ray-tracing</li> <li>• Unrestricted image resolution</li> <li>• OptiX 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 and larger scenes</li> <li>• Availability of Chaos V-Ray, Otoy Octane and Redshift third-party GPU renderers</li> </ul>	Single GPU Single Node
<b>Meshroom</b>	Czech Technical University (CTU)	Open source photogrammetry 3D software	<ul style="list-style-type: none"> <li>• CUDA-accelerated depth analysis</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>NX Ray Traced Studio</b>	Siemens Digital Industries Software	Embedded rendering feature for Siemens NX	<ul style="list-style-type: none"> <li>• Iray based</li> <li>• MDL</li> <li>• AI denoising</li> </ul>	Multi-GPU Single Node
<b>OctaneRender</b>	Otoy	CUDA-accelerated GPU renderer	<ul style="list-style-type: none"> <li>• GPU accelerated rendering</li> <li>• AI de-noising</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 AI de-noising</li> </ul>	Single GPU Single Node



<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>Trapcode</b>	Red Giant	Particle simulations and 3D effects for motion graphics and VFX. Now with Fluid Dynamics.	<ul style="list-style-type: none"> <li>GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>TurbulenceFD</b>	Jawset	Turbulence FD is a powerful simulation tool to create smoke, fire and explosion effects.	<ul style="list-style-type: none"> <li>GPU accelerated graphics, compute and simulation</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 and final-frame GPU rendering</li> </ul>	Multi-GPU Single Node
<b>vRt</b>	vRt	vRt is an open-source project aiming to offer Vulkan-based ray-tracing for modern graphics cards that offers a unified ray-tracing, cross-platform library built against Vulkan 1.1	<ul style="list-style-type: none"> <li>vRtC (compute-based, native, default, wide GPU support)</li> <li>vRtX (NVIDIA RTX only, more higher performance at now)</li> </ul>	Multi-GPU Single Node
<b>WispRenderer</b>	Bred University of Applied Sciences	General purpose high level rendering library with RTX, RTGI, HBAO+, and Ansel support.	<ul style="list-style-type: none"> <li>RTX, RTGI, HBAO+</li> <li>Ansel</li> </ul>	Multi-GPU Single Node

## COLOR CORRECTION AND GRAIN MANAGEMENT

APPLICATION NAME	COMPANYNAME	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>Dark Energy</b>	Cinnafilm	Application and plug-in for image enhancement	<ul style="list-style-type: none"> <li>Image de-noising and restoration</li> <li>Noise reduction, de-noise and de-grain</li> <li>Grain removal, image sharpening and texture management dust busting</li> <li>SDR to HDR upres</li> </ul>	Multi-GPU Single Node
<b>DaVinci Resolve</b>	Blackmagic Design	Color grading and editing	<ul style="list-style-type: none"> <li>Real-time color correction and de-noising</li> <li>RTX-accelerated AI features for re-timing and image enhancement</li> </ul>	Multi-GPU Single Node
<b>DeNoise AI</b>	Topaz Labs	DeNoise AI uses machine-learning to remove noise from your image while preserving detail for a crisp, clear result. Whether you are shooting with High ISO or in a low light scenario, DeNoise will correct your image without removing any important information or patterns in your image.	<ul style="list-style-type: none"> <li>GPU accelerated effects</li> </ul>	Single 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>FilmConvert Nitrate</b>	FilmConvert	A film emulation plugin that can take Log-based video and transform it into full color corrected media with a natural grain similar to that of commonly loved film stocks.	<ul style="list-style-type: none"> <li>GPU accelerated processing</li> <li>Cineon Log film emulations</li> <li>Full custom curve controls</li> <li>Advanced film grain controls</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>HDR Image Analyser</b>	aja	A 1RU waveform, histogram, vectorscope and Nit-level HDR monitoring solution for HD, UltraHD, 2K, and HD resolution with HDR and WCG content.	<ul style="list-style-type: none"> <li>• Precise, high quality UltraHD UI for native-resolution picture display</li> <li>• Advanced out of gamut and out of brightness detection with error intolerance</li> <li>• Support for SDR (Rec.709), ST2084/PQ and HLG analysis</li> <li>• CIE graph, Vectorscope, Waveform, Histogram</li> <li>• Out of gamut false color mode to easily spot out of gamut/out of brightness pixels</li> <li>• Data analyzer with pixel picker</li> <li>• Up to 4K/UltraHD 60p over 4x 3G-SDI inputs</li> <li>• SDI auto signal detection</li> <li>• File base error logging with timecode</li> <li>• Display and color processing look up table (LUT) support</li> <li>• Line mode to focus a region of interest onto a single horizontal or vertical line</li> <li>• Loop through output to broadcast monitors</li> <li>• Still store</li> <li>• Nit levels and phase metering</li> <li>• Built-in support for color spaces from ARRI, Canon, Panasonic, RED and Sony</li> </ul>	Single GPU Single Node
<b>Magic Bullet Colorista</b>	Red Giant	Real time, interactive, multi-layered masked color correction (video playback too!) with the Mercury Playback engine in Premiere Pro.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>Magic Bullet Looks</b>	Red Giant	Powerful looks and color correction for filmmakers.	<ul style="list-style-type: none"> <li>• GPU accelerated compute</li> </ul>	Single GPU Single Node
<b>Mist</b>	Marquise Technologies	Mastering tool for cinema, broadcast and over-the-top content	<ul style="list-style-type: none"> <li>• 100% CUDA-accelerated imaging pipeline for de-bayering, color grading, transcoding and image enhancement</li> <li>• Integrated Dolby Vision pipeline</li> </ul>	Multi-GPU Single Node
<b>Nucoda</b>	Digital Vision	Color grading	<ul style="list-style-type: none"> <li>• GPU-accelerated color grading</li> <li>• Accelerated scopes, playback and rendering</li> </ul>	Single GPU Single Node
<b>Pablo family</b>	Grass Valley	Color grading and finishing	<ul style="list-style-type: none"> <li>• Real time color correction</li> </ul>	Multi-GPU Single Node
<b>Pablo Rio</b>	Grass Valley	Pablo Rio is a color grading application that GV acquired when they purchased Snell.	<ul style="list-style-type: none"> <li>• CUDA-accelerated color grading</li> </ul>	Multi-GPU Single Node
<b>PFClean</b>	The Pixel Farm	Image restoration and remastering	<ul style="list-style-type: none"> <li>• CUDA-based image processing acceleration</li> </ul>	Multi-GPU Single Node
<b>RAW Converter</b>	ARRI	RAW de-Bayering and primary color grading	<ul style="list-style-type: none"> <li>• CUDA-accelerated de-bayering and primary grading</li> </ul>	Single GPU Single Node
<b>REDCINE-X PRO</b>	Red Digital Cinema	Primary color grading	<ul style="list-style-type: none"> <li>• CUDA-accelerated de-bayering and primary color grading</li> </ul>	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. Used by many color grading and video editing applications.	<ul style="list-style-type: none"> <li>• CUDA-accelerated wavelet decoding and de-bayering</li> </ul>	Single GPU Single Node
<b>Scratch</b>	Assimilate	Color grading and finishing	<ul style="list-style-type: none"> <li>• Accelerated de-bayering for real-time digital finishing</li> </ul>	Single GPU Single Node
<b>VFX Suite</b>	Red Giant	VFX Suite is a complete set of visual effects and motion graphics plugins for creating professional effects.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node

## COMPOSITING, FINISHING AND EFFECTS

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>After Effects</b>	Adobe	Motion graphics and effects	<ul style="list-style-type: none"> <li>• CUDA acceleration for up to 10x faster performance on key effects plus enhanced 3D ray tracing</li> </ul>	Single GPU Single Node
<b>Aura</b>	Rowbyte	Aura is a procedural plug-in for After Effects that creates elegant geometric shapes in 3D space. It's akin to a particle system but instead of rendering small particles all over the place, it generates vector like shapes (waves) that change over time much like the classic Radiowaves plug-in.	<ul style="list-style-type: none"> <li>• GPU-accelerated High Frequency Rendering</li> </ul>	Single GPU Single Node
<b>Clipster</b>	Rohde & Schwarz	Video and film player and DCI Packager	<ul style="list-style-type: none"> <li>• GPU-accelerated</li> <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	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Continuum</b>	Boris FX	Visual effects plug-in for creative effects, titling, and quick fixes.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>DE:Noise</b>	RE:Vision Effects	Reduce noise, dust, and artifacts with frame-to-frame motion tracking. Useful for low light shoots, CG renders with ray tracing sample artifacts, excessive film grain.	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>DEFlicker</b>	RE:Vision Effects	Reducing flicker and artifacts in high-frame-rate and time-lapse video.	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Element 3D</b>	Video Copilot	Advanced 3D object & particle render engine plugin for Adobe After Effects	<ul style="list-style-type: none"> <li>• GPU accelerated graphics and compute</li> </ul>	Single GPU Single Node
<b>Flame Premium</b>	Autodesk	Finishing and color grading	<ul style="list-style-type: none"> <li>• Integrated toolset for 3D VFX, editorial, and color grading</li> </ul>	Multi-GPU Single Node
<b>Flicker Free</b>	Digital Anarchy	Deflicker Time Lapse, Slow Motion, and Old Video. Flicker Free is a powerful, new way to deflicker video.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>Fusion</b>	Blackmagic Design	Effects and compositing	<ul style="list-style-type: none"> <li>• 3D tracking</li> <li>• Compositing</li> <li>• VR</li> </ul>	Single GPU Single Node
<b>HIERO</b>	Foundry	Multi-shot management tool that supports collaborative working, review and approval, quick production turnaround and delivery	<ul style="list-style-type: none"> <li>• Fluid, interactive playback</li> </ul>	Single GPU Single Node
<b>Imerge Pro</b>	FXhome	<p>Imerge Pro is layer-based image compositing software that is GPU accelerated, making performance astonishingly fast, even on high-resolution images.</p> <p>Create pro-level composites with unlimited layers and zero baked-in changes. Imerge Pro is the first photo editing software to keep your image data RAW and your layers self-contained.</p>	<ul style="list-style-type: none"> <li>• GPU-accelerated processing</li> </ul>	Single GPU Single Node
<b>Magic Bullet Denoiser</b>	Red Giant	Magic Bullet Denoiser III lets you reduce visible noise and grain in digital video produced by digital video cameras, camcorders, or film.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>Magic Bullet Film</b>	Red Giant	Gives digital footage the look of real film by emulating the entire photochemical process from the original film negative, to color grading, and finally to the print stock.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node

<b>Magic Bullet Suite</b>	Red Giant	Full suite of tools for color correction, finishing and film looks for filmmakers.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing and affects</li> </ul>	Single GPU Single Node
<b>Mamba FX</b>	SGO	High-end compositing	<ul style="list-style-type: none"> <li>• Faster keying, tracking, painting and restoration</li> </ul>	Single GPU Single Node
<b>MediaReactor</b>	Drastic Technologies	Debayering and processing of raw camera files.	<ul style="list-style-type: none"> <li>• GPU-accelerated compute</li> </ul>	Single GPU Single Node
<b>Mighty Bake</b>	Mighty Bake	A powerful, easy to use, all-in-one texture baking solution for any 3D artist	<ul style="list-style-type: none"> <li>• GPU accelerated processing</li> </ul>	Single GPU Single Node
<b>Mistika Ultima</b>	SGO	Color grading and finishing	<ul style="list-style-type: none"> <li>• Faster keying, tracking, painting and restoration, de-bayering</li> </ul>	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>• Export clips in many formats, including DPX and ProRes</li> </ul>	Single GPU Single Node
<b>Mocha Pro</b>	Boris FX	Mocha Pro is an award-winning planar tracking tool for motion tracking, rotoscoping, object removal, camera stabilization and general visual effects.	<ul style="list-style-type: none"> <li>• GPU accelerated planar tracking and object removal</li> </ul>	Single GPU Single Node
<b>Natron</b>	Natron	Natron is a free and open-source node-based compositing software application.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing and rendering</li> </ul>	Single GPU Single Node
<b>Neat Video</b>	Absoft	Digital filter with auto-profiling tool designed to reduce visible noise and grain found in footage.	<ul style="list-style-type: none"> <li>• GPU accelerated processing</li> </ul>	Single GPU Single Node
<b>NUKE</b>	Foundry	Compositing tool with 3D tracker	<ul style="list-style-type: none"> <li>• GPU-accelerated BLINK processing</li> <li>• Faster compositing and effects</li> </ul>	Single GPU Single Node
<b>Optics</b>	Boris FX	Optics is designed to simulate optical camera filters, specialized lenses, film stocks and grain, lens flares, optical lab processes, color correction as well as natural light and photographic effects. First collaborative product between Sapphire and Digital Film Tools. Plugin for Photoshop and Lightroom, also has a Windows and Mac standalone application.	<ul style="list-style-type: none"> <li>• GPU accelerated processing and affects</li> </ul>	Single GPU Single Node
<b>PFTrack</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>Plexus</b>	Rowbyte	Plexus is a plug-in designed to bring generative art closer to a non-linear program like After Effects. It lets you create, manipulate and visualize data in a procedural manner. Render the particles and create all sorts of interesting relationships between them based on various parameters using lines and triangles.	<ul style="list-style-type: none"> <li>• Plexus (interacts natively with AE's Camera)</li> <li>• High-quality, GPU-accelerated Depth of Field effects</li> </ul>	Single GPU Single Node
<b>Rotobot</b>	Kognat	An AI product for compositing packages which uses machine learning to generate mattes for machine-based rotoscoping.	<ul style="list-style-type: none"> <li>• CUDA accelerated AI rotoscoping</li> </ul>	Multi-GPU Single Node
<b>Sapphire</b>	Boris FX	The Sapphire suite is an all-in-one solution containing hundreds of effects, presets, and workflows that are aimed at taking professional video work to the next level.	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>SilhouetteFX</b>	Boris FX	Invaluable in post-production, Silhouette continues to bring best of class tools to the visual effects industry. As a fully featured GPU accelerated compositing system, its standout features are award winning rotoscoping and non-destructive paint as well as keying, matting, warping, morphing, and a total of 142 different nodes--all stereo enabled.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing and affects</li> </ul>	Single GPU Single Node

<b>Silhouette Paint</b>	Boris FX	Rotoscoping tool that allows for intensive VFX fixes, blemish cleanup, beauty effects, wire/object removal, style effects on video, and as an artistic paint tool. It is raster based so it has a smaller memory footprint (fastest paint plugin on the market), Integrated with Mocha Pro planar tracker	<ul style="list-style-type: none"> <li>• GPU accelerated processing and affects</li> </ul>	Single GPU Single Node
<b>Twixtor</b>	RE:Vision Effects	Optical flow tracking of pixel motion to synthesize new frames by warping & interpolating frames of the original sequence. Reduces artifacts & retime frames.	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node
<b>Video Essentials</b>	NewBlueFX	Comprehensive collection of titling, transitions and video effects.	<ul style="list-style-type: none"> <li>• Faster effects</li> </ul>	Single GPU Single Node

## (VIDEO) EDITING

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Blackmagic RAW SDK</b>	Blackmagic Design	Blackmagic RAW is a CPU and GPU-enabled SDK for decoding and debayering Blackmagic RAW files on MacOS, Windows and Linux	<ul style="list-style-type: none"> <li>• CUDA-accelerated de-coding and de-bayering</li> </ul>	Single GPU Single Node
<b>Catalyst Production Suite</b>	Sony Creative Software	4K, Sony RAW, and HD video editing. Includes 3 applications: Browse, Prepare, Edit	<ul style="list-style-type: none"> <li>• Faster effects, transitions and encoding</li> <li>• RAW camera de-bayering</li> </ul>	Single GPU Single Node
<b>CineMatch</b>	FilmConvert	CineMatch is a set of tools designed to help you match footage shot on different cameras to a baseline technical level - a seamless, matched timeline in Log or REC.709, ready for creative grading.	<ul style="list-style-type: none"> <li>• Real-time color matching conversions with CUDA</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 camera de-bayering</li> </ul>	Single GPU Single Node
<b>Filmora</b>	Wondershare	Filmora is an easy-to-use and trendy video editing software that lets you empower your story and be amazed at results, regardless of your skill level. With Filmora, you can get started with any new movie project by importing and editing your video, adding special effects and transitions, and sharing your final production on social media, mobile devices, or DVDs.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing</li> </ul>	Single GPU Single Node
<b>Gigapixel AI</b>	Topaz Labs	Photo up scaling by using AI to "fill in" and add new detail when enlarging photos.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>GPUSqueeze</b>	Multicamera Systems	GPUSqueeze is cross platform software library for multi-stream and ultra high speed video encoding, transcoding and processing using multi-GPU and distributed setups. The library uses highly optimized patent pending algorithms to achieve maximum speed, high hardware utilization and provides almost linear performance scaling with the increase of number of GPUs in the system.	<ul style="list-style-type: none"> <li>• GPU accelerated video encoding and decoding</li> </ul>	Multi-GPU Multi-Node
<b>HitFilm Pro</b>	FXhome	HitFilm Pro is an all-in-one video editor, compositor, and visual effects (VFX) software designed for filmmakers, professional video editors, and visual content producers.	<ul style="list-style-type: none"> <li>• GPU accelerated effects and decoding</li> </ul>	Single GPU Single Node
<b>Illustrator</b>	Adobe	Vector graphics software for creating logos, icons, drawings, typography, and illustrations for print, web, video, and mobile devices.	<ul style="list-style-type: none"> <li>• Entire canvas optimized for NVIDIA GPUs for faster pan &amp; zoom</li> <li>• Accelerated by CUDA-based NV Path Render technology</li> </ul>	Single GPU Single Node

<b>Lightroom Classic</b>	Adobe	Easily edits organizes, stores, and shares your photos.	<ul style="list-style-type: none"> <li>• GPU accelerated Develop module plus new Sensei features like "Enhance Details" with NVIDIA GPU AI optimization.</li> <li>• Up to 600% faster than integrated GPUs with controls like Texture, Dehaze, &amp; Sharpening</li> <li>• Improved editing in 1:1 view &amp; on hi-rez displays.</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>Live Planet</b>	Live Planet	Livestreaming, recording and delivery of stereoscopic 360 VR	<ul style="list-style-type: none"> <li>• Real time 360 3D capture and stitch</li> <li>• 4K</li> </ul>	Single GPU Single Node
<b>Luminar</b>	Skylum	Luminar is the world's first photo editor that adapts to your style & skill level. It is designed to make complex photo editing easy & enjoyable for everyone. Take advantage of over 300 powerful, yet simple photo editing tools that allow you to perform all kind of image editing tasks.	<ul style="list-style-type: none"> <li>• GPU accelerated processing and AI affects</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>Movavi Video Suite</b>	Movavi	An all-in-one video maker: an editor, converter, screen recorder, and more.	<ul style="list-style-type: none"> <li>• Faster conversion speed with NVIDIA CUDA</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 allowing remote GPU-accelerated production workflows</li> </ul>	Single GPU Single Node
<b>Photoshop</b>	Adobe	Photo editing to transform your images into anything you can imagine	<ul style="list-style-type: none"> <li>• Over 30 GPU accelerated features such as blur gallery, liquify, smart sharpen, perspective warp</li> </ul>	Single GPU Single Node
<b>Pinnacle Studio</b>	Corel	Video editing and sharing program.	<ul style="list-style-type: none"> <li>• GPU accelerated compute and effects</li> </ul>	Single GPU Single Node
<b>PowerDirector</b>	CyberLink	PowerDirector delivers professional-grade video editing and production for creators of all levels. Whether you are editing in 360 degrees, Ultra HD 4K or even the latest online media formats, PowerDirector remains the definitive Windows video editing solution for anyone, whether they are beginners or professionals.	<ul style="list-style-type: none"> <li>• GPU accelerated compute</li> </ul>	Single GPU Single Node
<b>Premiere Pro</b>	Adobe	Video editing software for film, TV, and the web.	<ul style="list-style-type: none"> <li>• Real-time video editing &amp; fast output rendering based on CUDA</li> </ul>	Multi-GPU Single Node
<b>Premiere Rush</b>	Adobe	Easy-to-use video editor for creating and sharing online videos.	<ul style="list-style-type: none"> <li>• CUDA</li> <li>• Real-time video editing</li> <li>• Fast output rendering</li> </ul>	Multi-GPU Single Node
<b>Sharpen AI</b>	Topaz Labs	Sharpening and shake reduction software that can tell difference between real detail and noise.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> <li>• Machine Learning</li> </ul>	Single GPU Single Node
<b>SmartCourtPro</b>	PlaySight	Sophisticated video and analytics training technology with the latest in AI, integrations and player development tools.	<ul style="list-style-type: none"> <li>• IVA</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>TotalFX</b>	NewBlueFX	Comprehensive collection of Titling, Compositing, Polishing and Styling tools.	<ul style="list-style-type: none"> <li>• GPU-accelerated affects</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> <li>• Uses NVENC to encode/decode H.264 and HEVC streams</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
<b>Video Enhance AI</b>	Topaz Labs	Trained on thousands of videos and combining information from multiple input video frames, Topaz Video Enhance AI will enlarge and enhance your footage up to 8K resolution with true details and motion consistency.	<ul style="list-style-type: none"> <li>• GPU accelerated AI inference and processing</li> </ul>	Single GPU Single Node
<b>Video Studio</b>	Corel	High quality tools that build, edit, and correct video skillfully.	<ul style="list-style-type: none"> <li>• GPU accelerated compute</li> </ul>	Single GPU Single Node
<b>VLC Media Player</b>	VideoLAN Organization	VLC is a free and open source cross-platform multimedia player and framework that plays most multimedia files as well as DVDs, Audio CDs, VCDs, and various streaming protocols.	<ul style="list-style-type: none"> <li>• NV Video Codec accelerated encoding and decoding</li> </ul>	Single GPU Single Node
<b>WonderLive</b>	Z Cam	Cinematic VR Camera with excellent image quality, stereoscopic 360 degrees; recording, and live streaming.	<ul style="list-style-type: none"> <li>• Up to 4K output resolution equirectangular image</li> <li>• Save live stitched video file</li> <li>• Preview live stitched video</li> <li>• RTMP live streaming output</li> <li>• Supports VRworks 360 video SDK</li> </ul>	Single GPU Single Node

## (IMAGE & PHOTO) EDITING

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Adjust AI</b>	Topaz Labs	Adjust AI is a one click application that leverages the power of machine learning to intelligently enhance photos.	<ul style="list-style-type: none"> <li>• GPU accelerated effects</li> </ul>	Single GPU Single Node
<b>Corel Draw</b>	Corel	Professional vector illustration, layout, photo editing and design tools	<ul style="list-style-type: none"> <li>• Faster processing of AI features</li> </ul>	Single GPU Single Node
<b>Corel Photo-Paint</b>	Corel	Corel PHOTO-PAINT is an advanced photo editing software that offers professional editing tools and support for PSD files, plus extensive RAW file support for over 300 types of cameras.	<ul style="list-style-type: none"> <li>• Faster processing of AI features</li> </ul>	Single GPU Single Node
<b>Fresco</b>	Adobe	Powerful painting and drawing app that let you create with realistic watercolors and oils	<ul style="list-style-type: none"> <li>• DirectX</li> </ul>	Single GPU Single Node
<b>JPEG to RAW AI</b>	Topaz Labs	AI powered conversion of JPEG to high-quality RAW for better editing. Prevent banding, remove compression artifacts, recover detail, and enhance dynamic range	<ul style="list-style-type: none"> <li>• GPU accelerated processing</li> </ul>	Single GPU Single Node
<b>Mask AI</b>	Topaz Labs	This is a AI-based masking tool for photography that lets creators automatically detect and remove objects from image.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing</li> </ul>	Single GPU Single Node
<b>Neat Image</b>	Absoft	Reduces noise, film grain, artifacts from photos.	<ul style="list-style-type: none"> <li>• GPU accelerated processing</li> </ul>	Single GPU Single Node
<b>ON1 Photo Raw</b>	ON1	Professional-grade photo organizer, raw processor, layered editor, and effects app, includes everything you need in one photography application.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing</li> </ul>	Single GPU Single Node
<b>PhotoLab</b>	DxO	PhotoLab is a photo editor with specializing in high-quality RAW processing and optical corrections for lens defect, along with powerful local image adjustment tools.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing and AI features</li> </ul>	Single GPU Single Node
<b>Topaz Studio</b>	Topaz Labs	Topaz Studio is an intuitive image effect toolbox with Topaz Labs' powerful acclaimed photo enhancement technology. It works a plugin within Lightroom, Photoshop, Affinity Photo, and others, as well as a standalone editor and host application for your other Topaz plugins.	<ul style="list-style-type: none"> <li>• GPU-accelerated processing</li> </ul>	Single GPU Single Node

## ENCODING AND DIGITAL DISTRIBUTION

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>4K Capture Utility for Windows</b>	ElGato	ElGato sells Capture Cards and offers a capture software with them. The ElGato 4K60 Pro Mk.II capture card includes an implementation of the Video Codec SDK (i.e. NVENC).	<ul style="list-style-type: none"> <li>• HDR recording over HEVC</li> <li>• HDR to SDR conversion</li> </ul>	Single GPU Single Node
<b>Alchemist on Demand</b>	Grass Valley	Video standards conversion	<ul style="list-style-type: none"> <li>• GPU-accelerated video processing 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 processing and encoding</li> </ul>	Single GPU Single Node
<b>Aurora</b>	Tektronix	Automated video quality measurement	<ul style="list-style-type: none"> <li>• CUDA-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</li> <li>• Real-time 4K 360 degree stitching from four camera inputs</li> <li>• Jetson TX-1</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 processing 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>Daniel2</b>	Cinegy	Resolution-independent, CUDA accelerated video codec.	<ul style="list-style-type: none"> <li>• 8K+ video playback faster than real time</li> <li>• 3D LUT color profiles supported</li> <li>• lossless 10-, 12-, 16-bit support</li> <li>• Adobe Premiere Pro plugin</li> </ul>	Single GPU Single Node
<b>Discord Go Live</b>	Discord	Broadcast feature that enables Discord users to broadcast their screen to a Discord channel	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	N/A
<b>DouYu App</b>	DouYu	Douyu's streaming application	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	Single 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</li> <li>• Color correction features and monitoring</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>FAST TICO-RAW</b>	intoPIX	The intoPIX TICO-RAW SDKs provide the highest quality, visually lossless codec for the optimization of your application's infrastructure. FastTICO-RAW SDKs are perfect for all professionals looking to deploy ultra-low latency, lossless RAW encoding over parts of their workflows.	<ul style="list-style-type: none"> <li>• CUDA GPU accelerated up to 10K decoding</li> <li>• Lossless and low latency</li> <li>• All operating systems</li> </ul>	Single GPU Single Node
<b>FAST TICO-XS</b>	intoPIX	The intoPIX FastTICO-XS SDKs provide the highest quality, lowest latency, visually lossless codec for the optimization of your application. FastTICO-XS SDKs are perfect for all professionals looking to deploy ultra-low latency, lossless encoding over their whole infrastructure and workflows.	<ul style="list-style-type: none"> <li>• CUDA GPU accelerated HD, UHD-4K and -8K encoding / decoding</li> <li>• Lossless and low latency</li> <li>• All operating systems</li> <li>• JPEG XS standard compliant</li> </ul>	Single GPU Single Node

<b>Handbrake</b>	Handbrake	HandBrake is an open-source, GPL-licensed, multiplatform, multithreaded video transcoder.	<ul style="list-style-type: none"> <li>• GPU accelerated encoding</li> </ul>	Single GPU Single Node
<b>HuYa App</b>	HuYa	Huya's streaming app	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	Single 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> <li>• Uses NVENC to encode/decode multiple H.264 and HEVC streams</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>Logitech Capture</b>	Logitech	Logitech's app to control their webcam	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	Single 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 encoding</li> </ul>	Single GPU Single Node
<b>Media Transcoding in the Cloud</b>	Ribbon Communications	<p>Industry-leading SBC media transcoding scaling capabilities in virtual and cloud deployments using NVIDIA GPUs to increase performance and decrease cost per transcoded session.</p> <p>Expanded SBC and PSX support for SIP Recording (SIPRec) allows enterprises and call centers to conduct up to four (4) simultaneous recordings of sessions via secure, encrypted technology.</p> <p>Expanded capabilities for Virtual Network Functions (VNF) instantiation with the ability to instantiate Ribbon PSX VNF aligned with the Open Network Automation Platform (ONAP) framework.</p> <p>Enhancements for operational efficiencies that allow CSPs to reduce configuration complexity and improve ease of use.</p> <p>Enhanced security across all products to deliver more restrictive access, reduction in possible network exposure and additional encryption.</p>	<ul style="list-style-type: none"> <li>• Ribbons Session Border Controller Release 7.0 now supports GPUs enabling greater performance and scale for media transcoding, at cost-effective price points, in cloud and virtualized environments.</li> <li>• Ribbons Centralized Policy and Routing (PSX) can be instantiated as a Virtual Network Function (VNF) aligned with the ONAP architecture.</li> <li>• Enterprises now have increased capacity for up to four (4) concurrent SIP Recording (SIPRec) sessions, enabling recorded data to be used for multiple purposes simultaneously such as real-time analytics for call center agents, recordings for corporate compliance and back-up, and lawful intercept</li> <li>• The Insight Element Management System (EMS) has an improved user interface for ease of use and offers improved provisioning and management processes</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>mxfsPEEDRAIL</b>	MOG Technologies	Baseband broadcast news and sports production video ingest product line that allows editing of growing files during ingest.	<ul style="list-style-type: none"> <li>• NVIDIA Video codec used for encoding for higher channel density</li> <li>• CUDA RAW de-coding, de-bayering, and video re-sizing and re-sampling</li> </ul>	Single GPU Single Node
<b>OBS Studio</b>	Open Broadcaster Software	Free and open source software for video recording and live streaming optimized for NVIDIA video encoder	<ul style="list-style-type: none"> <li>• NVENC</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) delivering 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>• De-grain/re-grain to any film look,</li> <li>• De-noise/re-texture 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>Skywatch</b>	MOG Technologies	Video and broadcast production management system for collecting audio/video usage and metadata.	<ul style="list-style-type: none"> <li>• NVIDIA Video codec used for encoding for higher channel density</li> <li>• CUDA RAW de-coding, de-bayering, and video re-sizing and re-sampling</li> </ul>	Single GPU Single Node
<b>Smart Render Editor</b>	Nablet	H.264 and HEVC video encoding using NV Video Codec	<ul style="list-style-type: none"> <li>• Accelerated, high-density video encoding</li> </ul>	Single GPU Single Node
<b>Smart Render SDK</b>	Nablet	Video de-noising, de-interlacing, JPEG 2000 encoding and video fingerprinting	<ul style="list-style-type: none"> <li>• CUDA accelerated video processing</li> <li>• NVIDIA Video codec</li> </ul>	Single GPU Single Node
<b>Speech Quality transformed using Neural Network Computing</b>	BabbleLabs	BabbleLabs has just launched broad production availability of our commercial speech API, web service, and phone mobile apps for iPhone and Android. These services clean up video and audio recordings to make the speech much easier to understand. The apps work on existing videos as well as new audio and video recorded inside the app.	<ul style="list-style-type: none"> <li>• Real time encoding/decoding of audio</li> <li>• Video signals</li> </ul>	Single GPU Single Node
<b>StreamLabs OBS</b>	StreamLabs	Branch of the OBS Studio project that adds a custom UI, integrates plugins, and a plugin store	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	Single GPU Single Node
<b>Tachyon</b>	Cinnafilm	Standards conversion	<ul style="list-style-type: none"> <li>• Video processing and frame rate conversion</li> <li>• Standards conversions and transcoding</li> <li>• SD to UHD, telecine correction, and frame rate normalization</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>Twitch Studio</b>	Twitch.tv	Broadcasting app focused on beginners	<ul style="list-style-type: none"> <li>• NVENC</li> <li>• Multi-video Codec support</li> </ul>	Single 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> <li>• GPU-accelerated noise, grain and dust detection/removal</li> </ul>	Multi-GPU Single Node
<b>Wormhole</b>	Cinnafilm	Time alteration	<ul style="list-style-type: none"> <li>• Retiming and motion compensation,</li> <li>• Super slow motion, and run length adjustment</li> <li>• Commercial insertion, audio retiming, and caption retiming</li> </ul>	Single 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
<b>XSplit Broadcaster</b>	SplitmediaLabs, Ltd.	Broadcast app for recording and streaming, now including a lightweight video editor	<ul style="list-style-type: none"> <li>• NVENC</li> <li>• Record</li> <li>• Stream</li> </ul>	N/A
<b>XSplit Gamecaster</b>	SplitmediaLabs, Ltd.	Simplified broadcast app for recording and streaming, now including a lightweight video editor	<ul style="list-style-type: none"> <li>• NVENC</li> </ul>	Single GPU Single Node

## ON-AIR GRAPHICS

APPLICATION NAME	COMPANYNAME	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> <li>• NVIDIA Video Codec for accelerated encoding and decoding HD and HEVC</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>Camino</b>	AJT Systems	Camino is a powerful 3D rendering system for live-to-air broadcast graphics, capable of up to 4K character generation. Camino's high end features, with excellent ease of use, combine to deliver an exceptional system for your broadcast graphics requirements.	<ul style="list-style-type: none"> <li>• Camino's real-time graphics overlay can be applied to tickertapes, scoreboards, schedule boards, program junctions, and TV show promotions</li> <li>• Graphics overlay may be done via predefined templates, which may then be populated with live data during playout</li> <li>• Makes real-time rendering of data-driven graphics possible in news and sports events.4K, 1080p, 720p and SD Support</li> <li>• NTSC and PAL Support</li> <li>• Graphics, Clips and 3D Objects Importer</li> <li>• 2D and 3D Primitives</li> <li>• Real-Time Key-Frame Animations</li> <li>• Real-Time 3D Scene Lighting</li> <li>• Timeline-Based Audio Support</li> <li>• Data Mapping to External Sources</li> <li>• Transition Logic</li> <li>• Automation Controller Support</li> <li>• Stereoscopic 3D 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>Click Effects PRIME</b>	ChyronHego	Click Effects PRIME is audiovisual content control and delivery solutions for live sports & entertainment productions.	<ul style="list-style-type: none"> <li>• Real-time graphics rendering</li> </ul>	Single GPU Single Node
<b>Cube</b>	Dalet	On-air Graphics	<ul style="list-style-type: none"> <li>• Real-time graphics rendering</li> </ul>	Single GPU Single Node
<b>Designer</b>	Disguise	Designer is the ultimate software to visualize, design, and sequence projects wherever you are, from concept all the way through to showtime.	<ul style="list-style-type: none"> <li>• Real-time graphics rendering</li> <li>• Synchronized video playback</li> <li>• Projection Mapping</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> <li>• RTX accelerated ray-tracing optional Epic Unreal Engine</li> </ul>	Single GPU Single Node
<b>InfinitySet</b>	Brainstorm	Realistic virtual sets	<ul style="list-style-type: none"> <li>• Real-time RTX ray tracing through UE4</li> <li>• HDR I/O</li> <li>• Physically-based rendering</li> <li>• RTX accelerated ray-tracing optional Epic Unreal Engine</li> </ul>	Single GPU Single Node
<b>KAIROS</b>	Panasonic	The IT/IP platform 'KAIROS' is a live video production platform developed based on a new concept and innovative architecture. It incorporates proprietary, ground-breaking software to maximize the CPU and GPU capacities for video processing.	<ul style="list-style-type: none"> <li>• Realtime playout</li> <li>• CUDA and NVEnc</li> <li>• Rivermax SMPTE 2110</li> <li>• GPU Accelerated Video</li> </ul>	Single GPU Single Node
<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>Multiviewers</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>Pixotope</b>	The Future Group	All-in-one, real-time virtual production system with integrated Unreal Engine photorealistic rendering. Open software-based solution for rapidly creating virtual studios, augmented reality (AR), and on-air graphics. Offers a real-time WYSIWYG editor, a virtual set auto-generation tool, its own powerful internal chroma keyer, and user-designed custom control panels.	<ul style="list-style-type: none"> <li>• Real-time rendering</li> <li>• RTX accelerated ray-tracing</li> </ul>	Single GPU Single Node
<b>PRIME</b>	ChyronHego	PRIME Graphics Platform is the next generation of pioneering real-time graphics solutions, helping broadcasters create engaging visuals for all types of programming.	<ul style="list-style-type: none"> <li>• Real-time graphics rendering</li> </ul>	Single GPU Single Node



<b>Reality Engine</b>	Zero Density	Photorealistic virtual studio solution in broadcast industry, powered by Epic Unreal Engine 4.24 Using Mellanox Rivermax API	<ul style="list-style-type: none"> <li>• RTX-accelerated ray-tracing with Unreal Engine</li> <li>• Node-based compositing system designed for real-time production</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	Create elegant video titles or 3D motion graphics.	<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 graphics 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 graphics 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 graphics rendering</li> </ul>	Single GPU Single Node

## ON-SET, REVIEW AND STEREO TOOLS

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>4kScope</b>	Drastic Technologies	4kScope software provides a real time, professional quality signal analysis tool for on set, production, post production, and research and development environments.	<ul style="list-style-type: none"> <li>• GPU accelerated effects and compute</li> </ul>	Single GPU Single Node
<b>8KScope</b>	Drastic Technologies	Real time, professional quality signal analysis tool for on set, production, post production, and research and development environments.	<ul style="list-style-type: none"> <li>• GPU-accelerated effects and compute</li> </ul>	Single GPU Single Node
<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>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>Net-X-Code</b>	Drastic Technologies	Net-X-Code is a distributed capture and conversion system: IP Capture, Control, Convert and Output for server level.	<ul style="list-style-type: none"> <li>• GPU accelerated compute</li> </ul>	Single GPU Single Node
<b>NewBlue Stream</b>	NewBlueFX	NewBlue Stream is a lightweight streaming and broadcast solution paired with dynamic, data-driven graphics	<ul style="list-style-type: none"> <li>• GPU-accelerated processing, encoding and decoding</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 review</li> <li>• NV Video Codec encoding and transcoding</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>VideoQC</b>	Drastic Technologies	videoQC is a suite of video and audio analysis and playback tools with both visual and automated quality checking tools. Takes the media coming into your facility and perform a series of automated tests on video, audio and metadata values against a template, then analyze the audio and video.	<ul style="list-style-type: none"> <li>• GPU accelerated effects and compute</li> </ul>	Single GPU Single Node
----------------	----------------------	--	---	---------------------------

## WEATHER GRAPHICS

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

## Medical Imaging

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>3D Slicer</b>	3D Slicer	3D Slicer is an open source software platform for medical image informatics, image processing, and three-dimensional visualization. Slicer brings free, powerful cross-platform processing tools to physicians, researchers, and the general public.	<ul style="list-style-type: none"> <li>• Multi organ: from head to toe</li> <li>• Support for multi-modality imaging including, MRI, CT, US, nuclear medicine, and microscopy</li> <li>• Bidirectional interface for devices</li> </ul>	Single GPU Single Node
<b>aidoc</b>	Aidoc Medical	AI based decision support software analyzing medical imaging to provide solutions for detecting acute abnormalities across the body, helping radiologists prioritize life threatening cases and expedite patient care. Agnostic to PACS and RIS systems	<ul style="list-style-type: none"> <li>• Classification and segmentation using deep learning on top of any PACS platform</li> </ul>	Single GPU Single Node
<b>AI-LAB</b>	American College of Radiology	ACR AI-LAB offers radiologists tools designed to help them learn the basics of AI and participate directly in the creation, validation and use of health care AI. It accelerates the development and adoption of artificial intelligence (AI) in clinical practice, empowering radiologists to create AI tools at their own institutions, to meet their own patient needs.	<ul style="list-style-type: none"> <li>• AI models for diagnostic imaging</li> <li>• AI models tailored to their local patient population</li> <li>• Patient data protection</li> </ul>	Single GPU Single Node
<b>deepflow</b>	Helmholtz Zentrum München	Deep learning tool for reconstructing cell cycle and disease progression using deep learning from flow cytometry data.	<ul style="list-style-type: none"> <li>• Tool will show that deep convolutional neural networks combined with nonlinear dimension reduction enable reconstructing biological processes based on raw image data</li> <li>• Tool will demonstrate this by reconstructing the cell cycle of Jurkat cells and disease progression in diabetic retinopathy. In further analysis of Jurkat cells</li> <li>• Tool will detect and separate a subpopulation of dead cells in an unsupervised manner and, in classifying discrete cell cycle stages</li> <li>• Tool will reach a sixfold reduction in error rate compared to a recent approach based on boosting on image features. In contrast to previous methods, deep learning based predictions are fast enough for on-the-fly analysis in an imaging flow cytometer</li> <li>• Uses MXNet, cv2, numpy, python3</li> </ul>	Single GPU Single Node

<b>Ibex Decision Support</b>	IBEX	IBEX run DL on prostate cancer digital pathology and to find any potential cancerous areas	<ul style="list-style-type: none"> <li>• Combines data from digitized glass slides and electronic medical records to reveal underlying patterns</li> <li>• Extracts valuable clinical insights that can transform how pathology and oncology are practiced and propel them into the information age</li> </ul>	Single GPU Single Node
<b>iIntuition</b>	Terarecon, Inc.	Intuition offers AI-driven advanced 3D and 4D medical imaging post-processing and visualization.	<ul style="list-style-type: none"> <li>• Volumetric Navigation, CT and MRI Suites</li> <li>• Interventional Radiology</li> <li>• EVAR / TAVR Planning</li> <li>• Body Fusion</li> <li>• Maxillo-Facial</li> <li>• iGENTLE noise reduction</li> <li>• Lung / Liver Segmentation</li> <li>• Mitral Valve (TMVR) Workflow</li> <li>• Lung Density Analysis-II</li> <li>• Intuition AI Adapter</li> <li>• Eureka Clinical AI Platform framework</li> <li>• Explorer UX/UI, and AI algorithm runtime licenses</li> </ul>	Multi-GPU Multi-Node
<b>LVO</b>	Viz.ai	Automatically identify suspected LVOs on CTA imaging in your network and to alert your on-call stroke physician within minutes	<ul style="list-style-type: none"> <li>• Real-Time Specialist Notifications</li> <li>• AI-Powered LVO Detection</li> <li>• Automated Maximum Intensity Projections (MIP)</li> </ul>	Single GPU Single Node
<b>MITK</b>	German Cancer Research Center	Free open-source software system for development of interactive medical image processing software	<ul style="list-style-type: none"> <li>• Interactive segmentation of slices in image volumes, including interactive region growing and easy correction, interpolation of missing slices, surface generation, and volumetry</li> <li>• Point based registration of medical image volumes allows to match two images based on two corresponding sets of points; Rigid registration of images by combination of the ITK registration objects (transforms, optimizers, metrics, etc.)</li> <li>• Measurement of distances and angles; Volume visualization, GPU-based, easy to modify transfer functions; Movie generation (Windows only)</li> <li>• Deformable Registration</li> </ul>	Single GPU Single Node
<b>PowerGrid</b>	University of Illinois Urbana-Champaign	Provides iterative non-cartesian MRI reconstruction	<ul style="list-style-type: none"> <li>• GPU accelerated implementations of the non-Uniform FFT and Discrete Fourier Transform</li> <li>• MPI is used to enable using multiple GPUs in one or several machines</li> <li>• Iterative reconstruction using physics-based model to correct for unwanted effects, such as field inhomogeneity and patient motion</li> </ul>	Multi-GPU Single Node
<b>Proprio</b>	Proprio	Proprio's multi-camera system, based on networked camera array, depth sensing, light field for surgeons to operate and access all the data they need. Offers training based in captured real cases in a safe and collaborative environment.	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	Single 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.)</li> <li>• 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)</li> <li>• Pen tool for freehand drawing</li> </ul>	Single GPU Single Node
<b>Radiology Assist</b>	Zebra Imaging	Receives imaging scans from various modalities and automatically analyzes them for a number of different clinical findings. Findings are provided in real time to radiologists or other physicians and hospital systems as needed.	<ul style="list-style-type: none"> <li>• Classification and segmentation on top of any PACS platform</li> </ul>	Single GPU Single Node
<b>Radlogics Virtual Resident</b>	RadLogics	Software platform imports any DICOM-compatible study directly from the modality or the PACS. The software platform provides APIs for image analysis algorithms to incorporate search, measurement, and other findings into the radiologist existing PACS and reporting system as a preliminary report.	<ul style="list-style-type: none"> <li>• Real time analytics on medical imaging</li> </ul>	Single GPU Single Node
<b>Vitrea®</b>	Vital Images	Vitrea provides advanced visualization tools to a range of medical specialists (including radiologists, cardiologists, oncologists and other specialists) so that they can visualize patient images and communicate with each other efficiently on a course of action. Vitrea is a crucial tool for clinical decision support and enabling physicians to communicate effectively about a common patient, and specialists rely on its detailed 2D, 3D and 4D images for confident analysis in critical scenarios.	<ul style="list-style-type: none"> <li>• Interface designed for viewing in the reading room</li> <li>• Improved clinical outcomes with clinical workflows and partner applications</li> <li>• Increased efficiency with a consistent user interface and experience for all modalities</li> <li>• Easy to deploy thin client solution does not require specialized software to reside on client computers.</li> </ul>	Multi-GPU Multi-Node
<b>XNAT ML</b>	Radiologics	XNAT is an open source imaging informatics platform developed by the Neuroinformatics Research Group at Washington University. It facilitates common management, productivity, and quality assurance tasks for imaging and associated data. XNAT is extensible and can be used to support a wide range of imaging-based projects.	<ul style="list-style-type: none"> <li>• Upload data using DICOM image data and metadata</li> <li>• Organize and share data within user-defined projects securely</li> <li>• Visualize and download using an embedded medical image viewer that supports a number of common medical imaging formats</li> <li>• Secure and manage access to data using a tiered architecture</li> <li>• Search and explore large data sets and create and share customized search patterns</li> <li>• Process data using pipelines that allow for the programming and automation of complex workflows</li> </ul>	Single GPU Single Node
<b>xvision</b>	Augmedics	Augmented reality guidance system for surgery, allows surgeons to see the patient's anatomy through skin and tissue as if they have 'x-ray vision' and to accurately guide instruments and implants during spine procedures	<ul style="list-style-type: none"> <li>• Transparent AR Display</li> <li>• Tracking system</li> </ul>	N/A

# Oil and Gas

APPLICATION NAME	COMPANYNAME	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>	Single GPU Single Node
<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</li> <li>• Scales to thousands of sensors across remote and geographically dispersed locations</li> <li>• 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 and earth modeling.	<ul style="list-style-type: none"> <li>• CUDA acceleration of fault extraction</li> </ul>	Multi-GPU Single Node
<b>Echelon</b>	Stone Ridge Technology	Full featured reservoir simulator designed from inception for GPU (Supported features)	<ul style="list-style-type: none"> <li>• Fully GPU-accelerated reservoir model</li> <li>• Dual-perm, dual porosity, pressure varying perm and porosity</li> <li>• Eclipse compatible input deck</li> </ul>	Multi-GPU Multi-Node
<b>GeoDepth</b>	Emerson	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</li> <li>• 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</li> <li>• 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>	Emerson	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</li> <li>• Pascal/Volta architecture</li> <li>• Multi-GPU</li> </ul>	Multi-GPU Multi-Node
<b>VoxelGeo</b>	Emerson	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

# Life Sciences

## BIOINFORMATICS

APPLICATION NAME	COMPANYNAME	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>AtacWorks</b>	NVIDIA	AtacWorks is a deep learning toolkit for coverage track denoising and peak calling from low-coverage or low-quality ATAC-Seq data.	<ul style="list-style-type: none"> <li>• Coverage track denoising</li> <li>• Retraining</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</li> <li>• 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. Makes 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</li> <li>• Kps-means</li> <li>• K-medoids</li> <li>• K-centers</li> <li>• Hierarchical clustering</li> <li>• Self-organizing map</li> </ul>	Multi-GPU Multi-Node
<b>Clara Genomics Analysis</b>	NVIDIA	Clara Genomics Analysis is a GPU-accelerated library for biological sequence analysis.	<ul style="list-style-type: none"> <li>• CUDA based libraries partial order alignment (cudapoa)</li> <li>• Global aligner (cudaligner)</li> <li>• Mapper (cudamapper)</li> </ul>	Multi-GPU Single 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>f5c</b>	University of New South Wales	An optimised re-implementation of the call-methylation and eventalign modules in Nanopolish. Given a set of basecalled Nanopore reads and the raw signals, f5c call-methylation detects the methylated cytosine and f5c eventalign aligns raw nanopore DNA signals (events) to the base-called read. f5c can optionally utilise NVIDIA graphics cards for acceleration.	<ul style="list-style-type: none"> <li>• Methylated cytosine base and frequency detection</li> <li>• Event alignment</li> </ul>	Single 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>• 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	MUMmer GPU is a 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</li> <li>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</li> <li>Comparable quality results</li> </ul>	Multi-GPU Single Node
<b>Parabricks</b>	NVIDIA	Parabricks provides 30-50 times faster secondary analysis of sequencer generated FASTQ files to variant call files (VCFs). Parabricks has accelerated the standard secondary analyses such as GATK4, Google's Deepvariant to generate equivalent results, while increasing throughput significantly.	<ul style="list-style-type: none"> <li>BWA-mem, Star, haplotype caller, CNVKit, Mutect2, Deep Variant, ImportGVCF, Select Variants, Genotype GVCF, Mark, Sort, BQSR, Merge, VQSR, Variant Filtration, CNNScore, and many quality checking tools.</li> </ul>	Multi-GPU Single Node
<b>PEANUT</b>	Open Source	Read mapper for DNA or RNA sequence that 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</li> <li>Reads 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
<b>Racon</b>	University of Zagreb, Faculty of Electrical Engineering and Computing	Racon is intended as a standalone consensus module to correct raw contigs generated by rapid assembly methods which do not include a consensus step. The goal of Racon is to generate genomic consensus which is of similar or better quality compared to the output generated by assembly methods which employ both error correction and consensus steps, while providing a speedup of several times compared to those methods.	<ul style="list-style-type: none"> <li>It supports data produced by both Pacific Biosciences and Oxford Nanopore Technologies. Racon can be used as a polishing tool after the assembly with either Illumina data or data produced by third generation of sequencing. The type of data inputted is automatically detected. Racon takes as input only three files: contigs in FASTA/FASTQ format, reads in FASTA/FASTQ format and overlaps/alignments between the reads and the contigs in MHAP/PAF/SAM format. Output is a set of polished contigs in FASTA format printed to stdout. All input files can be compressed with gzip (which will have impact on parsing time).</li> <li>Racon can also be used as a read error-correction tool. In this scenario, the MHAP/PAF/SAM file needs to contain pairwise overlaps between reads including dual overlaps.</li> </ul>	Single GPU Single Node
<b>racon-gpu</b>	Open Source	Racon is intended as a standalone consensus module to correct raw contigs generated by rapid assembly methods which do not include a consensus step. The goal of Racon is to generate genomic consensus which is of similar or better quality compared to the output generated by assembly methods which employ both error correction and consensus steps, while providing a speedup of several times compared to those methods. It supports data produced by both Pacific Biosciences and Oxford Nanopore Technologies.	<ul style="list-style-type: none"> <li>Racon can be used as a polishing tool after the assembly with either Illumina data or data produced by third generation of sequencing</li> <li>The type of data inputted is automatically detected.</li> <li>Racon takes as input only three files: contigs in FASTA/FASTQ format, reads in FASTA/FASTQ format and overlaps/alignments between the reads and the contigs in MHAP/PAF/SAM format. Output is a set of polished contigs in FASTA format printed to stdout. All input files can be compressed with gzip (which will have impact on parsing time).</li> <li>Racon can also be used as a read error-correction tool. In this scenario, the MHAP/PAF/SAM file needs to contain pairwise overlaps between reads including dual overlaps.</li> </ul>	Single GPU Single Node

<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</li> <li>• REML analysis</li> <li>• Regional Heritability (including multi-GPU)</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>Shasta</b>	University of California Santa Cruz	Shasta long read assembler is to rapidly produce accurate assembled sequence using as input DNA reads generated by Oxford Nanopore flow cells.	<ul style="list-style-type: none"> <li>• Uses a run-length representation of the read sequence to make the assembly process more resilient to errors in homopolymer repeat counts, which are the most common type of errors in Oxford Nanopore reads.</li> <li>• Using in some phases of the computation a representation of the read sequence based on markers, a fixed subset of short k-mers (k = 10)</li> </ul>	Single GPU Single Node
<b>SOAP3</b>	Genomics	GPU-based software for aligning short reads with a reference sequence. Finds 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</li> <li>• Reports all answers</li> </ul>	Multi-GPU Multi-Node
<b>SOAP3-dp</b>	The University of Hong Kong	SOAP3-dp is an ultra-fast GPU-based tool for short read alignment via index-assisted dynamic programming.	<ul style="list-style-type: none"> <li>• Borrows-Wheeler Transformation</li> <li>• Dynamic Programming</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	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ANNA-PALM</b>	Institut Pasteur	Accelerating Single Molecule Localization Microscopy with Deep Learning: ANNA-PALM is a computational method that can reconstruct super-resolution images from sparse single molecule localization data and/or widefield images. ANNA-PALM can produce high quality super-resolution images from data obtained in much shorter acquisition time than standard single molecule localization microscopy. By strongly reducing acquisition time, ANNA-PALM facilitates super-resolution imaging of large numbers of cells (high throughput imaging), large samples, and live cells.	<ul style="list-style-type: none"> <li>• Uses a much smaller number of low resolution frames than other methods</li> <li>• Processing by localization algorithms results in a sparse localization image using a neural network previously trained on conventional PALM images</li> <li>• Inputs sparse image and outputs a super-resolution image</li> <li>• Runs well on GPU due to acceleration available in Tensorflow</li> </ul>	Single GPU Single Node
<b>Appion</b>	New York Structural Biology Center	Appion is a “pipeline” for processing and analysis of EM images. Appion is integrated with Leginon data acquisition but can also be used stand-alone after uploading images (either digital or scanned micrographs) or particle stacks using a set of provided tools. Appion consists of a web based user interface linked to a set of python scripts that control several underlying integrated processing packages. All data input and output within Appion is managed using tightly integrated SQL databases. The goal is to have all control of the processing pipeline managed from a web based user interface and all output from the processing presented using web based viewing tools.	<ul style="list-style-type: none"> <li>• The underlying packages integrated into Appion include MotionCor2, Gctf, EMAN, Spider, Frealign, Imagic, XMIPP, IMOD, ProTomo, ACE, CTFFind and CTFTilt, findEM, DogPicker, TiltPicker, RMeasure, EM-BFACTOR, and Chimera.</li> </ul>	Single GPU Single Node
<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>crYOLO</b>	Max Planck Institute for Molecular Physiology	Novel automated particle picking software based on the deep learning object detection system ‘You Only Look Once’ (YOLO). CrYOLO is available as standalone program under <a href="http://sphire.mpg.de/">http://sphire.mpg.de/</a> and will be part of the image processing workflow in SPHIRE.	<ul style="list-style-type: none"> <li>• Part of the image processing workflow in SPHIRE.</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</li> <li>• Heterogeneous reconstruction</li> <li>• High-speed and high resolution refinement of 3D protein structures implemented on GPUs</li> <li>• Multiple simultaneous jobs on multiple GPUs</li> </ul>	Multi-GPU Multi-Node

<b>Dynamo</b>	Center for Cellular Imaging and Nano Analytics (C-CINA), Biozentrum, University of Basel	Dynamo is a software environment for subtomogram averaging of cryo-EM data.	<ul style="list-style-type: none"> <li>• Dynamo provides workflows all the way from tomograms to averages and classes.</li> <li>• In a full workflow, you would organize tomograms in catalogues, use them to pick particles and create alignment and classification projects to be run on different computing environments</li> <li>• Requires CUDA Toolkit of version 7.5 or higher and CUDA driver compatible with your actual GPU device</li> </ul>	Single GPU Single Node
<b>EMAN2</b>	Baylor College of Medicine	EMAN2 is the successor to EMAN1. It is a broadly based greyscale scientific image processing suite with a primary focus on processing data from transmission electron microscopes. EMAN's original purpose was performing single particle reconstructions (3-D volumetric models from 2-D cryo-EM images) at the highest possible resolution, but the suite now also offers support for single particle cryo-ET, and tools useful in many other subdisciplines such as helical reconstruction, 2-D crystallography and whole-cell tomography. Image processing in a suite like EMAN differs from consumer image processing packages like Photoshop in that pixels in images are represented as floating-point numbers rather than small (8-16 bit) integers. In addition, image compression is avoided entirely, and there is a focus on quantitative analysis rather than qualitative image display.	<ul style="list-style-type: none"> <li>• All EMAN2 programs, including GUI programs, are written in the easy-to-learn Python scripting language. This permits knowledgeable end-users to customize any of the code with unprecedented ease. If you aren't an advanced user, you can still make use of the integrated GUI and all of EMAN2's command-line programs.</li> </ul>	Single GPU Single Node
<b>emClarity</b>	Benjamin Himes	emClarity is a collection of gpu accelerated software developed to enable determination of biological structures at resolutions better than 1nm from heterogeneous specimen imaged by cryo-Electron Tomography.	<ul style="list-style-type: none"> <li>• Subtomogram averaging</li> <li>• Very high resolution single particle analysis</li> <li>• Hybrid electron microscopy.</li> </ul>	Multi-GPU Single Node
<b>Gautomatch</b>	MRC Laboratory of Molecular Biology	Gautomatch is a GPU accelerated program for accurate, fast, flexible and fully automatic particle picking from cryo-EM micrographs with or without templates.	<ul style="list-style-type: none"> <li>• Fast: typically, 1.5~2.0s with 15 templates, using a good GPU (e.g. GTX 980, Titan X)</li> <li>• Fully automatic with simple command on entire data sets</li> <li>• Convenient and easy to use</li> <li>• Flexible: with or without template, suitable for both basic or advanced users</li> <li>• Compatible with Relion/EMAN</li> <li>• Background correction: automatic correct the gradient background that affects the picking</li> <li>• Rejection of ice/carbon: automatically detect non-particle areas and reject them</li> <li>• Post-optimization: scripts available to re-filter the coordinates after picking within seconds</li> <li>• Accuracy: the user's satisfaction is the only 'gold standard' criterion</li> </ul>	Single GPU Single Node
<b>GCTF</b>	MRC Laboratory of Molecular Biology	Corrects contrast transfer function effects in electron microscope optics	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	Single 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>IMOD</b>	University of Colorado	IMOD is a set of image processing, modeling and display programs used for tomographic reconstruction and for 3D reconstruction of EM serial sections and optical sections. Contains tools for assembling and aligning data within multiple types and sizes of image stacks, viewing 3-D data from any orientation, and modeling and display of the image files.	<ul style="list-style-type: none"> <li>• ctfphaseflip : Corrects tilt series for microscope CTF by phase flipping</li> <li>• gputilttest : Test whether a GPU is reliable for computing reconstructions with the tilt program</li> <li>• 3dmod : Model editing and image display program. 3dmod can display three-dimensional graphic data sets in many views simultaneously, can model these data sets, and can display models and graphic data in 3-D. The views include a slice through the 3D volume, a projection of a sub-volume and orthogonal views with contour overlays.</li> <li>• xyzproj : Project 3-dimensional data at a series of tilts around the X, Y, or Z axis.</li> </ul>	Single GPU Single Node
<b>ITK</b>	Kitware	The National Library of Medicine Insight Segmentation and Registration Toolkit (ITK), or Insight Toolkit, is an open-source, cross-platform C++ toolkit for segmentation and registration. Segmentation is the process of identifying and classifying data found in a digitally sampled representation. Typically the sampled representation is an image acquired from such medical instrumentation as CT or MRI scanners. Registration is the task of aligning or developing correspondences between data. For example, in the medical environment, a CT scan may be aligned with a MRI scan in order to combine the information contained in both.	<ul style="list-style-type: none"> <li>• Library is used by Paraview, VTK, and many other software distributions</li> <li>• Many capabilities for multi-dimensional image processing and extraction tools</li> <li>• Most recent GPU acceleration of FFTs using cuFFT (cuFFTW) and matrix math accelerated through CUDA enabled Eigen3</li> </ul>	Single GPU Single Node
<b>Leginon</b>	New York Structural Biology Center	Leginon is a system designed for automated collection of images from a transmission electron microscope.	<ul style="list-style-type: none"> <li>• A Leginon application is image acquisition process that is built of several smaller pieces called 'nodes'</li> <li>• Nodes can be applications</li> <li>• Some of these are GPU accelerated applications such as Topaz, Relion, and MotionCor2</li> </ul>	Single 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>MotionCor2</b>	UCSF	A multi-GPU program that corrects beam-induced sample motion on dose fractionated movie stacks. Implements a robust iterative alignment algorithm that delivers precise measurement and correction of both global and non-uniform local motions at single pixel level across the whole frame. Suitable for both single-particle and tomographic images.	<ul style="list-style-type: none"> <li>• Overall, MotionCor2 is extremely robust, and sufficiently accurate at correcting local motions so that the very time-consuming and computationally-intensive particle polishing in RELION can be skipped. Importantly</li> <li>• Works on a wide range of data sets including cryo tomographic tilt series</li> </ul>	Multi-GPU Single Node

<b>PSSR</b>	Waitt Advanced Biophotonics Center Core	Deep Learning-Based Point-Scanning Super-Resolution Imaging allows point-scanning super-resolution (PSSR) imaging and facilitates point-scanning image acquisition with otherwise unattainable resolution, speed, and sensitivity.	<ul style="list-style-type: none"> <li>• Pre-trained models for</li> <li>• PSSR for Electron Microscopy (EM)</li> <li>• PSSR single frame (PSSR-SF) for mitoTracker</li> <li>• PSSR multiframe (PSSR-MF) for mitoTracker</li> <li>• PSSR for neuronal mitochondria</li> </ul>	Single GPU Single Node
<b>RELION</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>• 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
<b>Thunder</b>	Tsinghua University	THUNDER is a particle-filter algorithm based cryoEM image processing software for using THUNDER to analysis cryoEM images in purpose of achieving a 3D model.	<ul style="list-style-type: none"> <li>• Both image classification and highresolution 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 Multi-Node
<b>Tomviz</b>	Kitware	Tomviz enables 3D characterization of materials at the nano- and meso-scale, tailored for visualizing electron tomography data. It utilizes the large quantities of memory and processing resources required to render, manipulate, and analyze voluminous 3D tomograms.	<ul style="list-style-type: none"> <li>• 3D tomographic data processing, visualization, and analysis of</li> <li>• Python</li> <li>• Windows</li> <li>• Mac OS</li> <li>• Linux</li> </ul>	Single GPU Single Node
<b>Topaz</b>	Tristan Bepler	A pipeline for particle detection in cryo-electron microscopy images using convolutional neural networks trained from positive and unlabeled examples.	<ul style="list-style-type: none"> <li>• Deep learning for cryo EM data particle picking</li> <li>• Uses CUDA and pytorch</li> </ul>	Single GPU Single Node
<b>Warp</b>	Max Planck Institute for Biophysical Chemistry	Warp integrates novel algorithms for frame alignment, defocus estimation, particle picking and tomographic reconstruction in a rich user interface. Enables data quality monitoring in real time, data analysis at microscope level and obtains high-resolution structures before data collection is over.	<ul style="list-style-type: none"> <li>• CUDA enabled processing for electron microscopy</li> <li>• TensorFlow (v1.10)</li> <li>• CUDA kernels: backprojection, CTF, deconvolution, FFT, tomography refinement, and others</li> </ul>	Single GPU Single Node



## MOLECULAR DYNAMICS

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ACEMD</b>	Acellera 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>Colvars</b>	Temple University	<p>Software module for molecular simulation and analysis that provides a high-performance implementation of sampling algorithms defined on a reduced space of continuously differentiable functions (aka collective variables)</p> <p>The module itself implements a variety of functions and algorithms, including free-energy estimators based on thermodynamic forces, non-equilibrium work and probability distributions</p>	<ul style="list-style-type: none"> <li>• LAMMPS, NAMD, VMD</li> <li>• GPU support</li> </ul>	Multi-GPU Multi-Node
<b>Computational Crystallography Toolbox</b>	Lawrence Berkeley Laboratories	Open source component of the PHENIX system to advance automation of macromolecular structure determination. Useful for small-molecule crystallography and even general scientific applications	<ul style="list-style-type: none"> <li>• GPU acceleration for scattering and general purpose math via</li> <li>• CUDA and cuFFT</li> </ul>	Multi-GPU Single Node
<b>DeePMD-kit</b>	Princeton University	DeePMD-kit is a package written in Python/C++, designed to minimize the effort required to build deep learning based model of interatomic potential energy and force field and to perform molecular dynamics (MD). Addresses the accuracy-versus-efficiency dilemma in molecular simulations. Applications of DeePMD-kit span from finite molecules to extended systems and from metallic systems to chemically bonded systems.	<ul style="list-style-type: none"> <li>• TensorFlow</li> <li>• High-performance classical MD and quantum (path-integral) MD packages</li> <li>• Deep Potential series models</li> <li>• MPI and GPU support</li> </ul>	Multi-GPU Single Node
<b>DeepSite</b>	Acellera Ltd	DeepSite is a protein binding pocket predictor based on deep neural networks. Allows you to upload your structure on PDB format, monitor the progress of your job and visualize the results with our modern WebGL viewer.	<ul style="list-style-type: none"> <li>• Deep learning</li> <li>• Machine learning</li> <li>• Drug discovery in a web interface</li> </ul>	Single 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>FEP+</b>	Schrodinger, Inc.	Molecular Dynamics (MD) and Free Energy Perturbation (FEP) calculations occur on time scales that are computationally demanding to simulate. A key factor in determining whether a simulation will take days, hours, or minutes to run is the hardware being used. The advent of GPU computing, however, has opened the door to a new world of computationally intensive simulations that would not have been possible even a few years ago. Desmond's high-performance Molecular Dynamics code, together with continuously improving computer hardware technologies are helping scientists push the boundaries of discovery further than ever before. MD simulations to impact drug discovery has now been attained in FEP+, due to the confluence of hardware and software development along with the formulation of sufficiently accurate theoretical methods and models	<ul style="list-style-type: none"> <li>• Optimization of the FEP+ algorithm to take full advantage of the Desmond GPU MD engine enabling 2 to 4 ligands to be scored per day on a multi-GPU server.</li> </ul>	Multi-GPU Multi-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>Galamost</b>	CAS-CIAC	GALAMOST is a project of employing high-performance computational techniques to accelerate molecular simulation by fully utilizing the computational power of NVIDIA GPUs. Enables the investigation of polymeric systems in a large temporal and spatial scale at a very low cost.	<ul style="list-style-type: none"> <li>• Full Molecular Simulation on GPU</li> </ul>	Multi-GPU Multi-Node
<b>GALAMOST</b>	ChangChun CHINA	GALAMOST is a package of employing high-performance computational techniques on many-core processors to accelerate molecular dynamics simulations. The package is written with CUDA and C++ languages for particularly running on NVIDIA GPUs and focuses on the large scale simulations of soft matters.	<ul style="list-style-type: none"> <li>• General molecular dynamics</li> <li>• Dissipative particle dynamics (DPD)</li> <li>• Brownian dynamics (BD)</li> <li>• Coarse-graining molecular dynamics (CGMD)</li> <li>• Reaction model</li> <li>• Anisotropic particle models</li> <li>• MD-SCF</li> <li>• DNA 3SPN model</li> <li>• Rigid body method</li> <li>• Stretching method</li> </ul>	Single 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>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>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>	KTH Royal Institute of Technology	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>• ACEMD</li> <li>• PMEMD</li> <li>• NAMD</li> <li>• GROMACS</li> <li>• AMBER</li> <li>• 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</li> <li>• Gay-Berne</li> <li>• Tersoff</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>MOLECULAR OPERATING ENVIRONMENT</b>	Chemical Computing Group ULC	Calculate and Analyze pH-Dependent Protein Properties. MOEsaic Session Sharing and Project Customization. Determine Conformation Population from NMR NOE Data  Predict Relative Binding Energies with AMBER Thermodynamic Integration.	<ul style="list-style-type: none"> <li>• GPU Accelerated 3D Stereo Graphics</li> <li>• AMBER GPU accelerated support</li> </ul>	Single 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</li> <li>• 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</li> <li>• 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</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 for the Self Organized Polymer Model fully implemented on a GPU. 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</li> <li>• Multiple simulation trajectories can be performed simultaneously on a single GPU</li> <li>• Calpha and Calpha-Cbeta models</li> <li>• Simulations of protein forced unfolding</li> <li>• Novel simulations of nanoindentation in silico</li> <li>• Support for hydrodynamic interactions</li> <li>• Up to ~100 ms of simulation time per day,</li> <li>• 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	COMPANYNAME	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 4</b>	University of Florida	New SIA/aces4 development A new super instruction architecture with interface applications for quantum chemistry (aces4).	<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>ACES III</b>	University of Florida	ACES III takes the 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>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> </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>• Daubechies wavelets</li> </ul>	Multi-GPU Multi-Node
<b>BrianQC</b>	StreamNovation Ltd.	BrianQC is a software product in the field of quantum chemistry. It accelerates features of Q-Chem 5.0 or later. Optimized for simulating large molecules and tested up to 20,000 Cartesian Gaussian basis functions. Has full support of s, p, d, f and g-type orbitals. Full support for NVIDIA GPU architectures (Kepler, Maxwell, Pascal, Volta) with double precision accuracy on 64-bit Linux operation systems. Targets the speeds up of Q-Chem for every calculation that uses Coulomb or Exchange integrals over Gaussian basis functions or their first analytic derivative (including HF-SCF, DFT, SCF geom. opt, DFT geom. opt for most functionals, etc.)	<ul style="list-style-type: none"> <li>• The range of NVIDIA architectures supported by BrianQC has been expanded. In addition to GPUs powered by Kepler, Maxwell and Pascal, BrianQC now supports NVIDIA Tesla V100 GPU as well</li> <li>• Compatible with features of Q-Chem 5.0 or later</li> <li>• Optimized for simulating large molecules</li> <li>• Tested up to 20,000 Cartesian Gaussian basis functions</li> <li>• Full support of s, p, d, f and g-type orbitals</li> <li>• Full support for NVIDIA GPU architectures (Kepler, Maxwell, Pascal). Double precision accuracy</li> <li>• Runs on 64-bit Linux operation systems</li> <li>• Speeds up Q-Chem for every calculation that uses Coulomb or Exchange integrals over Gaussian basis functions or their first analytic derivative (including HF-SCF, DFT, SCF geom. opt, DFT geom. opt for most functionals, etc.)</li> </ul>	Multi-GPU Single 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>• (sslss) 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</li> <li>• Hartree-Fock</li> <li>• 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</li> <li>• 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</li> <li>• SP2 Algorithm</li> </ul>	Multi-GPU Single Node
<b>libxc</b>	TDDFT	Libxc is a library of exchange-correlation functionals for density-functional theory providing portable, well tested and reliable set of exchange and correlation functionals that can be used by all the ETSF codes and also other codes	<ul style="list-style-type: none"> <li>• GPU acceleration for quantum chemistry</li> <li>• LDA, GGA, hybrids and mGGA</li> <li>• Python 3 and C interfaces</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 capabilitiesTensor 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>MAPS</b>	Scienomics	MAPS CLASSICAL & MESOSCALE simulation toolkit contains world-class simulation engines such as LAMMPS, CHAMELEON, TOWHEE, NAMD. Includes a collection of ready-to-use workflows and a rich Force-Field library.	<ul style="list-style-type: none"> <li>• Typical calculations that can be executed include molecular dynamics simulations and Monte Carlo simulations, structure relaxation in periodic or molecular systems using both classical and quantum mechanics tools</li> <li>• Trajectory can be generated and then later analyzed using the appropriate tools</li> <li>• Additional simulations can be performed using PC-SAFT and related methods for thermodynamics modeling</li> </ul>	Single 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</li> <li>• Full diagonalization</li> <li>• 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.	<ul style="list-style-type: none"> <li>• Triples part of Reg-CCSD(T)</li> <li>• CCSD and EOMCCSD task schedulers</li> </ul>	Multi-GPU Single Node

<b>NWChemEX</b>	Pacific Northwest National Laboratories	NWChemEx targets developing high-performance computational models for the production of advanced biofuels and other bioproducts	<ul style="list-style-type: none"> <li>• GPU acceleration</li> <li>• libraries like libxc</li> </ul>	Single 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</li> <li>• Orthogonalization</li> <li>• Subspace diagonalization</li> <li>• Poisson solver</li> <li>• Time propagation</li> <li>• 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>QBox</b>	University of California Davis	Qbox is a C++/MPI scalable parallel implementation of first-principles molecular dynamics (FPMD) based on the plane-wave, pseudopotential formalism. Designed for operation on large parallel computers.	<ul style="list-style-type: none"> <li>• The availability of double precision graphics cards provides an opportunity to speed up electronic structure computations. We modify the Qbox code to utilize Fermi GPUs on the Keeneland platform</li> <li>• We use the CUFFT library to speed up Fourier transforms and perform asynchronous communication to cut down the cost of data transfers</li> <li>• The modified code is used in simulations of a 64-molecule water system with an 85 Ry plane wave energy cut off</li> <li>• Preliminary results show a 2-3 times speedup in the calculation of the charge density and in the application of the Hamiltonian operator to the wave function</li> <li>• We present these findings as well as further speedups measured in other parts of the code. <a href="http://eslab.ucdavis.edu/software/qbox">http://eslab.ucdavis.edu/software/qbox</a> <a href="http://keeneland.gatech.edu">http://keeneland.gatech.edu</a></li> </ul>	Single 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>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</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 (matix multiply)</li> <li>• Explicit computational kernels</li> <li>• 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</li> <li>• Supports s, p, d, f orbitals on energy calculation</li> <li>• HF gradient with s,p,d orbital support</li> <li>• GPU-based ERI generator</li> </ul>	Multi-GPU Single Node
<b>RESCU</b>	Hongzhiwei technology	RESCU is a KS-DFT calculation software that can study very large systems with only a small computer. Offers new, extremely powerful and parallel high efficiency KS-DFT self-consistent calculation method.	<ul style="list-style-type: none"> <li>• Parallel high efficiency processing-KS-DFT</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 and runs 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</li> <li>• Multipetaflops capable</li> <li>• Handles thousands of atoms with full DFT precision</li> <li>• Supports multiple GPUs per node</li> <li>• Fully open source</li> <li>• Installation support</li> <li>• Cray XE6/XK7</li> </ul>	Multi-GPU Single Node
<b>TAL-SH</b>	Oak Ridge National Lab	Tensor Algebra Library Routines for Shared Memory Systems accelerates three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	<ul style="list-style-type: none"> <li>• Tensor Algebra Library for Shared Memory Computers: Nodes equipped with multicore CPU, NVIDIA GPU, and Intel Xeon Phi (in progress)</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)</li> <li>• RMM-DIIS (ALGO = VERYFAST &amp; FAST)</li> <li>• K-Points and optimization for critical step in exact exchange calculations</li> </ul>	Multi-GPU Multi-Node

## (MOLECULAR) VISUALIZATION AND DOCKING

APPLICATION NAME	COMPANYNAME	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>AUTODOCK</b>	Scripps	The AutoDock Suite is a growing collection of methods for computational docking and virtual screening, for use in structure-based drug discovery and exploration of the basic mechanisms of biomolecular structure and function.	<ul style="list-style-type: none"> <li>• OpenCL-accelerated version of AutoDock4.2.6</li> <li>• AutoDock GPU</li> <li>• ADADELTA</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>	OpenEye 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>• High quality images and ease of interaction</li> <li>• Latest GPU computing acceleration techniques</li> <li>• Natural user interfaces such as Kinect and Wiimotes</li> </ul>	Single GPU Single Node
<b>MEGADOCK</b>	Akiyama_ Laboratory, Tokyo Institute of Technology	MEGADOCK is a fast protein-protein docking software when more acceleration is demanded for an interactome prediction, which is composed of millions of protein pairs.	<ul style="list-style-type: none"> <li>• MEGADOCK-GPU on 12 CPU cores</li> <li>• 3 GPU calculation speed 37.0 times faster than MEGADOCK on 1 CPU core</li> <li>• Novel docking software facilitating the application of docking techniques to assist large-scale protein interaction network analyses</li> </ul>	Multi-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</li> <li>• Pose evaluation</li> <li>• 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</li> <li>• Molecular surface values</li> </ul>	Single GPU Single Node
<b>VMD</b>	University of Illinois	Visualization and analyzation of large bio-molecular systems in 3-D graphics.	<ul style="list-style-type: none"> <li>• High quality rendering</li> <li>• Large structures (100M atoms)</li> <li>• Analysis and visualization tasks</li> <li>• Multiple GPU support for display of molecular orbitals</li> </ul>	Multi-GPU Single Node

## Research: Higher Education and Supercomputing

### NUMERICAL ANALYTICS

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>ArrayFire HPC</b>	ArrayFire	ArrayFire is a software development and consulting company with a passion for helping organizations develop high-performance computing solutions on modern computational platforms. Our core areas of expertise drive innovation in all areas of technical computing. We have extensive experience in CUDA and OpenCL programming, code acceleration and optimization, and software design. We also have specialized domain expertise in machine learning and computer vision. Our customers range from startups to Fortune 500 companies in a variety of industries, including public sector, finance, and media, and include government and academic research institutions.	<ul style="list-style-type: none"> <li>• ArrayFire contains hundreds of functions across various domains including:</li> <li>• Vector Algorithms</li> <li>• Image Processing</li> <li>• Computer Vision</li> <li>• Signal Processing</li> <li>• Linear Algebra</li> <li>• Statistics</li> <li>• and more...</li> </ul>	Multi-GPU Single Node
<b>Eigen</b>	Eigen	Eigen is a C++ template library for linear algebra: matrices, vectors, numerical solvers, and related algorithms.	<ul style="list-style-type: none"> <li>• CUDA enabled linear algebra</li> <li>• Eigen solver, reduction, random, etc.</li> </ul>	Single GPU Single Node
<b>Julia</b>	Julia Computing	Julia delivers dramatic improvements in simplicity, speed, scalability, capacity, and productivity to solve massive computational problems quickly and accurately, making it the preferred language for big data analytics.	<ul style="list-style-type: none"> <li>• Full support/integration of NVIDIA CUDA via Julia CUDA JIT plugin architecture</li> <li>• Free and open source (MIT licensed)</li> <li>• User-defined types are as fast and compact as built-ins</li> <li>• No need to vectorize code for performance; devectorized code is fast</li> <li>• Designed for parallelism and distributed computation</li> <li>• Lightweight "green" threading (coroutines)</li> <li>• Unobtrusive yet powerful type system</li> <li>• Elegant and extensible conversions and promotions for numeric and other types</li> <li>• Efficient support for Unicode, including but not limited to UTF-8</li> <li>• Call C functions directly (no wrappers or special APIs needed)</li> <li>• Powerful shell-like capabilities for managing other processes</li> <li>• Lisp-like macros and other metaprogramming facilities</li> </ul>	Multi-GPU Multi-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>MATLAB</b>	Mathworks	GPU acceleration for MATLAB (high-level technical computing language).	<ul style="list-style-type: none"> <li>• Acceleration for 200+ most used MATLAB functions</li> <li>• Acceleration of more than 500 most parallelizable MATLAB functions</li> <li>• Accelerated Signal Processing toolkit</li> <li>• Accelerated Image Processing toolkit</li> <li>• Accelerated Communications Systems toolkit</li> <li>• Available via an NGC container</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	COMPANYNAME	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.	<ul style="list-style-type: none"> <li>• Wilson-clover fermion linear solver</li> </ul>	Multi-GPU Single Node
<b>CADISHI</b>	Max Planck Institute	CADISHI is a software package that enables scientists to compute (Euclidean) distance histograms efficiently. Any sets of objects that have 3D Cartesian coordinates may be used as input, for example, atoms in molecular dynamics datasets or galaxies in astrophysical contexts.	<ul style="list-style-type: none"> <li>• Highly tuned CPU and GPU kernels</li> <li>• Python engine for throughput computing</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).	<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 and performs 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</li> <li>• Krylov solvers</li> <li>• 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 is 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>• CUDA is supported</li> <li>• 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>	Open Source	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</li> <li>• 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>GENE</b>	GENE	GENE (Gyrokinetic Electromagnetic Numerical Experiment) is an open source plasma microturbulence code which can be used to efficiently compute gyroradius-scale fluctuations and the resulting transport coefficients in magnetized fusion/astrophysical plasmas.	<ul style="list-style-type: none"> <li>• Basic Modeling</li> </ul>	Multi-GPU Multi-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>• Calculates average network density and velocity</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>• Calculates average network density and velocity</li> </ul>	Single GPU Single Node

<b>GTC</b>	University of California Irvine(UC Irvine)	Gyrokinetic Plasma Fusion for Modeling a Tokamak reactor	<ul style="list-style-type: none"> <li>• NVLINK</li> </ul>	Multi-GPU Multi-Node
<b>GTC Irvine</b>	University of California Irvine(UC Irvine)	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.	<ul style="list-style-type: none"> <li>• PUSHe, Collision and Poisson Solver</li> </ul>	Multi-GPU Multi-Node
<b>GTC-P</b>	Princeton Plasma Physics Lab	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.	<ul style="list-style-type: none"> <li>• Optimized with CUDA</li> <li>• OpenACC development underway</li> </ul>	Multi-GPU Single Node
<b>HACC</b>	HACC	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.	<ul style="list-style-type: none"> <li>• This code has been optimized with CUDA runs in full production mode</li> </ul>	Multi-GPU Single Node
<b>HAMR GPU</b>	HAMR	GPU accelerated General Relativistic Magneto Hydrodynamic application	<ul style="list-style-type: none"> <li>• Active galactic nuclei which assumes a radiatively inefficient sub-eddington rate torus</li> <li>• Axisymmetric ideal MHD</li> <li>• Viscosity and resistivity through use of Riemann solver (HLL)</li> <li>• Density floors to mass load the jet</li> <li>• Uses grids that can resolve the substructure of the jet over 5 orders of magnitude</li> </ul>	Multi-GPU Single Node
<b>MAESTRO</b>	MAESTRO	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.	<ul style="list-style-type: none"> <li>• Gravitational Field Solver</li> </ul>	Multi-GPU Single Node

<b>MILC</b>	USQCD	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.	<ul style="list-style-type: none"> <li>• Staggered fermions</li> <li>• Krylov solvers</li> <li>• Gauge-link fattening</li> </ul>	Multi-GPU Multi-Node
<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.	<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>ORB5</b>	EPFL	ORB5 is a global, gyrokinetic, Lagrangian, Particle-In-Cell (PIC), finite element, electromagnetic model	<ul style="list-style-type: none"> <li>• Plasma and background magnetic geometry</li> <li>• Axisymmetric ideal MHD equilibria</li> <li>• Computed with CHEASE code [9] kinetic electrons, or various approximate models: hybrid-trapped or adiabatic intra- and inter-species linearized collision operators electromagnetic perturbations, with the cancellation problem solved using enhanced control variates and a 'pullback' scheme</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</li> <li>• 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-particle acceleration and relativistic plasma physics</li> </ul>	Multi-GPU Multi-Node
<b>PPM</b>	PPM	Piecewise parabolic method is 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</li> <li>• Radiative transfer for reionization</li> <li>• Hydrodynamic solver using AMR</li> </ul>	Multi-GPU Multi-Node
<b>samadii/sciv</b>	Metariver Technology	Software for computing flow field in high vacuum condition using the DSMC(Direct Simulation with Monte Carlo) method. Simulating the interactions between gas and surfaces boundaries, the gas flow with molecular particles	<ul style="list-style-type: none"> <li>• DSMC simulator, gas dynamics solver</li> <li>• OLED &amp; Semiconductor deposition and etching analysis, Vacuum field analysis</li> <li>• PDL(Pixel Define Layer) growth analysis</li> <li>• Deposition mask toolkits, Wall growth, Chemical reaction</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	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<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>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>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>Inside Explorer</b>	Interspectral	An interactive and intuitive software with volumetric rendering and 3D-visualization of real captured data.	<ul style="list-style-type: none"> <li>• vGPU</li> </ul>	Single GPU Single 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>Pix4Dmapper</b>	Pix4D	This professional photogrammetry software uses images to generate point clouds, digital surface and terrain models, orthomosaics, textured models and more. It is most often used by geospatial professionals such as surveyors and civil engineers.	<ul style="list-style-type: none"> <li>• GPU accelerated processing</li> </ul>	Single GPU Single Node
<b>SPECFEM3D</b>	CIG	<p>There are two modules/apss in the SPECFEM family: GLOBE and CARTESIAN.</p> <p>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</p>	<ul style="list-style-type: none"> <li>• OpenCL and CUDA hardware accelerators, based on an automatic source-to-source transformation library</li> <li>• Simulates acoustic (fluid), elastic (solid), coupled acoustic/elastic, poroelastic or seismic wave propagation in any type of conforming mesh of hexahedra (structured or not).</li> </ul>	Multi-GPU Single Node
<b>Tecplot</b>	Tecplot	General purpose scientific visualization software for Aerodynamics, O&G, Internal Combustion and Geoscience applications	<ul style="list-style-type: none"> <li>• Rendering</li> </ul>	Single GPU Single Node
<b>VisIt</b>	LLNL	Scalable data anlysis and visualization application	<ul style="list-style-type: none"> <li>• Rendering and analysis tasks</li> </ul>	Multi-GPU Single Node
<b>vI3 (Argonne National Lab)</b>	Argonne National Lab	Large dataset visualization in cosmology, astrophysics, and biosciences fields.	<ul style="list-style-type: none"> <li>• Volume rendering of particles</li> </ul>	Multi-GPU Single Node

# Safety and Security

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>AI-NVR</b>	IronYun	Search in Video, Real time intrusion detection	<ul style="list-style-type: none"> <li>• Search amongst 1000s of videos for interesting activities or attributes.</li> </ul>	Single GPU Single Node
<b>Alert</b>	Irvine Sensors	Alert provides people counting and intrusion detection	<ul style="list-style-type: none"> <li>• People counting</li> <li>• Intrusion detection</li> </ul>	Single GPU Single Node
<b>Arvas</b>	VI Dimensions	ARVAS, is an Intelligent Video Analytics solution that uses advance statistical modelling based on deep machine learning technology to detect anomalies. This automated approach enables more accurate detection of complex risk pattern that would otherwise escape human analysts and caused high false alarm.	<ul style="list-style-type: none"> <li>• Abnormally Detection Features - Break-ins, robbery, rioting, floods, accidents, fights, arson, fire, maintenance and vandalism.</li> </ul>	Single GPU Single Node
<b>Better Tomorrow</b>	AnyVision	Face recognition for multiple industries	<ul style="list-style-type: none"> <li>• Face recognition</li> </ul>	Multi-GPU Single Node
<b>BioSurveillance NEXT, BioFinder</b>	Herta Security	Real time facial recognition and forensic alerts against multiple watchlists.	<ul style="list-style-type: none"> <li>• Supports crowded scenes and difficult lighting</li> <li>• Faster than real-time analysis</li> <li>• Partial face concealment</li> </ul>	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.	<ul style="list-style-type: none"> <li>• CUDA</li> </ul>	Multi-GPU Single Node
<b>Cylance</b>	Cylance	Advanced AI-based endpoint malware detection.	<ul style="list-style-type: none"> <li>• Endpoint malware detection solution</li> <li>• GPU deep learning technology</li> </ul>	Multi-GPU Single Node
<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</li> <li>• ALPR</li> <li>• Video analytics and pattern recognition,</li> <li>• Video processing and video enhancement</li> </ul>	Multi-GPU Multi-Node

<b>FindFace</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</li> <li>• TRT</li> <li>• nvenc</li> <li>• nvdec</li> </ul>	Multi-GPU Multi-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>• Facial Expression</li> <li>• Age Estimation</li> <li>• Gender</li> <li>• Ethnicity</li> <li>• Multi Face Tracking</li> <li>• Attention Time</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)</li> <li>• 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</li> <li>• Identifies 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>• Real-time video analysis and alerts</li> <li>• Video search and investigation</li> <li>• Big data analysis</li> <li>• Geospatial mapping and more</li> </ul>	Single GPU Single Node
<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</li> <li>• Optimized scalability using multithreading</li> <li>• Computationally efficient and compact face descriptors</li> <li>• Broad range of working conditions with domain-specific face descriptors</li> </ul>	Multi-GPU Single Node
<b>Nodeflux IVA</b>	Nodeflux	Nodeflux IVA products and services cover wide range of sector including but not limited to smart city, public sector and security, traffic management, toll management, store analytic (wholesale and retail), asset and facilities management, advertising, and transportation.	<ul style="list-style-type: none"> <li>• Face recognition</li> <li>• License plate recognition</li> <li>• Traffic violation detection</li> <li>• Traffic monitoring, and flood monitoring</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,</li> <li>• Incident Detection</li> <li>• Integrated Management</li> <li>• 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</li> <li>• Casino game table analytics</li> </ul>	Single GPU Single Node
<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</li> <li>• Search in Video</li> </ul>	Single GPU Single Node
<b>Tera, Tera+, Tera Vortex</b>	SmartCow	Embedded and Backend video analytics for real-time insights from your security and service-related monitoring systems.	<ul style="list-style-type: none"> <li>• Automatic number plate recognition</li> <li>• Traffic Management</li> <li>• Smart Car Parking Policy</li> <li>• Accident Detection</li> </ul>	Multi-GPU Single Node
<b>XIntelligence XHound XTransport</b>	XJERA LABS PTE.LTD	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</li> <li>• Face recognition</li> <li>• License plate recognition</li> </ul>	Single GPU Single Node
<b>XRVision, IoP</b>	XRVision	Face Recognition and Video Analytics for Uncontrolled, Crowded and In Motion Environments	<ul style="list-style-type: none"> <li>• Face Recognition and Video Analytics</li> <li>• Smart City, Public Safety, Transportation Analytics, Retail Analytics, Ordinance and Environment Safety</li> </ul>	Multi-GPU Single Node

# Tools and Management

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Altair Access</b>	Altair	A simple, powerful, and consistent portal for submitting and monitoring jobs on remote clusters and clouds, and for remote visualization. Brings high-end 3D visualization datacenter hardware right to the user.	<ul style="list-style-type: none"> <li>• 3D Remote Visualization</li> <li>• High-fidelity collaboration</li> <li>• Integrated with Altair PBS Professional for scheduling and control on GPU use and accounting</li> </ul>	Multi-GPU Multi-Node
<b>Altair PBS Professional</b>	Altair	Fast, powerful workload manager designed to improve productivity, optimize utilization and efficiency, and simplify administration for HPC clusters, clouds and supercomputers. Altair PBS Professional automates job scheduling, management, monitoring and reporting.	<ul style="list-style-type: none"> <li>• GPU auto discovery</li> <li>• Specify GPU count per CPU</li> <li>• Specify GPU type</li> <li>• GPU/CPU affinity</li> <li>• GPU awareness and equality in accounting, quotas, and fair share</li> <li>• GPU/CPU syntax/scheduling equivalence</li> <li>• Specify memory use per GPU</li> <li>• Add-on/integration project with support for NVIDIA Data Center GPU Management (DCGM) for GPU health checks &amp; accounting</li> <li>• Open source and commercial versions</li> </ul>	Multi-GPU Multi-Node
<b>Arm Forge (formerly Allinea)</b>	Arm	<p>Build reliable and optimized code for the right results on multiple Server and HPC architectures, from the latest compilers and C++ 11 standards including NVIDIA GPU hardware. Arm Forge combines Arm DDT, the leading debugger for time-saving high performance application debugging, Arm MAP, the trusted performance profiler for invaluable optimization advice across native and Python HPC codes, and Arm Performance Reports for advanced reporting capabilities.</p> <p>Arm Forge Professional (DDT &amp; MAP) providing 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 applications.</p>	<ul style="list-style-type: none"> <li>• Cross Platform: Moving to a new architecture or system is challenging enough without having to learn a new tool chain at the same time. Arm DDT, MAP and Performance Reports run everywhere - on your own laptop, the latest supercomputer, and tomorrow's upcoming architectures</li> <li>• Automatically detect memory bugs, profile behavior and see advanced performance metrics at all scales on Arm 64-bit, Intel Xeon, Intel Xeon Phi, NVIDIA GPUs, and OpenPOWER</li> <li>• Fast Debug: Arm DDT is the debugger of choice for developing of C++, C or Fortran parallel, and threaded applications on CPUs, GPUs and Intel Xeon Phi</li> <li>• Its powerful intuitive graphical interface helps you easily detect memory bugs and divergent behavior at all scales, making Arm DDT the number one debugger in research, industry and academia.</li> <li>• Low-overhead Profiling: Profile your code without distorting application behavior. Arm MAP is Arm Forge's scalable low-overhead profiler of C++, C, Fortran and Python with no instrumentation or code changes required. It helps developers accelerate their code by revealing the causes of slow performance.</li> <li>• From multicore Linux workstations to the largest supercomputers, you can profile realistic test cases with typically less than 5% runtime overhead.</li> </ul>	Multi-GPU Multi-Node

<b>Arm Forge (formerly Allinea) continued</b>	Arm		<ul style="list-style-type: none"> <li>• Short Learning Curve: Arm DDT offers a powerful intuitive GUI that sets the standard for multi-process and multi-threaded debugging</li> <li>• Complex software debugging is made simple whether you're working on a PC or offline, with the help of zero-click variable comparisons, built-in memory debugging, and powerful array visualizations - for today's increasingly parallel processors, clusters, and supercomputers.</li> <li>• Wide Issue Coverage: Arm MAP exposes a wide set of performance indicators, including MPI metrics, PAPI counters, IO metrics, energy metrics and even your own custom metrics</li> <li>• Profile computation (with self and child and call tree representations over time), thread activity (to identify over-subscribed cores and sleeping threads that waste CPU time for OpenMP and pthreads), instruction types, as well as synchronization and I/O performance.</li> <li>• Single and Multi Threaded Profiling: Arm MAP profiles parallel, multithreaded, and single threaded C, C++, Fortran, F90 and Python codes, providing in-depth analysis and bottleneck pinpointing to the source line</li> <li>• Unlike most profilers , it can profile pthreads, OpenMP or MPI for parallel and threaded code, including communication and workload imbalance issues for MPI and multi-process codes</li> </ul>	Multi-GPU Multi-Node
<b>Artec Leo</b>	Artec 3D	A smart 3D scanner that enables you to see your object projected in 3D directly on the HD display.	<ul style="list-style-type: none"> <li>• Jetpack</li> <li>• Tx2</li> </ul>	Single GPU Single 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</li> <li>• Powerful Cluster Management Shell as alternative user interface</li> <li>• Full Support for NVIDIA libraries</li> <li>• CUDA</li> <li>• OpenCL</li> <li>• OpenACC</li> <li>• CUDA-aware libraries</li> <li>• NCCL</li> <li>• CUB</li> <li>• Comprehensive monitoring of GPUs</li> <li>• Brings in GPU resources from public (AWS, Azure) and private (OpenStack) clouds within minutes</li> <li>• Automated scaling of the cluster based on pre-defined policies</li> <li>• Supports several popular Linux distributions: RHEL and derivatives, SUSE SLES and Ubuntu LTS</li> <li>• GPU-enabled Docker containers</li> <li>• Offers a complete deep learning stack</li> <li>• Deployment for popular HPC filesystems and management of fast interconnects</li> <li>• Scales up with multiple NVIDIA DGX systems</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. Controls the software compilation process using simple platform and compiler independent configuration files, and generates native makefiles and workspaces that can be used in the compiler environment of your choice.	<ul style="list-style-type: none"> <li>• Color output for make</li> <li>• Progress output for make</li> <li>• Incremental linking support with vs 8,9 and manifests</li> <li>• Supports out-of-tree builds</li> <li>• Auto-rerun of cmake if any cmake input files change (works with vs 8, 9 using ide macros)</li> <li>• Auto depend information for C++, C, and Fortran</li> <li>• Graphviz output for visualizing dependency trees</li> <li>• Full support for library versions</li> <li>• Full cross platform install system</li> <li>• Generate project files for major IDEs: Visual Studio, Xcode,</li> <li>• Eclipse, KDevelop not tied to make, other portable generators like ant possible</li> <li>• Ability to add custom rules and targets</li> <li>• Compute link depend information, and chaining of dependent libraries</li> <li>• Works with parallel make and is fast, can build very large projects like KDE on build farms</li> </ul>	N/A

<b>ELPA</b>	Max Planck Institute	The publicly available ELPA library provides highly efficient and highly scalable direct eigensolvers for symmetric matrices. Though especially designed for use for PetaFlop/s applications solving large problem sizes on massively parallel supercomputers, ELPA eigensolvers have proven to be also very efficient for smaller matrices.	<ul style="list-style-type: none"> <li>• Improved one-step ScaLAPACK-type solver ELPA1</li> <li>• Novel two-step solver ELPA2</li> </ul>	Multi-GPU Multi-Node
<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. HPCToolkit provides accurate measurements of a program's work, resource consumption, and inefficiency, correlates these metrics with the program's source code, works with multilingual, fully optimized binaries, has very low measurement overhead, and scales to large parallel systems. HPCToolkit's measurements provide support for analyzing a program execution cost, inefficiency, and scaling characteristics both within and across nodes of a parallel system.	<ul style="list-style-type: none"> <li>• Collects accurate and precise calling-context-sensitive performance measurements for unmodified fully optimized applications at very low overhead (1-5%)</li> <li>• Uses asynchronous sampling triggered by system timers and performance monitoring unit events to drive collection of call path profiles and optionally traces</li> <li>• To associate calling-context-sensitive measurements with source code structure, hpcstruct analyzes fully optimized application binaries and recovers information about their relationship to source code</li> <li>• Relates object code to source code files, procedures, loop nests, and identifies inlined code</li> <li>• Overlays call path profiles and traces with program structure computed by hpcstruct and correlates the result with source code</li> <li>• Handles thousands of profiles from a parallel execution by performing this correlation in parallel</li> <li>• hpcprof and hpcprof/mpi generate a performance database that can be explored using the hpcviewer and hpctraceviewer user interfaces</li> <li>• Is a graphical user interface that interactively presents performance data in three complementary code-centric views (top-down, bottom-up, and flat), as well as a graphical view that enables one to assess performance variability across threads and processes</li> <li>• Designed to facilitate rapid top-down analysis using derived metrics that highlight scalability losses and inefficiency rather than focusing exclusively on program hot spots</li> <li>• Presents a hierarchical, time-centric view of a program execution. The tool can rapidly render graphical views of trace lines for thousands of processors for an execution tens of minutes long even a laptop</li> <li>• hpctraceviewer's hierarchical graphical presentation is quite different than that of other tools - it renders execution traces at multiple levels of abstraction by showing activity over time at different call stack depths</li> </ul>	Multi-GPU Multi-Node

<b>IBM Spectrum LSF</b>	IBM Corporation	IBM Spectrum LSF is a highly scalable and highly available HPC workload manager that features intelligent, policy driven scheduling, superior resource utilization, and comprehensive support for GPUs.	<ul style="list-style-type: none"> <li>• Enforcement of GPU allocations via cgroups</li> <li>• Exclusive allocation and round robin shared mode allocation</li> <li>• CPU-GPU affinity</li> <li>• Boost control</li> <li>• Power management</li> <li>• Multi-Process Server (MPS) support</li> <li>• NVIDIA Volta and DCGM support</li> </ul>	Multi-GPU Multi-Node
<b>Magma</b>	ICL - University of Tennessee Knoxville	MAGMA provides a dense linear algebra library similar to LAPACK but for heterogeneous/hybrid architectures, starting with current "Multicore+GPU" systems.	<ul style="list-style-type: none"> <li>• Linear system solvers</li> <li>• Eigenvalue problem solvers</li> <li>• Auxiliary BLAS</li> <li>• Batched LA</li> <li>• Sparse LA</li> <li>• CPU/GPU Interface</li> <li>• Multiple precision support</li> <li>• Non-GPU-resident factorizations</li> <li>• Multicore and multi-GPU support</li> <li>• MAGMA Analytics/DNN</li> <li>• LAPACK testing</li> <li>• Linux</li> <li>• Windows</li> <li>• Mac OS</li> </ul>	Multi-GPU Single Node
<b>PAPI</b>	ICL - University of Tennessee Knoxville	PAPI provides the tool designer and application engineer with a consistent interface and methodology that enables software engineers to see, in near real time, the relation between software performance and processor events.	<ul style="list-style-type: none"> <li>• The Performance API (PAPI) project specifies a standard application programming interface (API) for accessing hardware performance counters available on most modern microprocessors</li> <li>• These counters exist as a small set of registers that count Events, occurrences of specific signals related to the processor's function</li> <li>• Monitoring these events facilitates correlation between the structure of source/object code and the efficiency of the mapping of that code to the underlying architecture</li> </ul>	Multi-GPU Multi-Node

<b>Parallelware Trainer</b>	Appentra Solutions	<p>Parallelware Trainer is an interactive, real-time code editor with features that facilitate the learning, usage, and implementation of parallel programming by understanding how and why sections of code can be parallelized.</p> <p>Users are actively involved in learning parallel programming through observation, comparison, and hands-on experimentation.</p> <p>Parallelware Trainer provides support for widely used parallel programming strategies using OpenMP and OpenACC with execution on multicore processors and GPUs.</p>	<ul style="list-style-type: none"> <li>• Interactive, real-time editor GUI that shows you how and where to implement parallelism.</li> <li>• Assists in the parallelization of code using OpenMP and OpenACC.</li> <li>• Transparent, local/ remote, execution and benchmarking.</li> <li>• Support for the C programming language. Full Fortran support coming soon.</li> <li>• Detailed report of opportunities for parallelism discovered in your code.</li> <li>• Support for multiple compilers including GCC, Intel and PGI.</li> <li>• Benefits: <ul style="list-style-type: none"> <li>&gt; Faster, more effective learning.</li> <li>&gt; Reduced learning curve.</li> <li>&gt; All-in-one learning tool for parallel programming.</li> <li>&gt; Immediate use of parallel programming.</li> <li>&gt; Support for multicore processors and GPUs.</li> </ul> </li> </ul>	N/A
<b>SLURM</b>	SchedMD	<p>SLURM is a highly configurable open source workload and resource manager that 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 workflows.</p>	<ul style="list-style-type: none"> <li>• First-class GPU support</li> <li>• Scales to millions of cores and tens of thousands of GPGPUs</li> <li>• Military grade security</li> <li>• Heterogenous platform support allowing users to take advantage of GPGPUs.</li> <li>• Flexible plugin framework enables Slurm to meet complex customization requirements</li> <li>• Topology aware job scheduling for maximum system utilization</li> <li>• Extensive scheduling options including advanced reservations, suspend/resume, backfill, fair-share and preemptive scheduling for critical jobs</li> <li>• No single point of failure</li> </ul>	Multi-GPU Multi-Node
<b>STRIVR</b>	StriVR	<p>STRIVR offers an end-to-end Immersive Learning platform that revolutionizes the way people and businesses train, learn, and perform.</p>	<ul style="list-style-type: none"> <li>• VRWorks 360 Video</li> </ul>	Single GPU Single Node

<b>TAU - Tuning and Analysis Utilities</b>	University of Oregon	<p>TAU Performance System is a portable profiling and tracing toolkit for performance analysis of parallel programs written in Fortran, C, C++, UPC, Java, Python.</p> <p>TAU (Tuning and Analysis Utilities) is capable of gathering performance information through instrumentation of functions, methods, basic blocks, and statements as well as event-based sampling. All C++ language features are supported including templates and namespaces. The API also provides selection of profiling groups for organizing and controlling instrumentation. The instrumentation can be inserted in the source code using an automatic instrumentor tool based on the Program Database Toolkit (PDT), dynamically using DyninstAPI, at runtime in the Java Virtual Machine, or manually using the instrumentation API.</p> <p>TAU's profile visualization tool, paraprof, provides graphical displays of all the performance analysis results, in aggregate and single node/context/thread forms. The user can quickly identify sources of performance bottlenecks in the application using the graphical interface. In addition, TAU can generate event traces that can be displayed with the Vampir, Paraver or JumpShot trace visualization tools.</p>	<ul style="list-style-type: none"> <li>• Instrumentation</li> <li>• PerfDMF</li> <li>• Paraprof</li> <li>• Load Profiles</li> <li>• Metric Window</li> <li>• Thread Windows</li> <li>• Communication Matrix</li> <li>• 3D Visualization</li> <li>• Derived Metrics</li> <li>• Selective Instrumentation</li> <li>• PerfExplorer</li> <li>• Cluster Analysis</li> <li>• Correlation Analysis</li> <li>• Scalability Chart</li> <li>• Preset Charts</li> <li>• Custom Charts</li> <li>• Visualizations</li> <li>• Eclipse Introduction</li> <li>• Selective Instrumentation</li> <li>• Instrumenting Java</li> <li>• Configuration Manager</li> </ul>	Multi-GPU Multi-Node
<b>Torque Moab</b>	Adaptive Computing	<p>Moab HPC Suite is a workload and resource orchestration platform that automates the scheduling, managing, monitoring, and reporting of HPC workloads on massive scale.</p> <p>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.</p>	<ul style="list-style-type: none"> <li>• Requests and schedules gpus based on gpu location in NUMA systems</li> <li>• Collects and report smetrics and status information</li> <li>• Sets gpu mode at job run time</li> </ul>	Multi-GPU Multi-Node
<b>Totalview</b>	Perforce	<p>TotalView is the leading dynamic analysis and debugging tool designed to handle complex CPU and GPU based multi-threaded, multi-process and multi-node cluster applications. TotalView supports the latest CUDA SDK's, NVIDIA GPU hardware, Linux x86-64, Arm64, and OpenPower platforms and applications utilizing MPI and OpenMP technologies.</p>	<ul style="list-style-type: none"> <li>• OpenACC directives</li> <li>• CUDA running directly on NVIDIA latest GPUs</li> <li>• Linux and GPU device thread visibility</li> <li>• CUDA function calls, host pinned memory regions and CUDA contexts</li> <li>• Handling CUDA functions inline and on the stack</li> <li>• Command line interface (CLI) commands for CUDA functions</li> <li>• MPI applications on CUDA-accelerated clusters</li> </ul>	Multi-GPU Multi-Node

<b>Univa Grid Engine</b>	Univa	<p>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.</p>	<ul style="list-style-type: none"> <li>• Manage Nvidia CUDA</li> <li>• OpenACC</li> <li>• OpenCL plus MPI hybrid apps</li> <li>• Optimizes scheduling with resource-mapped GPUs</li> <li>• Manages GPU apps within or without Docker containers</li> <li>• Obtain visibility with CUDA-specific metrics for GPU monitors and reports</li> <li>• Extend on-premise deployments to incorporate cloud-based GPU instances</li> </ul>	Multi-GPU Multi-Node
<b>Vampir</b>	TU Dresden	<p>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.</p> <p>The combined handling and visualization of instrumented and sampled event traces generated by Score-P enables an outstanding performance analysis capability of highly-parallel applications. Current developments also include the analysis of memory and I/O behavior that often impacts an application's performance.</p>	<ul style="list-style-type: none"> <li>• Powerful zooming and scrolling in all displays</li> <li>• Adaptive statistics for user selected time ranges</li> <li>• Filtering of processes, functions, messages, collective operations</li> <li>• Hierarchical grouping of threads, processes, and nodes</li> <li>• Support of source code locations</li> <li>• Integrated snapshot and printing for publishing</li> <li>• Customizable displays</li> <li>• VampirServer</li> <li>• Ultra scalable re-design of established Vampir functionality</li> <li>• Distributed performance data visualization</li> <li>• Increased scalability compared to sequential approach</li> <li>• Responsive performance data browsing from remote sites</li> <li>• Includes support for: NVIDIA CUDA, CUPTI, CUDA libraries</li> <li>• Performance Analysis Framework</li> <li>• Easy to use performance analysis framework for parallel programs</li> <li>• Graphical data representation enables detailed understanding of dynamic processes on massively parallel systems</li> <li>• In-depth event based analysis of parallel run-time behavior and interprocess communication</li> <li>• Identification of performance problems and bottlenecks</li> </ul>	Multi-GPU Multi-Node



## Agriculture

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Taranis</b>	Taranis	Taranis provides a platform for discovering various crop health issues, helping farmers take care of both land and crops and making sure they get the best of their yield.	<ul style="list-style-type: none"> <li>• report plant population to farmers</li> <li>• detect when a weed emerges in field and constitutes a potential threat</li> <li>• calculate amounts of nutrients in vegetation, water content in the soil, plant temperature</li> <li>• identify and categorize the top relevant diseases for prevalent crops</li> </ul>	Multi-GPU Multi-Node

## Business Process Optimization

APPLICATION NAME	COMPANYNAME	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
<b>Automated checkout</b>	Focal Systems	Focal's Product Recognition eliminates barcode scanning entirely at the cashier and achieves 99% accuracy on thousands of products.	<ul style="list-style-type: none"> <li>• cuDNN</li> <li>• TensorRT</li> </ul>	Multi-GPU Single Node
<b>DataX.AI</b>	CrowdANALYTIX	Cloud-based crowd-sourced analytics services that create an online retail product catalog, on-boarding SKU in minutes instead of the manual process of tagging and provide produce info and removing human error involved.	<ul style="list-style-type: none"> <li>• cuDNN</li> </ul>	Single GPU Single Node
<b>Helix</b>	Maxerience	CPG product training platform: creates digital copies of products right at the production line in a matter of minutes, and creates an AI model in less than 30 minutes!	<ul style="list-style-type: none"> <li>• TensorRT</li> </ul>	Single GPU Single Node
<b>Part Finder Kiosk</b>	Slyce	A visual search and image recognition solution for retailers and brands	<ul style="list-style-type: none"> <li>• Real time scan item and direct customer to item's location in store</li> <li>• Find a replacement or additional info</li> <li>• Feature Jetpack</li> </ul>	Single GPU Single Node
<b>Peak Trading Out Of Stock</b>	BeMyEye	Out of Stock (OOS) and Almost OOS (AOOS) crowd sourcing solutions for retailers	<ul style="list-style-type: none"> <li>• Product recognition on the cloud</li> </ul>	Single GPU Single Node
<b>Perfect Shelf</b>	BeMyEye	Track Hypermarkets, Supermarkets, Discounters, Managed Convenience and Chemists, using unique blend of IR technologies and crowdsourcing, to provide you with on-shelf sales fundamental data across an entire category	<ul style="list-style-type: none"> <li>• Real time inferencing on the cloud</li> <li>• SKU recognition</li> </ul>	Single GPU Single Node
<b>Predictive Pricing</b>	Evo Pricing	Market-driven optimal prices based on demand, competition, product features and customer feedback	<ul style="list-style-type: none"> <li>• GPU on the cloud</li> </ul>	Multi-GPU Single Node
<b>Third Wave Automation</b>	Third Wave Automation	Automation cloud robotics and machine learning technology to material handling forklift automation in a warehouse	<ul style="list-style-type: none"> <li>• Geforce 2080 Ti</li> </ul>	Single GPU Single Node

For more information on GPU-accelerated applications please visit, [www.nvidia.com/testlaapps](http://www.nvidia.com/testlaapps)

