



GPU-ACCELERATED APPLICATIONS



GPU-ACCELERATED APPLICATIONS

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

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



CONTENTS

- 1 Computational Finance
- 2 Climate, Weather and Ocean Modeling
- 2 Data Science and Analytics
- 5 Artificial Intelligence
 - DEEP LEARNING AND MACHINE LEARNING
- 10 Federal, Defense and Intelligence
- 12 Design for Manufacturing/Construction: CAD/CAE/CAM
 - CFD (MFG)
 - CFD (RESEARCH DEVELOPMENTS)
 - COMPUTATIONAL STRUCTURAL MECHANICS
 - DESIGN AND VISUALIZATION
 - ELECTRONIC DESIGN AUTOMATION
 - INDUSTRIAL INSPECTION
- 23 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
- 34 Medical Imaging
- 35 Oil and Gas
- 36 Life Sciences
 - BIOINFORMATICS
 - MICROSCOPY
 - MOLECULAR DYNAMICS
 - QUANTUM CHEMISTRY
 - (MOLECULAR) VISUALIZATION AND DOCKING
 - AGRICULTURE
- 49 Research: Higher Education and Supercomputing
 - NUMERICAL ANALYTICS
 - PHYSICS
 - SCIENTIFIC VISUALIZATION
- 54 Safety and Security
- 57 Tools and Management
- 65 Unclassified Categories

Computational Finance

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--------------------------------------|----------------------------|---|--|---------------------------|
| Accelerated Computing Engine | 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"> • Web-like API with Native bindings for Python, R, Scala, C • Custom models and data streams | Multi-GPU Single Node |
| Adaptiv Analytics | 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"> • Codes in C# supported transparently, with minimal code changes • Supports multiple backends including CUDA and OpenCL • Switches transparently between multiple GPUs and CPUS depending on the deal support and load factors. | Multi-GPU Single Node |
| Alea.cuBase F# | QuantAlea's | F# package enabling a growing set of F# capability to run on a GPU. | <ul style="list-style-type: none"> • F# for GPU accelerators | Multi-GPU Single Node |
| Esther | 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"> • High quality models not admitting closed form solutions • Efficient solvers based on full matrix linear algebra powered by GPUs and Monte Carlo algorithms | Multi-GPU Single Node |
| Global Risk | MISYS | Regulatory compliance and enterprise wide risk transparency package. | <ul style="list-style-type: none"> • Risk analytics | Multi-GPU Single Node |
| Hybridizer C# | Altimesh | Multi-target C# framework for data parallel computing. | <ul style="list-style-type: none"> • C# with translation to GPU • Multi-Core Xeon | Multi-GPU Single Node |
| MACS Analytics Library | Murex | Analytics library for modeling valuation and risk for derivatives across multiple asset classes. | <ul style="list-style-type: none"> • Market standard models for all asset classes paired with the most efficient resolution methods (Monte Carlo simulations and Partial Differential Equations) | Multi-GPU Single Node |
| MiAccLib 2.0.1 | 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"> • Text Processing: Exact Match, Approximate\Similarity Text, Wild Card, MultiKeyword and MultiColumnMultiKeyword, etc • Data Security: Accelerated Encryption/Description for AES-128 • Video Analytics: Accelerated Intrusion Detection Algorithm | Multi-GPU Single Node |
| NAG | Numerical Algorithms Group | Random number generators, Brownian bridges, and PDE solvers | <ul style="list-style-type: none"> • Monte Carlo and PDE solvers | Single GPU Single Node |
| O-Quant options pricing | O-Quant | Offering for risk management and complex options and derivatives pricing using GPUs. | <ul style="list-style-type: none"> • Cloud-based interface to price complex derivatives representing large baskets of equities | Multi-GPU Multi-Node |
| Oneview | Numerix | Numerix introduced GPU support for Forward Monte Carlo simulation for Capital Markets and Insurance. | <ul style="list-style-type: none"> • Equity/FX basket models with BlackScholes/Local Vol models for individual equities and FX • Algorithms: AAD (Automatic Algebraic Differential) • New approaches to AAD to reduce time to market for fast Price Greeks and XVA Greeks | Multi-GPU Multi-Node |
| Pathwise | Aon Benfield | Specialized platform for real-time hedging, valuation, pricing and risk management. | <ul style="list-style-type: none"> • Spreadsheet-like modeling interfaces • Python-based scripting environment • Grid middleware | Multi-GPU Single Node |
| SciFinance | SciComp, Inc | Derivative pricing (SciFinance) | <ul style="list-style-type: none"> • Monte Carlo and PDE pricing models | Single GPU Single Node |
| Synerscope Data Visualization | SynerScope | Visual big data exploration and insight tools | <ul style="list-style-type: none"> • Graphical exploration of large network datasets including geo-spatial and temporal components | Single GPU Single Node |

| | | | | |
|---------------------|--------------------|--|---|--------------------------|
| Volera | Hanweck Associates | Real-time options analytical engine (Volera) | <ul style="list-style-type: none"> • Real-time analytics | Multi-GPU Single Node |
| Xcelerit SDK | 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"> • C++ programming language, cross-platform (back-end generates CUDA and optimized CPU code) • Supports Windows and Linux operating systems | Multi-GPU Single Node |

Climate, Weather and Ocean Modeling

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------|------------------|--|--|-------------------------|
| COSMO | COSMO Consortium | Regional numerical weather prediction and climate research model | <ul style="list-style-type: none"> • Radiation only in the trunk release • All features in the MCH branch used for operational weather forecasting | Multi-GPU Multi-Node |
| E3SM-EAM | US DOE | Global atmospheric model used as component to E3SM global coupled climate model. | <ul style="list-style-type: none"> • Dynamics and most physics | Multi-GPU Multi-Node |
| Gales | KNMI, TU Delft | Regional numerical weather prediction model | <ul style="list-style-type: none"> • Full Model | Multi-GPU Multi-Node |
| WRF AceCAST-WRF | 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"> • ARW dynamics • 19 physics options including enough to run the full WRF model on GPUs | Multi-GPU Multi-Node |

Data Science and Analytics

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------------|-------------|---|---|--------------------------|
| Anaconda Distribution | Anaconda | <p>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, enabling individual data scientists to:</p> <ul style="list-style-type: none"> • Quickly download 1,500+ Python/R data science packages • Manage libraries, dependencies, and environments with Conda • Develop and train machine learning and deep learning models with scikit-learn, TensorFlow, and Theano • Analyze data with scalability and performance with Dask, NumPy, pandas, and Numba • Visualize results with Matplotlib, Bokeh, Dash, and Holoviews <p>Anaconda is the leading Python package manager, that is the lead contributor to several open source data science libraries. Anaconda includes Numba, a Python-to-GPU compiler that compiles easy-to-read Python code to many-core and GPU architectures. Also includes single-line install of key deep learning packages for GPUs such as pytorch. Anaconda has been downloaded over 15M times and is used for AI & ML data science workloads using TensorFlow, Theano, Keras, Caffe, Neon, Lasagne, NLTK, spaCY.</p> | <ul style="list-style-type: none"> • Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND • Sorts algorithms from the CUB and Modern GPU libraries • Includes Numba (JIT python compiler) and Dask (python scheduler) • Includes single-line install of numerous DL frameworks such as pytorch | Multi-GPU Single Node |

| | | | | |
|-------------------------------------|--------------------|---|--|--------------------------|
| ArgusSearch | Planet AI | Deep Learning driven document search tool. | <ul style="list-style-type: none"> • Fast full text search engine • Searches hand-written and text documents, including PDF • Allows almost any arbitrary requests (Regular Expressions are supported) • Provides a list of matches sorted by confidence | Multi-GPU Single Node |
| Automatic Speech Recognition | Capio | In-house and Cloud-based speech recognition technologies | <ul style="list-style-type: none"> • Real-time and offline (batch) speech recognition • Exceptional accuracy for transcription of conversational speech • Continuous Learning (System becomes more accurate as more data is pushed to the platform) | Multi-GPU Single Node |
| BlazingSQL | BlazingDB | GPU-accelerated SQL Engine for analytics available on all major CSP and on-premise deployment. | <ul style="list-style-type: none"> • Modern SQL Engine • Supports petabyte scale applications • Supports traditional big data format | Multi-GPU Single Node |
| BrytlytDB | Brytlyt | In-GPU-memory database built on top of PostgreSQL | <ul style="list-style-type: none"> • GPU-Accelerated joins, aggregations, scans, etc. on PostgreSQL • Visualization platform bundled with database is called SpotLyt. | Multi-GPU Multi-Node |
| CuPy | Preferred Networks | CuPy (https://github.com/cupy/cupy) is a GPU-accelerated scientific computing library for Python with a NumPy compatible interface. | <ul style="list-style-type: none"> • CUDA • multi-GPU support | Multi-GPU Single Node |
| Datalogue | Datalogue | AI powered pipelines that automatically prepare any data from any source for immediate & compliant use. | <ul style="list-style-type: none"> • Data transformation • Ontology mapping • Data standardization • Data augmentation | Multi-GPU Single Node |
| DeepGram | DeepGram | Voice processing solution for call centers, financials and other scenarios. | <ul style="list-style-type: none"> • Speech to text and phonetic search using GPU deep learning | Multi-GPU Single Node |
| Driverless AI | 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"> • Automated machine learning and feature extraction • Automated statistical visualization • Interpretability toolkit for machine learning models | Multi-GPU Single Node |

| | | | | |
|-------------------------|----------------------------------|--|---|--------------------------|
| GPUdb | 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"> • Query against big data in real time • No pre-indexing allows for complex, ad-hoc query chains • Interactively explore large, streaming data sets | Multi-GPU Single Node |
| H2O4GPU | 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"> • Available algorithms include Gradient Boosting Machines (GBM's) • Generalized Linear Models (GLM's) • K-Means Clustering • SVD • PCA • K-means • XGBoost. • It can be used as a drop-in replacement for scikit-learn with support for GPUs on selected (and ever-growing) algorithms. • 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. | Multi-GPU Single Node |
| IntelligentVoice | 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"> • Advanced Speech Recognition across large data sets • JumpTo Technology, for data visualisation • E-Discovery • Extraction from phone calls • IM & Email defining key phrases and emotional analysis • Compliance, defining key conversations and interactions | Multi-GPU Single Node |
| Jedox | 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"> • This database holds all relevant data in GPU memory • Tesla K40 & 12 GB on-board RAM • Scales up with multiple GPUs • Keeps close to 100 GB of compressed data in GPU memory on a single server system • Fast analysis, reporting, and planning | Multi-GPU Single Node |
| Labellio | 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"> • Neural net fine-tuning for image data • Data crawling and data browsing • Drag-and-drop style data cleansing backed by AI support | Multi-GPU Single Node |
| Numba | 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"> • On-the-fly code generation (at import time or runtime, at the user's preference) • Native code generation for the CPU (default) and GPU hardware • Integration with the Python scientific software stack (enabled via Numpy) • JIT compilation of Python functions for execution on various targets (including CUDA) | Multi-GPU Single Node |

| | | | | |
|-----------------------------|------------|--|--|---------------------------|
| OmniSci | 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"> • Uses LLVM's nvtx backend to generate CUDA kernels • OpenGL- (EGL) based rendering • Can run in a docker container using NVIDIA-docker | Multi-GPU Single Node |
| Polymatica | Polymatica | Analytical OLAP and Data Mining Platform | <ul style="list-style-type: none"> • Visualization, Reporting, OLAP in-memory with GPU acceleration • Data Mining • Machine Learning • Predictive Analytics | Multi-GPU Multi-Node |
| Scream DB | Scream | GPU accelerated SQL database engine for big data analytics. Scream speeds SQL analytics by 100X by translating SQL queries into highly parallel algorithms run on the GPU. | <ul style="list-style-type: none"> • Up to 100TB of raw data can be stored and queried in a standard 2U server • Inserts and analyzes hundreds of billions of records in seconds • No indexes required • No changes to SQL code or data science paradigms required | Multi-GPU Single Node |
| SynerScope | 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"> • Real-time Interaction with data | Single GPU Single Node |
| ZX Lib (Fuzzy Logic) | Tanay | Financial analytics and data mining library | <ul style="list-style-type: none"> • Monte Carlo simulations • Pricing of vanilla and exotic options • Fixed income analytics • Data mining | Multi-GPU Single Node |

Artificial Intelligence

DEEP LEARNING AND MACHINE LEARNING

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-------------------|-------------|---|---|--------------------------|
| AiFi Nano | AiFi | Cashier-free (like Amazon grab and go solution) and stock out retail software | <ul style="list-style-type: none"> • cuDNN • TensorRT • DeepStream | Multi-GPU Single Node |
| AlphaSense | AlphaSense | PaaS for Financial analysis based on public corporate information. Geared at financial analysts within financial services.. Allows very fast searches of public corporate information, and allows questing answering format ("the Google for Analyst research") | <ul style="list-style-type: none"> • 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") | Multi-GPU Single Node |

| | | | | |
|--|------------------|--|--|--------------------------|
| Anaconda Enterprise | Anaconda | Anaconda Enterprise combines core AI technologies, governance, and cloud-native architecture. Each pieceâ€”core AI, governance, and cloud nativeâ€”are critical components to enabling organizations to automate AI at speed and scale. | <ul style="list-style-type: none"> • Access 1,500+ secure Python and R data science packages and libraries from Anaconda. • Curate a private package repository controlled by IT. • Craft package policies by blacklisting and white listing license types and versions. • Leverage code and GPU-specific Conda packages designed to accelerate computation and train models. • Share centralized GPU clusters across teams, using custom resource profiles to establish resource limits. • Create Anaconda installers with custom sets of packages for Windows, Mac, and Linux. • Easily distribute your own proprietary packages to share code, algorithms, and models. AE5.3 v2.19 Use Open Source Software Securely • Quickly and easily share notebooks with others, using your preferred IDE. • Grant or restrict access to individual notebooks by user or group. • Benefit from automated version control in data science projects. • Connect to Hadoop/Spark clusters and other data sources for distributed workloads. • Create custom resource profiles by role to efficiently allocate resources across teams. | Multi-GPU Single Node |
| Apache Mahout | Apache Mahout | Mahout is building an environment for quickly creating scalable performant machine learning applications. | <ul style="list-style-type: none"> • Extremely easy to add new algorithms • Distributed instead of single machine | Multi-GPU Multi-Node |
| Artificial Intelligence Radio Transceiver (AIR-T) | 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"> • The AIR-T is designed to be an edge-compute inference engine for deep learning algorithms. | N/A |
| ARYA.ai | 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"> • Deep learning • TensorFlow. | Multi-GPU Multi-Node |
| Avitas Systems - Inspection as a Service | 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"> • Drone based data capture • RGB Camera, Laser and Infrared sensing • Deep learning driven Object detection for Inspection • Detect corrosion levels, damaged/missing parts, encroaching vegetation volumes. • AI workbench • Photogrammetry | Multi-GPU Multi-Node |

| | | | | |
|--|--------------------------|--|---|--------------------------|
| BIDMach - | UC Berkeley | The fastest machine learning library available. Holds the record for many common machine learning algorithms. | <ul style="list-style-type: none"> • Written in Scala and supports Scala and Java interfaces • Supports linear regression, logistic regression, SVM, LDA, K-Means and other operations | Multi-GPU Single Node |
| Bons.ai | 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"> • Easy to use programming interface. Bons.ai • Novel programming language called Inkling • Primary focus on reinforcement learning | Multi-GPU Single Node |
| C3 Fraud Detection | C3 IoT | C3 IoT is a Platform-as-a-Service for industrial customers including utilities, manufacturing, retail, finance, and healthcare. GPUs consumed exclusively through Amazon AWS. | <ul style="list-style-type: none"> • Deep learning models, including RNNs | Multi-GPU Single Node |
| Caffe2 | 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"> • GPU cluster processing • Mass image data | Multi-GPU Single Node |
| CatBoost | Yandex | CatBoost is an open-source gradient boosting library with categorical features support. | <ul style="list-style-type: none"> • Extremely fast learning on GPU • Multi-GPU • Multi-Node | Multi-GPU Multi-Node |
| Chainer | Preferred Networks, Inc. | DL framework that makes the construction of neural networks (NN) flexible and intuitive. | <ul style="list-style-type: none"> • Dynamic NN construction, which makes debugging easier • CPU/GPU-agnostic coding, which is promoted by CuPy, partially NumPy-compatible multidimensional array library for CUDA • Data-dependent NN construction, which fully exploits the control flows of Python without magic | Multi-GPU Multi-Node |
| Clarifai | 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. | <ul style="list-style-type: none"> • GPU-based training and inference • Recognizes and indexes images with predefined classifiers, or with custom classifiers | Multi-GPU Single Node |
| CNTK | Microsoft | 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"> • Speech Recognition • Machine Translation • Image Recognition • Image Captioning • Text Processing and Relevance • Language Understanding • Language Modeling | Multi-GPU Single Node |
| Databricks Unified Analytics Platform | Databricks | Databricks provides a cloud-based platform designed to make big data and machine learning simple. | <ul style="list-style-type: none"> • GPU instances available with CUDA drivers included • GPU support provided by Spark scheduler • Integration of TensorFlow, Keras • TensorFrames data connector • Deep learning pipelines/workflows • Transfer learning and image loading | Multi-GPU Multi-Node |
| DeepInstinct | DeepInstinct | Zero day end point malware detection solution offered to enterprise markets. | <ul style="list-style-type: none"> • Zero-day threats & APT attack detection on endpoints, servers and mobile devices | Multi-GPU Single Node |

| | | | | |
|--------------------------|----------------------------|--|--|---------------------------|
| Deeplearning4j | 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"> Integrates with Hadoop and Spark to run distributed Java and Scala APIs Composable framework that facilitates building your own nets Includes ND4J, the Numpy for Java. | Multi-GPU Single Node |
| Dessa | Dessa | Deep Learning Platform based on TensorFlow. Allows end-to-end workflows. Targets consumer banking and insurance industries. | <ul style="list-style-type: none"> Deep learning workflows can be built Based on TensorFlow Use cases in consumer banking and Insurance | Multi-GPU Multi-Node |
| Dextro | Axon | Dextro's API uses deep learning systems to analyze and categorize videos in real-time. | <ul style="list-style-type: none"> Object and scene detection Machine transcription for audio Motion and movement detection | Multi-GPU Single Node |
| Dr. Retail | SkyREC | Instore data analytics | <ul style="list-style-type: none"> TensorRT 5.1 nvJPEG NVEnc NVDec | Single GPU Single Node |
| Gridspace | 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"> Speech-to-text transcription Compliance Call grading Call topic modeling Customer service enhancement Customer churn prediction | N/A |
| Keras | 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"> cuDNN version (depends on the version of TensorFlow and Theano installed with Keras) Supported Interfaces: Python | Multi-GPU Single Node |
| MatConvNet | 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"> Building Blocks Simple CNN wrapper DagNN wrapper cuDNN implemented | Multi-GPU Single Node |
| Matroid | 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"> Matroid is multi-cloud and allows it customers to easily switch between AWS, Azure and Google Cloud. | Multi-GPU Multi-Node |
| MetaMind | 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"> GPU-based training and inference Recognizes image and analyzes text Creates and trains classifiers with tooling for uploading and managing datasets | Multi-GPU Single Node |
| Mobiliya ThirdEye | 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"> CUDA/cuDNN TensorRT Deepestream Jetpack | Single GPU Single Node |
| MXNet | 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"> MXnet supports cuDNN v5 for GPU acceleration | Multi-GPU Multi-Node |

| | | | | |
|-------------------------------|----------------------|---|--|---------------------------|
| Neon | 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"> • Training, inference and deployment of deep learning models • Processes over 442M images per day on a Titan X | Multi-GPU Single Node |
| NVCaffe | Berkeley AI Research | The Caffe deep learning framework makes implementing state-of-the-art deep learning easy. | <ul style="list-style-type: none"> • Process over 40M images per day with a single NVIDIA K40 or Titan GPU | Single GPU Single Node |
| out of stock detection | 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"> • On-Shelf Availability Analytics per hour • Real-time Alerts on your “never be outs” | Multi-GPU Single Node |
| PaddlePaddle | 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"> • Optimized math operations through SSE/AVX intrinsics, BLAS libraries (e.g. MKL, ATLAS, cuBLAS) or customized CPU/GPU kernels • Highly optimized recurrent networks which can handle variable-length sequence without padding • Optimized local and distributed training for models with high dimensional sparse data | Multi-GPU Single Node |
| SAS | SAS | SAS Machine Learning. SAS Viya Visual Data Mining and Visualization suites now leverage GPU deep learning | <ul style="list-style-type: none"> • Volta V100 with tensor cores • TensorRT for inference on the NVIDIA Jetson TX2 box • RNN • Multiple GPUs on a single SMP node • Homogeneous and heterogeneous MPP with synchronized Stochastic Gradient Descent | Multi-GPU Multi-Node |
| Sentient | Sentient | Sentient is an AI platform company with special focus on digital marketing, ecommerce and finance trading applications. | <ul style="list-style-type: none"> • Sentient is using GPU deep learning in its commercially available ecommerce, digital marketing and financial trading applications • Studio.ml is a new project designed to make AI development easier by hiding most of the complexity • Studio.ml runs on-premise and in the cloud | Single GPU Single Node |
| Smart Skin | 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) - automatic retopology of 3D and 4D data using machine learning - photogrammetry : noise-reduction and hole-patching using machine learning - realistic lip-sync using 4D-trained neural network | <ul style="list-style-type: none"> • CUDA • Hairworks • PhysX • cuDNN • OptiX | Multi-GPU Multi-Node |
| SpaceKnow PaaS | SpaceKnow | PaaS for deep learning extraction of satellite data information targeted at Financial Services and Defense Intelligence. Tracks macro/micro-economic activity by applying deep learning to satellite images. | <ul style="list-style-type: none"> • Extracts economic activity from satellite images using deep learning • Provides batch mode extraction | Multi-GPU Multi-Node |

| | | | | |
|--------------------------------|-------------|--|--|---------------------------|
| Tensorflow | 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"> TensorFlow is flexible, portable and performant creating an open standard for exchanging research ideas and putting machine learning in products | Multi-GPU Single Node |
| Theano | LISA Lab | Theano is a symbolic expression compiler that powers large-scale computationally intensive scientific investigations. | <ul style="list-style-type: none"> Abstract expression graphs for transparent GPU acceleration | Multi-GPU Single Node |
| Torch7 | Open Source | Torch7 is an interactive development environment for machine learning and computer vision. | <ul style="list-style-type: none"> Computational back-ends for multicore GPUs | Multi-GPU Single Node |
| UETorch | Facebook | UETorch provides an embedded Torch environment within the powerful Unreal Engine 4. This allows one to have deep learning models directly interact with the game world, and paves way for powerful research. | <ul style="list-style-type: none"> Game interaction and physics CUDA-optimized deep learning and neural networks CuDNN supported | Multi-GPU Single Node |
| Unify.ID | Unify.ID | Behavioral user authentication service | <ul style="list-style-type: none"> Identifies individuals based on unique factors such as the way they walk, type and sit | Multi-GPU Single Node |
| ViMo | Motionloft | Video analytics using Tx1/Tx2, people counting, queue management, bounce rate, gender, age, heatmap and path tracing. | <ul style="list-style-type: none"> Jetpack TensorRT | Single GPU Single Node |
| Visual Intelligence API | 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"> 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"> Curate and organize visual content Search and recommend visually Get insights and analytics visually | Single GPU Single Node |
| Zippin | 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"> Jetpack | Single GPU Single Node |

Federal, Defense and Intelligence

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------------|--------------|---|--|--------------------------|
| Advanced Ortho Series | DigitalGlobe | Geospatial visualization | <ul style="list-style-type: none"> Image orthorectification | Multi-GPU Single Node |
| ArcGIS Pro | ESRI | <p>Viewshed2 determines the raster surface locations visible to a set of observer features, using geodesic methods.</p> <p>Aspect - Determines the compass direction that the downhill slope faces for each location</p> <p>Slope - Determines the slope (gradient or steepness) from each cell of a raster</p> | <ul style="list-style-type: none"> Viewshed2 - transforms the elevation surface into a geocentric 3D coordinate system and runs 3D sightlines to each transformed cell center 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. Slope - The output slope raster can be calculated in two types of units, degrees or percent (percent rise). | Multi-GPU Single Node |
| Blaze Terra | Eternix | Geospatial visualization tool | <ul style="list-style-type: none"> 3D visualization of geospatial data | Multi-GPU Single Node |

| | | | | |
|-----------------------------------|----------------------|--|--|---------------------------|
| Elcomsoft | Elcomsoft | High-performance distributed password recovery software with NVIDIA GPU acceleration and scalability to over 10,000 workstations. | <ul style="list-style-type: none"> • GPU acceleration for password recovery • 10-100x speedup for password recovery | Multi-GPU Single Node |
| ENVI | Harris | Image Processing and Analytics | <ul style="list-style-type: none"> • Deep Learning training • Deep learning inferencing • Image orthorectification • Image transformation • Atmospheric correction • Panchromatic co-occurrence texture filter | Multi-GPU Single Node |
| Geomatics GXL | PCI | Image processing | <ul style="list-style-type: none"> • Image orthorectification • Additional image processing | Multi-GPU Single Node |
| GeoWeb3d Desktop | Geoweb3d | Geospatial visualization of 3D and 2D data, mensuration and mission planning | <ul style="list-style-type: none"> • 3D visualization and analysis of geospatial data | Multi-GPU Single Node |
| Graphistry | Graphistry | Graphistry is the first visual investigation platform to handle increasing enterprise-scale workloads. | <ul style="list-style-type: none"> • Graph reasoning • GPU-accelerated visual analytics • Visual pivoting • Rich investigation templating | Multi-GPU Single Node |
| Ikena ISR | 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"> • Real-time super-resolution-based video enhancement on live streams • Geospatial visualization • Target detection and tracking • Fast 2-D mapping | Multi-GPU Single Node |
| LuciadLightspeed | Hexagon Gespatial | Geospatial visualization and analysis | <ul style="list-style-type: none"> • Geospatial situational awareness | Single GPU Single Node |
| Manifold Systems | Manifold Systems | Full-featured GIS, vector/raster processing & analysis | <ul style="list-style-type: none"> • Manifold surface tools | Multi-GPU Single Node |
| OmniSIG | 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"> • Operates in a real-time streaming fashion • Ingests radio samples from many common radio interfaces • Make use of packet formats like VITA49 or SDDS. • Can be used from any device with a browser, including mobile handsets • OmniSIG software also provides its metadata output stream in JSON form for use by other applications | Multi-GPU Single Node |
| SNEAK | OpCoast | Electromagnetic signals propagation modeling for complex urban and terrain environments. | <ul style="list-style-type: none"> • Ray tracing, DTED and remote sensing inputs | Multi-GPU Single Node |
| SocetGXP | 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"> • Automated 3D feature extraction | Multi-GPU Single Node |
| Terrabuilder PhotoMesh | 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"> • 3D model building from imagery • Building texture generation | Multi-GPU Single Node |

Design for Manufacturing/Construction: CAD/CAE/CAM

CFD (MFG)

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--|----------------------|---|--|---------------------------|
| ADS Flow Solver - Code LEO | ADSCFD, Inc. | CFD simulation for turbochargers, turbines, and compressors | <ul style="list-style-type: none"> • URANS, structured/unstructured solver, explicit density based | Multi-GPU Multi-Node |
| Altair AcuSolve | 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"> • Linear equation solver | Multi-GPU Single Node |
| Altair nanoFluidX | 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"> • Extremely fast • Single and Multiphase Flows • Arbitrary motion definition • Time-dependent acceleration • Inlets/outlets • Surface tension and adhesion • Steady-state thermal solutions through coupling | Multi-GPU Multi-Node |
| Altair ultraFluidX | 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"> • CUDA-accelerated high-fidelity flow field computations based on the Lattice Boltzmann method • CUDA-aware MPI support for multi-GPU and multi-node usage • Efficient implementation of tailor-made automotive features, including rotating wheels, belt systems, boundary layer suction and porous media support | Multi-GPU Multi-Node |
| ANSYS Fluent | ANSYS | General purpose CFD software | <ul style="list-style-type: none"> • Linear equation solver • Radiation heat transfer model • Discrete Ordinate Radiation model | Multi-GPU Multi-Node |
| ANSYS Icepak | ANSYS | CFD software for electronics thermal management | <ul style="list-style-type: none"> • Linear Equation Solver | Multi-GPU Multi-Node |
| ANSYS Polyflow | ANSYS | CFD software for the analysis of polymer and glass processing | <ul style="list-style-type: none"> • Direct Solvers | Multi-GPU Single Node |
| CPFD Barracuda-VR and Barracuda | CPFD | Fluidized bed modeling software | <ul style="list-style-type: none"> • Linear equation solver • Particle calculations | Single GPU Single Node |
| DYVERSO | Realflow | 3D modeling, animation, and rendering | <ul style="list-style-type: none"> • Fluid solver (DY-SPH, DY-PBD) | Single GPU Single Node |
| EXN/Aero | Envenio | On-demand HPC-cloud CFD solver | <ul style="list-style-type: none"> • Multiphase, heat transfer, buoyancy, multi-grid, concurrent single/double precision, ideal gas, incompressible & compressible flows, RANS/LES/DES, conjugate heat transfer | Multi-GPU Multi-Node |
| FFT Actran | 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"> • Discontinuous Galerkin Method (DGM) solver | Multi-GPU Single Node |
| Fine/Open | 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"> • Incompressible, low and high speed flows • Efficient preconditioned compressible solver with fast agglomerated multigrid acceleration and adaptation techniques to combine completely unstructured hexahedral grids | Multi-GPU Multi-Node |

| | | | | |
|-----------------------|--------------------------|---|---|---------------------------|
| FINE/Turbo | Numeca International | Structured, multi-block, multi-grid CFD solver targeting the turbo machinery industry | <ul style="list-style-type: none"> Multi-grid solver | Multi-GPU Multi-Node |
| GeoPlat-RS | 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"> CUDA Spectral Decomposition with CUFFT library | Multi-GPU Single Node |
| HiFUN | 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"> HiFUN imbibes most recent CFD technologies; many of them home grown HiFUN exhibits highly scalable parallel performance with its ability to scale upto several thousand processors on massively parallel computing platforms Capable of handling complex geometries and flow physics arising in high lift flows | Multi-GPU Single Node |
| JSCAST | 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"> Basic module includes pre-/post processors, solvers and material property database. Optional modules include models for mold filling, solidification, casting deformation due to macro-shrinkage or due to the influence of back-pressure | Single GPU Single Node |
| midas NFX(CFD) | Midas | General purpose CFD software based on FEM | <ul style="list-style-type: none"> Linear equation solver (Iterative Solver and AMG Preconditioner) | Single GPU Single Node |
| MIKE 21 | DHI | 2D hydrological modelling of coast and sea | <ul style="list-style-type: none"> Hydrodynamics Advection-dispersion Sand and mud transport Coupled modelling Particle tracking Oil spill Ecological modelling Agent based modelling Various wave models | Multi-GPU Single Node |
| MIKE 3 | DHI | 3D Modeling of Coast and Sea | <ul style="list-style-type: none"> Hydrodynamics Advection-dispersion 8 Agent based modeling Underwater acoustic simulator, sand transport, mud transport, particle tracking, oil spill, ecological modeling, agent based modeling | Multi-GPU Multi-Node |
| MIKE FLOOD | DHI | 1D & 2D urban, coastal, and riverine flood modelling | <ul style="list-style-type: none"> Hydrodynamics | Multi-GPU Single Node |
| Numerix | Zeus | Simulation of flow around buildings | <ul style="list-style-type: none"> Discrete computational technique | Multi-GPU Single Node |

| | | | | |
|--------------------------|---------------------------------|---|---|---------------------------|
| Pacefish | Numeric Systems GmbH | CFD application for ground transportation and building aerodynamics | <ul style="list-style-type: none"> • Lattice-Boltzmann Method for single-phase flows • Transient simulation • Isothermal modeling • Integrated fast and robust pre-processor for complex geometries • Local grid refinement • uRANS (K-Omega-SST), hybrid uRANS-LES (SST-DDES & SST-IDDES) • LES (Smagorinsky) turbulence modeling | Multi-GPU Single Node |
| Particleworks | 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"> • Viscosity model, viscosity/pressure term solution, turbulence model, airflow, surface tension model, rigid body, thermal properties, external force and aeration. • Boundary conditions: particle wall, polygon wall, inflow and outflow boundary, simulation domain, and pump. | Multi-GPU Multi-Node |
| PowerViz | Dassault Systèmes SIMULIA Corp. | Industry proven, modern post-processing app for EXA POWERFLOW CFD | <ul style="list-style-type: none"> • Rendering • Ray tracing | Multi-GPU Single Node |
| Siemens STAR-CCM+ | Siemens Digital Industries | Post-processing for CFD-focused multiphysics simulation | <ul style="list-style-type: none"> • Rendering | Single GPU Single Node |
| Simcenter 3D | Siemens PLM Software | Industry proven, modern pre- & post-processing app for multidiscipline CAE | <ul style="list-style-type: none"> • Rendering • Raytracing | Multi-GPU Single Node |
| Speed IT FLOW | Vratis | Incompressible single-phase CFD software | <ul style="list-style-type: none"> • Finite-volume solver: Simple and piso, incompressible single-phase flows with k-OmegaSST turbulence | Single GPU Single Node |
| Turbostream | Turbostream Ltd. | CFD software for turbomachinery flows | <ul style="list-style-type: none"> • Explicit solver | Multi-GPU Multi-Node |
| zCFD | Zenotech Simulation Unlimited | General purpose CFD solver | <ul style="list-style-type: none"> • Turbulent flow (RANS, URANS, DDES or LES) including automatic scalable wall functions | Multi-GPU Single Node |

CFD (RESEARCH DEVELOPMENTS)

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------|--|--|--|--------------------------|
| ALYA | 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"> • Incompressible Flows • Compressible Flows • Non-linear Solid Mechanics • Species transport equations • Excitable Media • Thermal Flows • N-body collisions | Multi-GPU Multi-Node |
| DualSPHysics | University of Manchester | SPH-based CFD software | <ul style="list-style-type: none"> • SPH model | Multi-GPU Single Node |
| HiPSTAR | University of Southampton and University of Melbourne - Sandberg | CFD software for compressible reacting flows | <ul style="list-style-type: none"> • Explicit solver | Multi-GPU Single Node |
| PyFR | Imperial College - Vincent | General purpose CFD software for compressible flows | <ul style="list-style-type: none"> • High-order explicit solver based on flux reconstruction method | Multi-GPU Multi-Node |
| RAPTOR | US DOE | CFD formulation of turbulent combustion for fuel injector and other engine applications | <ul style="list-style-type: none"> • Flow solver | Multi-GPU Multi-Node |
| S3D | Sandia and Oak Ridge NL | Direct numerical solver (DNS) for turbulent combustion | <ul style="list-style-type: none"> • Chemistry model | Multi-GPU Multi-Node |

COMPUTATIONAL STRUCTURAL MECHANICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------------|----------------------|---|---|---------------------------|
| Adams | MSC Software | Multi-Body Dynamics simulation software | <ul style="list-style-type: none"> • Rendering | Single GPU Single Node |
| Altair HyperWorks | 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"> • OpenGL v3.2 / OpenCL v2.0 support • Anti-aliasing | Single GPU Single Node |
| Altair OptiStruct | 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"> • Direct solver (BCS) • Eigenvalue solvers (AMSES and Lanczos) • Iterative solver (PCG) | Single GPU Single Node |
| ANSYS Mechanical | ANSYS | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Direct and iterative solvers | Multi-GPU Multi-Node |
| EDEM | DEM Solutions, Ltd. | Market-leading Discrete Element Method (DEM) software for bulk material simulation | <ul style="list-style-type: none"> • DEM solver | Single GPU Single Node |
| GranuleWorks | 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"> • Size distribution, contact force model, rolling resistance model, liquid bridge force model, van der Waals force model, heat transfer and external force. • Boundary conditions: polygon wall, inflow and outflow boundary, and simulation domain. • Coupling with Particleworks MPS solver: support for aeration and pumps | Multi-GPU Multi-Node |
| Helyx PEM | 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"> • Polyhedral Elements Method solver | Single GPU Single Node |
| Impetus Afea | Impetus Afea | Predicts large deformations of structures and components exposed to extreme loading conditions | <ul style="list-style-type: none"> • Non-linear Explicit Finite-Element Solver | Multi-GPU Single Node |
| Irazu | Geomechanica Inc. | Simulation and analysis tool for rock mechanics, involving large deformations, fracturing and multi-physics phenomena. | <ul style="list-style-type: none"> • Explicit 2D and 3D FEM and FDEM solvers • Coupled hydraulic, mechanical, transport, thermal and fracture processes | Single GPU Single Node |
| LS-Dyna Implicit | LSTC | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Linear equation solver | Multi-GPU Single Node |
| Marc | MSC Software | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Direct sparse solver | Multi-GPU Single Node |
| MatDEM | 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"> • Full product support on GPU | Multi-GPU Single Node |
| midas GTS NX | Midas | Simulation tool for geo-technical analysis | <ul style="list-style-type: none"> • Linear equation solver(Multi Frontal Solver) | Single GPU Single Node |
| midas NFX(Structural) | Midas | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Linear equation solver(Multi Frontal Solver) | Single GPU Single Node |
| MSC Nastran | MSC Software | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Direct sparse solver | Multi-GPU Single Node |
| NX Nastran | Siemens PLM Software | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Linear and nonlinear equation solver | Multi-GPU Multi-Node |

| | | | | |
|--------------------------------|------------------------------------|--|--|---------------------------|
| PERMAS-XPU | INTES GmbH | General purpose structural simulation software | <ul style="list-style-type: none"> • Linear Equation Solver | Single GPU Single Node |
| RecurDyn | FunctionBay, Inc. | Multi-Flexible Body Dynamics simulation software | <ul style="list-style-type: none"> • Rendering | Single GPU Single Node |
| Rocky DEM | Rocky DEM | Discrete Element Modeling (DEM)-based particle simulation software | <ul style="list-style-type: none"> • Explicit DEM solver (dry/sticky contact rheologies) • 1-way & 2-way coupling with ANSYS Fluent and ANSYS Mechanical | Multi-GPU Single Node |
| SIMULIA 3DEXPERIENCE | Dassault Systèmes SIMULIA Corp. | Realistic simulation solution (Uses Abaqus Standard for GPU computing) | <ul style="list-style-type: none"> • Direct sparse solver | Single GPU Single Node |
| SIMULIA Abaqus/Standard | Dassault Systèmes SIMULIA Corp. | Simulation and analysis tool for structural mechanics | <ul style="list-style-type: none"> • Direct sparse solver • AMS Solver • Steady State Dynamics | Multi-GPU Multi-Node |

DESIGN AND VISUALIZATION

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-----------------------------|------------------------------------|--|--|---------------------------|
| 3DEXCITE DeltaGen | Dassault Systèmes | High-end 3D visualization and realtime interaction to help increase visual quality, speed, and flexibility. | <ul style="list-style-type: none"> • Interactive ray tracing and global illumination. • Integration with Siemens TeamCenter. • Cluster support Realtime & Offline Production Process Integration and scene building. • Scene Analysis, Xplore DeltaGen, SDK for DeltaGen. | Multi-GPU Single Node |
| Abaqus/CAE | Dassault Systèmes SIMULIA Corp. | Complete solution for Abaqus finite element modeling, visualization, and process automation | <ul style="list-style-type: none"> • Rendering | Multi-GPU Single Node |
| Accelerad | 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"> • Up to forty times faster using OptiX • Renderings with large numbers of ambient bounces • Calculations over many thousands of sensor points • Fast simulation of annual climate-based daylighting metrics • 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. | N/A |
| Additive Mfg Toolkit | 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"> • CUDA | N/A |
| Allplan | Nemetschek | Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction. | <ul style="list-style-type: none"> • OpenGL and DirectX based GPU rendering | Single GPU Single Node |
| ANSA | BETA CAE Systems | Industry proven, modern pre-processing app for CAE | <ul style="list-style-type: none"> • OpenGL | Single GPU Single Node |
| ANSYS Discovery Live | 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"> • OpenGL-based visualization • CUDA-based Structural Stress, Modal, Fluid Dynamics, Thermal, Electrical Conduction and Coupled Multi-Physics simulations | Single GPU Single Node |

| | | | | |
|--|---------------------------|---|---|---------------------------|
| ANSYS SPEOS | ANSYS | Physically accurate optical simulation software dedicated to predictive illumination and optical performance of systems. SPEOS combines lighting performance modelling with extensive dedicated libraries and optimizer capabilities. Extensive set of specialized features, such as optical part design, optical sensors test, HUD design and analysis and infrared modelling. SPEOS Live Preview allows to quickly visualise the final aspect of your product integrated into a whole system and environment. High-fidelity visualization of the final result, based on unique human vision algorithm. | <ul style="list-style-type: none"> • SPEOS Live Preview • 360 degrees for immersive or observer view | Single GPU Single Node |
| ANSYS VRXPERIENCE for Headlamps and Sensors | 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. Speed up the engineering process at an early stage of development on digital test tracks. Drive your future car with realistic traffic conditions, including various weather, oncoming vehicles, and pedestrians scenarios, anticipating your vehicle's reaction to any critical situation. | <ul style="list-style-type: none"> • Multispectral Physics-based real time lighting simulation with multi-display capabilities (driving simulator). | Multi-GPU Multi-Node |
| ANSYS VRXPERIENCE for HMI and Perceived Quality | 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. Scalable rendering capabilities, ranging from rasterization to fully GPU raytraced SPEOS Live Preview. Test and validate the full cockpit design for HMIs, including virtual displays and actuators, combining accurate rendering and finger tracking. Full HMI evaluation for next-generation vehicles, using virtual reality and driving simulation. Take into account human factors analyses and cognitive workload with direct interaction with reactive virtual interfaces, from touchscreens to switches. Evaluate the relevance of the displayed information, in real time, for a safer drive. | <ul style="list-style-type: none"> • Physics-based real time lighting simulation with VR capabilities from HMD to CAVEs (multi-GPU, multi-node) • SPEOS Live Preview (raytracing) based on CUDA/OptiX benefiting from RTX architecture (single GPU) | Multi-GPU Multi-Node |
| ANSYS Workbench | ANSYS | Industry proven, modern pre- & post-processing app for CAE | <ul style="list-style-type: none"> • Rendering | Multi-GPU Single Node |
| Apex | MSC Software | Industry proven, modern pre- & post-processing app for CAE | <ul style="list-style-type: none"> • Rendering | Single GPU Single Node |
| ArchiCAD | Nemetschek | Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction. | <ul style="list-style-type: none"> • OpenGL and DirectX based GPU rendering | Single GPU Single Node |
| AutoCAD | Autodesk | 2D and 3D CAD designing, drafting, modeling, architectural drawing, and engineering software. | <ul style="list-style-type: none"> • Surface, mesh and solid modeling tools, model documentation tools, parametric drawing capabilities • Open GL • Native DWG support • GRID Support. | Single GPU Single Node |
| Avatar VR | NeuroDigital Technologies | Haptic VR gloves for training design or remote operation. | <ul style="list-style-type: none"> • PhysX | Single GPU Single Node |

| | | | | |
|---|---------------------------|--|--|---------------------------|
| CATIA | 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"> • GPU performance scaling • VR native integration with HTC Vive | Single GPU Single Node |
| CATIA Live Rendering | Dassault Systèmes | Realistic 3D Rendering on full CATIA 3D CAD model. | <ul style="list-style-type: none"> • Physically Based Rendering with no data preparation thanks to native NVIDIA Iray Photoreal integration and interactive realistic rendering using NVIDIA Iray IRT | Multi-GPU Single Node |
| COMSOL | COMSOL | General Purpose CAE simulation software | <ul style="list-style-type: none"> • OpenGL | Multi-GPU Single Node |
| Creo Parametric | 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"> • GPU accelerated real-time engineering simulation with Creo Simulation Live • Full scene anti-aliasing • Order independent transparency • Better lighting and enhanced shaded-with-edges mode • Immersive design environment with realistic materials | Single GPU Single Node |
| Easy 3D Scan | 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"> • OpenGL | Single GPU Single Node |
| FEMAP | Siemens PLM | Industry proven, modern pre- & post-processing app for CAE | <ul style="list-style-type: none"> • Rendering | Single GPU Single Node |
| IC.IDO | 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"> • NV Pro Pipeline (RiX) for OpenGL rendering • VRWorks SPS and VR SLI (NVLink support) • DesignWorks, including VR Occlusion Culling open source sample and OptiX | Multi-GPU Single Node |
| Inspire Studio/Render (formerly known as Evolve) | 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"> • NURBS modeling • PolyNURBS modeling • OpenGL 4.5 Core • OpenGL-based real-time high-quality rendering • Interactive high-quality rendering using Thea Render • Production rendering using Thea Render • Integrated "dark room" environment to manage render queue and post-processing of rendered images | Single GPU Single Node |
| Inventor | Autodesk | 3D mechanical design, documentation, and product simulation. | <ul style="list-style-type: none"> • Uses BIM for intelligent building components to improve design accuracy | Single GPU Single Node |
| Iray | NVIDIA | A ready-to-integrate, physically-based, photorealistic rendering solution. | <ul style="list-style-type: none"> • Iray Interactive • Iray Photoreal • Iray Server • Fast interactive ray tracing • Physically-based, global-illumination rendering • Distributed cluster rendering. | Multi-GPU Multi-Node |
| Iray for 3ds Max | Siemens Industry Software | A physically-based renderer plugin for Autodesk 3ds Max | <ul style="list-style-type: none"> • Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support and AI based denoising | Multi-GPU Multi-Node |

| | | | | |
|-----------------------|--------------------------------|--|--|---------------------------|
| Iray for Maya | Ox1 Software & Consulting GmbH | A physically-based renderer plugin for Autodesk Maya. | Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising | Multi-GPU Multi-Node |
| Iray for Rhino | migenius Pty Ltd | Iray plugin for Rhino | <ul style="list-style-type: none"> Iray Photoreal and Iray Interactive support VCA clustering Cloud rendering MDL support. | Single GPU Single Node |
| Iray Server | migenius Pty Ltd | The scaling solution for any Iray based application | Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support and AI-based denoising | Multi-GPU Multi-Node |
| LensMechanix | 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"> Optical product teams need an easier and faster way to get from design to manufacture LensMechanix is the answer LensMechanix is software for mechanical engineers who design housing for optical products in CAD With LensMechanix, mechanical engineers can access the complete design data of optical systems designed in OpticStudio and start designing the mechanical envelope right away They can then validate their mechanical design and fix issues before building a physical prototype | Single GPU Single Node |
| META | BETA CAE Systems | High-performance post-processing software for CAE | <ul style="list-style-type: none"> OpenGL | Single GPU Single Node |
| NX | Siemens Digital Industries | Siemens PLM Software premium design app with full Iray integration, supporting multi-gpu rendering. Still CPU bound for most tasks otherwise | Iray, MDL - see NX Ray Traced Studio Database Entry | Multi-GPU Multi-Node |
| OpticStudio | 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"> Share designs between OpticStudio and CAD packages as native files, giving mechanical engineers full access to the optical coordinate system and all critical dimensions there is no need for file format conversions which can cause loss of design data Simulate the impact of mechanical components on optical performance to uncover any issues and make informed design decisions Check for, and resolve errors, before building costly physical prototypes | N/A |
| Painter | Corel | Raster-based digital art application for drawing, sketching and painting. | <ul style="list-style-type: none"> GPU accelerated brushes | Single GPU Single Node |
| Patran | MSC Software | Industry proven, modern pre- & post-processing app for CAE | <ul style="list-style-type: none"> Rendering | Single GPU Single Node |
| Quark VR | 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"> CUDA | Single GPU Single Node |
| Recap PRO | 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"> Generation of 3D meshed models from laser scans or photos of an object GPU accelerated photogrammetry process from 2D to 3D 3D model display accelerated by GPU for smooth navigation of converted models in all display modes | Multi-GPU Single Node |

| | | | | |
|-----------------------------|----------------------------|--|---|---------------------------|
| RETOMO | BETA CAE Systems | New software for the generation of 3D-tesellated models from CT-scan images | <ul style="list-style-type: none"> • OpenGL | Single GPU Single Node |
| Review | PiXYZ | Imports any CAD data to prepare and experience your content with VR. | <ul style="list-style-type: none"> • Large CAD file support with NVIDIA Pascal Single Pass Stereo extension integration | Single GPU Single Node |
| Revit | Autodesk | Building Information Modeling (BIM) for architecture, engineering and construction. | <ul style="list-style-type: none"> • Modeling (BIM) to design, build, and maintain higher-quality, more energy-efficient buildings • GRID support | Single GPU Single Node |
| Rhino | McNeel & Assoc. | General purpose conceptual/ industrial design software for AEC and Manufacturing industries, including a real-time ray-traced display mode that is CUDA-based. | <ul style="list-style-type: none"> • Fast, scalable OpenGL pipeline leverages latest NVIDIA GPUs • GPU computed shaders and memory optimizations • Rhino v6 has real-time ray traced viewport mode • Photorealistic rendering leverages NVIDIA RTX technology | Single GPU Single Node |
| Siemens STAR-CCM+ VR | Siemens Digital Industries | Immersive VR for CFD results visualization | <ul style="list-style-type: none"> • HTC Vive virtual reality headset | Single GPU Single Node |
| Simpleware | Synopsys | 3D image data visualization, analysis and model generation software | <ul style="list-style-type: none"> • OpenGL | Single GPU Single Node |
| SketchUp | Trimble | SketchUp, formerly Google 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"> • OpenGL support | Single GPU Single Node |
| SolidEdge | Siemens Digital Industries | Mid range CAD option from Siemens | n/a | Single GPU Single Node |
| SOLIDWORKS | 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"> • High performance in Shaded, Shaded w/ Edges, and RealView modes, FSAA for sharp edges, Order Independent Transparency • Real time photorealistic renderings with SOLIDWORKS Visualize, an Iray-based application. | Single GPU Single Node |
| SOLIDWORKS Visualize | Dassault Systèmes | Easy to use photorealistic rendering software based on NVIDIA Iray | <ul style="list-style-type: none"> • Iray-based ray-tracing • Animation support • Network rendering • OptiX-based Artificial Intelligence denoiser | Single GPU Single Node |
| Spotscale | Spotscale | 3D reconstruction algorithms are tailored for buildings and urban environments. Using drones to captured data. | <ul style="list-style-type: none"> • cuDNN | Multi-GPU Single Node |
| Studio | PiXYZ | Interactively prepare & optimize any CAD data before using your favorite staging tool. | <ul style="list-style-type: none"> • Large scale CAD format • Support for multi-CAD file standard, prepare, optimize and heal your geometry before experiencing it in VR | Single GPU Single Node |
| Substance Designer | Allegorithmic | Material shader edition and market reference for procedural texture creation. | <ul style="list-style-type: none"> • Iray rendering including textures/ substances and bitmap texture export to render in any Iray powered compatible with MDL | Multi-GPU Single Node |
| Substance Painter | Allegorithmic | Intuitive interactive 3D painting software with physics and particle support. | <ul style="list-style-type: none"> • Iray rendering to enhance all artwork released with the software | Multi-GPU Single Node |
| Sunata | Atlas 3D | 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"> • Thermal simulation | Multi-GPU Single Node |

| | | | | |
|--------------------|----------------------------|---|--|---------------------------|
| T-FLEX CAD | Top Systems | 3D and 2D parametric design, simulation, photorealistic rendering | <ul style="list-style-type: none"> • High performance visualization • Real time photorealistic rendering • CUDA | Multi-GPU Single Node |
| TeamCenter | Siemens Digital Industries | Product lifecycle management solutions from design to simulation to production to service. | <ul style="list-style-type: none"> • Design software, NX, and PLM viewer applications • TcVis and Active Workspace • GRID support | Single GPU Single Node |
| UE4 | 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"> • GPU Accelerated Rendering on OpenGL, DirectX and Vulkan • Phys-X implemented | Multi-GPU Single Node |
| Vectorworks | Nemetschek | Building Information Modeling (BIM) enabled design software for the Architecture, Landscape, and Entertainment industries. | <ul style="list-style-type: none"> • OpenGL based GPU rendering | Multi-GPU Single Node |
| VRED | 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"> • Enhanced geometry behavior • Automotive product interoperability • Navigation in a scene • Import Alias layer structure • Asset Manager improvements • Integrated file converter • Analytic rendering modes • Gap Analysis tool • Oculus Rift support • Animation module • Multiple rendering modes • Subsurface scattering • Displacement mapping | Multi-GPU Single Node |
| WYSIWYG | 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"> • GPU accelerated Shaded Views and Virtual Views | Multi-GPU Single Node |
| ZLVE | Zerolight | Immersive customer experience with VR or web GPU streaming | VRS and foveated rendering for VR and 3D experience through AWS GPU streaming | Multi-GPU Single Node |

ELECTRONIC DESIGN AUTOMATION

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------|-------------|--|---|--------------------------|
| Altair Feko | Altair | Comprehensive computational electromagnetics (CEM) code used widely in the telecommunications, automobile, space and defense industries to solve high-frequency problems. | <ul style="list-style-type: none"> • FDTD solver • MoM solver • RL-GO solver • CMA Solver | Multi-GPU Single Node |
| ANSYS HFSS | 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"> • Transient solver • FEM solver • OpenGL rendering | Multi-GPU Single Node |
| ANSYS HFSS SBR+ | ANSYS | Simulation tool for installed antenna performance and antenna-to-antenna coupling. | <ul style="list-style-type: none"> • High-frequency solver • OpenGL rendering | Multi-GPU Single Node |
| ANSYS Maxwell | 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"> • Eddy Current Solver | Multi-GPU Single Node |

| | | | | |
|----------------------------------|------------------------------------|--|--|---------------------------|
| ANSYS Nexxim | 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"> • AMI analysis | Single GPU Single Node |
| CDP | 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"> • Simulation-based processing | Multi-GPU Multi-Node |
| CST MPHYSICS STUDIO | 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"> • Conjugated Heat Transfer Solver | Single GPU Single Node |
| CST STUDIO SUITE | 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"> • Transient Solver • Integral Equation Solver • Asymptotic Solver • Multilayer Solver | Multi-GPU Multi-Node |
| EMPro | KeySight | Modeling and simulation environment for analyzing 3D EM effects of high speed and RF/Microwave components. | <ul style="list-style-type: none"> • FDTD solver | Multi-GPU Single Node |
| JMAG | JMAG | FEA software for electromechanical design. Fast solver High quality mesh Advanced modeling technologies. | <ul style="list-style-type: none"> • EM transient solver • EM time harmonic solver • EM static solver | Multi-GPU Single Node |
| KeySight ADS | KeySight | Simulation tool for design of RF, microwave and high speed digital circuits | <ul style="list-style-type: none"> • Transient Convolution simulation with BSIM4 models | Single GPU Single Node |
| REMCOM XFDTD | REMCOM | 3D EM Simulation | <ul style="list-style-type: none"> • FDTD Solver | Multi-GPU Multi-Node |
| SEMCAD-X | SPEAG | 3D EM modeling and simulation | <ul style="list-style-type: none"> • FDTD solver | Single GPU Single Node |
| Serenity | Lucernhammer | EM Simulation (RCS) tool | <ul style="list-style-type: none"> • MoM solver | Multi-GPU Single Node |
| Sim4Life | ZMT Zurich MedTech AG | 3D Electromagnetics & Acoustic modeling and simulation | <ul style="list-style-type: none"> • FDTD and Acoustics Solvers | Multi-GPU Single Node |
| TrueMask MDP | D2S | GPU-accelerated simulation and data preparation for mask writing. | <ul style="list-style-type: none"> • Simulation-based processing | Multi-GPU Multi-Node |
| TrueModel | D2S | GPU-accelerated simulation and geometric checking of curvilinear shapes. | <ul style="list-style-type: none"> • Simulation-based processing | Multi-GPU Multi-Node |
| Virtuoso - Cadence | Cadence Design Systems | EDA design simulation. Primary app is Allegro | <ul style="list-style-type: none"> • Visualization and acceleration for EDA and CAD design software. | Multi-GPU Multi-Node |
| VSim for Electromagnetics | Tech-X Corporation | Physics Simulation and modeling software for EM | <ul style="list-style-type: none"> • FDTD solver | Single GPU Single Node |
| WIPL-D 2D Solver | WIPL-D | 2D EM modeling and simulation | <ul style="list-style-type: none"> • MoM Solver | Multi-GPU Single Node |
| WIPL-D Pro | WIPL-D | 3D EM modeling and simulation | <ul style="list-style-type: none"> • MoM (Method of Moments) Solver • DDS (Domain Decomposition Solver) | Multi-GPU Multi-Node |
| Wireless InSite | REMCOM | Uses Optix 4.1 for Ray-tracing and Propagation prediction | <ul style="list-style-type: none"> • X3D Ray Tracer | Multi-GPU Single Node |

INDUSTRIAL INSPECTION

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------------|-------------|---|---|---------------------------|
| Cognex VisionPro ViDi | 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"> • Feature localization and identification • Segmentation and defect detection • Object and scene classification • Text & character recognition | Single GPU Single Node |

| | | | | |
|----------------------------|----------------|--|---|---------------------------|
| HALCON | 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"> • Deep learning - pre-trained networks optimized for latency or precision • HALCON also provides an IDE for training neural networks • Sub-pixel detection, edge detection, counting, OCR, barcode reading, 3D reconstruction from stereo | Single GPU Single Node |
| IBM Visual Insights | IBM | 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"> • Cloud-based DL training, deployment on (spec'ed) edge server | Multi-GPU Single Node |

Media and Entertainment

ANIMATION, MODELING AND RENDERING

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------------|-----------------|---|--|---------------------------|
| 3ds Max | Autodesk | 3D modeling, animation, and rendering | <ul style="list-style-type: none"> • Faster interactive graphics • Availability of Arnold with AI denoising • Availability of Chaos V-Ray, Otoy Octane, Redshift, cebas finalRender third-party GPU renderers | Multi-GPU Single Node |
| Agisoft PhotoScan | 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"> • CUDA-accelerated photogrammetry solution • RTX opportunity | Multi-GPU Single Node |
| Altair Thea Render | 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"> • GPU-accelerated hybrid renderer • Advanced material layering system with subsurface scattering, displacement mapping, physical sun-sky and IES support | Multi-GPU Single Node |
| Beauty Box | Digital Anarchy | Automatic masking and skin retouching. | <ul style="list-style-type: none"> • GPU accelerated graphics and compute | Single GPU Single Node |
| Blender | Blender Inst | 3D modeling, rendering and animation | <ul style="list-style-type: none"> • GPU-accelerated viewport | Single GPU Single Node |
| Blender Cycles | Blender Inst | GPU renderer | <ul style="list-style-type: none"> • CUDA-accelerated rendering | Multi-GPU Single Node |
| Character Animator | Adobe | Character animator that uses your expressions & movements to animate characters in real-time | <ul style="list-style-type: none"> • Auto lip syncing • Deep Learning | Single GPU Single Node |
| Cinema 4D | Maxon | 3D modeling, animation, and rendering | <ul style="list-style-type: none"> • Increased model complexity at interactive rates • Support for Chaos V-Ray, Otoy Octane and Redshift third-party GPU renderers • Accelerated ProRender GPU rendering | Single GPU Single Node |
| Daz Studio | 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"> • GPU accelerated compute • Rendering via NVIDIA IRAY and Optix | Multi-GPU Single Node |
| Dimension CC | Adobe | 3D design tool enabling graphic designers to compose, adjust, and render photorealistic images. | <ul style="list-style-type: none"> • Accelerated graphics & MDL (Material Definition Language) | Single GPU Single Node |

| | | | | |
|-----------------------------|----------------------------|--|---|---------------------------|
| finalRender | Cebas | GPU renderer | <ul style="list-style-type: none"> • CUDA-based GPU rendering | Multi-GPU Single Node |
| HIERO Player | Foundry | Shot management, conform and review timeline | <ul style="list-style-type: none"> • Fluid, interactive playback | Single GPU Single Node |
| Houdini | SideFX | Procedural 3D modeling, animation and rendering | <ul style="list-style-type: none"> • Faster simulations | Multi-GPU Single Node |
| Indigo | Glare Technology | Unbiased, physically-based renderer. | <ul style="list-style-type: none"> • GPU-accelerated rendering | Multi-GPU Single Node |
| KATANA | Foundry | Powerful look development and lighting tool | Faster interactive graphics | Single GPU Single Node |
| Lightwave 3D | NewTek | 3D modeling, animation, and rendering | <ul style="list-style-type: none"> • Increased model complexity at interactive rates | Single GPU Single Node |
| LuxRender | LuxRender | GPU 3D Renderer | <ul style="list-style-type: none"> • GPU-accelerated ray tracing | Single GPU Single Node |
| Mantra | SideFX | Houdini Mantra renderer | <ul style="list-style-type: none"> • Much faster interactive rendering using OptiX AI de-noising | Single GPU Single Node |
| MARI | Foundry | 3D paint tool that allows painting directly onto 3D models | <ul style="list-style-type: none"> • Faster interactive painting | Single GPU Single Node |
| Massive | Massive | Simulation and visualization tools for autonomous agent driven animation for film, games, television, architecture and transportation. | <ul style="list-style-type: none"> • GPU accelerated effects | Single GPU Single Node |
| Maxwell | Next Limit | CUDA-accelerated interactive and final-frame renderer | <ul style="list-style-type: none"> • Unrestricted image resolution • Network rendering • De-noising | Multi-GPU Single Node |
| Maya | Autodesk | 3D modeling, animation, and rendering | <ul style="list-style-type: none"> • Increased model complexity and larger scenes • Availability of Chaos V-Ray, Otoy Octane and Redshift third-party GPU renderers | Single GPU Single Node |
| MODO | Foundry | 3D modeling, animation and rendering | <ul style="list-style-type: none"> • Increased model complexity, larger scenes | Single GPU Single Node |
| Motion Builder | Autodesk | Character animation and motion capture | <ul style="list-style-type: none"> • Increased model complexity at interactive rates | Single GPU Single Node |
| Mudbox | Autodesk | 3D sculpting | <ul style="list-style-type: none"> • Increased model complexity at interactive rates | Single GPU Single Node |
| NX Ray Traced Studio | Siemens Digital Industries | Embedded rendering application for Siemens NX | <ul style="list-style-type: none"> • iRay based • MDL • AI denoising | Multi-GPU Single Node |
| Octane Render | Otoy | CUDA-accelerated GPU renderer | <ul style="list-style-type: none"> • GPU rendering | Multi-GPU Single Node |
| Realflow | Next Limit | Fluid simulation system | <ul style="list-style-type: none"> • GPU-accelerated simulation | Single GPU Single Node |
| RealityCapture | Capturing Reality | Photogrammetry | <ul style="list-style-type: none"> • CUDA-accelerated, fast photogrammetry | Multi-GPU Single Node |
| Redshift Renderer | Redshift | GPU-accelerated, biased renderer | <ul style="list-style-type: none"> • CUDA-based GPU final-frame rendering • Mac and Windows supported | Multi-GPU Single Node |
| Sculptris | Pixologic | 3D sculpting | <ul style="list-style-type: none"> • Increased model complexity at interactive rates | Single GPU Single Node |
| Trapcode | Red Giant | Particle simulations and 3D effects for motion graphics and VFX. Now with Fluid Dynamics. | <ul style="list-style-type: none"> • GPU accelerated effects | Single GPU Single Node |
| TurbulenceFD | Jawset | Turbulence FD is a powerful simulation tool to create smoke, fire and explosion effects. | <ul style="list-style-type: none"> • GPU accelerated graphics, compute and simulation | Single GPU Single Node |
| V-Ray GPU | Chaos Group | GPU renderer with CPU Hybrid rendering | <ul style="list-style-type: none"> • CUDA interactive and final-frame GPU rendering | Multi-GPU Single Node |

COLOR CORRECTION AND GRAIN MANAGEMENT

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------------------|-----------------------|--|--|---------------------------|
| ARRI de-bayering SDK | ARRI | RAW de-bayering SDK | <ul style="list-style-type: none"> De-bayering of ARRI RAW and primary color grading. | Single GPU Single Node |
| Baselight | FilmLight | Color grading | <ul style="list-style-type: none"> Real-time color correction | Multi-GPU Single Node |
| Cinema RAW SDK | Canon | RAW de-bayering | <ul style="list-style-type: none"> GPU-accelerated de-bayering | Single GPU Single Node |
| Dark Energy | Cinnafilm | Application and plug-in for image enhancement | <ul style="list-style-type: none"> Image de-noising and restoration | Multi-GPU Single Node |
| DaVinci Resolve | Blackmagicdesign | Color grading and editing | <ul style="list-style-type: none"> Real-time color correction and de-noising | Multi-GPU Single Node |
| DeNoise AI | 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"> GPU accelerated effects | Single GPU Single Node |
| Diamant-Film Restoration | HS-Art | Film cleanup and restoration | <ul style="list-style-type: none"> CUDA accelerated optical flow, de-flicker, in-painting and over 30 filters | Multi-GPU Single Node |
| Grain and Noise Reducer | Wavelet Beam | Video noise reduction | <ul style="list-style-type: none"> CUDA-accelerated grain and noise reduction | Multi-GPU Single Node |
| HDR Image Analyser | 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"> Precise, high quality UltraHD UI for native-resolution picture display Advanced out of gamut and out of brightness detection with error intolerance Support for SDR (Rec.709), ST2084/PQ and HLG analysis CIE graph, Vectorscope, Waveform, Histogram Out of gamut false color mode to easily spot out of gamut/out of brightness pixels Data analyzer with pixel picker Up to 4K/UltraHD 60p over 4x 3G-SDI inputs SDI auto signal detection File base error logging with timecode Display and color processing look up table (LUT) support Line mode to focus a region of interest onto a single horizontal or vertical line Loop through output to broadcast monitors Still store Nit levels and phase metering Built-in support for color spaces from ARRI, Canon, Panasonic, RED and Sony | Single GPU Single Node |
| Magic Bullet Colorista | 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"> GPU accelerated effects | Single GPU Single Node |
| Magic Bullet Looks | Red Giant | Powerful looks and color correction for filmmakers. | <ul style="list-style-type: none"> GPU accelerated compute | Single GPU Single Node |
| Mist | Marquise Technologies | Mastering tool for cinema, broadcast and over-the-top content | <ul style="list-style-type: none"> CUDA-accelerated de-bayering, color grading, transcoding and image enhancement | Multi-GPU Single Node |

| | | | | |
|-----------------------------------|--------------------|---|--|---------------------------|
| Nucoda | Digital Vision | Color grading | • GPU-accelerated color grading | Single GPU Single Node |
| Pablo family | Grass Valley | Color grading and finishing | • Real time color correction | Multi-GPU Single Node |
| PFClean | The Pixel Farm | Image restoration and remastering | • CUDA-based image processing acceleration | Multi-GPU Single Node |
| RAW Converter | ARRI | RAW de-Bayering and primary color grading | • CUDA-accelerated de-bayering and primary grading | Single GPU Single Node |
| Red Digital Cinema R3D SDK | Red Digital Cinema | Red Digital Cinema camera SDK that decodes and de-bayers Red RAW camera data, and allows primary color grading. Used by many color grading and video editing applications. | • CUDA-accelerated wavelet decoding and de-bayering | Single GPU Single Node |
| REDCINE-X PRO | Red Digital Cinema | Primary color grading | • CUDA-accelerated de-bayering and primary color grading | Single GPU Single Node |
| Scratch | Assimilate | Color grading and finishing | • Accelerated debayering for real-time digital finishing | Single GPU Single Node |
| SpeedGrade CC | Adobe | Color grading for editors, filmmakers, colorists and visual effects artists. As of March 2019, SpeedGrade is being phased out in favor of Lumetri Color tools in Premiere Pro | • GPU accelerations for real-time grading and finishing with Lumetri Deep Color Engine | Single GPU Single Node |
| VFX Suite | Red Giant | VFX Suite is a complete set of visual effects and motion graphics plugins for creating professional effects. | • GPU accelerated effects | Single GPU Single Node |

COMPOSITING, FINISHING AND EFFECTS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------------|-------------------|---|--|---------------------------|
| After Effects | Adobe | Motion graphics and effects | • CUDA acceleration for up to 10x faster performance on key effects plus enhanced 3D ray tracing | Multi-GPU Single Node |
| Clipster | Rohde & Schwarz | Video and film player and DCI Packager | • GPU-accelerated • Video scaling • Color space conversion • Data format conversion | Multi-GPU Single Node |
| Continuum 2019 | Boris FX | Visual effects plug-in for creative effects, titling, and quick fixes. | • GPU accelerated effects | Single GPU Single Node |
| DE:Noise | 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. | • Faster effects | Single GPU Single Node |
| Element 3D | Video CoPilot | Advanced 3D object & particle render engine plugin for Adobe After Effects | • GPU accelerated graphics and compute | Single GPU Single Node |
| Flame Premium | Autodesk | Finishing and color grading | • Integrated toolset for 3D VFX, editorial, and color grading | Multi-GPU Single Node |
| Flicker Free | Digital Anarchy | Deflicker Time Lapse, Slow Motion, and Old Video. Flicker Free is a powerful, new way to deflicker video. | • GPU accelerated effects | Single GPU Single Node |
| Fusion | Blackmagicdesign | Effects and compositing | • Faster effects | Single GPU Single Node |
| HIERO | Foundry | Multi-shot management tool that supports collaborative working, review and approval, quick production turnaround and delivery | • Fluid, interactive playback | Single GPU Single Node |
| Magic Bullet Denoiser | Red Giant | Magic Bullet Denoiser III lets you reduce visible noise and grain in digital video produced by digital video cameras, camcorders, or film. | • GPU accelerated effects | Single GPU Single Node |

| | | | | |
|--------------------------|----------------------|--|---|---------------------------|
| Magic Bullet Film | 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"> • GPU accelerated effects | Single GPU Single Node |
| Mamba FX | SGO | High-end compositing | <ul style="list-style-type: none"> • Faster keying, tracking, painting and restoration | Single GPU Single Node |
| MediaReactor | Drastic Technologies | Debayering and processing of raw camera files. | <ul style="list-style-type: none"> • GPU-accelerated compute | Single GPU Single Node |
| Mistika Ultima | SGO | Color grading and finishing | <ul style="list-style-type: none"> • Faster keying, tracking, painting and restoration, de-bayering | Single GPU Single Node |
| Mistika VR | SGO | Near real-time optical flow stitching | <ul style="list-style-type: none"> • GPU-accelerated video stitching with manual controls • Compatible with most camera rigs • Stitch, review and improve results in seconds • Export clips in many formats, including DPX and ProRes | Single GPU Single Node |
| Mocha Pro 2019 | 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"> • GPU accelerated planar tracking and object removal | Single GPU Single Node |
| Neat Video | Absoft | Digital filter with auto-profiling tool designed to reduce visible noise and grain found in footage. | <ul style="list-style-type: none"> • GPU accelerated processing | Single GPU Single Node |
| NUKE | Foundry | Compositing tool with 3D tracker | <ul style="list-style-type: none"> • GPU-accelerated BLINK processing • Faster compositing and effects | Single GPU Single Node |
| PFTTrack | The Pixel Farm | 3D scene creation and tracking | <ul style="list-style-type: none"> • CUDA-accelerated tracking | Multi-GPU Single Node |
| Sapphire 2019 | 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"> • Faster effects | Single GPU Single Node |
| Twixtor | 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"> • Faster effects | Single GPU Single Node |
| Video Essentials | NewBlueFX | Comprehensive collection of titling, transitions and video effects. | <ul style="list-style-type: none"> • Faster effects | Single GPU Single Node |

(VIDEO) EDITING

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|----------------------------------|--------------|---|--|---------------------------|
| Catalyst Production Suite | Sony | 4K, Sony RAW, and HD video editing. Includes 3 applications: Browse, Prepare, Edit | <ul style="list-style-type: none"> • Faster effects, transitions and encoding • RAW camera de-bayering | Single GPU Single Node |
| Edius Pro | Grass Valley | Video editing | <ul style="list-style-type: none"> • Faster effects • RAW de-bayering | Single GPU Single Node |
| Final Cut Pro | Apple | Video editing | <ul style="list-style-type: none"> • Faster effects | Single GPU Single Node |
| Gigapixel AI | Topaz Labs | Photo up scaling by using AI to "fill in" and add new detail when enlarging photos. | <ul style="list-style-type: none"> • GPU accelerated effects | Single GPU Single Node |
| HitFilm Pro | 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"> • GPU accelerated effects and decoding | Multi-GPU Multi-Node |

| | | | | |
|--------------------------|------------------------|---|--|---------------------------|
| Illustrator CC | 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"> • Entire canvas optimized for NVIDIA GPUs for faster pan & zoom. Accelerated by CUDA-based NV Path Render technology | Single GPU Single Node |
| Lightroom Classic | Adobe | Easily edit, s organizes, stores, and shares your photos. | <ul style="list-style-type: none"> • GPU accelerated Develop module plus new Sensei features like "Enhance Details" with NVIDIA GPU AI optimization | Single GPU Single Node |
| Lightworks | EditShare | Video editing | <ul style="list-style-type: none"> • Faster effects • CUDA-accelerated de-bayering | Single GPU Single Node |
| Live Planet | Live Planet | Livestreaming, recording and delivery of stereoscopic 360 VR | <ul style="list-style-type: none"> • Real time 360 3D capture and stitch • 4K | Single GPU Single Node |
| Media Composer | Avid | Video editing | <ul style="list-style-type: none"> • Faster video effects, unique stereo 3D capabilities | Single GPU Single Node |
| MXF | Film Partners | Collaborative editing system supporting Avid Media Composer, Adobe Premiere Pro, Grass Valley Edius and Blackmagic Resolve | <ul style="list-style-type: none"> • NVIDIA Video Codec allowing remote GPU-accelerated production workflows | Single GPU Single Node |
| Photoshop | Adobe | Photo editing to transform your images into anything you can imagine | <ul style="list-style-type: none"> • Over 30 GPU accelerated features such as blur gallery, liquify, smart sharpen, & perspective warp | Single GPU Single Node |
| Pinnacle Studio | Corel | Video editing and sharing program. | <ul style="list-style-type: none"> • GPU accelerated compute and effects | Single GPU Single Node |
| Pixvana | Pixvana | Pixvana is a Seattle software startup building a video creation and delivery platform for the emerging mediums of virtual, augmented, and mixed reality (XR). | <ul style="list-style-type: none"> • vGPU | Multi-GPU Multi-Node |
| PowerDirector | CyberLink | PowerDirector delivers professional-grade video editing and production for creators of all levels. Whether you're 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"> • GPU accelerated compute | Single GPU Single Node |
| Premiere Pro | Adobe | Video editing software for film, TV, and the web. | <ul style="list-style-type: none"> • Real-time video editing & fast output rendering based on CUDA | Multi-GPU Single Node |
| Premiere Rush | Adobe | Easy-to-use video editor for creating and sharing online videos. | <ul style="list-style-type: none"> • CUDA • Real-time video editing • fast output rendering | Multi-GPU Single Node |
| Qube | Grass Valley | Broadcast video editing | <ul style="list-style-type: none"> • Faster video effects • Unique stereo 3D capabilities | Single GPU Single Node |
| Sharpen AI | Topaz Labs | Sharpening and shake reduction software that can tell difference between real detail and noise. | <ul style="list-style-type: none"> • GPU accelerated effects • Machine Learning | Single GPU Single Node |
| SmartCourtPro | PlaySight | Sophisticated video and analytics training technology with the latest in AI, integrations and player development tools. | <ul style="list-style-type: none"> • IVA | Single GPU Single Node |
| Smoke | Autodesk | Finishing and editing | <ul style="list-style-type: none"> • Faster effects | Single GPU Single Node |
| Vegas Pro | Magix | Video editing | <ul style="list-style-type: none"> • Faster video effects and encoding | Single GPU Single Node |
| Velocity | Imagine Communications | Video editing | <ul style="list-style-type: none"> • Faster effects | Single GPU Single Node |
| Video Studio | Corel | High quality tools that build, edit, and correct video skillfully. | <ul style="list-style-type: none"> • GPU accelerated compute | Single GPU Single Node |

| | | | | |
|-------------------|-------|--|--|---------------------------|
| WonderLive | Z Cam | Cinematic VR Camera with excellent image quality, stereoscopic 360 Degrees; recording, and live streaming. | <ul style="list-style-type: none"> • Up to 4K output resolution equirectangular image • Save live stitched video file • Preview live stitched video • RTMP live streaming output • Supports VRworks 360 video SDK | Single GPU Single Node |
|-------------------|-------|--|--|---------------------------|

(IMAGE & PHOTO) EDITING

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-------------------|-------------|--|--|---------------------------|
| Adjust AI | 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"> • GPU accelerated effects | Single GPU Single Node |
| Corel Draw | Corel | Professional vector illustration, layout, photo editing and design tools. | <ul style="list-style-type: none"> • Faster effects | Single GPU Single Node |
| Neat Image | Absoft | Reduces noise, film grain, artifacts from photos. | <ul style="list-style-type: none"> • GPU accelerated processing | Single GPU Single Node |

ENCODING AND DIGITAL DISTRIBUTION

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------------------|--------------|--|---|---------------------------|
| Alchemist on Demand | Grass Valley | Video standards conversion | <ul style="list-style-type: none"> • GPU-accelerated video processing and encoding | Multi-GPU Single Node |
| Amberfin | Dalet | Transcoding and video quality analysis | <ul style="list-style-type: none"> • GPU-accelerated video processing and encoding | Single GPU Single Node |
| Aurora | Tektronix | Automated video quality measurement | <ul style="list-style-type: none"> • GPU-accelerated video quality assessment | Single GPU Single Node |
| AW-360C10 | Panasonic | 360-degree Live Camera designed for live sporting events, concerts and stadium events | <ul style="list-style-type: none"> • Low-latency, real-time 4K 360 degree stitching from four camera inputs using Jetson TX-1 • Control from tablet or over wi-fi from PC • Automatic exposure and white balance adjustment | Single GPU Single Node |
| Content Agent | Root6 | Automated transcoding and workflow management | <ul style="list-style-type: none"> • GPU-accelerated video processing and encoding | Multi-GPU Single Node |
| Core | ArcVideo | Video processing and transcoding Live | <ul style="list-style-type: none"> • Accelerated transcoding and encoding | Multi-GPU Single Node |
| Elemental Live | Elemental | Live streaming video processing and encoding | <ul style="list-style-type: none"> • Video encoding and video processing | Multi-GPU Single Node |
| Elemental Server | Elemental | File-based video processing and encoding | <ul style="list-style-type: none"> • Video encoding and video processing | Multi-GPU Single Node |
| Fast CinemaDNG Processor | Fastvideo | RAW video debayering, denoising and color correction completely on GPU side | <ul style="list-style-type: none"> • High-quality GPU-based RAW video processing up to 160 fps • Wavelet, realtime de-noising • Color correction features and monitoring • Export to 16-bit TIF or 10-bit ProResFull-sized video processing • Realtime 4K, 6K, and 8K playback supported | Multi-GPU Single Node |
| JPEG2000 Codec | Comprimato | JPEG2000 encoding and decoding for DCP, IMF, video editing, broadcast contribution, and archiving. | <ul style="list-style-type: none"> • Faster-than-real-time UltraHD / 4K • Lossy and mathematically lossless • High-bit-depth (HDR) | Multi-GPU Single Node |

| | | | | |
|---------------------------------------|-----------------------|--|---|---------------------------|
| Lightspeed Live | 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"> • Video processing and transcoding | Multi-GPU Single Node |
| Live | ArcVideo | High-density, real-time video processing and encoding. | <ul style="list-style-type: none"> • Accelerated broadcast encoding with NVIDIA CUDA and NVENC | Multi-GPU Single Node |
| Media Transcoding in the Cloud | 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"> • 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. • Ribbons Centralized Policy and Routing (PSX) can be instantiated as a Virtual Network Function (VNF) aligned with the ONAP architecture. • 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 • The Insight Element Management System (EMS) has an improved user interface for ease of use and offers improved provisioning and management processes | Single GPU Single Node |
| Medialooks SDK | Medialooks | MFormats SDK provides complete control over the video pipeline | <ul style="list-style-type: none"> • NVIDIA Video Codec used for accelerated encoding and encoding | Single GPU Single Node |
| Multiplatform Transcoder | ERLAB | Video processing and encoding software | <ul style="list-style-type: none"> • Pre-processing encoding, decoding, post-processing and delivery | Single GPU Single Node |
| mxSPEEDRAIL | 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"> • NVIDIA Video codec used for encoding for higher channel density • CUDA RAW de-coding, de-bayering, and video re-sizing and re-sampling | Single GPU Single Node |
| Piko TV | Kizil Elektronik | Linear broadcast encoder | <ul style="list-style-type: none"> • H.264 and HEVC 4K encoding for broadcast channels | Single GPU Single Node |
| PixelStrings | 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"> • Motion-compensated frame rate conversion • High-quality de-interlacing • Texture-aware scaling • Degrain/regrain to any film look, • Denoise/retexture to limit banding • Reverse telecine/pulldown pattern correction • Interlace artifact and dust removal • Runtime retiming | Multi-GPU Single Node |

| | | | | |
|--|-----------------------|---|---|---------------------------|
| Skywatch | MOG Technologies | Video and broadcast production management system for collecting audio/video usage and metadata. | <ul style="list-style-type: none"> • NVIDIA Video codec used for encoding for higher channel density • CUDA RAW de-coding, de-bayering, and video re-sizing and re-sampling | Single GPU Single Node |
| Smart Render Editor | Nablet | H.264 and HEVC video encoding using NV Video Codec | <ul style="list-style-type: none"> • Accelerated, high-density video encoding | Single GPU Single Node |
| Smart Render SDK | Nablet | Video de-noising, de-interlacing, JPEG 2000 encoding and video fingerprinting | <ul style="list-style-type: none"> • CUDA accelerated video processing • NVIDIA Video codec | Single GPU Single Node |
| Speech Quality transformed using Neural Network Computing | 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"> • Real time encoding/decoding of audio • Video signals | Single GPU Single Node |
| Tachyon | Cinnafilm | Standards conversion | <ul style="list-style-type: none"> • Video processing and frame rate conversion | Multi-GPU Single Node |
| Tornado | Marquise Technologies | Transcoding engine for IMF and DCP facilities | <ul style="list-style-type: none"> • Image re-sizing up to 8K • Color space conversion: 601/709, REC 2020, DCI XYZ, ACES 1.0 • De-bayering: ARRIRAW, DNG, RED R3D, SONY F65, F55 RAW, Phantom flex 4K, Canon C500 • Mezzanine: ProRes 444, Avid DNxHD 444, XDCAM, AVC Intra, AS-11 DPP, IMF • Uncompressed: DPX, TIFF, OpenEXR | Single GPU Single Node |
| Transkoder | Colorfront | Encoding and transcoding for DCP, and IMF mastering | <ul style="list-style-type: none"> • JPEG2000 encoding and decoding • 32-bit floating point processing on multiple GPUs • MXF wrapping, accelerated checksums and AES encryption and decryption, • IMF/IMP and DCI/DCP package authoring, editing, transwrapping | Multi-GPU Single Node |
| Vantage LightSpeed | 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"> • Video transcoding and processing | Multi-GPU Single Node |
| Viarte | Isovideo | Video standards conversion | <ul style="list-style-type: none"> • CUDA-accelerated video processing and encoding | Multi-GPU Single Node |
| VidiCert | Joanneum Research | Video and film quality assurance | <ul style="list-style-type: none"> • CUDA accelerated video quality analysis | Multi-GPU Single Node |
| Wowza Streaming Engine Transcoder | Wowza | H.264 video encoding | <ul style="list-style-type: none"> • NVENC accelerated video encoding | Single GPU Single Node |

ON-AIR GRAPHICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-----------------------------|------------------------|---|--|---------------------------|
| Air | Cinegy | Broadcast play-out server | <ul style="list-style-type: none"> Real-time on-air graphics | Single GPU Single Node |
| Brodcaast Dscript 3D | Monarch | 3D on-air graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |
| Camino | 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"> Camino's real-time graphics overlay can be applied to tickertapes, scoreboards, schedule boards, program junctions, and TV show promotions Graphics overlay may be done via predefined templates, which may then be populated with live data during payout Makes real-time rendering of data-driven graphics possible in news and sports events.4K, 1080p, 720p and SD Support NTSC and PAL Support Graphics, Clips and 3D Objects Importer 2D and 3D Primitives Real-Time Key-Frame Animations Real-Time 3D Scene Lighting Timeline-Based Audio Support Data Mapping to External Sources Transition Logic Automation Controller Support Stereoscopic 3D rendering | Single GPU Single Node |
| Capture | Cinegy | Video ingest | <ul style="list-style-type: none"> Uses NVENC to encode/decode multiple H.264 and HEVC streams | Single GPU Single Node |
| Clarity | Pixel Power | On-air graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |
| Cube | Dalet | On-air Graphics | <ul style="list-style-type: none"> Real-time graphics rendering | Single GPU Single Node |
| eStudio | Brainstorm | Virtual sets and motion graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |
| InfinitySet | Brainstorm | Realistic virtual sets | <ul style="list-style-type: none"> Real-time RTX ray tracing through UE4 HDR I/O Physically-based rendering | Single GPU Single Node |
| Livebook GFX | 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"> Graphics solution for compact live sports productions | Multi-GPU Single Node |
| Mosaic | ChyronHego | On-air graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |
| Multiviewer | Evertz | Broadcast multiviewer | <ul style="list-style-type: none"> Uses NVENC H.264 and HEVC encoding and decoding | Single GPU Single Node |
| Nexio Channelbrand | Imagine Communications | On-air graphics | <ul style="list-style-type: none"> Real-time rendering | Multi-GPU Single Node |
| Nexio G8 | Imagine Communications | On-air graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |
| Nexio TitleOne | Imagine Communications | On-air graphics | <ul style="list-style-type: none"> Real-time rendering | Single GPU Single Node |

| | | | | |
|-------------------------------|--------------|---|--|---------------------------|
| Reality Virtual Studio | Zero Density | Photorealistic virtual studio solution in broadcast industry, powered by Epic Unreal Engine 4 | <ul style="list-style-type: none"> • Node-based compositing system designed for real-time production • Same content can be used for broadcast and VR • 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 | Single GPU Single Node |
| Titler Pro | NewBlueFX | Create elegant video titles or 3D motion graphics. | <ul style="list-style-type: none"> • GPU-accelerated graphics | Single GPU Single Node |
| tOG | RT Software | On-air graphics | <ul style="list-style-type: none"> • Real-time rendering | Single GPU Single Node |
| Type | Cinegy | On-air Graphics | <ul style="list-style-type: none"> • Real-time graphics rendering | Single GPU Single Node |
| Vertigo | Grass Valley | On-air Graphics | <ul style="list-style-type: none"> • Real-time rendering | Single GPU Single Node |
| Virtuoso | Monarch | Virtual sets and motion graphics | <ul style="list-style-type: none"> • Real-time rendering | Single GPU Single Node |
| Viz Engine | vizrt | On-air graphics and virtual sets | <ul style="list-style-type: none"> • Real-time graphics rendering | Single GPU Single Node |
| Wasp3D - CG | Wasp3D | On-air graphics and virtual sets | <ul style="list-style-type: none"> • Real-time graphics rendering | Single GPU Single Node |

ON-SET , REVIEW AND STEREO TOOLS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------------|-----------------------|--|---|---------------------------|
| 4kScope | 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"> • GPU accelerated effects and compute | Single GPU Single Node |
| Cortex Dailies | MTI Film | Review, color grading and transcoding on set | <ul style="list-style-type: none"> • CUDA accelerated grading and transcoding | Multi-GPU Single Node |
| Fluid 4K Review | BlueFish444 | Review and approval of 4K content | <ul style="list-style-type: none"> • Real-time video review | Single GPU Single Node |
| ICE | Marquise Technologies | IMF reference video player | <ul style="list-style-type: none"> • RAW data support for ARRIRAW, DNG, RED R3D, SONY F65, F55 RAW, Phantom flex 4K and Canon C500 • HDR content encoded in Dolby Vision, HDR10, HDR10+ or HLG • Uncompressed formats support: DPX, TIFF and OpenEXR | Single GPU Single Node |
| Net-X-Code | 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"> • GPU accelerated compute | Single GPU Single Node |
| On-Set Dailies | Colorfront | Review, color grading and transcoding on set | <ul style="list-style-type: none"> • Real-time | Multi-GPU Single Node |
| Previzion | Lightcraft | On-set virtual production | <ul style="list-style-type: none"> • Real-time, virtual set production | Single GPU Single Node |
| VideoQC | 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"> • GPU accelerated effects and compute | Single GPU Single Node |

WEATHER GRAPHICS

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

Medical Imaging

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-----------------------|-------------------------------|---|--|---------------------------|
| aidoc | 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"> Classification and segmentation using deep learning on top of any PACS platform | Single GPU Single Node |
| deepflow | 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"> Tool will show that deep convolutional neural networks combined with nonlinear dimension reduction enable reconstructing biological processes based on raw image data Tool will demonstrate this by reconstructing the cell cycle of Jurkat cells and disease progression in diabetic retinopathy. In further analysis of Jurkat cells Tool will detect and separate a subpopulation of dead cells in an unsupervised manner and, in classifying discrete cell cycle stages 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 Uses MXNet, cv2, numpy, python3 | Single GPU Single Node |
| Ibex Decision Support | IBEX | IBEX run DL on prostate cancer digital pathology and to find any potential cancerous areas | <ul style="list-style-type: none"> Combines data from digitized glass slides and electronic medical records to reveal underlying patterns Extracts valuable clinical insights that can transform how pathology and oncology are practiced and propel them into the information age | Single GPU Single Node |
| MITK | German Cancer Research Center | Free open-source software system for development of interactive medical image processing software | <ul style="list-style-type: none"> Interactive segmentation of slices in image volumes, including interactive region growing and easy correction, interpolation of missing slices, surface generation, and volumetry 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.) Measurement of distances and angles; Volume visualization, GPU-based, easy to modify transfer functions; Movie generation (Windows only) Deformable Registration | Single GPU Single Node |

| | | | | |
|-------------------------|---|---|---|---------------------------|
| PowerGrid | University of Illinois Urbana-Champaign | Provides iterative non-cartesian MRI reconstruction | <ul style="list-style-type: none"> • GPU accelerated implementations of the non-Uniform FFT and Discrete Fourier Transform • MPI is used to enable using multiple GPUs in one or several machines • Iterative reconstruction using physics-based model to correct for unwanted effects, such as field inhomogeneity and patient motion | Multi-GPU Single Node |
| Proprio | 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"> • CUDA | Single GPU Single Node |
| RadiAnt | Medixant | RadiAnt DICOM Viewer provides basic tools for the manipulation and measurement of images | <ul style="list-style-type: none"> • Fluid zooming and panning, Brightness and contrast adjustments, negative mode, Preset window settings for Computed Tomography (lung, bone, etc.) • Ability to rotate (90, 180 degrees) or flip (horizontal and vertical) images, Segment length, Mean, minimum and maximum parameter values (e.g. density in Hounsfield Units in Computed Tomography) within circle/ellipse and its area, Angle value (normal and Cobb angle) • Pen tool for freehand drawing | Single GPU Single Node |
| Radiology Assist | 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"> • Classification and segmentation on top of any PACS platform | Single GPU Single Node |

Oil and Gas

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--------------------------|------------------------|---|--|---------------------------|
| 6X | Ridgeway Kite | Reservoir Simulation on Tesla | <ul style="list-style-type: none"> • CUDA Simulation Parallelization | Single GPU Single Node |
| AISight for SCADA | BRS Labs | Proactive integrity management and real-time precursor alerts for enhanced SCADA operations in oil and gas. | <ul style="list-style-type: none"> • 24/7 real-time analysis and alerting • Scales to thousands of sensors across remote and geographically dispersed locations • Historical analysis and trend reports | Multi-GPU Single Node |
| AxRTM | Acceleware | Reverse Time Migration Software | <ul style="list-style-type: none"> • CUDA accelerated libraries for building RTM software | Multi-GPU Multi-Node |
| DecisionSpace | Halliburton (Landmark) | E&P platform for geoscience, well planning, drilling and earth modeling. | <ul style="list-style-type: none"> • CUDA acceleration of fault extraction | Multi-GPU Single Node |
| Echelon | Stoneridge Technology | Reservoir simulator | <ul style="list-style-type: none"> • Fully GPU-accelerated reservoir model • Dual-perm, dual porosity, pressure varying perm and porosity • Eclipse compatible input deck | Multi-GPU Multi-Node |
| GeoDepth | Emerson | Seismic Interpretation Suite | <ul style="list-style-type: none"> • CUDA-accelerated RTM | Multi-GPU Multi-Node |
| Geoteric | Geoteric | Seismic interpretation | <ul style="list-style-type: none"> • Attributes calculations • Geobodies extraction | Multi-GPU Single Node |

| | | | | |
|----------------------------|--------------------------|---|--|--------------------------|
| Graydient S (SCADA) | Giant Grey | Machine learning anomaly detection for large scale industrial data. | <ul style="list-style-type: none"> • Proactive integrity management and real-time precursor alerts for enhanced SCADA operations in oil and gas • 24/7 real-time analysis and alerting scaling to thousands of sensors across remote and geographically dispersed location | Multi-GPU Single Node |
| HUESpace | Bluware | Library SDK toolkit for creating applications for seismic compression and seismic/geospatial imaging and interpretation. | <ul style="list-style-type: none"> • CUDA acceleration for compression • Large-scale visualization | Multi-GPU Single Node |
| InsightEarth | CGG | Seismic Interpretation Suite | <ul style="list-style-type: none"> • OpenCL acceleration for AFE • 3D Curvature attributes | Multi-GPU Single Node |
| Omega2 RTM | Schlumberger | Seismic processing | <ul style="list-style-type: none"> • Multiple algorithms (RTM, etc) | Multi-GPU Multi-Node |
| PumaFlow IFP | Beicip-Franlab | Reservoir simulation | <ul style="list-style-type: none"> • GPU-accelerated linear solver | Multi-GPU Single Node |
| Roxar RMS | Emerson | Reservoir modeling | <ul style="list-style-type: none"> • Multi GPU capabilities via HUESpace | Multi-GPU Single Node |
| RTM | Tsunami | Seismic processing | <ul style="list-style-type: none"> • RTM algorithm | Multi-GPU Multi-Node |
| Seismic City RTM | Seismic City | RTM Seismic Processing | <ul style="list-style-type: none"> • CUDA acceleration | Multi-GPU Multi-Node |
| SKUA | Emerson | Reservoir modeling | <ul style="list-style-type: none"> • Faults, Horizons and Flow Simulation Grid | Multi-GPU Single Node |
| tNavigator | 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"> • CUDA • Pascal/Volta architecture • Multi-GPU | Multi-GPU Multi-Node |
| VoxelGeo | Emerson | Seismic Interpretation Package | <ul style="list-style-type: none"> • Multi-GPU volume rendering • Horizon-flattening • Attribute calculations | Multi-GPU Single Node |

Life Sciences

BIOINFORMATICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-------------------|---|--|---|--------------------------|
| Arioc | Johns Hopkins University | High-throughput read alignment with GPU-accelerated exploration of the seed-and-extend search space. | <ul style="list-style-type: none"> • Single-end alignment, paired-end alignment • Output in SAM or database-ready binary formats • Multiple GPU implementation | Multi-GPU Single Node |
| BarraCUDA | University of Cambridge Metabolic Research Labs | Sequence mapping software | <ul style="list-style-type: none"> • Alignment of short sequencing reads • Alignment of indels with gap openings and extensions. | Multi-GPU Multi-Node |
| BEAGLE-lib | 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"> • Evaluation of likelihood for sequence evolution on trees and Arbitrary models (e.g. nucleotide, amino acid, codon) • Speed-ups (over CPU only version): nucleotide model = up to 25x, codon model = up to 50x. | Multi-GPU Single Node |

| | | | | |
|--------------------|---|--|--|---------------------------|
| Campaign | SimTK | An open-source library of GPU-accelerated data clustering algorithms and tools. | <ul style="list-style-type: none"> • K-means • Kps-means • K-medoids • K-centers • Hierarchical clustering • Self-organizing map | Multi-GPU Multi-Node |
| CUDASW++ | Open Source | Open source software for Smith-Waterman protein database searches on GPUs. | <ul style="list-style-type: none"> • Parallel search of Smith-Waterman database. | Multi-GPU Single Node |
| CUSHAW | Open Source | Parallelized short read aligner | <ul style="list-style-type: none"> • Parallel, accurate long read aligner for large genomes | Multi-GPU Single Node |
| G-BLASTN | Hong Kong Baptist University | GPU-accelerated nucleotide alignment tool based on the widely used NCBI-BLAST. | <ul style="list-style-type: none"> • Blastn and megablast modes of NCBI-BLAST | Single GPU Single Node |
| GHOST-Z GPU | Akiyama_Laboratory, Tokyo Institute of Technology | Sequence homology search tool. | <ul style="list-style-type: none"> • Shotgun Metagenome Analysis. | Multi-GPU Multi-Node |
| GPU-Blast | Carnegie Mellon University | Local search with fast k-tuple heuristic | <ul style="list-style-type: none"> • Protein alignment according to BLASTP | Single GPU Single Node |
| mCUDA-MEME | Open Source | Ultrafast scalable motif discovery algorithm based on MEME . | <ul style="list-style-type: none"> • Scalable motif discovery algorithm based on MEME | Multi-GPU Single Node |
| MUMmer GPU | Open Source | MUMmer GPU is a high-throughput local sequence alignment program | <ul style="list-style-type: none"> • Aligns multiple query sequences against reference sequence in parallel | Single GPU Single Node |
| NVBIO | 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"> • Data structures, algorithms • Utility routines useful for building complex computational genomics applications on CPU-GPU systems | Multi-GPU Single Node |
| NVBowtie | Open Source | A largely complete implementation of the Bowtie2 aligner on top of NVBIO. | <ul style="list-style-type: none"> • Good coverage of Bowtie2 features • Comparable quality results | Multi-GPU Single Node |
| PEANUT | Open Source | Read mapper for DNA or RNA sequence that reads to a known reference genome. | <ul style="list-style-type: none"> • Achieves supreme sensitivity and speed compared to current state of the art • Reads mappers like BWA MEM, Bowtie2 and RazerS3 • PEANUT reports both only the best hits or all hits | Single GPU Single Node |
| racon-gpu | 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"> • 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). • 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. | Single GPU Single Node |

| | | | | |
|------------------------|-------------------------------------|---|--|--------------------------|
| REACTA | 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"> • GRM creation • REML analysis • Regional Heritability (including multi-GPU) | Multi-GPU Single Node |
| Sanjeevani | SCFBio-IITD | Sanjeevani3 software suite has been developed as a computational pathway paving the way expressly towards automating lead candidate design: screening potential molecules for drug likeness, optimizing their geometry, determining partial atomic charges, assigning force field parameters, locating potential active sites on the target and docking the candidates in the active site of a given biological target. This is followed by a rigorous analysis of the binding free energies and further optimization of the candidate molecules. | <ul style="list-style-type: none"> • This code is able to sample 10 times more ligand conformations leading to increase in accuracy. The ACC ported ParDOCK code is now able to predict correct pose of a protein-ligand interaction from 96.8% percent times, compared to 94.3% earlier (for poses under 1) and 89.9% times compared to 86.7% earlier (for poses under 0.5) | Multi-GPU Multi-Node |
| SeqNFind | 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"> • Hardware and software for reference assembly, blast, SW, HMM, de novo assembly | Multi-GPU Single Node |
| SOAP3 | 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"> • Short read alignment tool that is not heuristic based • Reports all answers | Multi-GPU Multi-Node |
| SOAP3-dp | 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"> • Borrows-Wheeler Transformation • Dynamic Programming | Multi-GPU Single Node |
| Synomics Studio | Row Analytics | Multi-Omics Biomarker Network Discovery and ValidationSynomics 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"> • Multi-SNP association studies (GWAS studies with up to 30 SNPs/SNVs in combination) • 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) • Representative performance for 15,000 case:controls, 200,000 SNPs • 2 SNP associations found and validated in 12 mins on single 20 core IBM POWER8NVL with 4x Tesla P100 GPU • 17 SNP associations found and validated in 6 days on single 20 core IBM POWER8NVL with 4x Tesla P100 GPU | Multi-GPU Single Node |
| UGene | Unipro | Open source Smith-Waterman for SSE/CUDA, Suffix array based repeats finder and dotplot. | <ul style="list-style-type: none"> • Fast short read alignment | Multi-GPU Single Node |
| WideLM | Open Source | Fits numerous linear models to a fixed design and response. | <ul style="list-style-type: none"> • Parallel linear regression on multiple similarly-shaped models | Multi-GPU Single Node |

MICROSCOPY

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------|--|--|---|---------------------------|
| ANNA-PALM | 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"> • Uses a much smaller number of low resolution frames than other methods • Processing by localization algorithms results in a sparse localization image using a neural network previously trained on conventional PALM images • Inputs sparse image and outputs a super-resolution image • Runs well on GPU due to acceleration available in Tensorflow | Single GPU Single Node |
| Appion | New York Structural Biology Center | Appion is a "pipeline" for processing and analysis of EM images. Appion is integrated with Legion 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"> • 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. | Single GPU Single Node |
| BioEM | Max Planck Institute | GPU-accelerated computing of Bayesian inference of electron microscopy images. | <ul style="list-style-type: none"> • BioEM can use CUDA for the cross-correlation step, which essentially consists of an image multiplication in Fourier space and a Fourier back-transformation | Multi-GPU Single Node |
| crYOLO | 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 http://sphire.mpg.de/ and will be part of the image processing workflow in SPHIRE. | <ul style="list-style-type: none"> • Part of the image processing workflow in SPHIRE. | Multi-GPU Single Node |
| cryoSPARC | 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"> • Ab-initio reconstruction • Heterogeneous reconstruction • High-speed and high resolution refinement of 3D protein structures implemented on GPUs • Multiple simultaneous jobs on multiple GPUs | Multi-GPU Multi-Node |
| Dynamo | 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"> • Dynamo provides workflows all the way from tomograms to averages and classes. • 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 • Requires CUDA Toolkit of version 7.5 or higher and CUDA driver compatible with your actual GPU device | Single GPU Single Node |

| | | | | |
|-------------------|-------------------------------------|---|--|---------------------------|
| emClarity | 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"> • Subtomogram averaging • Very high resolution single particle analysis • Hybrid electron microscopy. | Multi-GPU Single Node |
| Gautomatch | 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"> • Fast: typically, 1.5~2.0s with 15 templates, using a good GPU (e.g. GTX 980, Titan X) • Fully automatic with simple command on entire data sets • Convenient and easy to use • Flexible: with or without template, suitable for both basic or advanced users • Compatible with Relion/EMAN • Background correction: automatic correct the gradient background that affects the picking • Rejection of ice/carbon: automatically detect non-particle areas and reject them • Post-optimization: scripts available to re-filter the coordinates after picking within seconds • Accuracy: the user's satisfaction is the only "gold standard" criterion | Single GPU Single Node |
| Huygens | Scientific Volume Imaging | Huygens Products: Greatly improve your microscope images | <ul style="list-style-type: none"> • Deconvolution of volumetric images and time series from widefield, confocal, light sheet, super-resolution STED microscopes and more • Chromatic aberration and cross-talk correction, image stabilization and stitching • Visualization, tracking, colocalization and object analysis • Multi-GPU and cluster support | Multi-GPU Single Node |
| IMOD | 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"> • ctphaseflip : Corrects tilt series for microscope CTF by phase flipping • gputilttest : Test whether a GPU is reliable for computing reconstructions with the tilt program • 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. • xyzproj : Project 3-dimensional data at a series of tilts around the X, Y, or Z axis. | Single GPU Single Node |

| | | | | |
|----------------------|-------------------------------------|--|--|---------------------------|
| ITK | 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"> • Library is used by Paraview, VTK, and many other software distributions • Many capabilities for multi-dimensional image processing and extraction tools • Most recent GPU acceleration of FFTs using cuFFT (cuFFTW) and matrix math accelerated through CUDA enabled Eigen3 | Single GPU Single Node |
| Leginon | 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"> • A Leginon application is image acquisition process that is built of several smaller pieces called 'nodes' • Nodes can be applications • Some of these are GPU accelerated applications such as Topaz, Relion, and MotionCor2 | Single GPU Single Node |
| Microvolution | Microvolution | Nearly instantaneous 3D deconvolution & up to 200 times faster. | <ul style="list-style-type: none"> • 3D deconvolution for fluorescence microscopy • Written for use only on GPUs • Multi-GPU support | Single GPU Single Node |
| MotionCor2 | 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"> • 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 • Works on a wide range of data sets including cryo tomographic tilt series | Multi-GPU Single Node |
| RELION | 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"> • Image classification and high resolution refinement accelerated up to 40-fold • Template-based particle selection accelerated almost 1000-fold • Reduced memory requirements • High-resolution cryo-EM structure determination in a matter of day on a single workstation | Multi-GPU Single Node |
| Thunder | 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"> • Both image classification and highresolution refinement accelerated up to 40-fold • Template-based particle selection accelerated almost 1000-fold • Reduced memory requirements • High-resolution cryo-EM structure determination in a matter of day on a single workstation | Multi-GPU Multi-Node |
| Topaz | 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"> • Deep learning for cryo EM data particle picking • Uses CUDA and pytorch | Single GPU Single Node |

| | | | | |
|-------------|--|--|---|---------------------------|
| Warp | 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"> • CUDA enabled processing for electron microscopy • TensorFlow (v1.10) • CUDA kernels: backprojection, CTF, deconvolution, FFT, tomography refinement, and others | Single GPU Single Node |
|-------------|--|--|---|---------------------------|

MOLECULAR DYNAMICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------|---|---|--|--------------------------|
| ACEMD | Acellera Ltd | GPU simulation of molecular mechanics force fields, implicit and explicit solvent | <ul style="list-style-type: none"> • Written for use only on GPUs. | Multi-GPU Multi-Node |
| AMBER | University of California at San Francisco | Suite of programs to simulate molecular dynamics on biomolecule. | <ul style="list-style-type: none"> • PMEMD Explicit Solvent and GB Implicit Solvent | Multi-GPU Single Node |
| CHARMM | Harvard University | MD package to simulate molecular dynamics on biomolecule. | <ul style="list-style-type: none"> • Implicit (5x) • Explicit (2x) • Solvent via OpenMM, now ported natively to GPUs | Multi-GPU Single Node |
| DESMOND | David E. Shaw Research | High-speed molecular dynamics simulations of biological systems. | <ul style="list-style-type: none"> • The code uses novel parallel algorithms and numerical techniques to achieve high performance and accuracy | Multi-GPU Single Node |
| ESPreSo | 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"> • Hydrodynamic Electrokinetic forces • P3M electrostatics. | Multi-GPU Single Node |
| FEP+ | 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"> • 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. | Multi-GPU Multi-Node |
| Folding@Home | Stanford University | A distributed computing project that studies protein folding, misfolding, aggregation, and related diseases. | <ul style="list-style-type: none"> • Powerful distributed computing molecular dynamics system • Implicit solvent and folding | Multi-GPU Single Node |
| Galamost | 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"> • Full Molecular Simulation on GPU | Multi-GPU Multi-Node |

| | | | | |
|--|------------------------------|---|---|---------------------------|
| GALAMOST | 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"> • General molecular dynamics • Dissipative particle dynamics (DPD) • Brownian dynamics (BD) • Coarse-graining molecular dynamics (CGMD) • Reaction model • Anisotropic particle models • MD-SCF • DNA 3SPN model • Rigid body method • Stretching method | Single GPU Single Node |
| Genesis | 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"> • Powerful parallelization for hybrid (CPU+GPU) systems • Full electrostatics with PME • Large (1-100 million atoms) biological systems | Multi-GPU Single Node |
| GENESIS | RIKEN | GENESIS (GENERALized-Ensemble Simulation System) is a software package for molecular dynamics simulations and trajectory analyses. | <ul style="list-style-type: none"> • Powerful parallelization for hybrid (CPU+GPU) systems • Full electrostatics with PME • Large (1-100 million atoms) biological systems | Multi-GPU Single Node |
| GPUgrid.net | Acellera Ltd | A distributed computing project that uses GPUs for molecular simulations. | <ul style="list-style-type: none"> • High-performance all-atom biomolecular simulations • Explicit solvent and binding | Multi-GPU Single Node |
| GROMACS | GROMACS | Simulation of biochemical molecules with complicated bond interactions. | <ul style="list-style-type: none"> • Implicit (5x) • Explicit (2x) Solvent | Multi-GPU Single Node |
| HALMD | HALMD | Large-scale simulations of simple and complex liquids. | <ul style="list-style-type: none"> • Simple fluids and binary mixtures (pair potentials, high-precision NVE and NVT, dynamic correlations) | Single GPU Single Node |
| HOOMD-Blue | University of Michigan | Particle dynamics package written grounds up for GPUs. | <ul style="list-style-type: none"> • Written for use only on GPUs | Multi-GPU Single Node |
| HTMD | Acellera Ltd | High throughput molecular dynamics simulations. | <ul style="list-style-type: none"> • Available via Conda and github • ACEMD • PMEMD • NAMD • GROMACS • AMBER • CHARMM force fields • Adaptive sampling, Markov State Models, visualization, protein preparation and ligand parameterization | Multi-GPU Single Node |
| LAMMPS | Sandia National Lab | Classical molecular dynamics package | <ul style="list-style-type: none"> • Lennard-Jones • Gay-Berne • Tersoff | Multi-GPU Multi-Node |
| MELD | University of Calgary | OpenMM plugin written for GPUs. | <ul style="list-style-type: none"> • Integrative approach to combine physics and information • Orders of magnitude faster protein folding than brute force MD | Multi-GPU Single Node |
| MOLECULAR OPERATING ENVIRONMENT | 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"> • GPU Accelerated 3D Stereo Graphics • AMBER GPU accelerated support | Single GPU Single Node |

| | | | | |
|-----------------|--|--|---|---------------------------|
| myPresto | N2PC/AIST/ JBIC, Japan | Open Source Computational Drug Discovery Suite. | <ul style="list-style-type: none"> • High performance virtual screening by MD binding • Free energy calculation. | Multi-GPU Multi-Node |
| NAMD | University of Illinois at Champaign Urbana | Designed for high-performance simulation of large molecular systems. | <ul style="list-style-type: none"> • Full electrostatics with PME and most simulation features • 100M atom capable | Multi-GPU Single Node |
| OpenMM | Stanford University | Library and application for molecular dynamics for HPC with GPUs. | <ul style="list-style-type: none"> • Molecular Dynamics toolkit • Extensible and growing • Implicit and explicit solvent, custom forces | Multi-GPU Single Node |
| PolyFTS | University of California at Santa Barbara | Classical molecular simulation code for studying polymer self-assembly and thermodynamics. | <ul style="list-style-type: none"> • Uses auxiliary fields as the fundamental simulation degrees of freedom • Uses cuFFT extensively (~ 80%) • CUDA code is ~20% • Multi CPU or single GPU per job • 1x = Ivy Bridge E5-2690 CPU all 10 cores • 3-8X on K40 or K80 (utilizing 1/2 of the K80) | Single GPU Single Node |
| SOP-GPU | 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"> • Langevin dynamics simulations using the coarse-grained Self Organized Polymer (SOP) model • Multiple simulation trajectories can be performed simultaneously on a single GPU • Calpha and Calpha-Cbeta models • Simulations of protein forced unfolding • Novel simulations of nanoindentation in silico • Support for hydrodynamic interactions • Up to ~100 ms of simulation time per day, • Systems of up to 1,000,000 amino-acids (on GPUs with 6GB or great memory) | Single GPU Single Node |

QUANTUM CHEMISTRY

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------|------------------------------------|---|---|--------------------------|
| Abinit | 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"> • Local Hamiltonian • Non-local Hamiltonian • LOBPCG algorithm • Diagonalization/ orthogonalization. | Multi-GPU Single Node |
| ACES 4 | 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"> • Integrating scheduling GPU into SIAL programming language and SIP runtime environment | Multi-GPU Single Node |
| ACES III | 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"> • Integrating scheduling GPU into SIAL programming language and SIP runtime environment. | Multi-GPU Multi-Node |
| ADF | Software for Chemistry & Materials | Density Functional Theory (DFT) software package that enables first-principles electronic structure calculations. | <ul style="list-style-type: none"> • Geometry optimizations and frequency calculations with GGA functionals. | Multi-GPU Single Node |
| BigDFT | BigDFT | Implements density functional theory by solving the Kohn-Sham equations describing the electrons in a material. | <ul style="list-style-type: none"> • Daubechies wavelets | Multi-GPU Multi-Node |

| | | | | |
|------------------|---------------------------------------|--|--|--------------------------|
| BrianQC | 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"> • 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 • Compatible with features of Q-Chem 5.0 or later • Optimized for simulating large molecules • Tested up to 20,000 Cartesian Gaussian basis functions • Full support of s, p, d, f and g-type orbitals • Full support for NVIDIA GPU architectures (Kepler, Maxwell, Pascal). Double precision accuracy • Runs on 64-bit Linux operation systems • 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.) | Multi-GPU Single Node |
| CP2K | CP2K | Program to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems. | <ul style="list-style-type: none"> • DBCSR (space matrix multiply library) | Multi-GPU Multi-Node |
| GAMESS-UK | 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"> • (ss ss) type integrals within calculations using Hartree-Fock ab initio methods and density functional theory • Supports organics and inorganics. | Multi-GPU Multi-Node |
| GAMESS-US | Ames Laboratory/Iowa State University | Computational chemistry suite used to simulate atomic and molecular electronic structure. | <ul style="list-style-type: none"> • Libqc with Rys Quadrature Algorithm • Hartree-Fock • MP2 and CCSD | Multi-GPU Multi-Node |
| Gaussian | Gaussian, Inc. | Predicts energies, molecular structures, and vibrational frequencies of molecular systems. | <ul style="list-style-type: none"> • Joint NVIDIA • PGI and Gaussian collaboration | Multi-GPU Single Node |
| GPAW | GPAW | Real-space grid DFT code written in C and Python. | <ul style="list-style-type: none"> • Electrostatic poisson equation • Orthonormalizing of vectors • Residual minimization method (rmm-diis) | Multi-GPU Multi-Node |
| gWL-LSMS | ORNL | Materials code for investigating the effects of temperature on magnetism. | <ul style="list-style-type: none"> • Generalized Wang-Landau method | Multi-GPU Multi-Node |
| LATTE | Open Sourcee | Density matrix computations | <ul style="list-style-type: none"> • CU_BLAS • SP2 Algorithm | Multi-GPU Single Node |
| LSDalton | 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"> • (T) correction to the CCSD energy • RI-MP2 energy/gradient (in development) • CCSD energy (in development) • GPU-based ERI generator (in development) | Multi-GPU Single Node |

| | | | | |
|------------------|--------------------------------|--|--|---------------------------|
| MAPS | 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"> • 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 • Trajectory can be generated and then later analyzed using the appropriate tools • Additional simulations can be performed using PC-SAFT and related methods for thermodynamics modeling | Single GPU Single Node |
| MOLCAS | MOLCAS | Methods for calculating general electronic structures in molecular systems in both ground and excited states. | <ul style="list-style-type: none"> • CU_BLAS | Multi-GPU Single Node |
| MOPAC2012 | MOPAC | Semiempirical Quantum Chemistry | <ul style="list-style-type: none"> • Pseudodiagonalization • Full diagonalization • Density matrix assembling via Magma libraries | Single GPU Single Node |
| NWChem | 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"> • Triples part of Reg-CCSD(T) • CCSD and EOMCCSD task schedulers | Multi-GPU Single Node |
| Octopus | Harvard University | Used for ab initio virtual experimentation and quantum chemistry calculations. | <ul style="list-style-type: none"> • Full GPU support for ground-state, real-time calculations • Kohn-Sham Hamiltonian • Orthogonalization • Subspace diagonalization • Poisson solver • Time propagation • DFT application | Single GPU Single Node |
| PEtot | Lawrence Berkeley Laboratories | First principles materials code that computes the behavior of the electron structures of materials. | <ul style="list-style-type: none"> • Density functional theory (DFT) plane wave pseudopotential calculations | Multi-GPU Single Node |
| Q-CHEM | Q-Chem Inc. | Computational chemistry package designed for HPC clusters. | <ul style="list-style-type: none"> • Various features including RI-MP2 | Single GPU Single Node |

| | | | | |
|-------------------------|---------------------------------|--|--|---------------------------|
| QBox | 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"> • 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 • We use the CUFFT library to speed up Fourier transforms and perform asynchronous communication to cut down the cost of data transfers • The modified code is used in simulations of a 64-molecule water system with an 85 Ry plane wave energy cut off • 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 • We present these findings as well as further speedups measured in other parts of the code. http://eslab.ucdavis.edu/software/qbox http://keeneland.gatech.edu | Single GPU Single Node |
| QMCPACK | 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"> • Main features | Multi-GPU Multi-Node |
| Quantum Espresso | 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"> • PWscf package: linear algebra (matix multiply), explicit computational kernels, 3D FFTs | Multi-GPU Multi-Node |
| QUICK | Michigan State University | QUICK is a GPU-enabled ab initio quantum chemistry software package. | <ul style="list-style-type: none"> • Running Hartree-Fock and DFT energy on GPU • Supports s, p, d, f orbitals on energy calculation • HF gradient with s,p,d orbital support • GPU-based ERI generator | Multi-GPU Single Node |
| RESCU | 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"> • Parallel high efficiency processing- KS-DFT | Multi-GPU Single Node |
| RMG | 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"> • Supports 10k+ GPU nodes • Multipetaflops capable • Handles thousands of atoms with full DFT precision • Supports multiple GPUs per node • Fully open source • Installation support • Cray XE6/XK7 | Multi-GPU Single Node |
| TAL-SH | 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"> • Tensor Algebra Library for Shared Memory Computers: Nodes equipped with multicore CPU, NVIDIA GPU, and Intel Xeon Phi (in progress) | Multi-GPU Multi-Node |
| TeraChem | PetaChem LLC | Quantum chemistry software designed to run on NVIDIA GPU. | <ul style="list-style-type: none"> • Full GPU-based solution; Performance compared to GAMESS CPU version | Multi-GPU Single Node |
| VASP | 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"> • Blocked Davidson (ALGO = NORMAL & FAST) • RMM-DIIS (ALGO = VERYFAST & FAST) • K-Points and optimization for critical step in exact exchange calculations | Multi-GPU Single Node |

(MOLECULAR) VISUALIZATION AND DOCKING

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--|---|---|---|---------------------------|
| Amira | Thermo fisher Scientific | A multifaceted software platform for visualizing, manipulating, and understanding Life Science and bio-medical data. | <ul style="list-style-type: none"> • 3D visualization of volumetric data and surfaces | Single GPU Single Node |
| BINDSURF | Universidad Catolica de Murcia | A virtual screening methodology that uses GPUs to determine protein binding sites. | <ul style="list-style-type: none"> • Allows fast processing of large ligand databases | Single GPU Single Node |
| BUDE | Bristol University Docking Station | Molecular docking program | <ul style="list-style-type: none"> • Empirical Free Energy Force field | Single GPU Single Node |
| FastROCS | OpenEye Scientific Software, Inc. | Molecule shape comparison application | <ul style="list-style-type: none"> • Real-time shape similarity searching/ comparison | Multi-GPU Multi-Node |
| Interactive Molecule Visualizer | University of Illinois | Experimental interactive molecule visualizer based on a ray-tracing engine. | <ul style="list-style-type: none"> • High quality images and ease of interaction • Latest GPU computing acceleration techniques • Natural user interfaces such as Kinect and Wiimotes | Single GPU Single Node |
| MEGADOCK | 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"> • MEGADOCK-GPU on 12 CPU cores • 3 GPU calculation speed 37.0 times faster than MEGADOCK on 1 CPU core • Novel docking software facilitating the application of docking techniques to assist large-scale protein interaction network analyses | Multi-GPU Single Node |
| Molegro Virtual Docker 6 | QIAGEN | Method for performing high accuracy flexible molecular docking. | <ul style="list-style-type: none"> • Energy grid computation • Pose evaluation • Guided differential evolution | Single GPU Single Node |
| PIPER Protein Docking | Boston University | Protein-protein docking program | <ul style="list-style-type: none"> • Molecule docking | Single GPU Single Node |
| PyMol | Schrodinger, Inc. | User-sponsored molecular visualization system on an open-source foundation. | <ul style="list-style-type: none"> • Lines: 460% increase • Cartoons: 1246% increase • Surface: 1746% increase • Spheres: 753% increase • Ribbon: 426% increase | Single GPU Single Node |
| VEGA ZZ | University of California, San Francisco | Molecular Modeling Toolkit | <ul style="list-style-type: none"> • Virtual logP • Molecular surface values | Single GPU Single Node |
| VMD | University of Illinois | Visualization and analyzation of large bio-molecular systems in 3-D graphics. | <ul style="list-style-type: none"> • High quality rendering • Large structures (100M atoms) • Analysis and visualization tasks • Multiple GPU support for display of molecular orbitals | Multi-GPU Single Node |

AGRICULTURE

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------|-------------|--|--|-------------------------|
| Taranis | 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"> • Report plant population to farmers (regardless of growth stage of the crop) • Detect when a weed emerges in the field and constitutes a potential threat to the yield and then classifies it • Calculate the amounts of nutrients in vegetation, the water content in the soil, plant temperature • Identify and categorize the top relevant diseases for prevalent crops | Multi-GPU Multi-Node |

Research: Higher Education and Supercomputing

NUMERICAL ANALYTICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|----------------------|-----------------|---|--|---------------------------|
| ArrayFire HPC | 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 defense, finance, and media, and include government and academic research institutions. | <ul style="list-style-type: none"> • ArrayFire contains hundreds of functions across various domains including: <ul style="list-style-type: none"> • Vector Algorithms • Image Processing • Computer Vision • Signal Processing • Linear Algebra • Statistics • and more... | Multi-GPU Single Node |
| Eigen | Eigen | Eigen is a C++ template library for linear algebra: matrices, vectors, numerical solvers, and related algorithms. | <ul style="list-style-type: none"> • CUDA enabled linear algebra • eigen solver, reduction, random, etc. | Single GPU Single Node |
| Julia | 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"> • Full support/integration of NVIDIA CUDA via Julia CUDA JIT plugin architecture • Free and open source (MIT licensed) • User-defined types are as fast and compact as built-ins • No need to vectorize code for performance; devectorized code is fast • Designed for parallelism and distributed computation • Lightweight "green" threading (coroutines) • Unobtrusive yet powerful type system • Elegant and extensible conversions and promotions for numeric and other types • Efficient support for Unicode, including but not limited to UTF-8 • Call C functions directly (no wrappers or special APIs needed) • Powerful shell-like capabilities for managing other processes • Lisp-like macros and other metaprogramming facilities | Multi-GPU Multi-Node |
| Mathematica | Wolfram | A symbolic technical computing language and development environment. | <ul style="list-style-type: none"> • Development environment for CUDA and OpenCL • GPU acceleration for Wolfram Finance Platform | Multi-GPU Single Node |
| MATLAB | Mathworks | GPU acceleration for MATLAB (high-level technical computing language). | <ul style="list-style-type: none"> • Acceleration for 200+ most used MATLAB functions • Acceleration of more than 500 most parallelizable MATLAB functions • Accelerated Signal Processing toolkit • Accelerated Image Processing toolkit • Accelerated Communications Systems toolkit | Multi-GPU Single Node |

| | | | | |
|----------------------|-------|--|---|---------------------------|
| NMath Premium | 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"> Automatically offloads computations to the GPU. | Single GPU Single Node |
|----------------------|-------|--|---|---------------------------|

PHYSICS

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|------------------|----------------------|---|--|---------------------------|
| AWP | 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"> 3D Finite Difference Computation | Single GPU Single Node |
| BQCD | USQCD | Lattice quantum chromodynamics application, used for nuclear ad high energy physics calculations. | <ul style="list-style-type: none"> Wilson-clover fermion linear solver | Multi-GPU Single Node |
| CADISHI | 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"> Highly tuned CPU and GPU kernels Python engine for throughput computing | Multi-GPU Single Node |
| CASTRO | 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"> Gravitational Field Solver | Multi-GPU Single Node |
| Changa | 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"> Gravitational Model has been accelerated using CUDA | Single GPU Single Node |
| Chemora | 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"> Chemora embeds the equations' computational kernels into dynamically compiled loop nests shaped for input size and GPU structure | Multi-GPU Single Node |
| Cholla | Cholla | Computational Hydrodynamics On ParaLLeL Architectures for Astrophysics | <ul style="list-style-type: none"> Models the Euler equations on a static mesh and evolves the fluid properties of thousands of cells simultaneously using GPUs 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 (>256³) on a single device; calculations can be extended onto multiple devices with nearly ideal scaling beyond 64 GPUs | Multi-GPU Single Node |
| Chroma | USQCD | Lattice Quantum Chromodynamics (LQCD) | <ul style="list-style-type: none"> Wilson-clover fermions Krylov solvers Domain-decomposition | Multi-GPU Multi-Node |
| CPS | USQCD | Lattice quantum chromodynamics application, used for nuclear ad high energy physics calculations. | <ul style="list-style-type: none"> Wilson, domain-wall and Mbius fermion linear solvers | Multi-GPU Single Node |

| | | | | |
|----------------------------|------------------------------------|---|--|---------------------------|
| CPS (GRID) | 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"> • CUDA is supported • The GRID code from Edinburgh is currently being optimized. | Multi-GPU Multi-Node |
| CST PARTICLE STUDIO | Dassault Systèmes SIMULIA Corp. | Self-consistent simulation of charged particles in electromagnetic fields | <ul style="list-style-type: none"> • Particle-in-Cell Solver | Multi-GPU Multi-Node |
| GADGET | Max Planck Institute | A code for cosmological simulations of structure formation. | <ul style="list-style-type: none"> • MPI | Multi-GPU Multi-Node |
| GAMER | 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"> • Adaptive mesh refinement (AMR). Hydrodynamics with self-gravity • A variety of GPU-accelerated hydrodynamic and Poisson solvers • Hybrid OpenMP/MPI/GPU parallelization • Concurrent CPU/GPU execution for performance optimization. Hilbert space-filling curve for load balance | Multi-GPU Single Node |
| GENE | 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"> • Basic Modeling | Multi-GPU Multi-Node |
| GPU-AH | 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"> • Calculates average network density and velocity | Single GPU Single Node |
| GPUwalls | 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"> • calculates average network density and velocity | Single GPU Single Node |
| GTC Irvine | GTC | 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"> • PUSHe, Collision and Poisson Solver | Multi-GPU Multi-Node |

| | | | | |
|-----------------|------------------------------|---|---|--------------------------|
| GTC-P | 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"> • Optimized with CUDA • OpenACC development underway | Multi-GPU Single Node |
| HACC | 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"> • This code has been optimized with CUDA runs in full production mode | Multi-GPU Single Node |
| HAMR GPU | HAMR | GPU accelerated General Relativistic Magneto Hydrodynamic application | <ul style="list-style-type: none"> • Active galactic nuclei which assumes a radiatively inefficient sub-eddington rate torus • Axisymmetric ideal MHD • Viscosity and resistivity through use of Riemann solver (HLL) • Density floors to mass load the jet • Uses grids that can resolve the substructure of the jet over 5 orders of magnitude | Multi-GPU Single Node |
| MAESTRO | 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"> • Gravitational Field Solver | Multi-GPU Single Node |
| MILC | USCQD | 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"> • Staggered fermions • Krylov solvers • Gauge-link fattening | Multi-GPU Multi-Node |
| NekCEM | 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"> • 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 | Multi-GPU Multi-Node |
| OSIRIS | UCLA Plasma Physics Group | Simulates Plasma Physics including Laser interaction | <ul style="list-style-type: none"> • 2 dimensions of the particle push have been optimized with CUDA • Additional optimization is being planned with OpenACC | Multi-GPU Single Node |
| PIConGPU | 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"> • Simulation of laser-wakefield acceleration of electrons. | Multi-GPU Single Node |

| | | | | |
|---------------|-------|---|--|---------------------------|
| PPM | 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"> Turbulent, compressible mixing of gases in the context of stars near the ends of their lives and also in inertial confinement fusion | Single GPU Single Node |
| QUDA | USQCD | Library for Lattice QCD calculations using GPUs. | <ul style="list-style-type: none"> QUDA supports the following fermion formulations: Wilson, Wilson-clover, Twisted mass, Improved staggered (asqtad or HISQ) and Domain wall | Multi-GPU Single Node |
| RAMSES | CEA | Simulates astrophysical problems on different scales (e.g. star formation, galaxy dynamics, cosmological structure formation). | <ul style="list-style-type: none"> GPU acceleration Radiative transfer for reionization Hydrodynamic solver using AMR | Multi-GPU Multi-Node |
| XGC | PPPL | Simulates edge effects for MHD plasma physics | <ul style="list-style-type: none"> The particle push portion has been optimized with CUDA and is being fully optimized with OpenACC and CUDA | Multi-GPU Multi-Node |

SCIENTIFIC VISUALIZATION

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|-----------------------------------|-------------------------|---|---|---------------------------|
| Animator | GNS | Industry proven, modern post-processing app for CAE | <ul style="list-style-type: none"> Rendering | Multi-GPU Single Node |
| ANSYS EnSight | ANSYS | Industry proven post-processing app for CAE | <ul style="list-style-type: none"> Rendering Ray tracing | Multi-GPU Single Node |
| FieldView | IntelligentLight | Visualization application for CFD | <ul style="list-style-type: none"> Rendering | Single GPU Single Node |
| HVR (LCSE, U of Minnesota) | University of Minnesota | Interactive volume rendering application | <ul style="list-style-type: none"> Volume rendering | Multi-GPU Single Node |
| Index | NVIDIA | Interactive distributed volumetric compute and visualization framework. | <ul style="list-style-type: none"> Parallel distributed 3D rendering of dense or sparse volumes Accurate ray casting or ray tracing at high resolution of full size datasets Plug-in to ParaView also available. | Multi-GPU Multi-Node |
| ParaView | Kitware | Scalable data analysis and visualization application. One of the main vis tools at HPC sites. | <ul style="list-style-type: none"> Rendering and analysis tasks Plugin for NVIDIA Index OptiX rendering backend CUDA accelerated filters (data transformation routines) | Multi-GPU Multi-Node |
| SPECFEM3D | CIG | There are two modules/apss in the SPECFEM family: GLOBE and CARTESIAN. The global model is the former Gordon Bell Awardee code. Used for global inversion. Also part of the CAAR effort (although, that one is mostly focused on workflow, rather than the actual model). The regional model is CARTESIAN and it is the app used for seismic simulations, earthquake models, submarine acoustics etc. In addition to being used as a community app, Specfem3D is also use as a proxy app for proprietary codes | <ul style="list-style-type: none"> OpenCL and CUDA hardware accelerators, based on an automatic source-to-source transformation library Simulates acoustic (fluid), elastic (solid), coupled acoustic/elastic, poroelastic or seismic wave propagation in any type of conforming mesh of hexahedra (structured or not). | Multi-GPU Single Node |
| Tecplot | Tecplot | General purpose scientific visualization software for Aerodynamics, O&G, Internal Combustion and Geoscience applications | <ul style="list-style-type: none"> Rendering | Single GPU Single Node |
| VisIt | LLNL | Scalable data anlysis and visualization application | <ul style="list-style-type: none"> Rendering and analysis tasks | Multi-GPU Single Node |
| v13 (Argonne National Lab) | Argonne National Lab | Large dataset visualization in cosmology, astrophysics, and biosciences fields. | <ul style="list-style-type: none"> Volume rendering of particles | Multi-GPU Single Node |

Safety and Security

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--|----------------|--|---|---------------------------|
| AI-NVR | IronYun | Search in Video, Real time intrusion detection | <ul style="list-style-type: none"> • Search amongst 1000s of videos for interesting activities or attributes. | Single GPU Single Node |
| Alert | Irvine Sensors | Alert provides people counting and intrusion detection | <ul style="list-style-type: none"> • People counting • Intrusion detection | Single GPU Single Node |
| Arvas | 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"> • Abnormally Detection Features - Break-ins, robbery, rioting, floods, accidents, fights, arson, fire, maintenance and vandalism. | Single GPU Single Node |
| Better Tomorrow | Anyvision | Face recognition for multiple industries | <ul style="list-style-type: none"> • Face recognition | Multi-GPU Single Node |
| BioSurveillance NEXT, BioFinder | Herta Security | Real time facial recognition and forensic alerts against multiple watchlists. | <ul style="list-style-type: none"> • Supports crowded scenes and difficult lighting • Faster than real-time analysis • Partial face concealment | Multi-GPU Single Node |
| Cezurity EVO | 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"> • CUDA | Multi-GPU Single Node |
| Cylance | Cylance | Advanced AI-based endpoint malware detection. | <ul style="list-style-type: none"> • Endpoint malware detection solution • GPU deep learning technology | Multi-GPU Single Node |
| FaceControl | 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"> • Non-cooperative biometrical facial recognition system • ALPR • Video analytics and pattern recognition, • Video processing and video enhancement | Multi-GPU Multi-Node |

| | | | | |
|--|-------------------------------------|---|--|---------------------------|
| FindFace | 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"> • CUDA • TRT • nvenc • nvdec | Multi-GPU Multi-Node |
| Glueck Media; Glueck Analytics | 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"> • Facial Expression • Age Estimation • Gender • Ethnicity • Multi Face Tracking • Attention Time | Multi-GPU Single Node |
| Ikena Forensic, Ikena Spotlight | 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"> • Multi-filter, render-less video reconstruction (super-resolution, stabilization, light/color correction) • Automatic tracking for redaction video from body cameras, CCTV and other sources | Multi-GPU Single Node |
| iMotionFocus | 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"> • GPU accelerated machine learning • Identifies abnormal activity within video streams | Multi-GPU Single Node |
| innovi | 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"> • Real-time video analysis and alerts • Video search and investigation • Big data analysis • Geospatial mapping and more | Single GPU Single Node |
| LUNA | 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"> • Face detection, face alignment, facial descriptor extraction, face matching, facial attribute classification and face spoofing prevention • Optimized scalability using multithreading • Computationally efficient and compact face descriptors • Broad range of working conditions with domain-specific face descriptors | Multi-GPU Single Node |
| Nodeflux IVA | Nodeflux | Nodeflux IVA products and services cover wide range of sector including but not limited to smart city, defense and security, traffic management, toll management, store analytic (wholesale and retail), asset and facilities management, advertising, and transportation. | <ul style="list-style-type: none"> • Face recognition • License plate recognition • Traffic violation detection • Traffic monitoring, and flood monitoring | Multi-GPU Single Node |

| | | | | |
|--|-----------------|---|---|---------------------------|
| OpenALPR | OpenALPR | Automatic license plate and vehicle make/model/year recognition software applied to video streams from IP cameras. | <ul style="list-style-type: none"> • High accuracy license plate character recognition spanning North America, Europe, United Kingdom, Australia, Korea, Singapore and Brazil • APIs and source code available for embedded applications and web services | Multi-GPU Single Node |
| Recotraffic; Recosecure; Recohospital | 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"> • Traffic Data Collection, • Incident Detection • Integrated Management • Vehicle Classification and supporting related application | Multi-GPU Single Node |
| SenDISA Platform | 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"> • Intelligent Transportation - parking enforcement • Casino game table analytics | Single GPU Single Node |
| Syndex Pro | 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"> • Review hours of video in minutes • Search in Video | Single GPU Single Node |
| Tera, Tera+, Tera Vortex | SmartCow | Embedded and Backend video analytics for real-time insights from your security and service-related monitoring systems. | <ul style="list-style-type: none"> • Automatic number plate recognition • Traffic Management • Smart Car Parking Policy • Accident Detection | Multi-GPU Single Node |
| XIntelligence XHound XTransport | Xjera Labs | AI-based image and video analytics solution. This solution is ideal for people counting and recognition and vehicle counting for various commercial applications, with proven accuracy, high-level customization, and robust security. | <ul style="list-style-type: none"> • People counting • Face recognition • License plate recognition | Single GPU Single Node |
| XRVision, IoP | XRVision | Face Recognition and Video Analytics for Uncontrolled, Crowded and In Motion Environments | <ul style="list-style-type: none"> • Face Recognition and Video Analytics • Smart City, Public Safety, Transportation Analytics, Retail Analytics, Ordinance and Environment Safety | Multi-GPU Single Node |

Tools and Management

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|--------------------------------|-------------|---|---|-------------------------|
| Altair Access | 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"> • 3D Remote Visualization • High-fidelity collaboration • Integrated with Altair PBS Professional for scheduling and control on GPU use and accounting | Multi-GPU Multi-Node |
| Altair PBS Professional | Altair | Fast, powerful workload manager designed to improve productivity, optimize utilization & 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"> • GPU auto discovery • Specify GPU count per CPU • Specify GPU type • GPU/CPU affinity • GPU awareness and equality in accounting, quotas, and fair share • GPU/CPU syntax/scheduling equivalence • Specify memory use per GPU • Add-on/integration project with support for NVIDIA Data Center GPU Management (DCGM) for GPU health checks & accounting | Multi-GPU Multi-Node |

| | | | | |
|---|-----|--|---|-------------------------|
| Arm Forge (formerly Allinea) | 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 & 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"> • 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 • 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 • 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 • 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. • 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 • From multicore Linux workstations to the largest supercomputers, you can profile realistic test cases with typically less than 5% runtime overhead. • Short Learning Curve: Arm DDT offers a powerful intuitive GUI that sets the standard for multi-process and multi-threaded debugging • 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. • 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 • 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. • 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 • 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 | Multi-GPU Multi-Node |
|---|-----|--|---|-------------------------|

| | | | | |
|-------------------------------|------------------|---|---|---------------------------|
| Artec Leo | 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"> • Jetpack • Tx2 | Single GPU Single Node |
| Bright Cluster Manager | 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"> • Intuitive web app provides comprehensive view of GPU and cluster metrics • Powerful Cluster Management Shell as alternative user interface • Full Support for NVIDIA libraries <ul style="list-style-type: none"> • CUDA • OpenCL • OpenACC • CUDA-aware libraries • NCCL • CUB • Comprehensive monitoring of GPUs • Brings in GPU resources from public (AWS, Azure) and private (OpenStack) clouds within minutes • Automated scaling of the cluster based on pre-defined policies • Supports several popular Linux distributions: RHEL and derivatives, SUSE SLES and Ubuntu LTS • GPU-enabled Docker containers • Offers a complete deep learning stack • Deployment for popular HPC filesystems and management of fast interconnects • Scales up with multiple NVIDIA DGX systems | Multi-GPU Multi-Node |
| CMake | 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"> • Color output for make • Progress output for make • Incremental linking support with vs 8,9 and manifests • Supports out-of-tree builds • Auto-rerun of cmake if any cmake input files change (works with vs 8, 9 using ide macros) • Auto depend information for C++, C, and Fortran • Graphviz output for visualizing dependency trees • Full support for library versions • Full cross platform install system • Generate project files for major IDEs: Visual Studio, Xcode, Eclipse, KDevelop not tied to make, other portable generators like ant possible • Ability to add custom rules and targets • Compute link depend information, and chaining of dependent libraries • Works with parallel make and is fast, can build very large projects like KDE on build farms | N/A |

| | | | | |
|-------------------|----------------------|---|--|-------------------------|
| ELPA | 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"> • Improved one-step ScaLAPACK-type solver ELPA1 • Novel two-step solver ELPA2 | Multi-GPU Multi-Node |
| HPCToolkit | 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"> • Collects accurate and precise calling-context-sensitive performance measurements for unmodified fully optimized applications at very low overhead (1-5%) • Uses asynchronous sampling triggered by system timers and performance monitoring unit events to drive collection of call path profiles and optionally traces • 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 • Relates object code to source code files, procedures, loop nests, and identifies inlined code • Overlays call path profiles and traces with program structure computed by hpcstruct and correlates the result with source code • Handles thousands of profiles from a parallel execution by performing this correlation in parallel • hpcprof and hpcprof/mpi generate a performance database that can be explored using the hpcviewer and hpctraceviewer user interfaces • 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 • Designed to facilitate rapid top-down analysis using derived metrics that highlight scalability losses and inefficiency rather than focusing exclusively on program hot spots • 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 • 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 | Multi-GPU Multi-Node |

| | | | | |
|-------------------------|---|---|--|--------------------------|
| IBM Spectrum LSF | 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"> • Enforcement of GPU allocations via cgroups • Exclusive allocation and round robin shared mode allocation • CPU-GPU affinity • Boost control • Power management • Multi-Process Server (MPS) support • NVIDIA Volta and DCGM support | Multi-GPU Multi-Node |
| Magma | 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"> • Linear system solvers • Eigenvalue problem solvers • Auxiliary BLAS • Batched LA • Sparse LA • CPU/GPU Interface • Multiple precision support • Non-GPU-resident factorizations • Multicore and multi-GPU support • MAGMA Analytics/DNN • LAPACK testing • Linux • Windows • Mac OS | Multi-GPU Single Node |
| PAPI | 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"> • The Performance API (PAPI) project specifies a standard application programming interface (API) for accessing hardware performance counters available on most modern microprocessors • These counters exist as a small set of registers that count Events, occurrences of specific signals related to the processor's function • 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 | Multi-GPU Multi-Node |

| | | | | |
|-----------------------------|--------------------|---|---|---------------------------|
| Parallelware Trainer | 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"> • Interactive, real-time editor GUI that shows you how and where to implement parallelism. • Assists in the parallelization of code using OpenMP and OpenACC. • Transparent, local/ remote, execution and benchmarking. • Support for the C programming language. Full Fortran support coming soon. • Detailed report of opportunities for parallelism discovered in your code. • Support for multiple compilers including GCC, Intel and PGI. • Benefits: <ul style="list-style-type: none"> • Faster, more effective learning. • Reduced learning curve. • All-in-one learning tool for parallel programming. • Immediate use of parallel programming. • Support for multicore processors and GPUs. | N/A |
| SLURM | 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"> • First-class GPU support • Scales to millions of cores and tens of thousands of GPGPUs • Military grade security • Heterogenous platform support allowing users to take advantage of GPGPUs. • Flexible plugin framework enables Slurm to meet complex customization requirements • Topology aware job scheduling for maximum system utilization • Extensive scheduling options including advanced reservations, suspend/ resume, backfill, fair-share and preemptive scheduling for critical jobs • No single point of failure | Multi-GPU Multi-Node |
| STRIVR | 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"> • VRWorks 360 Video | Single GPU Single Node |

| | | | | |
|--|----------------------|---|---|-------------------------|
| TAU - Tuning and Analysis Utilities | 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"> • Instrumentation • PerfDMF • Paraprof • Load Profiles • Metric Window • Thread Windows • Communication Matrix • 3D Visualization • Derived Metrics • Selective Instrumentation • PerfExplorer • Cluster Analysis • Correlation Analysis • Scalability Chart • Preset Charts • Custom Charts • Visualizations • Eclipse Introduction • Selective Instrumentation • Instrumenting Java • Configuration Manager | Multi-GPU Multi-Node |
| Torque Moab | 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"> • Requests and schedules gpus based on gpu location in NUMA systems • Collects and report smetrics and status information • Sets gpu mode at job run time | Multi-GPU Multi-Node |

| | | | | |
|--------------------------|----------|--|---|-------------------------|
| Totalview for HPC | Perforce | <p>TotalView for HPC is a debugging software designed for high-performance computing (HPC) environments. TotalView enables faster fault isolation, improved memory optimization, and dynamic visualization for your high-scale HPC apps. TotalView understands high-scale parallel and multicore applications, with unprecedented control over processes and thread execution and visibility into program states and data.</p> | <ul style="list-style-type: none"> • Sessions Manager for managing and loading debugging sessions, "Sessions Manager" • Graphical User Interface with powerful data visualization capabilities, "The GUI" • Command Line Interface (CLI) for scripting and batch environments, "The CLI" • Stepping commands and specialized breakpoints that provide fine-grained control, "Stepping and Breakpoints" • Examining complex data sets, "Data Display and Visualization" • Controlling threads and processes, "Tools for Multi-Threaded and Parallel Applications" • Automatic batch debugging, "Batch and Automated Debugging" • Running TotalView remotely, "Remote Display" • Debugging CUDA code running on the host system and the NVIDIA® GPU, "CUDA Debugger" • Debugging remote programs, "Debugging on a Remote Host" • Memory debugging capabilities integrated into the debugger, "Memory Debugging" • Recording and replaying running programs, "Reverse Debugging" <p>Compilers: Versions are given as a range, from the earliest supported version to the latest supported version, which is usually the current version. All versions within the range are supported.</p> <p>Version information first lists compilers that support both C/C++ and Fortran, followed by compilers specific to one language or the other.</p> <ul style="list-style-type: none"> • Operating Systems: Specific supported versions are listed. If a whole number is given, all minor versions of that whole number are supported. • MPI Products: No versions are given. The rule is: if a product version can be compiled with a supported compiler, that product version is supported. | Multi-GPU Multi-Node |
| Univa Grid Engine | 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"> • Manage Nvidia CUDA • OpenACC • OpenCL plus MPI hybrid apps • Optimizes scheduling with resource-mapped GPUs • Manages GPU apps within or without Docker containers • Obtain visibility with CUDA-specific metrics for GPU monitors and reports • Extend on-premise deployments to incorporate cloud-based GPU instances | Multi-GPU Multi-Node |

| | | | | |
|---------------|------------|---|---|-------------------------|
| Vampir | 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"> • Powerful zooming and scrolling in all displays • Adaptive statistics for user selected time ranges • Filtering of processes, functions, messages, collective operations • Hierarchical grouping of threads, processes, and nodes • Support of source code locations • Integrated snapshot and printing for publishing • Customizable displays • VampirServer • Ultra scalable re-design of established Vampir functionality • Distributed performance data visualization • Increased scalability compared to sequential approach • Responsive performance data browsing from remote sites • Includes support for: NVIDIA CUDA, CUPTI, CUDA libraries • Performance Analysis Framework • Easy to use performance analysis framework for parallel programs • Graphical data representation enables detailed understanding of dynamic processes on massively parallel systems • In-depth event based analysis of parallel run-time behavior and interprocess communication • Identification of performance problems and bottlenecks | Multi-GPU Multi-Node |
|---------------|------------|---|---|-------------------------|

Unclassified Categories

| APPLICATION NAME | COMPANYNAME | PRODUCT DESCRIPTION | SUPPORTED FEATURES | GPU SCALING |
|---------------------------|---------------|---|---|---------------------------|
| Automated checkout | 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"> • cuDNN • TensorRT | Multi-GPU Single Node |
| DataX.AI | CrowdANALYTIX | Aloud-based crowd-sourced analytics services that create an online retail product catalog, onboarding 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"> • cuDNN | Single GPU Single Node |
| Helix | 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"> • TensorRT | Single GPU Single Node |

For more information on GPU-accelerated applications please visit, www.nvidia.com/teslaapps

