




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.

CONTENTS

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

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



Computational Finance

APPLICATION NAME	COMPANY/DEVELOPER	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.	• Web-like API with Native bindings for Python, R, Scala, C • Custom models and data streams are easy to add	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.	• Existing models code 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	• 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.	• 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	• Risk analytics	Multi-GPU Single Node
Hybridizer C#	Altimesh	Multi-target C# framework for data parallel computing.	• C# with translation to GPU or Multi-Core Xeon	Multi-GPU Single Node
MACS Analytics Library	Murex	Analytics library for modeling valuation and risk for derivatives across multiple asset classes	• 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 respectively.	• 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	• Monte Carlo and PDE solvers	Single GPU Single Node
O-Quant options pricing	O-Quant	Offering for risk management and complex options / derivatives pricing using GPU	• Cloud-based interface to price complex derivatives representing large baskets of equities	Multi-GPU Multi-Node
Onewiew	Numerix	Numerix introduced GPU support for Forward Monte Carlo simulation for Capital Markets and Insurance	• 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	• Spreadsheet-like modeling interfaces, Python-based scripting environment, and Grid middleware	Multi-GPU Single Node
SciFinance	SciComp, Inc	Derivative pricing (SciFinance)	• Monte Carlo and PDE pricing models	Single GPU Single Node
Synerscope Data Visualization	Synerscope	Visual big data exploration and insight tools	• Graphical exploration of large network datasets including geo-spatial and temporal components	Single GPU Single Node

> Indicates new application

Volera	Hanweck Associates	Real-time options analytical engine (Volera)	Real-time options analytics engine	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.	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	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
ACME-Atm	US DOE	Global atmospheric model used as component to ACME global coupled climate model	Dynamics only	Multi-GPU Multi-Node
COSMO	COSMO Consortium	Regional numerical weather prediction and climate research model	Radiation only in the trunk release, all features in the MCH branch used for operational weather forecasting	Multi-GPU Multi-Node
Gales	KNMI, TU Delft	Regional numerical weather prediction model	Full Model	Multi-GPU Multi-Node
WRF AceCAST-WRF	TempoQuest Inc.	WRF model from NCAR, but now commercialized by TQI. Used for numerical weather prediction and regional climate studies. All popular aspects of WRF model are GPU developed.	All popular aspects of WRF model are GPU developed: - 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	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
ANACONDA	Anaconda	Anaconda is the leading python package manager, that is the lead contributor to several open source data science libraries. Anaconda includes Numba, a Python-to-GPU compiler that compiles easy-to-read Python code to many-core and GPU architectures. Also includes single-line install of key deep learning packages for GPUs such as pytorch. Anaconda has been downloaded over 15M times and is used for AI & ML data science workloads using TensorFlow, Theano, Keras, Caffe, Neon, Lasagne, NLTK, spaCY.	Includes Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND, and sorting algorithms from the CUB and Modern GPU libraries Includes Numba (JIT python compiler) and Dask (python scheduler) Includes single-line install of numerous DL frameworks such as pytorch	Multi-GPU Single Node
ArgusSearch	Planet AI	Deep Learning driven document search	- 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	- 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
Blazegraph GPU	Blazegraph	First and fastest GPU accelerated platform for graph query. It provides drop-in acceleration for existing RDF/ Sparql and Tinkerpop/ Blueprints graph applications. It provides high-level graph database APIs with transparent GPU acceleration for graph query.	- GPU-accelerated SPARQL graph query - Data Management using the RDF interchange model - Tinkerpop/Blueprints Graph Support - Billions of edges on a single multi-GPU node - SaaS and Appliance models available.	Multi-GPU Single Node

BlazingDB	BlazingDB	GPU-accelerated relational database for data warehousing scenarios available for AWS and on-premise deployment.	Modern data warehousing application supporting petabyte scale applications	Multi-GPU Single Node
BrytlytDB	Brytlyt	In-GPU-memory database built on top of PostgreSQL	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.	- 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	- Data transformation - Ontology mapping - Data standardization - Data augmentation	Multi-GPU Single Node
DeepGram	DeepGram	Voice processing solution for call centers, financials and other scenarios	Speech to text and phonetic search using GPU deep learning	Multi-GPU Single Node
Driverless AI	H2O.ai	Automated Machine Learning with Feature Extraction. Essentially BI for Machine Learning and AI, with accuracy very similar to Kaggle Experts.	- 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.	- 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
Gunrock	UC Davis	Gunrock is a library for graph processing on the GPU. Gunrock achieves a balance between performance and expressiveness by coupling high performance GPU implementations with a high-level programming model, that requires minimal GPU programming knowledge.	- Direction-optimizing BFS, SSSP, PageRank, Connected Components, Betweenness-centrality, triangle counting - Multi-GPU support for frontier-based methods	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.	- Currently supporting tree based methods (GBM & Random Forest), GLM, Kmeans and are working on a bunch of other algorithms that are coming soon - Supports TensorFlow, Caffe and MXNet	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.	- 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.	- This database holds all relevant data in GPU memory and is thus an ideal application to utilize the Tesla K40 & 12 GB on-board RAM - Scale that up with multiple GPUs and keep close to 100 GB of compressed data in GPU memory on a single server system for fast analysis, reporting, and planning.	Multi-GPU Single Node
Labellio	KYOCERA Communication Systems Co	The world's easiest deep learning web service for computer vision, which allows everyone to build own image classifier with only web browser.	Neural net fine-tuning for image data, data crawling, data browsing as well as drag-and-drop style data cleansing backed by AI support	Multi-GPU Single Node

Numba	Anaconda	JIT compilation of Python functions for execution on various targets (including CUDA)	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.	- Uses LLVM's nvptx backend to generate CUDA kernels. OpenGL- (EGL) based rendering is not open source. - Can run in a docker container using NVIDIA-docker.	Multi-GPU Single Node
Polymatica	Polymatica	Analytical OLAP and Data Mining Platform	Visualization, Reporting, OLAP in-memory with GPU acceleration, Data Mining, Machine Learning, Predictive Analytics	Multi-GPU Multi-Node
Sqream DB	Sqream	GPU accelerated SQL database engine for big data analytics. Sqream speeds SQL analytics by 100X by translating SQL queries into highly parallel algorithms run on the GPU.	- 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.	Real-time Interaction with data	Single GPU Single Node
ZX Lib (Fuzzy Logic)	Tanay	Financial analytics and data mining library	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	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
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")	- 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 takes Anaconda to the next level and makes it easy, secure, and manageable to scale powerful analytics workflows from the laptop to the server and then scaled out to your cluster, while also incorporating collaboration, publishing, security, and Hadoop-optimized deployment.	- Anaconda Enterprise opens up the full capabilities of your GPU or multi-core processor to the Python programming language. Common operations like linear algebra, random number generation, FFT and Monte Carlo simulation run faster, and take advantage of multiple cores - Identify and remedy performance bottlenecks easily with data, code and in-notebook profilers - Includes Bindings to CUDA libraries: cuBLAS, cuFFT, cuSPARSE, cuRAND, and sorting algorithms from the CUB and Modern GPU libraries	Multi-GPU Single Node
Apache Mahout	Apache Mahout	Mahout is building an environment for quickly creating scalable performant machine learning applications.	Designed to make it extremely easy to add new algorithms. Designed to be 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 AIR-T is equipped with three signal processors including a 256 core NVIDIA Jetson TX2, a field programmable gate array (FPGA), and dual embedded CPUs.	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. Incorporates TensorFlow. Focusing on consumer banking and insurance industries.	Deep learning platform with end-to-end workflows for Enterprise. Incorporates TensorFlow.	Multi-GPU Multi-Node
Avitas Systems - Inspection as a Service	Avitas Systems	Inspection solution offering enhanced, robotic-based autonomous inspection, advanced predictive analytics, digital inspection data warehousing, and intelligent inspection planning recommendations available in a web-based interface	DL for visual inspection	
BIDMach -	UC Berkeley	The fastest machine learning library available. Holds the record for many common machine learning algorithms. Both BIDMach and its sister library BIDmat were originated at UC Berkley.	- 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.	Easy to use programming interface. Bons.ai has its own 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. Their first DL use cases targets fraud detection from time series power consumption data.	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). This allows data-parallel via MPI.	Using the GPU cluster processing mass image data	Multi-GPU Single Node
CatBoost	Yandex	CatBoost is an open-source gradient boosting library with categorical features support	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.	- 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. Clarifai uses GPUs to train large neural networks to solve practical problems in advertising, media, and search across a wide variety of industries.	- GPU-based training and inference - Recognizes and indexes images with predefined classifiers, or with custom classifiers	Multi-GPU Single Node

CNTK	Microsoft	Microsoft's Computational Network Toolkit (CNTK) is a unified computational network framework that describes deep neural networks as a series of computational steps via a directed graph.	• Supports many applications, including Speech Recognition, Machine Translation, Image Recognition, Image Captioning, Text Processing and Relevance, Language Understanding, Language Modeling	Multi-GPU Single Node
Databricks cloud	Databricks	Databricks cloud is GPU-accelerated PaaS offering build on top of AWS.	• 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
DeepBench	Baidu Research	The primary purpose of DeepBench is to benchmark operations that are important to deep learning on different hardware platforms.	• DeepBench consists of a set of basic operations (dense matrix multiplies, convolutions and communication) as well as some recurrent layer types • Both forward and backward operations are tested • This first version of the benchmark will focus on training performance in 32-bit floating-point arithmetic	Multi-GPU Single Node
DeepInstinct	DeepInstinct	Zero day end point malware detection solution offered to enterprise markets.	• 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.	• 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. Targeting consumer banking and insurance industries.	• 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.	• Object and scene detection • Machine transcription for audio • Motion and movement detection.	Multi-GPU Single Node
Gridspace	Gridspace	Voice analytics to turn your streaming speech audio into useful data and service metrics. Instrument your contact call center and work communications today with powerful deep learning-driven voice analytics	• 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, and capable of running on top of either TensorFlow or Theano. Keras was developed with a focus on enabling fast experimentation.	• cuDNN version depends on the version of TensorFlow and Theano installed with Keras • Supported Interfaces: Python	Multi-GPU Single Node

> Krisp	2Hz	2Hz, Inc is building deep learning based technologies that improve voice audio quality in real-time communications. 2Hz's Noise Suppression technology doesn't have multiple-microphone dependency and operates on any audio stream bi-directionally. With this innovation, Telcos and Cloud Communication Service Providers have the opportunity to move the ;voice enhancement; workloads to the network (as opposed to running on the user device) and be in the direct control of the quality of voice services provided.	The benchmark results illustrate that the NVIDIA Turing T4 platform is very promising for the voice processing use case given its price and energy efficiency. As expected, the Tesla V100 provides the most scalability per single GPU.	Single GPU Single Node
MatConvNet		CNNs for MathWorks MATLAB, allows you to use MATLAB GPU support natively rather than writing your own CUDA code	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. For example, video detectors can be trained to identify Steve Jobs and then search for Steve Jobs appearance anywhere in the video.	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.	- 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
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.	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. Neon features a rich set of example and pre-trained models for image, video, text, deep reinforcement learning and speech applications.	Training, inference and deployment of deep learning models. Process 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.	Process over 40M images per day with a single NVIDIA K40 or Titan GPU	Single 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.	- 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
Sentient	Sentient	Sentient is an AI platform company with special focus on digital marketing, ecommerce and finance trading applications.	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.	

SpaceKnow PaaS	SpaceKnow	PaaS for deep learning extraction of satellite data information, targeted at Financial Services and Defense Intelligence Can track macro/micro-economic activity by applying deep learning to satellite images.	• Extracts economic activity from satellite images using deep learning • Can provide 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.	• 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.	• 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.	• Computational back-ends for multicore GPUs.	Multi-GPU Single Node
UETorch	Facebook	It provides an embedded Torch environment within the powerful Unreal Engine 4. This allows one to have deep learning models directly interact with the game world, and paves way for powerful research. An example of doing AI Research using UETorch is for a neural network to learn physics and intuition about the real world.	• 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	• Identifies individuals based on unique factors, such as the way they walk, type and sit.	Multi-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.	• Visual Intelligence API allows leader enterprises in verticals like e-commerce and online auctions, media and entertainment and retailers, to analyze content related with faces, brands and context tags to perform actions like: Curate and organize visual content Search and recommend visually Get insights and analytics visually	

Federal, Defense and Intelligence

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Advanced Ortho Series	DigitalGlobe	Geospatial visualization	• Image orthorectification	Multi-GPU Single Node
ArcGIS Pro	ESRI	Viewshed2 - Determines the raster surface locations visible to a set of observer features, using geodesic methods. Aspect - Determines the compass direction that the downhill slope faces for each location Slope - Determines the slope (gradient or steepness) from each cell of a raster	• 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	• 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.	• GPU acceleration for password recovery • 10-100x speedup for password recovery	Multi-GPU Single Node
ENVI	Harris	Image Processing and Analytics	• Image orthorectification • Image transformation • Atmospheric correction • Panchromatic co-occurrence texture filter	Multi-GPU Single Node
Geomatics GXL	PCI	Image processing	• Image orthorectification and additional image processing	Multi-GPU Single Node
GeoWeb3d Desktop	Geoweb3d	Geospatial visualization of 3D and 2D data; mensuration; mission planning	• 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.	• Graph reasoning • GPU-accelerated visual analytics, visual pivoting, and 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 for intelligence analysts	• Real-time super-resolution-based video enhancement on live streams, geospatial visualization, target detection and tracking, and fast 2-D mapping	Multi-GPU Single Node
LuciadLightspeed	Luciad	Geospatial visualization and analysis	• Geospatial situational awareness	Single GPU Single Node
Manifold Systems	Manifold Systems	Full-featured GIS, vector/raster processing & analysis	• 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.	• The OmniSIG sensor typically operates in a real-time streaming fashion, ingests radio samples from many common radio interfaces, and can make use of packet formats like VITA49 or SDDS. The web-based user interface can be used from any device with a browser, including mobile handsets, and the OmniSIG software also provides its metadata output stream in JSON form for use by other applications.	Multi-GPU Single Node
SNEAK	OpCoast	Electromagnetic signals propagation modeling for complex urban and terrain environments.	• 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, which can extract 3-D objects from stereo images. ASM can extract dense 3-D point clouds from stereo images, and extract accurate building edges and corners from stereo images with high resolution, large overlaps, and high dynamic range.	• 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.	• 3D model building from imagery; building texture generation.	Multi-GPU Single Node

Design for Manufacturing/Construction: CAD/CAE/CAM

COMPUTATIONAL FLUID DYNAMICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
ADS Flow Solver - Code LEO	ADSCFD, Inc.	CFD simulation for turbochargers, turbines, and compressors	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.	Linear equation solver	Multi-GPU Single Node
ANSYS Fluent	ANSYS	General purpose CFD software	- Linear equation solver - Radiation heat transfer model - Discrete Ordinate Radiation model	Multi-GPU Multi-Node
ANSYS Icepak	ANSYS	CFD software for electronics thermal management	Linear Equation Solver	Multi-GPU Multi-Node
ANSYS Polyflow	ANSYS	CFD software for the analysis of polymer and glass processing	Direct Solvers	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. It can be used to predict the oiling in powertrain systems with rotating shafts gears and analyze forces and torques on individual components of the system.	- 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.	- 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
CPFD Barracuda-VR and Barracuda	CPFD	Fluidized bed modeling software	Linear equation solver, particle calculations	Single GPU Single Node
> DYVERSO	Realflow	3D modeling, animation, and rendering	Fluid solver (DY-SPH, DY-PBD)	Single GPU Single Node
EXN/Aero	Envenio	On-demand HPC-cloud CFD solver	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.	Discontinuous Galerkin Method (DGM) solver	Multi-GPU Single Node
FINE/Turbo	Numeca International	Structured, multi-block, multi-grid CFD solver targeting the turbo machinery industry	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. Current stage of development: Implementation of BlackOil model; Creation of pre- and post-processing modules (BlackOil model). Implementation of compositional model; Creation of pre- and post-processing modules (BlackOil model and compositional model).	• 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. The code is ported using OpenACC directives for NVIDIA GPU	• 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	• 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	• Linear equation solver (Iterative Solver and AMG Preconditioner)	Single GPU Single Node
> MIKE 3	DHI	3D Modeling of Coast and Sea	• Hydrodynamics, Advection-dispersion, Agent based modeling, underwater acoustic simulator, sand transport, mud transport, particle tracking, oil spill, ecological modeling, agent based modeling.	Multi-GPU Multi-Node
MIKE 21	DHI	2D hydrological modelling of coast and sea	• 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 FLOOD	DHI	1D & 2D urban, coastal, and riverine flood modelling	• Hydrodynamics	Multi-GPU Single Node
Moldflow	Autodesk	Plastic mold injection software	• Linear equation solver	Single GPU Single Node
Numerix	Zeus	Simulation of flow around buildings	• Discrete computational technique	Multi-GPU Single Node

> Pacefish	Numeric Systems GmbH	CFD application for ground transportation and building aerodynamics	• 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.	• 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	Exa	Industry proven, modern post-processing app for EXA POWERFLOW CFD	• Rendering • Ray tracing	Multi-GPU Single Node
> Siemens STAR-CCM+	Siemens PLM	Post-processing for CFD-focused multiphysics simulation	• Rendering	Single GPU Single Node
Simcenter 3D	Siemens PLM Software	Industry proven, modern pre- & post-processing app for multidiscipline CAE	• Rendering, Raytracing	Multi-GPU Single Node
Speed IT FLOW	Vratis	Incompressible single-phase CFD software	• 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	• Explicit solver	Multi-GPU Multi-Node
zCFD	Zenotech Simulation Unlimited	General purpose CFD solver		Multi-GPU Single Node

CFD (Research Developments)

APPLICATION NAME	COMPANY/DEVELOPER	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.	• 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	• SPH model	Multi-GPU Single Node
GIN3D	Boise St - Senocak	General purpose CFD software for incompressible flows	• Implicit solver	Multi-GPU Single Node
HiFILES	Stanford - Jameson	General purpose CFD software for compressible flows.	• Explicit solver	Multi-GPU Multi-Node
HiPSTAR	University of Southampton and University of Melbourne - Sandberg	CFD software for compressible reacting flows	• Explicit solver	Multi-GPU Single Node
JENRE, Propel (NRL)	US Naval Research Lab	CFD software for compressible flows	• Explicit solver	Multi-GPU Single Node

MUST	ITP Aero Engines	An edge-based CFD solver for RANS eqns on unstructured grid. It is proprietary to IPT, a division of Rolls Royce, and used in-house. GPU-based version of the MUST solver is used in production for the design of turbines and compressors: http://dx.doi.org/10.1080/10618562.2017.1294686	Flux computations over the edges using a multigrid solver	Multi-GPU Single Node
PyFR	Imperial College - Vincent	General purpose CFD software for compressible flows.	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	Flow solver	Multi-GPU Multi-Node
S3D	Sandia and Oak Ridge NL	Direct numerical solver (DNS) for turbulent combustion	Chemistry model	Multi-GPU Multi-Node

COMPUTATIONAL STRUCTURAL MECHANICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Adams	MSC Software	Multi-Body Dynamics simulation software	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.	- Rendering - Anti-aliasing - Ambient Occlusion	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.	- Direct solver (BCS) - Eigenvalue solvers (AMSES and Lanczos) - Iterative solver (PCG)	Single GPU Single Node
Altair Radioss Implicit	Altair	Leading structural analysis solver for highly non-linear problems under dynamic loadings. It is used across all industries worldwide to improve the crashworthiness, safety, and manufacturability of structural designs.	Direct and iterative solvers	Multi-GPU Single Node
ANSYS Mechanical	ANSYS	Simulation and analysis tool for structural mechanics	Direct and iterative solvers	Multi-GPU Multi-Node
EDEM	DEM Solutions, Ltd.	Market-leading Discrete Element Method (DEM) software for bulk material simulation	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.	- 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)	Polyhedral Elements Method solver	Single GPU Single Node
Impetus Afea	Impetus Afea	Predicts large deformations of structures and components exposed to extreme loading conditions	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	- 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	• Linear equation solver	Multi-GPU Single Node
Marc	MSC Software	Simulation and analysis tool for structural mechanics	• Direct sparse solver	Multi-GPU Single Node
> MatDEM	Nanjing University	Using innovative GPU matrix calculation method and 3D contact algorithm, MatDEM realizes 14 million 3D unit motion calculations per second (2D 4 million), and the calculation unit number and calculation speed are more than 30 times that of foreign commercial software (3 million 3D Unit, 10 million two-dimensional unit). The software implements automatic stacking modeling, layered material, joint surface and load settings, rich post-processing functions and secondary development. Graduate students can complete large-scale discrete element simulation of geology and geotechnical engineering through simple learning.	• All	Single GPU Single Node
midas GTS NX	Midas	Simulation tool for geo-technical analysis	• Linear equation solver(Multi Frontal Solver)	Single GPU Single Node
midas NFX(Structural)	Midas	Simulation and analysis tool for structural mechanics	• Linear equation solver(Multi Frontal Solver)	Single GPU Single Node
MSC Nastran	MSC Software	Simulation and analysis tool for structural mechanics	• Direct sparse solver	Multi-GPU Single Node
NX Nastran	Siemens PLM Software	Simulation and analysis tool for structural mechanics	• Linear and nonlinear equation solver	Multi-GPU Multi-Node
PERMAS-XPU	INTES GmbH	General purpose structural simulation software		Single GPU Single Node
RecurDyn	FunctionBay, Inc.	Multi-Flexible Body Dynamics simulation software	• Rendering	Single GPU Single Node
Rocky DEM	Rocky DEM	Discrete Element Modeling (DEM)-based particle simulation software	• 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)	• Direct sparse solver	Single GPU Single Node
SIMULIA Abaqus/Standard	Dassault Systèmes SIMULIA Corp.	Simulation and analysis tool for structural mechanics	• Direct sparse solver • AMS Solver • Steady State Dynamics	Multi-GPU Multi-Node

DESIGN AND VISUALIZATION

APPLICATION NAME	COMPANY/DEVELOPER	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.	• 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	• 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. It was developed by Nathaniel Jones at the MIT Sustainable Design Lab and modeled after the popular Radiance software suite developed by Greg Ward at Lawrence Berkeley National Laboratory.	• 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
Allplan	Nemetschek	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction	• OpenGL and DirectX based GPU rendering	Single GPU Single Node
ANSA	BETA CAE Systems	Industry proven, modern pre-processing app for CAE	• 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	• OpenGL-based visualization & CUDA-based Structural Stress, Modal, Fluid Dynamics and Thermal simulations	Single GPU Single Node
ANSYS Workbench	ANSYS	Industry proven, modern pre- & post-processing app for CAE	• Rendering	Multi-GPU Single Node
Apex	MSC Software	Industry proven, modern pre- & post-processing app for CAE	• Rendering	Single GPU Single Node
ArchiCAD	Nemetschek	Complete Building Information Modeling (BIM) for Architecture, Engineering, and Construction	• OpenGL and DirectX based GPU rendering	Single GPU Single Node
AutoCAD	Autodesk	2D and 3D CAD design, drafting, modeling, architectural drawing, and engineering software. Supports Open GL. Native DWG support.	• Surface, mesh, and solid modeling tools, model documentation tools, parametric drawing capabilities • Native DWG support • GRID Support.	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.	• 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	• 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	• OpenGL	Multi-GPU Single Node
Creo Parametric	PTC	Parametric design solution suite.	• Anti-aliasing, better lighting and enhanced shaded-with-edges mode • Immersive design environment with realistic materials. GRID Support.	Single GPU Single Node
FEMAP	Siemens PLM	Industry proven, modern pre- & post-processing app for CAE	• 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.	• NV Pro Pipeline (RiX) for OpenGL rendering, VRWorks SPS and VR SLI (NVLink support), and DesignWorks, including VR Occlusion Culling open source sample and OptiX	Multi-GPU Single Node
Inventor	Autodesk	3D mechanical design, documentation, and product simulation.	• 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.	• Iray Interactive • Iray Photoreal • Iray Server • Fast interactive ray tracing • Physically-based, global-illumination rendering • Distributed cluster rendering.	Multi-GPU Single Node
Iray for 3ds Max	Siemens Industry Software	Iray physically-based renderer plugin for Autodesk 3ds Max	• Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising	Multi-GPU Multi-Node
Iray for Maya	0x1 Software & Consulting GmbH	Iray 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	Iray plugin for Rhino.	• Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support.	
Iray Server	migenius	The scaling solution for any Iray based application	• Iray Photoreal and Iray Interactive support, VCA clustering, Cloud rendering, MDL support, AI based denoising	Multi-GPU Multi-Node
> META	BETA CAE Systems	High-performance post-processing software for CAE	• OpenGL	Single GPU Single Node
NX	Siemens PLM	Siemens PLM Software premium design app, has full Iray integration and support multi-gpu rendering. Still CPU bound for most tasks otherwise	• Iray, MDL	Multi-GPU Multi-Node
Optis HIM	ANSYS	Human Ergonomics Evaluation software for MRO (Maintenance Repair Overall) with large model dynamic review in VR with accurate physics behavior	• PhysX	
Optis Speos	ANSYS	Light simulation and ray tracing software		Multi-GPU Single Node
Optis Theia-RT	ANSYS	Light simulation stand-alone validation tool with OpenGL PBR engine, deterministic ray tracer, Monte-Carlo progressive GPU renderer (beta) with a strong physics focus for engineering decision making.	• Fast real-time engine integrating complex precomputed light effects	Multi-GPU Single Node
Optis VRX	ANSYS	Driving, headlamp, ADAS/AD simulators leveraging Theia-RT render engine and dynamic environment details (traffic, pedestrians...). VRX connects to Mathlab/ Simulink and ingests camera, IR, Lidar, Ultrasound and Radar sensor simulation.	• VR simulation with HMD and CAVE support	Multi-GPU Single Node
PLM Software NX and Teamcenter	Siemens	Product lifecycle management solutions from design to simulation to production to service.	• Design software, NX, and PLM viewer applications, TcVis and Active Workspace • GRID support	Single GPU Single Node
Patran	MSC Software	Industry proven, modern pre- & post-processing app for CAE	• Rendering	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 that can be cleaned up, fixed, edited, scaled, measured, re-topologized, decimated, aligned, compared and optimized for downstream workflows entirely in ReMake.	• 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's for smooth navigation of converted models in all display modes	Multi-GPU Single Node
Review	PiXYZ	Import any CAD data to prepare and experience your content with VR	• Large CAD file support with NVIDIA Pascal Single Pass Stereo extension integration	Single GPU Single Node

> Indicates new application

Revit	Autodesk	Building Information Modeling (BIM) for architecture, engineering, and construction.	- 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.		Single GPU Single Node
> Siemens STAR-CCM+ VR	Siemens PLM	Immersive VR for CFD results visualization	HTC Vive virtual reality headset	Single GPU Single Node
Simpleware	Synopsys	3D image data visualization, analysis and model generation software	OpenGL	Single GPU Single Node
SolidEdge	Siemens PLM	Mid range CAD option from Siemens	n/a	Single GPU Single Node
SOLIDWORKS	Dassault Systèmes	Covers all aspects of product development process with a seamless, integrated workflow; design, verification, sustainable design, communication and data management.	- 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	Iray-based ray-tracing, animation support, network rendering.	Single GPU Single Node
Studio	PiXYZ	Interactively prepare & optimize any CAD data before using your favorite staging tool	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, market reference for procedural texture creation.	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.	Iray rendering to enhance all artwork released with the software	Multi-GPU Single Node
T-FLEX CAD	Top Systems	3D and 2D parametric design, simulation, photorealistic rendering	High performance visualization, real time photorealistic rendering, CUDA	Multi-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	- 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	OpenGL based GPU rendering	Multi-GPU Single Node
VRED	Autodesk	VRED 3D visualization software helps automotive designers and engineers create product presentations, design reviews, and virtual prototypes. Use Digital Prototyping to quickly visualize ideas and evaluate designs.	- 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	The WYSIWYG software products, designed specifically for lighting professionals, offers a range of solutions to meet the needs of designers, assistants, electricians, console operators, teachers, and students.	The speed of wysiwyg's Shaded Views depends entirely on GPU, the GPU will have an easier time rendering ten risers consolidated into one Mesh, than rendering them as individual risers, Wysiwys also support NVIDIA SLI technologies	Multi-GPU Single Node
Zerolight	Zerolight	Immersive design review for POS		Multi-GPU Single Node

ELECTRONIC DESIGN AUTOMATION

APPLICATION NAME	COMPANY/DEVELOPER	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.	- 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.	- Transient solver - FEM solver	Multi-GPU Single Node
ANSYS HFSS SBR+	ANSYS	Simulation tool for installed antenna performance and antenna-to-antenna coupling.	High-frequency solver	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.	Eddy Current Solver	Multi-GPU Single Node
ANSYS Nexxim	ANSYS	Circuit simulation engine for RF/analog/mixed-signal IC design; IBIS-AMI analysis speedup with GPU computing.	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	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.	Conjugated Heat Transfer Solver	Multi-GPU Multi-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	- 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.	FDTD solver	Multi-GPU Single Node
JMAG	JMAG	FEA software for electromechanical design. Fast solver High quality mesh Advanced modeling technologies.	- 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	Transient Convolution simulation with BSIM4 models	Single GPU Single Node
REMCOM XFtdt	REMCOM	3D EM Simulation	FDTD Solver	Multi-GPU Multi-Node
SEMCAD-X	SPEAG	3D EM modeling and simulation	FDTD solver	Single GPU Single Node
Serenity	Lucernhammer	EM Simulation [RCS] tool	MoM solver	Multi-GPU Single Node

Sim4Life	ZMT Zurich MedTech AG	3D Electromagnetics & Acoustic modeling and simulation	• FDTD and Acoustics Solvers	Multi-GPU Single Node
TrueMask MDP	D2S	GPU-accelerated simulation and data preparation for mask writing	• Simulation-based processing	Multi-GPU Multi-Node
TrueModel	D2S	GPU-accelerated simulation and geometric checking of curvilinear shapes	• Simulation-based processing	Multi-GPU Multi-Node
Virtuoso - Cadence	Cadence Design Systems	EDA design simulation. Primary app = Allegro	• 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	• FDTD solver	Single GPU Single Node
WIPL-D 2D Solver	WIPL-D	2D EM modeling and simulation	• MoM Solver	Multi-GPU Single Node
WIPL-D Pro	WIPL-D	3D EM modeling and simulation	• MoM Solver	Multi-GPU Multi-Node
Wireless InSite	REMCOM	Uses Optix 4.1 for Ray-tracing and Propagation prediction	• X3D Ray Tracer	Multi-GPU Single Node

INDUSTRIAL INSPECTION

APPLICATION NAME	COMPANY/DEVELOPER	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.	• Feature localization & identification Segmentation & defect detection Object & 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.	• Deep learning - pre-trained networks optimized for latency or precision. HALCON also provides an IDE for training neural networks. Features include sub-pixel detection, edge detection, counting, OCR, barcode reading, 3D reconstruction from stereo	Single GPU Single Node
IBM Visual Insights	IBM	IBM Visual Insights uses cognitive capabilities to review and analyze parts, components, and products and identify defects by matching patterns to images of defects that it has previously analyzed and classified. Deploy models to edge computing on production lines to facilitate rapid image capture by camera and cognitive identification of defects. Quickly assess quality inspection metrics across manufacturing processes.	• Cloud-based DL training, deployment on (spec'ed) edge server	Multi-GPU Single Node

Media & Entertainment

ANIMATION, MODELING AND RENDERING

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
3ds Max	Autodesk	3D modeling, animation, and rendering	• Faster interactive graphics • Availability of Arnold with AI denoising • Availability of Chaos V-Ray, Otroy Octane, Redshift, cebas finalRender third-party GPU renderers	Multi-GPU Single Node
Agisoft PhotoScan	AGI Soft	Agisoft PhotoScan is a stand-alone software product that performs photogrammetric processing of digital images and generates 3D spatial data to be used in GIS applications, cultural heritage documentation, and visual effects production as well as for indirect measurements of objects of various scales.	• CUDA-accelerated photogrammetry solution	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	• GPU-accelerated hybrid renderer • Advanced material layering system with subsurface scattering, displacement mapping, physical sun-sky and IES support	Multi-GPU Single Node
Blender	Blender Inst	3D modeling, rendering and animation	• GPU-accelerated viewport	Single GPU Single Node
Blender Cycles	Blender Inst.	GPU renderer	• CUDA-accelerated rendering	Multi-GPU Single Node
Cinema 4D	Maxon	3D modeling, animation, and rendering	• Increased model complexity at interactive rates • Support for Chaos V-Ray, Ottoy 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. Whether you are a novice or proficient 3D artist or 3D animator, Daz Studio enables you to create amazing 3D art.	• Iray Interactive • Iray Photoreal • MDL support	Multi-GPU Single Node
finalRender	Cebas	GPU renderer	• CUDA-based GPU rendering	Multi-GPU Single Node
FurryBall	AAA Studio	GPU renderer	• NVIDIA OptiX and DirectX GPU rendering	Single GPU Single Node
HIERO Player	Foundry	Shot management, conform and review timeline	• Fluid, interactive playback	Single GPU Single Node
Houdini	SideFX	Procedural 3D modeling, animation and rendering	• Faster simulations	Multi-GPU Single Node
Indigo	Glare Technology	Unbiased, physically-based renderer.	• GPU-accelerated rendering.	Multi-GPU Single Node
KATANA	The Foundry	Powerful look development and lighting tool	• Faster interactive graphics	Single GPU Single Node
Kilton/Megaton	Blastcode	Physics-based simulation plug in	• Faster simulation	Single GPU Single Node
Lightwave	NewTek	3D modeling, animation, and rendering	• Increased model complexity at interactive rates	Single GPU Single Node
LuxRender	LuxRender	GPU 3D Renderer	• GPU-accelerated ray tracing	Single GPU Single Node
MARI	The Foundry	3D paint tool allows painting directly onto 3D models	• Faster interactive painting	Single GPU Single Node
Mantra	SideFX	• Houdini Mantra renderer	• Much faster interactive rendering using OptiX AI de-noising	
Maxwell	Next Limit	CUDA-accelerated interactive and final-frame renderer	• unrestricted image resolution • network rendering • de-noising	Multi-GPU Single Node
Maya	Autodesk	3D modeling, animation, and rendering	• Increased model complexity, larger scenes • Availability of Chaos V-Ray, Ottoy Octane and Redshift third-party GPU renderers	Single GPU Single Node
MODO	Foundry	3D modeling, animation and rendering	• Increased model complexity, larger scenes	Single GPU Single Node
Motion Builder	Autodesk	Character animation and motion capture	• Increased model complexity at interactive rates	Single GPU Single Node
Mudbox	Autodesk	3D sculpting	• Increased model complexity at interactive rates	Single GPU Single Node
Octane Render	Otoy	CUDA-accelerated GPU renderer	• GPU rendering	Multi-GPU Single Node

Realflow	Next Limit	Fluid simulation system	• GPU-accelerated simulation	Single GPU Single Node
RealityCapture	Capturing Reality	Photogrammetry	• CUDA-accelerated, fast photogrammetry	Multi-GPU Single Node
Redshift Renderer	Redshift	GPU-accelerated, biased renderer	• CUDA-based GPU final-frame rendering • Mac and Windows supported	Multi-GPU Single Node
Sculptris	Pixologic	3D sculpting	• Increased model complexity at interactive rates	Single GPU Single Node
TurbulenceFD	Jawset	Voxel-based gaseous fluid dynamics plug-in	• 12X performance boost using NVIDIA GPUs	Single GPU Single Node
Twinmotion	Abvent	AEC project review/phasing/marketing showcase including VR viewer with mulit CAD/BIM/AEC import and a simple to use interface but powered by UE4	• UE4 PBR material VR experience	Single GPU Single Node
V-Ray GPU	Chaos Group	GPU renderer with CPU Hybrid rendering	• CUDA interactive and final-frame GPU rendering	Multi-GPU Single Node

COLOR CORRECTION AND GRAIN MANAGEMENT

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
ARRI de-bayering SDK	ARRI	RAW de-bayering SDK	• De-bayering of ARRI RAW and primary color grading.	Single GPU Single Node
Baselight	FilmLight	Color grading	• Real-time color correction	Multi-GPU Single Node
Cinema RAW SDK	Canon	RAW de-bayering	• GPU-accelerated de-bayering	Single GPU Single Node
DaVinci Resolve	Blackmagic Design	Color grading and editing	• Real-time color correction and de-noising	Multi-GPU Single Node
Dark Energy	Cinnafilm	Application and plug-in for image enhancement	• Image de-noising and restoration	Multi-GPU Single Node
Diamant-Film Restoration	HS-Art	Film cleanup and restoration	• CUDA accelerated optical flow, de-flicker, in-painting and over 30 filters	Multi-GPU Single Node
Effects Suite	Red Giant	Visual effects plug-in	• Faster effects	Single GPU Single Node
Grain and Noise Reducer	Wavelet Beam	Video noise reduction	• 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.	• 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	
Magic Bullet Looks	Red Giant	Color and finishing tools	• Faster effects	Single GPU Single Node
Mist	Marquise Technologies	Mastering tool for cinema, broadcast and over-the-top content	• CUDA-accelerated de-bayering, color grading, transcoding and image enhancement	Multi-GPU Single Node

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
REDCINE-X PRO	Red Digital Cinema	Primary color grading	• CUDA-accelerated de-bayering and primary color grading	Single GPU Single Node
Red Digital Cinema R3D SDK	Red Digital Cinema	Red Digital Cinema camera SDK decodes and de-bayers Red RAW camera data and allows primary color grading. It is used by many color grading and video editing applications.	• CUDA accelerated wavelet decoding and de-bayering	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	• Real-time grading and finishing with Lumetri Deep Color Engine	Single GPU Single Node

COMPOSITING, FINISHING AND EFFECTS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
After Effects CC	Adobe	Motion graphics and effects	• Faster effects plus 3D ray tracing based on NVIDIA OptiX and up to 10x faster perf on about 10% effects • More effects dev & optimization in progress • Moving the effects pipeline on top of CUDA for the modular pipeline that will also be used as an effect in Premiere Pro	Multi-GPU Single Node
Clipster	Rohde & Schwarz	Video and film player and DCI Packager	• Video scaling • Color space conversion • Data format conversion	Multi-GPU Single Node
Complete	CoreMelt	Visual effects plug-in	• Faster effects	Single GPU Single Node
Continuum Complete	Boris FX	Visual effects plug-in	• Faster effects	Single GPU Single Node
Element 3D	Video Copilot	Advanced 3D object and particle render engine. Plugin for Adobe After Effects	• OpenCL - Targeting RTX ray tracing in mid 2019	Single GPU Single Node
FilmTouch 2.0	Pixelan	Video effects plug-in	• Faster effects	Single GPU Single Node
Flame Premium	Autodesk	Finishing and color grading	• Integrated toolset for 3D VFX, editorial, and color grading	Multi-GPU Single Node
Fusion	Blackmagic Design	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
Mamba FX	SGO	High-end compositing	• Faster keying, tracking, painting and restoration	Single GPU Single Node
Mistika Ultima	SGO	Color grading and finishing	• Faster keying, tracking, painting and restoration, de-bayering	Single GPU Single Node

Mistika VR	SGO	Near real-time optical flow stitching	• 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	Planar tracking - GPU accelerated object removal - must-have feature for the 360 video creators who need to remove camera rigs on stereo 8k projects https://borisfx.com/videos/mocha-pro-2019-improved-remove-module/		
Monsters GT	Boris FX	Visual effects plug-in	• Faster effects	Single GPU Single Node
NUKE	Foundry	Compositing tool with 3D tracker	• GPU-accelerated BLINK processing • Faster compositing and effects	Single GPU Single Node
Open FX	Neat Video	Video noise reduction plug-in	• Faster effects	Single GPU Single Node
PFTTrack	The Pixel Farm	3D scene creation and tracking	• CUDA-accelerated tracking	Multi-GPU Single Node
ROBUSKEY	Robuskey	Chroma keyer plug-in	• Faster effects	Single GPU Single Node
Sapphire 2019	Boris FX	Visual effects plug-in	• Faster effects	Single GPU Single Node
Twixtor	RE:Vision Effects	Visual effects plug-in	• Faster effects	Single GPU Single Node
Video Essentials	NewBlueFX	Video effects plug-in	• Faster effects	Single GPU Single Node

EDITING

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
> A.I. Gigapixel	Topaz Labs	Photo upscaling by up to 600% based on AI inference on NVIDIA GPUs	• AI	
Catalyst Browse/Prepare/Edit	Sony	4K, Sony RAW, and HD video editing	• Faster effects, transitions and encoding	Single GPU Single Node
Edius Pro	Grass Valley	Video editing	• Faster effects • RAW de-bayering	Single GPU Single Node
Final Cut Pro	Apple	Video editing	• Faster effects	Single GPU Single Node
Illustrator CC	Adobe	Vector-based digital design	• Entire canvas optimized for NVIDIA based on NV Path Render for faster pan and zoom - Approx 30% faster than standard OpenGL optimizations on Adobe-built Mac version	Single GPU Single Node
Lightroom CC (cloud service)	Adobe	Cloud-based version of Lightroom. AI features available online based on NVIDIA-enabled deep learning.		
> Lightroom Classic CC	Adobe	Photo editing	• Develop module is GPU accelerated speed up is only modest Minimal AI features like auto-tagging and intelligent image adjustments added • Recent AI inference added via "Enhance Details" feature	Single GPU Single Node
Lightworks	EditShare	Video editing	• Faster effects • CUDA-accelerated de-bayering	Single GPU Single Node

> Live Planet	Live Planet	Capture footage at the highest possible visual quality, stereoscopically, with no dizziness or nausea so that viewers can dwell in content experiences for long periods of time. Automatically stitch video footage in real time on the capture device, to such a high degree of perfection (no stitch lines, no distortion) that the content would be ready to be consumed live as an application may require. Reliably deliver content whether live or recorded over dynamic network conditions in high quality video and audio to the myriad of VR and platforms available, and do so with push-button simplicity on the part of the publisher.	• NVIDIA Tegra	Single GPU Single Node
> Luminar 3	Skylum	AI-based photo editor. Unique focus on sky effects		
Media Composer	Avid	Video editing	• 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	• NVIDIA Video Codec allows remote GPU-accelerated production workflows	Single GPU Single Node
Photoshop CC	Adobe	Image editing	• Natural canvas OpenGL accelerated, Blur galleries OpenCL accelerated	Single GPU Single Node
Premiere Pro CC	Adobe	Video editing	• Real-time video editing & accelerated output rendering via Mercury Playback Engine - CUDA / OpenCL Summer 2018 - Video Codec SDK - both encode & decode (including HEVC for 8K)	Multi-GPU Single Node
> Premiere Rush CC	Adobe	Streamlined version of Premiere Pro for quick turn editing projects. Simplified Premiere Pro for YouTube content creator types. Targeted for users to "learn in 3 minutes" 5-10x PPRO mkt size to about 15-20Mu TAM	• CUDA / OpenCL Video Codec SDK STILL SPECULATIVE - both encode & decode (H.264 & also HEVC for 8K)	
> 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).	• vGPU	Multi-GPU Multi-Node
Qube	Grass Valley	Broadcast video editing	• 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. Using AI inference on NVIDIA GPUs Standalone or Plugin for both Photoshop & Lightroom	• AI inference	
Skybox 360/VR Tools	Mettle	VR design plugin for Premiere Pro	• Fast VR processing via OpenGL and CUDA acceleration	Single GPU Single Node
Skybox Studio	Mettle	VR design plugin for After Effects	• Fast VR processing via OpenGL and CUDA acceleration	Single GPU Single Node
Smoke	Autodesk	Finishing and editing	• Faster effects	Single GPU Single Node
Vegas Pro	Magix	Video editing	• Faster video effects and encoding	Single GPU Single Node
Velocity	Imagine Communications	Video editing	• Faster effects	Single GPU Single Node
> Z Cam V1 Pro	Z Cam	Cinematic VR Camera with excellent image quality, stereoscopic recording, and live streaming	• Z CAM is the first company to integrate the NVIDIA VRWorks 360 degree Video SDK into the new V1 Pro and their WonderStitch and WonderLive stitching applications.	Single GPU Single Node

ENCODING AND DIGITAL DISTRIBUTION

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Alchemist on Demand	Grass Valley	Video standards conversion	GPU-accelerated video processing and encoding	Multi-GPU Single Node
Amberfin	Dalet	Transcoding and video quality analysis	GPU-accelerated video processing and encoding	Single GPU Single Node
Aurora	Tektronix	Automated video quality measurement	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	- 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	GPU-accelerated video processing and encoding	Multi-GPU Single Node
Core	ArcVideo	Video processing and transcoding Live	Accelerated transcoding and encoding	Multi-GPU Single Node
Elemental Live	Elemental	Live streaming video processing and encoding	Video encoding and video processing	Multi-GPU Single Node
Elemental Server	Elemental	File-based video processing and encoding	Video encoding and video processing	Multi-GPU Single Node
Fast CinemaDNG Processor	Fastvideo	RAW video debayering, denoising and color correction completely on GPU side	- 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.	- 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	Video processing and transcoding	Multi-GPU Single Node
Live	ArcVideo	High-density, real-time video processing and encoding.	Accelerated broadcast encoding with NVIDIA CUDA and NVENC	Multi-GPU Single Node
Media Encoder CC	Adobe	Output aggregator and encoder for Premiere Pro & After Effects	Faster output rendering based on Mercury Playback Engine	Multi-GPU Single Node
Medialooks SDK	Medialooks	MFormats SDK provides complete control over the video pipeline	NVIDIA Video Codec used for accelerated encoding and decoding	Single GPU Single Node

> Media Transcoding in the Cloud	Ribbon Communications	Solutions Benefits: Industry-leading SBC media transcoding scaling capabilities in virtual and cloud deployments using NVIDIA GPUs to increase performance and decrease cost per transcoded session. 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. Expanded capabilities for Virtual Network Functions (VNF) instantiation with the ability to instantiate Ribbon's PSX VNF aligned with the Open Network Automation Platform (ONAP) framework. Enhancements for operational efficiencies that allow CSPs to reduce configuration complexity and improve ease of use. Enhanced security across all products to deliver more restrictive access, reduction in possible network exposure and additional encryption.	Ribbon's 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. Ribbon's 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.	
Multiplatform Transcoder	ERLAB	Video processing and encoding software	Pre-processing encoding, decoding, post-processing and delivery	Single GPU Single Node
> mxfsPEEDRAIL	MOG Technologies	Baseband broadcast news and sports production video ingest product line for that allows editing of growing files during ingest.	- 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	H.264 and HEVC 4K encoding for broadcast channels	Single GPU Single Node
PixelStrings	Cinnafilm	Cloud-based image processing Platform-as-a-Service (PaaS) delivers the high-quality, automated video conversion and frame optimization	- 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
Server 2	Sorenson Media	Video transcoding for server app	H.264 video encoding and video processing	Multi-GPU Single Node
> Skywatch	MOG Technologies	Video and broadcast production management system for collecting audio/video usage and metadata.	- NVIDIA Video codec used for encoding for higher channel density - CUDA RAW de-coding, de-bayering, and video re-sizing and re-sampling	
> 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.	Real time encoding/decoding of audio, video signals	
Squeeze Desktop 7	Sorenson Media	Video transcoding application and plug-In	H.264 video encoding and video processing	Multi-GPU Single Node
Tachyon	Cinnafilm	Standards conversion	Video processing and frame rate conversion	Multi-GPU Single Node

Tornado	Marquise Technologies	Transcoding engine for IMF and DCP facilities	• 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	• 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	• Video transcoding and processing	Multi-GPU Single Node
Viarte	Isovideo	Video standards conversion	• CUDA-accelerated video processing and encoding	Multi-GPU Single Node
VidiCert	Joanneum Research	Video and film quality assurance	• CUDA accelerated video quality analysis	Multi-GPU Single Node
Wowza Streaming Engine Transcoder	Wowza	H.264 video encoding	• NVENC accelerated video encoding	Single GPU Single Node

ON-AIR GRAPHICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Air	Cinegy	Broadcast play-out server	• Real-time on-air graphics	Single GPU Single Node
Broadcast Dscript 3D	Monarch	3D on-air graphics	• Real-time rendering	Single GPU Single Node
Capture	Cinegy	Video ingest	• Uses NVENC to encode/decode multiple H.264 and HEVC streams	Single GPU Single Node
Clarity	Pixel Power	On-air graphics	• Real-time rendering	Single GPU Single Node
Cube	Dalet	On-air Graphics	• Real-time graphics rendering	Single GPU Single Node
eStudio	Brainstorm	Virtual sets and motion graphics	• Real-time rendering	Single GPU Single Node
GS2 Graphics Engine	ChyronHego	On-air graphics	• Real-time rendering	Single GPU Single Node
InfinitySet	Brainstorm			
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.	• Graphics solution for compact live sports productions	Multi-GPU Single Node
Mosaic	ChyronHego	On-air graphics	• Real-time rendering	Single GPU Single Node

Multiviewer	Evertz	Broadcast multiviewer	• Uses NVENC H.264 and HEVC encoding and decoding	Single GPU Single Node
Nexio Channelbrand	Imagine Communications	On-air graphics	• Real-time rendering	Multi-GPU Single Node
Nexio G8	Imagine Communications	On-air graphics	• Real-time rendering	Single GPU Single Node
Nexio TitleOne	Imagine Communications	On-air graphics	• 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	• 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	video titling	• GPU-accelerated graphics	Single GPU Single Node
tOG	RT Software	On-air graphics	• Real-time rendering	Single GPU Single Node
Type	Cinegy	On-air Graphics	• Real-time graphics rendering	Single GPU Single Node
Vertigo	Grass Valley	On-air Graphics	• Real-time rendering	Single GPU Single Node
Virtuoso	Monarch	Virtual sets and motion graphics	• Real-time rendering	Single GPU Single Node
Viz Engine	Vizrt	On-air graphics and virtual sets	• Real-time graphics rendering	Single GPU Single Node
Wasp3D - CG	Wasp3D	On-air graphics and virtual sets	• Real-time graphics rendering	Single GPU Single Node

ON-SET, REVIEW AND STEREO TOOLS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Cortex Dailies	MTI Film	Review, color grading and transcoding on set	• CUDA accelerated grading and transcoding	Multi-GPU Single Node
Dimension	Blackmagic Design	3D stereoscopic workflow	• Real-time	Single GPU Single Node
Fluid 4K Review	BlueFish444	Review and approval of 4K content	• Real-time video review	Single GPU Single Node
ICE	Marquise Technologies	IMF reference video player	• 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
On-Set Dailies	Colorfront	Review, color grading and transcoding on set	• Real-time	Multi-GPU Single Node
Previzion	Lightcraft	On-set virtual production	• Real-time, virtual set production	Single GPU Single Node
RV	Autodesk	Review and approval of 4K content	• Real-time	Single GPU Single Node

WEATHER GRAPHICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Cinemative HD	Accuweather	Weather graphics	• Real-time graphics	Single GPU Single Node
Metacast	ChyronHego	Weather graphics	• Real-time graphics	Single GPU Single Node
MeteoEarth	MeteoGraphics	Weather graphics	• Real-time graphics	Single GPU Single Node
Max Weather	WSI	Weather graphics	• Real-time graphics	Single GPU Single Node
Storyteller	Accuweather	Weather graphics	• Real-time graphics	Single GPU Single Node

Medical Imaging

APPLICATION NAME	COMPANY/DEVELOPER	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	• classification and segmentation using deep learning on top of any PACS platform	
Centricity Universal Viewer	General Electric	PACS Featuring a single image repository across 2D and 3D studies, Centricity Universal Viewer intuitively brings together the tools needed by radiologists, cardiologists and other clinicians to provide enterprise-wide access on a single desktop.	• Intelligent productivity tools, Advanced visualization applications, An advanced mammography workflow, Cross Enterprise Display, Advanced Cardiology tools, Access from anywhere1 Image enable your EMR.	
> deepflow	Helmholtz Zentrum München	Deep learning tool for reconstructing cell cycle and disease progression using deep learning from flow cytometry data.	• 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	Through digitalized pathology - they run DL on prostate cancer pathology and update if there was anything cancerous found. combine data from digitized glass slides and electronic medical records to reveal underlying patterns and extract valuable clinical insights that can transform how pathology and oncology are practiced and propel them into the information age.		Single GPU Single Node
IntelliSpace	Phillips	PACS		
iNtution iEMV	Terarecon, Inc.	Advanced Image Post-Processing		

JiveX	VISUS Health IT	free DICOM Viewer provides comparable functionality as the JiveX Review Client such as measuring tools, zoom, etc. The hanging protocols, which are less relevant outside the clinical setting, are not available.		
PowerGrid	University of Illinois Urbana-Champaign	Advanced MRI reconstruction modeling	• Discrete Fourier Transform	Multi-GPU Single Node
RadiAnt	Medixant	RadiAnt DICOM Viewer provides basic tools for the manipulation and measurement of images	• 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	
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.	• Classification and segmentation on top of any PACS platform	
SYNAPSE	Fujifilm Corporation			
Visage 7	Pro Medicus, Ltd.	DICOM viewer, also referred to as: enterprise viewer, a universal viewer (UniViewer), or an archive neutral viewer.		
Vitreia Vision	Vital Images	Advanced Image Post-Processing		

Oil and Gas

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
6X	Ridgeway Kite	Reservoir Simulation on Tesla	• CUDA Simulation Parallelization	
AlSight for SCADA	BRS Labs	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 locations including historical analysis and trend reports.	Multi-GPU Single Node
AxRTM	Acceleware	Reverse Time Migration Software	• CUDA accelerated libraries for building RTM software	Multi-GPU Multi-Node
DecisionSpace	Halliburton (Landmark)	E&P platform for geoscience, well planning, drilling, earth modeling	• CUDA acceleration of fault extraction	Multi-GPU Single Node
Echelon	Stoneridge Technology	Reservoir simulator	• Fully GPU-accelerated reservoir model, including dual-perm, dual porosity, pressure varying perm and porosity • Eclipse compatible input deck	Multi-GPU Multi-Node
GeoDepth	Emerson	Seismic Interpretation Suite	• CUDA-accelerated RTM	Multi-GPU Multi-Node
Geoteric	Geoteric	Seismic interpretation	• Attributes calculations • Geobodies extraction	Multi-GPU Single Node
Graydient S (SCADA)	Giant Grey	Machine learning anomaly detection for large scale industrial data.	• 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	• CUDA acceleration for compression and large-scale visualization	Multi-GPU Single Node
InsightEarth	CGG	Seismic Interpretation Suite	• OpenCL acceleration for AFE and 3D Curvature attributes	Multi-GPU Single Node
Omega2 RTM	Schlumberger	Seismic processing	• Multiple algorithms (RTM, etc)	Multi-GPU Multi-Node
PumaFlow IFP	Beicip-Franlab	Reservoir simulation	• GPU-accelerated linear solver	Multi-GPU Single Node
Roxar RMS	Emerson	Reservoir modeling	• Multi GPU capabilities via HUESpace	Multi-GPU Single Node
RTM	Tsunami	Seismic processing	• RTM algorithm	Multi-GPU Multi-Node
Seismic City RTM	Seismic City	RTM Seismic Processing	• CUDA acceleration	Multi-GPU Multi-Node
SKUA	Emerson	Reservoir modeling	• 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.	• CUDA, Pascal/Volta architecture, Multi-Node GPU	Multi-GPU Multi-Node
VoxelGeo	Emerson	Seismic Interpretation Package	• Multi-GPU volume rendering • Horizon-flattening • Attribute calculations	Multi-GPU Single Node

Research: Higher Education and Supercomputing

COMPUTATIONAL CHEMISTRY AND BIOLOGY

Bioinformatics

APPLICATION NAME	COMPANY/DEVELOPER	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	• 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	• 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. It can make use of highly-parallel processors such as those in graphics cards (GPUs) found in many PCs.	• 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
> BQSR	Parabricks	Base Quality Score Recalibration is a data pre-processing step that detects systematic errors made by the sequencer when it estimates the quality score of each base call.		Single GPU Single Node
> BWA-Mem	Parabricks			

Campaign	SimTK	An open-source library of GPU-accelerated data clustering algorithms and tools.	• K-means (and Kps-means, a K-means variant for GPUs with parallel sorting for improved performance) • K-medoids • K-centers (a K-medoids variant in which medoids are placed only once according to a heuristic) • Hierarchical clustering and Self-organizing map	Multi-GPU Multi-Node
CUDASW++	Open Source	Open source software for Smith-Waterman protein database searches on GPUs.	• Parallel search of Smith-Waterman database.	Multi-GPU Single Node
CUSHAW	Open Source	Parallelized short read aligner	• 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.	• Blastn and megablast modes of NCBI-BLAST	Single GPU Single Node
GHOST-Z GPU	Akiyama_Laboratory, Tokyo Institute of Technology	Sequence homology search tool.	• Good for Shotgun Metagenome Analysis.	Multi-GPU Multi-Node
GPU-Blast	Carnegie Mellon University	Local search with fast k-tuple heuristic	• Protein alignment according to BLASTP	Single GPU Single Node
mCUDA-MEME	Open Source	Ultrafast scalable motif discovery algorithm based on MEME .	• Scalable motif discovery algorithm based on MEME	Multi-GPU Single Node
MUMmer GPU	Open Source	High-throughput local sequence alignment program	• 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.	• Data structures, algorithms, and 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.	• Good coverage of Bowtie2 features and comparable quality results	Multi-GPU Single Node
PEANUT	Open Source	Read mapper for DNA or RNA sequence reads to a known reference genome.	• Achieves supreme sensitivity and speed compared to current state of the art read mappers like BWA MEM, Bowtie2 and RazerS3 • PEANUT reports both only the best hits or all hits	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.	• GRM creation, REML analysis, Regional Heritability (including multi-GPU)	Multi-GPU Single Node
SOAP3	Genomics	GPU-based software for aligning short reads with a reference sequence. It can find all alignments with k mismatches, where k is chosen from 0 to 3.	• 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: Ultra-fast GPU-based tool for short read alignment via index-assisted dynamic programming.	• Borrows-Wheeler Transformation, Dynamic Programming	Multi-GPU Single 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.	• Hardware and software for reference assembly, blast, SW, HMM, de novo assembly	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.	- 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.	Fast short read alignment	Multi-GPU Single Node
WideLM	Open Source	Fits numerous linear models to a fixed design and response.	Parallel linear regression on multiple similarly-shaped models	Multi-GPU Single Node

Microscopy

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
BioEM	Max Planck Institute	GPU-accelerated computing of Bayesian inference of electron microscopy images	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
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.	- Ab-initio reconstruction, heterogeneous reconstruction, and high-speed highresolution refinement of 3D protein structures implemented on GPUs - Lean memory usage: 768 X 768 X 768 box size on a 12GB GPU for refinement - Multiple simultaneous jobs on multiple GPUs	Multi-GPU Single Node
> emClarity	Benjamin Himes	emClarity stands for “enhanced Macro-molecular Classification and Alignment for highResolution In situ Tomography”. It 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.	Subtomogram averaging for very high resolution single particle analysis using hybrid electron microscopy.	Multi-GPU Single 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.	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

Huygens	Scientific Volume Imaging	Huygens Products: Greatly improve your microscope images	• 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. The package 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. IMOD was developed primarily by David Mastronarde, Rick Gaudette, Sue Held, Jim Kremer, Quanren Xiong, and John Heumann at the University of Colorado.	• ctfphaseflip : 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.	• 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
Microvolution	Microvolution	Nearly instantaneous 3D deconvolution & up to 200 times faster.	• 3D deconvolution for fluorescence microscopy • Written for use only on GPUs • Multi-GPU support	Single GPU Single Node
> MotionCor2	UCSF			
Phasefocus II Box	phasefocus	The Phasefocus II box product implements data processing aspects of Phasefocus's imaging methods that are known collectively as the Phasefocus Virtual Lens or & Ptychography.	• Computational diffractive imaging engine (ptychography) • Multi-GPU support	Single 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).	• Both 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. Here, we describe a protocol for using THUNDER to analysis cryoEM images in purpose of achieving a 3D model. A JSON file is used for setting parameters in THUNDER. Meaning of each attribute is discussed in this protocol. The .thu file format is defined for storing information of each particle image, including CTF parameters, rotation, translation, defocus adjustment and grading weight. The definition of this .thu file is introduced in this protocol. Finally, the procedure of running THUNDER and examining the output of THUNDER is contained in this protocol.	- 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.	Deep learning for cryo EM data particle picking. Uses CUDA 8 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. Thanks to its on-the-fly processing mode and integration with SPA tools like cryoSPARC and RELION, you can monitor data quality in real time, analyze data at the microscope, and obtain high-resolution structures even before you data collection is over.	CUDA enabled processing for electron microscopy includes numerous accelerated components in addition to using TensorFlow (v1.10). CUDA kernels include backprojection, CTF, deconvolution, FFT, tomography refinement, and others.	Multi-GPU Single Node

Molecular Dynamics

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
ACEMD	Acellera Ltd	GPU simulation of molecular mechanics force fields, implicit and explicit solvent	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.	PMEMD Explicit Solvent and GB Implicit Solvent	Multi-GPU Single Node
CHARMM	Harvard University	MD package to simulate molecular dynamics on biomolecule.	- 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.	The code uses novel parallel algorithms and numerical techniques to achieve high performance and accuracy	Multi-GPU Single Node
ESPResSo	ESPResSo	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.	- Hydrodynamic Electrokinetic forces - P3M electrostatics.	Multi-GPU Single Node
Folding@Home	Stanford University	A distributed computing project that studies protein folding, misfolding, aggregation, and related diseases.	- 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 utilize the computational power of NVIDIA GPUs. This project is under the direction of Prof. Zhong-Yuan Lu and Prof. Zhao-Yan Sun, and the package is developed and maintained by Dr. You-Liang Zhu. In addition to high-performance computation, a lot of advanced coarse-graining methods and models have been incorporated in this package (listed below). By the boost of GPU, GALAMOST enable not only us but also other researchers to investigate polymeric systems in a great large temporal and spatial scale, but at a very low cost.	• (1) General molecular dynamics includes Lennard-Jones (LJ), WeeksChandlerAndersen (WCA) potentials and Nos-Hoover, Berendsen, Andersen thermostats and Andersen, Berendsen barostats etc. (2) Dissipative particle dynamics (DPD) is a stochastic simulation technique for simulating the dynamic and rheological properties of simple and complex fluids. (3) Brownian dynamics (BD) describes the Langevin dynamics in the motion of particles in solution. (4) Coarse-graining molecular dynamics (CGMD) can be implemented by reading the numerical potential derived from iterative Boltzmann inversion (IBI) method which is developed by Muller-Plathe or other structure-based bottom-up coarse-graining methods. (5) Reaction model changes the topological connections between particles according to certain probability which is derived from real reaction rates. This model which is developed by Hong Liu can be applied to the chain-growth polymerization, step-growth polymerization, "Graft to" polymerization, polymeric ligand exchanged, the molecular mobility of polymeric graft, and so on. (6) Anisotropic particle models describe the rigid ellipsoidal particles by Gay-Berne model and describe the soft patchy particles by a soft anisotropic particle developed by Zhan-Wei Li. (7) MD-SCF is a hybrid particle-field molecular dynamics technique which combines the self-consistent field (SCF) theory and molecular dynamics (MD). It is developed by Giuseppe Millano. It will largely speed up some slowly evolving collective processes in MD simulations, such as microphase separation and self-assembly of polymeric systems. (8) DNA 3SPN model is a coarse-grained three-site-per-nucleotide model of DNA and reduce the complexity of a nucleotide to three interactions sites, one each for the phosphate, sugar, and base. (9) Rigid body method describes the translational and rotational motion of a rigid body which consists of a group of particles. (10) Stretching method imposes a non-equilibrated simulation of extension on polymeric systems to calculate the mechanical properties of materials.	Multi-GPU Multi-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.	• 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.	• 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.	• High-performance all-atom biomolecular simulations • Explicit solvent and binding	Multi-GPU Single Node
GROMACS	GROMACS	Simulation of biochemical molecules with complicated bond interactions.	• Implicit (5x) • Explicit (2x) Solvent	Multi-GPU Single Node
HALMD	HALMD	Large-scale simulations of simple and complex liquids.	• Simple fluids and binary mixtures (pair potentials, high-precision NVE and NVT, dynamic correlations)	Single GPU Single Node
HOOBMD-Blue	University of Michigan	Particle dynamics package written grounds up for GPUs.	• Written for use only on GPUs	Multi-GPU Single Node
HTMD	Acellera Ltd	High throughput molecular dynamics simulations	• Available via Conda and github • Support ACEMD, PMEMD, NAMD, GROMACS • AMBER and 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	• Lennard-Jones, Gay-Berne, Tersoff, and dozens more potentials	Multi-GPU Multi-Node
MELD	University of Calgary	OpenMM plugin written for GPUs	• Integrative approach to combine physics and information • Orders of magnitude faster protein folding than brute force MD	Multi-GPU Single Node
myPresto	N2PC/AIST/JBIC, Japan	Open Source Computational Drug Discovery Suite.	• 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.	• 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.	• Molecular Dynamics toolkit that is the heart of Folding@Home and is used by serveral/a growing number of other applications. • 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.	• 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, where SOP stands for the Self Organized Polymer Model fully implemented on a GPU, is a scientific software package designed to perform Langevin Dynamics Simulations of the mechanical or thermal unfolding, and mechanical indentation of large biomolecular systems in the experimental subsecond (millisecond-to-second) timescale.	• Langevin dynamics simulations using the coarse-grained Self Organized Polymer (SOP) model, Multiple simulation trajectories can be performed simultaneously on a single GPU, Calpha and Calpha-Cbeta models are supported, Simulations of protein forced unfolding, Novel simulations of nanoindentation in silico, Support for hydrodynamic interactions, Up to ~100 ms of simulation time per day, Systems of up to 1,000,000 amino-acids (on GPUs with 6GB or great memory)	Single GPU Single Node

Quantum Chemistry

APPLICATION NAME	COMPANY/DEVELOPER	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.	- Local Hamiltonian - Non-local Hamiltonian - LOBPCG algorithm - Diagonalization/ orthogonalization.	Multi-GPU Single Node
ACES III	University of Florida	Takes best features of parallel implementations of quantum chemistry methods for electronic structure.	Integrating scheduling GPU into SIAL programming language and SIP runtime environment.	Multi-GPU Multi-Node
ACES 4	University of Florida	New SIA/aces4 development A new super instruction architecture with interface applications for quantum chemistry (aces4) is now under development. Details and downloads can be found on the project's git hub site.	Integrating scheduling GPU into SIAL programming language and SIP runtime environment.	Multi-GPU Single Node
ADF	Software for Chemistry & Materials	Density Functional Theory (DFT) software package that enables first-principles electronic structure calculations.	- Geometry optimizations and frequency calculations with GGA functionals. https://www.scm.com/doc/ADF/Input/Technical_Settings.html?highlight=GPU?	Multi-GPU Single Node
BigDFT	BigDFT	Implements density functional theory by solving the Kohn-Sham equations describing the electrons in a material.	DFT; Daubechies wavelets, part of Abinit	Multi-GPU Multi-Node
> BrianQC	StreamNovation Ltd.	“Drug discovery and development is a very complex and multidisciplinary field connected to biology, chemistry and medical sciences. The process of discovery and development takes a 10-15 years long timespan and huge financial effort to see from start to success. During this period millions of possible drug candidates are examined. It costs several billion EUR and 5 to 10 years to introduce a new drug on average. According to our aiming R&D costs could be reduced thanks to the unique simulation tool developed by our team. We have already developed a unique algorithm, which makes high complexity calculations possible and effectively on GPU. Further on we have finished the integration of a GPU based and thus massively parallel integrator module. This module calculates integrals for quantum chemistry (QC) on parallel systems effectively. Our software is also able to simulate such large-scale molecules in reasonable time with high accuracy, which has not been possible until now. We address to open new research areas, as we can accurately calculate not only the features of large molecules but the attachment of active substances, which was managed by approximations so far.”	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.)	Single GPU Single Node

Chameleon	Scienomics	MAPS CLASSICAL & MESOSCALE simulation toolkit contains world-class simulation engines such as LAMMPS, CHAMELEON, TOWHEE, NAMD. In addition an important collection of ready-to-use workflows and a rich Force-Field library are included.	LAMMPS-ATOMISTIC PLUGIN A comprehensive package for Force-Field based molecular dynamics and mechanics simulations of bulk, interfacial and transport properties of liquid, solid, or gaseous systems. LAMMPS handles a variety of boundary conditions and is suitable for any atomic, polymeric, biological, metallic, or granular system as well as reactions. A wide variety of Force-Field parameter files is included such as Amber, CHARMM, CVFF, Dreiding, EAM, Martini, PCFF, ReaxFF, SciPCFF, TraPPE and UFF potentials. LAMMPS-DPD PLUGIN Enables MAPS users to set up, execute and analyze Dissipative Particle Dynamics simulations with LAMMPS. Dissipative Particle Dynamics simulations can be used to predict properties of mesoscopic systems. CHAMELEON PLUGIN SCIENOMICS proprietary Monte-Carlo simulation engine allowing the relaxation of complex, realistic polymeric systems at the atomistic or mesoscale level. CHAMELEON implements connectivity altering techniques that allows the fast and smart configurational space sampling and succeeds where standard molecular dynamics fails. TOWHEE PLUGIN Multi-purpose Monte Carlo simulation code for the prediction of fluid phase equilibria, adsorption behavior in porous materials and for relaxing complex systems using atom-based Force-Fields. MAPS OPTIMIZER Performs Force-Field based geometry optimization of both finite and periodic systems. It supports many different types of Force-Fields and includes several standard and advanced convergence algorithms. NAMD PLUGIN A molecular dynamics and mechanics simulations package for studying dynamic phenomena and bulk properties of large molecular systems. It is designed for high-performance computational tasks such as geometry optimization, molecular dynamics (within several ensembles) for periodic and non-periodic systems. CONFORMER SEARCH PLUGIN Conformer generation based on a systematic approach applying a set of torsion rules. It provide a set of low-energy conformers according to a diversity criterion. FHMIXING PLUGIN A Monte Carlo based simulation tool for binary mixtures which generates ensembles of pairs of molecules using the Molecular Silverware algorithm. These are then analyzed using the lattice based Flory Huggins theory. Miscibility, mixing energies, phase diagrams and many other properties can be computed very rapidly for quick screening of molecules. TEAMFF Force field database containing a large number of high quality force field parameters which can be easily extended. It contains multiple Force-Fields that can be combined for successful assignment.	Single GPU Single Node
-----------	------------	--	---	---------------------------

CP2K	CP2K	Program to perform atomistic and molecular simulations of solid state, liquid, molecular and biological systems.	• 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.	• (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.	• 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.	• Joint NVIDIA, PGI and Gaussian collaboration	Multi-GPU Single Node
GPAW	GPAW	Real-space grid DFT code written in C and Python	• 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.	• Generalized Wang-Landau method	Multi-GPU Multi-Node
LATTE	Open Sourcee	Density matrix computations	• 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.	• (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
MOLCAS	MOLCAS	Methods for calculating general electronic structures in molecular systems in both ground and excited states.	• CU_BLAS	Multi-GPU Single Node
MOPAC2012	MOPAC	Semiempirical Quantum Chemistry	• Pseudodiagonalization, full diagonalization, and 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.Tensor Algebra Library Routines for Shared Memory Systems which is being used to GPU accelerate three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	• 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.	• 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.	• Density functional theory (DFT) plane wave pseudopotential calculations	Multi-GPU Single Node
Q-CHEM	Q-Chem Inc.	Computational chemistry package designed for HPC clusters.	• Various features including RI-MP2	Single GPU Single Node

QBox	Open Source	Qbox is a C++/MPI scalable parallel implementation of first-principles molecular dynamics (FPMD) based on the plane-wave, pseudopotential formalism. Qbox is designed for operation on large parallel computers.		Single GPU Single Node
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.	• Main features	Multi-GPU Multi-Node
Quantum Espresso/ PWscf	Quantum Espresso Foundation	An integrated suite of computer codes for electronic structure calculations and materials modeling at the nanoscale.	• PWscf package: linear algebra (matrix 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.	• 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. RESCU is the abbreviation of Real space Electronic Structure Calculator. Its core is a new, extremely powerful and parallel high efficiency KS-DFT self-consistent calculation method.		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, it has been run successfully on systems with thousands of nodes (including GPU nodes) and hundreds of thousands of CPU cores.	• Supports 10k+ GPU nodes, multipetaflops capable • Handles thousands of atoms with full DFT precision • Supports multiple GPUs per node • Fully open source, with installation support, Downloads, documentation, forums www.rmgdft.org • Cray XE6/XK7	Multi-GPU Single Node
TAL-SH	Oak Ridge National Lab	Tensor Algebra Library Routines for Shared Memory Systems which is being used to accelerated three (3) CAAR codes; NWChem, LSDALTON and DIRAC.	• TAL-SH: Tensor Algebra Library for Shared Memory Computers: Nodes equipped with multicore CPU, NVIDIA GPU, and Intel Xeon Phi (in progress). Author: Dmitry I. Lyakh (Liakh): quant4me@gmail.com, liakhdi@ornl.gov Copyright (C) 2014-2016 Dmitry I. Lyakh (Liakh) Copyright (C) 2014-2016 Oak Ridge National Laboratory (UT-Battelle) LICENSE: GNU Lesser General Public License v.3 API reference manual: DOC/TALSH_manual.txt BUILD: Modify the header of the Makefile accordingly and run make.	Multi-GPU Multi-Node
TeraChem	PetaChem LLC	Quantum chemistry software designed to run on NVIDIA GPU.	• 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.	• 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

Visualization and Docking

APPLICATION NAME	COMPANY/DEVELOPER	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.	• 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.	Allows fast processing of large ligand databases	Single GPU Single Node
BUDE	Bristol University Docking Station	Molecular docking program	Empirical Free Energy Force field	Single GPU Single Node
FastROCS	OpenEye Scientific Software, Inc.	Molecule shape comparison application	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.	Targeting high quality images and ease of interaction, IMV uses the latest GPU computing acceleration techniques, combined with natural user interfaces such as Kinect and Wiimotes	Single GPU Single Node
Molegro Virtual Docker 6	QIAGEN	Method for performing high accuracy flexible molecular docking.	Energy grid computation, pose evaluation and guided differential evolution	Single GPU Single Node
PIPER Protein Docking	Boston University	Protein-protein docking program	Molecule docking	Single GPU Single Node
PyMol	Schrodinger, Inc.	User-sponsored molecular visualization system on an open-source foundation	- 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	Virtual logP, molecular surface values	Single GPU Single Node
VMD	University of Illinois	Visualization and analyzing large bio-molecular systems in 3-D graphics	High quality rendering, large structures (100M atoms), analysis and visualization tasks, multiple GPU support for display of molecular orbitals	Multi-GPU Single Node

NUMERICAL ANALYTICS

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Accelereyes-ArrayFire	ArrayFire	Comprehensive GPU function library	Hundreds of functions for math, signal/ image processing, 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.	CUDA enabled linear algebra including eigen solver, reduction, random, etc. Still considered an experimental feature	Single GPU Single Node
HiPLAR		3High Performance Linear Algebra in R	Supports GPU and multi-core platforms, compatible with legacy R code, no new data types or operators, auto-tuning, support for R Matrix package	Multi-GPU Single Node
MATLAB	Mathworks	GPU acceleration for MATLAB (high-level technical computing language).	Acceleration for 200+ of most used -- and more than 500 -- MATLAB functions (incl. Parallel Computing Toolkit, Signal Processing, Image Processing, Communications Systems, etc)	Multi-GPU Single Node
Mathematica	Wolfram	A symbolic technical computing language and development environment.	- Development environment for CUDA and OpenCL - GPU acceleration for Wolfram Finance Platform.	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.	Automatically offloads computations to the GPU.	Single GPU Single Node

PHYSICS

APPLICATION NAME	COMPANY/DEVELOPER	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.	3D Finite Difference Computation	Single GPU Single Node
BQCD	USQCD	Lattice quantum chromodynamics application, used for nuclear and high energy physics calculations. Most usage is concentrated in EMEA	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.	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). The approach uses a nested hierarchy of logically-rectangular grids with simultaneous refinement in both space and time.	Gravitational Field Solver	Multi-GPU Single Node
Changa	CHANGA	Astrophysics code performs collisionless N-body simulations. It can perform cosmological simulations with periodic boundary conditions in comoving coordinates or simulations of isolated stellar systems.	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.	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	- 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 ($>256^3$) 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)	Wilson-clover fermions, Krylov solvers, Domain-decomposition	Multi-GPU Multi-Node
CPS	USQCD	Lattice quantum chromodynamics application, used for nuclear and high energy physics calculations.	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 QCDQC, IBM Blue Gene machines, and builds for scalar machines or parallel machines with QMP.	• QUDA 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	• Particle-in-Cell Solver	Multi-GPU Multi-Node
GADGET	Max Planck Institute	A code for cosmological simulations of structure formation	• MPI	Multi-GPU Multi-Node
GAMER	OpenSource	A GPU-accelerated Adaptive Mesh Refinement Code for astrophysical applications. Currently the code solves the hydrodynamics with self-gravity.	• 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
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.	• Besides evolving the network timestep-by-timestep, it also calculates the average network density and velocity. These quantities confirm that that the simulation reproduces the behaviour of local string networks when compared to literature results.	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.	• Besides evolving the network timestep-by-timestep, it also calculates the average network density and velocity. These quantities confirm that that the simulation reproduces the behaviour of wall networks in the literature up to machine precision (compared to our earlier CPU version of the code).	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.	• Key accelerated features are the 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.	• 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.	• 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	• 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.	• Gravitational Field Solver	Multi-GPU Single Node
MILC	USQCD	Lattice Quantum Chromodynamics (LQCD) codes simulate how elemental particles are formed and bound by the strong force to create larger particles like protons and neutrons.	• 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. The code is actively developed at Mathematics and Computer Science Division of Argonne National Laboratory.	• 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	• 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.	• Simulation of laser-wakefield acceleration of electrons.	Multi-GPU Single Node
PPM	PPM	Piecewise parabolic method, a higher-order extension of Godunov's method which uses spatial interpolation and allows for a steeper representation of discontinuities, particularly contact discontinuities.	• 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.	• 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).	• GPU acceleration is applied for radiative transfer for reionization, and the hydrodynamic solver using AMR	Multi-GPU Multi-Node
XGC	PPPL	Simulates edge effects for MHD plasma physics	• 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	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
3D Slicer	Open Source	Medical visualization & segmentation	• Rendering, image processing	Single GPU Single Node
Animator	GNS	Industry proven, modern post-processing app for CAE	• Rendering	Multi-GPU Single Node
ANSYS EnSight	ANSYS	Industry proven post-processing app for CAE	• Rendering • Ray tracing	Multi-GPU Single Node
FieldView	IntelligentLight	Visualization application for CFD	• Rendering	Single GPU Single Node
FluoRender (SCI, U of Utah)	University of Utah	Interactive rendering tool for confocal microscopy data visualization.	• Multi-channel volume rendering	Single GPU Single Node
HVR (LCSE, U of Minnesota)	University of Minnesota	Interactive volume rendering application	• Volume rendering	Multi-GPU Single Node
ImageVis3D	SCI, University of Utah	Simple, scalable, and interactive volume rendering application.	• Out-of-core volume rendering	Single GPU Single Node
Index	NVIDIA	Interactive distributed volumetric compute and visualization framework.	• 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.	• Rendering and analysis tasks • Plugin for NVIDIA Index • OptiX rendering backend • CUDA accelerated filters (data transformation routines)	Multi-GPU Multi-Node
Seg3D (SCI, U of Utah)	SCI, University of Utah	Segmentation application for medical data	• Rendering, image processing	Single GPU Single 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	• 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	• Rendering	
VisIt	LLNL	Scalable data anlysis and visualization application	• Rendering and analysis tasks	Multi-GPU Single Node
Visualzation Toolkit (VTK)	Open Source	Data anlysis and visualization toolkit	• Rendering	Single GPU Single Node
vl3 (Argonne National Lab)		Large dataset visualization in cosmology, astrophysics, and biosciences fields.	• Volume rendering of particles	Multi-GPU Single Node

Safety and Security

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
AI-NVR	IronYun	Search in Video	• Search in Video	
Alert	Irvine Sensors	Alert provides people counting and intrusion detection	• People counting, Intrusion detection	
Arvas	VI Dimensions			
Better Tomorrow	Anyvision	Face recognition for multiple industries	• Face recognition	Multi-GPU Single Node
BioSurveillance NEXT, BioFinder	Herta Security	Real time facial recognition and forensic alerts against multiple watchlists.	• Supports crowded scenes, difficult lighting, faster than real-time analysis, partial face concealment	Multi-GPU Single Node
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.	• CUDA	Multi-GPU Single Node
Cylance	Cylance	Advanced AI-based end point malware detection	• End point malware detection solution build using GPU deep learning technology	Multi-GPU Single Node
Discover Experiencer	Viisights			
Face++	Megvii	Face Recognition for many verticals	• Face Recognition	
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.	• Non-cooperative biometrical facial recognition system, ALPR, video analytics and pattern recognition, video processing and video enhancement	Multi-GPU Single Node
FindFace Enterprise Server SDK	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.	• CUDA, NVDec, DeepStream (testing)	Multi-GPU Single 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.	Specific features supported includes; Facial Expression, Age Estimation, Gender, Ethnicity, Multi Face Tracking, Attention Time	Multi-GPU Single Node
Graydient V (Video)	Giant Grey	Machine learning anomaly detection for enhanced video analytics.	- Proactive event detection and real-time alerts for safety, unauthorized access prevention, and loss prevention 24/7 real-time analysis and alerting scaling to thousands of video streams across remote and geographically dispersed locations	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	Multi-filter, render-less video reconstruction (super-resolution, stabilization, light/color correction), and automatic tracking for redaction video from body cameras, CCTV and other sources	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.	GPU accelerated machine learning to identify 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	The solution provides real-time video analysis and alerts, video search and investigation, big data analysis, geospatial mapping and more	
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.	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's 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.	Specific product feature includes, face recognition, license plate recognition, traffic violation detection, Traffic monitoring, and flood monitoring. More such IVA Functions are being added with customer needs.	Multi-GPU Single Node
OpenALPR	OpenALPR	Automatic license plate and vehicle make/model/year recognition software applied to video streams from IP cameras.	- 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.	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.	Intelligent Transportation - parking enforcement Casino game table analytics	
SenseFace	Sensetime	Intelligent engine of visual surveillance system with ex ante, concurrent & ex post technical support service. Applicable to security, missing people & suspects search.	Face Recognition	
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.	Review hours of video in minutes, Search in Video	
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.	People counting, face recognition, license plate recognition	
Xpose	Lentix			

Tools and Management

APPLICATION NAME	COMPANY/DEVELOPER	PRODUCT DESCRIPTION	SUPPORTED FEATURES	GPU SCALING
Allinea Forge	Allinea now owned by ARM Ltd	Allinea Forge Professional provides all you will need to debug, profile and optimize for high performance - from single threads through to complex parallel HPC and scientific codes with MPI, OpenACC, OpenMP, threads or NVIDIA CUDA	CUPTI, cudagdb	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. PBS Professional automates job scheduling, management, monitoring and reporting.	Commonly requested capabilities that are supported: - 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 NVIDIA's Data Center GPU Management (DCGM); GPU health checks & detailed accounting	Multi-GPU Multi-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.	• Intuitive web app provides comprehensive view of GPU and cluster metrics • Powerful Cluster Management Shell as alternative user interface • Fully Support for NVIDIA libraries, CUDA, OpenCL, OpenACC, CUDA-aware libraries, NCCL, and CUB • Comprehensive monitoring of GPU • Brings in GPU resources from public (AWS, Azure) and private (OpenStack) clouds within minutes • 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	Multi-GPU Multi-Node
cmake	Kitware	CMake is an open-source, cross-platform family of tools designed to build, test and package software. CMake is used to control the software compilation process using simple platform and compiler independent configuration files, and generate native makefiles and workspaces that can be used in the compiler environment of your choice. The suite of CMake tools were created by Kitware in response to the need for a powerful, cross-platform build environment for open-source projects.		N/A
> 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.	• improved one-step ScaLAPACK-type solver ELPA1 and 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.	• CUPTI	Multi-GPU Multi-Node

IBM Spectrum LSF	IBM Corporation	IBM Spectrum LSF is a complete workload management solution for demanding HPC environments. Featuring intelligent, policy-driven scheduling, it helps organizations to improve competitiveness by accelerating research and design while controlling costs through superior resource utilization and ease of use. Building on over 20 years of experience, IBM Spectrum LSF features a highly scalable and available architecture designed to address the challenge of aligning compute resources with business priorities. With the ability to detect, monitor and schedule GPU enabled workloads to the appropriate resources, IBM Spectrum LSF enables users to easily take advantage of the benefits provided by GPUs.	• Enforcement of GPU allocations via cgroups • Exclusive allocation and round robin shared mode allocation • CPU-GPU affinity • Boost control • Power management/li> • Multi-Process Server (MPS) support • NVIDIA Volta and DCGM support	Multi-GPU Multi-Node
Magma	ICL - University of Tennessee Knoxville	The MAGMA project aims to develop a dense linear algebra library similar to LAPACK but for heterogeneous/hybrid architectures, starting with current "Multicore+GPU" systems. The MAGMA research is based on the idea that, to address the complex challenges of the emerging hybrid environments, optimal software solutions will themselves have to hybridize, combining the strengths of different algorithms within a single framework. Building on this idea, we aim to design linear algebra algorithms and frameworks for hybrid manycore and GPU systems that can enable applications to fully exploit the power that each of the hybrid components offers.		Multi-GPU Single Node
open SpeedShop	Krell Institute	Open SpeedShop (O SS) is an open source multi-platform performance tool enabling performance analysis of HPC applications running on both single node and large scale platforms. O SS gathers and displays several types of information to aid in solving performance problems, including: program counter sampling for a quick overview of the applications performance, call path profiling to add caller/callee context and locate critical time consuming paths, access to the machine hardware counter information, input/output tracing for finding I/O performance problems, MPI function call tracing for MPI load imbalance detection, memory analysis, POSIX thread tracing, NVIDIA CUDA analysis, and OpenMP analysis. O SS offers a command-line interface (CLI), a graphical user interface (GUI) and a python scripting API user interface.	• CUPTI	Multi-GPU Multi-Node
PAPI	ICL - University of Tennessee Knoxville	PAPI provides the tool designer and application engineer with a consistent interface and methodology for use of the performance counter hardware found in most major microprocessors. PAPI enables software engineers to see, in near real time, the relation between software performance and processor events. In addition, PAPI provides access to a collection of components that expose performance measurement opportunities across the hardware and software stack.	• CUPTI	Multi-GPU Multi-Node

Parallware Trainer	Appentra Solutions	Parallware Trainer is the new interactive, real-time editor with GUI features to facilitate the learning, usage, and implementation of parallel programming. Users are actively involved in learning parallel programming through observation, comparison, and hands-on experimentation. Parallware Trainer provides support for widely used parallel programming strategies using OpenACC and OpenMP with execution on multicore processors and GPUs.		N/A
SLURM	SchedMD	SLURM is a highly configurable open source workload and resource manager. In its simplest configuration, Slurm can be installed and configured in a few minutes. Use of optional plugins provides the functionality needed to satisfy the needs of demanding HPC centers with diverse job types, policies and work flows. Advanced configurations use plug-ins to provide features like accounting, resource limit management, by user or bank account, and support for sophisticated scheduling algorithms.	• 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
> Taranis	Taranis	Our deep learning engine use bleeding edge math models and hardware platforms on the cloud and has been trained by over 60 expert agronomists providing more than 1,000,000 examples of crop health issues	• 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
TAU - Tuning and Analysis Utilities	University of Oregon	TAU is a program and performance analysis tool framework. TAU provides a suite of static and dynamic tools to form an integrated analysis environment for parallel Fortran, C++, C, Java, and Python applications. Also, recent advancements in TAU's code analysis capabilities have allowed new static tools to be developed, such as an automatic instrumentation tool	• CUPTI	Multi-GPU Multi-Node
Torque Moab	Adaptive Computing	Moab HPC Suite is a workload and resource orchestration platform that automates the scheduling, managing, monitoring, and reporting of HPC workloads on massive scale. TORQUE provides control over batch jobs and distributed computing resources. It is an advanced open-source product based on the original PBS project* and incorporates the best of both community and professional development.	• Request/schedule gpus based on gpu location in NUMA systems • Collect and report metrics and status information • Set gpu mode at job run time	Multi-GPU Multi-Node

Totalview for HPC	Rogue Wave Software	TotalView for HPC allows simultaneous debug many processes and threads in a single window. Work backwards from failure through reverse debugging, isolating the root cause faster by eliminating repeated restarts of the application. Reproduce difficult problems that occur in concurrent programs that use threads, OpenACC, OpenMP, MPI and CUDA	• CUDA debug API	Multi-GPU Multi-Node
Univa Grid Engine	Univa	The Univa Grid Engine suite is a leading workload management system. The solution maximizes the use of shared resources in a data center and applies advanced management policy enforcement to deliver results faster, more efficiently, and with lower overall costs. The product suite can be deployed in any technology environment, including containers: on-premise, hybrid or in the cloud.	• Manage NVIDIA CUDA, OpenACC, OpenCL plus MPI hybrid apps • Optimize scheduling with resource-mapped GPUs • Manage GPU apps within or without Docker containers • Obtain visibility with CUDA-specific metrics for GPU monitors and reports • Extend on-premise deployments to incorporate cloud-based GPU instances	Multi-GPU Multi-Node
Vampir	TU Dresden	Vampir provides an easy-to-use framework that enables developers to quickly display and analyze arbitrary program behavior at any level of detail. The tool suite implements optimized event analysis algorithms and customizable displays that enable fast and interactive rendering of very complex performance monitoring data.	• CUPTI	Multi-GPU Multi-Node

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

