



2018 North American Machine Learning
Technology Innovation Award



2018
BEST PRACTICES
AWARDS

Contents

Background and Company Performance	3
<i>Industry Challenges</i>	3
<i>Technology Attributes and Future Business Value</i>	3
<i>Conclusion</i>	6
Significance of Technology Innovation	7
Understanding Technology Innovation	7
<i>Key Benchmarking Criteria</i>	8
Best Practices Award Analysis for Wave Computing	8
<i>Decision Support Scorecard</i>	8
<i>Technology Attributes</i>	9
<i>Future Business Value</i>	9
<i>Decision Support Matrix</i>	10
Best Practices Recognition: 10 Steps to Researching, Identifying, and Recognizing Best Practices	11
The Intersection between 360-Degree Research and Best Practices Awards.....	12
<i>Research Methodology</i>	12
About Frost & Sullivan	12

Background and Company Performance

Industry Challenges

Traditional computing architectures, such as central processing unit (CPU) and graphics processing unit (GPU), are more functionally oriented for control flow applications. These architectures are accelerating most hardware through parallel processing and floating-point arithmetic capabilities. With the onset of deep learning programs and software applications, the conventional processor architecture has become inefficient in meeting the needs of the machine learning industry due to the time-consuming master/slave process. Machine learning-based computing industries, especially in the data center arena, need an architecture that can process huge amounts of unstructured data in a parallel fashion.

In a conventional master/slave process control architecture, repetitive bidirectional data transfers as well as parallel-serial-parallel bottlenecks between the system's CPU and GPU create an overhead. Bottlenecks created by the traditional CPU architecture result in performance inefficiency, especially in handling Big Data applications that have large volumes of unstructured data from social networks, search engines, and smartphone sensors. GPUs decrease the problem through parallel processing; however, they require a significant amount of power.

Frost & Sullivan ongoing research shows that another major challenge faced by the machine learning stakeholder ecosystem is the training period of deep neural networks (DNNs). Training DNNs using traditional architecture solutions takes a week or more, with limitations in performance and scalability. Frost & Sullivan points out that, enterprises clearly need an innovative technology that can reduce the growing training times for the DNN model and provide new compute accelerators, with improved performance in an IoT ecosystem.

Technology Attributes and Future Business Value

Industry Impact

Headquartered in Campbell, California, Wave Computing developed its innovative, patented dataflow computing technology architecture to address performance- and scalability-related issues of the machine learning-based computing industry. Wave Computing's dataflow-based compute appliances are accelerating the performance and scalability of both shallow neural networks and DNNs. The team at Wave Computing developed this architecture to compile the deep learning network dataflow graph directly into an optimized wave graph. This optimization enables dynamic instruction reconfiguration that is required in training neural networks with on-chip data persistence.

Frost & Sullivan appreciates Wave Computing's leadership initiative in transforming the conventional computer chip architecture and deep learning industries, and bringing them back to the mainstream. With increasing volumes of unstructured data, Frost & Sullivan believes that Wave Computing's dataflow technology will have a significant impact on machine learning applications in the short term.

Product Impact

Wave Computing's native dataflow architecture is empowering the company's machine learning appliances and solutions and making them quite ideal for public and private enterprises. The dataflow technology enables Wave Computing's compute appliances to deliver high-speed performance for data center and on-premise based neural network training. Wave Computing's solutions and appliances leverage the dataflow architecture to exploit data and model parallelisms that are present in deep learning models, such as convolutional and recurrent neural networks.

For instance, to train a shallow recurrent neural network, a Wave machine learning compute appliance uses a dataflow processing units (DPUs). The compute appliance leverages the dataflow technology's parallel framework functionality, which enables Wave Computing's solution to deliver training more than 600 times faster for Word2vec, a two-layered shallow neural network model, than current hardware acceleration systems.

According to Frost & Sullivan analysis, the dataflow architecture's ability to provide the simple development and deployment of high-speed deep learning models in a machine learning environment provides Wave Computing with a distinct competitive edge over its other competitors.

Scalability

Wave Computing's research team developed dataflow systems that utilize DPUs comprising 16,000 interconnected dataflow processing elements (PEs). The interconnected dataflow PEs enable enterprises to scale the performance of deep learning training and inference models for various frameworks, such as TensorFlow.

Wave Computing's dataflow architecture empowers data scientists with complete software solutions and stacks, including the WaveFlow SDK, WaveFlow agent library, WaveFlow execution engine, and Wave machine learning framework interface. These stacks provide developers and data scientists with localized communication for executing dataflow graphs and for delivering faster results. Wave Computing leveraged its dataflow technology to design new processing units to support and scale the performance of frameworks such as Google TensorFlow and Microsoft Cognitive Toolkit. Wave Computing's deep learning computers based on the dataflow technology deliver up to hundreds of times the performance improvement depending upon the type of network being trained/inferenced compared to traditional CPU, GPU and field-programmable gate array (FPGA) compute accelerator-based systems.

Frost & Sullivan research concludes that the architecture's enhanced capability to scale up to four appliances within a single data center node is making Wave Computing an ideal choice among machine learning enterprises and hyper scalers. The network congestion-reducing ability within the data center enables significant improvement in performance and scalability for training machine learning networks, thereby setting Wave Computing apart from its other competitors. For instance, Google's Inception v3 model can be trained

in as little as four hours by Wave Computing's dataflow solution in a single data center node.

Visionary Innovation

Wave Computing developed its DPU technology to meet enterprises' future demands in the Big Data arena by providing a scalable architecture that improves the machine learning performance. The technology can accelerate a new generation of dataflow-based deep learning frameworks. Initially optimized for the data center, each Wave Computing's 3U compute appliance form factor delivers up to 2.9 PetaOps per second of performance, more than 250,000 processing elements, over 2 terabytes (TB) of high-speed memory, and up to 16 TB of SSD storage.

The major differentiating factor for Wave Computing, compared to its competitors, is its scalable architecture that is designed to improve functional density and computational efficiency with power-reducing mechanisms, while processing elements in DPUs.

According to Frost & Sullivan analysis, Wave Computing's strategy of leveraging data reuse, data parallelism, and model parallelism to accelerate the training of DNN models is enabling enterprises to utilize these models in a more cost-effective way and to gain exciting opportunities in newer markets and application areas.

Customer Acquisition

Wave Computing successfully taped out its commercial 16 nanometer (nm) DPU chip and raised more than \$56 million of additional funds in its Series D round to support the completion and market introduction of its compute appliances. Since inception, Wave Computing has raised more than \$115 million. For 2018, the company is specifically focusing on the direct sales of its compute appliances and solutions. In the future, Wave Computing may broaden its offerings by becoming a chip supplier and providing a range of DPU chips that are focused on deep learning model training and inferencing applications, as well as on low power consumption applications. Today, with the DPU chip and dataflow technology combination, Wave Computing is expanding its customer base worldwide, including in Europe and Asia.

In April 2017, Wave Computing introduced an Early Access Program (EAP) for data scientists and developers to provide them with cloud-based access to a Wave Computing compute appliance prototype before the actual sale kicks off, which is scheduled for late December 2017.

Frost & Sullivan expects that Wave Computing's proprietary technology and focus on developing dataflow-enabled servers for data centers and on-premise environments will help the company achieve significant market share in the near term, thus further providing it with a competitive edge. Wave Computing is strengthening its product portfolio by introducing a deskside appliance for on-premise environments in 2018.

Human Capital and Growth Potential

Wave Computing's CEO, Derek Meyer, holds a Master of Science in Astrophysics from Florida Atlantic University. Mr. Meyer has over 20 years of experience in leading the

semiconductor and computer systems industries and specializes in scaling start-ups to new heights. His industry expertise brings immense knowledge to Wave Computing, enabling tremendous growth potential for the company in terms of technology development.

With a dedicated team focused on developing cutting-edge research in the machine learning industry and holding a suite of patents that offers a strong competitive advantage, Wave Computing has positioned itself as a solid participant with dataflow technology for the machine learning industry. The company's patented native dataflow technology is outperforming all other machine learning products in the industry, thereby attracting new employees and creating a healthy environment for them to excel and to accelerate the functionality of deep learning applications beyond the data center. Wave Computing is aggressively hiring, and has plans to double its headcount by the summer of this year.

Furthermore, Wave Computing has an intellectual property (IP) portfolio of over 70 patents and a successful track record of executing customized chip development for its strategic partners. The company's agile team has brought tremendous expertise in developing the dataflow technology and enhancing its functionality to meet the machine learning industry's increasing demands.

Frost & Sullivan analysis finds that the machine learning industry remains in the emerging stage of development. At Wave Computing, research and development (R&D) is highly active and is bolstering technology development and commercialization activities. In the long term, the company's strong IP portfolio and leadership team will help it thrive in this competitive environment.

Conclusion

The increasing demand for deep learning applications and the exponential growth of unstructured data are clearly creating challenges for the machine learning industry. Through its patented dataflow processing technology, Wave Computing is solving the performance and scalability challenges imposed by traditional solutions used in current advanced machine learning applications.

The company's dataflow-based compute appliances will enable enterprise and public sector customers to improve their business processes through insights gained from the growing volume of data, compared to CPU or co-processor-based offerings. The DPU's high-speed performance, scalability, and energy-efficiency features will drive the adoption of Wave Computing's technology at a faster rate in the short, medium, and long terms.

With its strong overall performance, Wave Computing has earned the 2018 Frost & Sullivan Technology Innovation Award.

Significance of Technology Innovation

Ultimately, growth in any organization depends upon finding new ways to excite the market by maintaining a long-term commitment to innovation. At its core, technology innovation, or any other type of innovation, can only be sustained with leadership in three key areas: understanding demand, nurturing the brand, and differentiating from the competition.



Understanding Technology Innovation

Technology innovation begins with a spark of creativity that is systematically pursued, developed, and commercialized. This spark can result from a successful partnership, a productive in-house innovation group, or a bright-minded individual. Regardless of the source, the success of any new technology is ultimately determined by its innovativeness and its impact on the business as a whole.

Key Benchmarking Criteria

For the Technology Innovation Award, Frost & Sullivan analysts independently evaluated two key factors—Technology Attributes and Future Business Value—according to the criteria identified below.

Technology Attributes

- Criterion 1: Industry Impact
- Criterion 2: Product Impact
- Criterion 3: Scalability
- Criterion 4: Visionary Innovation
- Criterion 5: Application Diversity

Future Business Value

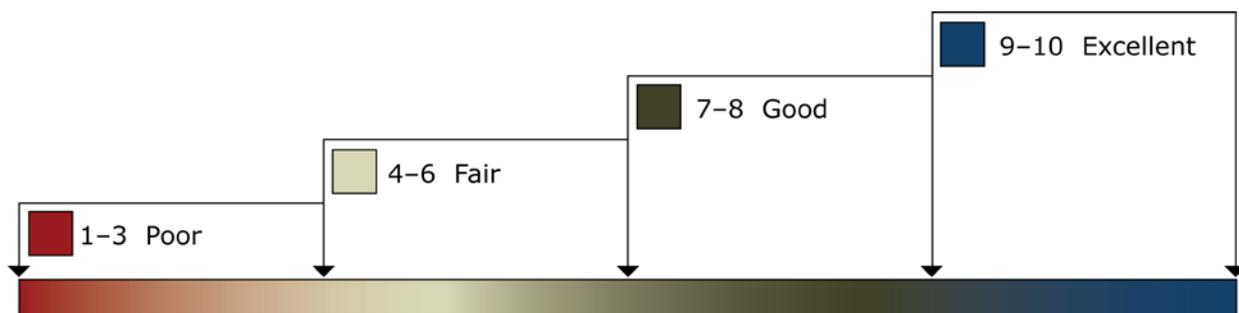
- Criterion 1: Financial Performance
- Criterion 2: Customer Acquisition
- Criterion 3: Technology Licensing
- Criterion 4: Brand Loyalty
- Criterion 5: Human Capital

Best Practices Award Analysis for Wave Computing

Decision Support Scorecard

To support its evaluation of best practices across multiple business performance categories, Frost & Sullivan employs a customized Decision Support Scorecard. This tool allows our research and consulting teams to objectively analyze performance, according to the key benchmarking criteria listed in the previous section, and to assign ratings on that basis. The tool follows a 10-point scale that allows for nuances in performance evaluation. Ratings guidelines are illustrated below.

RATINGS GUIDELINES



The Decision Support Scorecard is organized by Technology Attributes and Future Business Value (i.e., these are the overarching categories for all 10 benchmarking criteria; the definitions for each criterion are provided beneath the scorecard.). The research team confirms the veracity of this weighted scorecard through sensitivity analysis, which confirms that small changes to the ratings for a specific criterion do not lead to a significant change in the overall relative rankings of the companies.

The results of this analysis are shown below. To remain unbiased and to protect the interests of all organizations reviewed, we have chosen to refer to the other key participants as Competitor 2 and Competitor 3.

<i>Measurement of 1–10 (1 = poor; 10 = excellent)</i>			
Technology Innovation	Technology Attributes	Future Business Value	Average Rating
Wave Computing	9.5	9.0	9.2
Competitor 2	8.0	8.0	8.0
Competitor 3	7.5	7.0	7.2

Technology Attributes

Criterion 1: Industry Impact

Requirement: Technology enables the pursuit of groundbreaking ideas, contributing to the betterment of the entire industry.

Criterion 2: Product Impact

Requirement: Specific technology helps enhance features and functionalities of the entire product line for the company.

Criterion 3: Scalability

Requirement: Technology is scalable, enabling new generations of products over time, with increasing levels of quality and functionality.

Criterion 4: Visionary Innovation

Requirement: Specific new technology represents true innovation based on a deep understanding of future needs and applications.

Criterion 5: Application Diversity

Requirement: New technology serves multiple products, multiple applications, and multiple user environments.

Future Business Value

Criterion 1: Financial Performance

Requirement: Potential is high for strong financial performance in terms of revenues, operating margins, and other relevant financial metrics.

Criterion 2: Customer Acquisition

Requirement: Specific technology enables acquisition of new customers, even as it enhances value to current customers.

Criterion 3: Technology Licensing

Requirement: New technology displays great potential to be licensed across many sectors and applications, thereby driving incremental revenue streams.

Criterion 4: Brand Loyalty

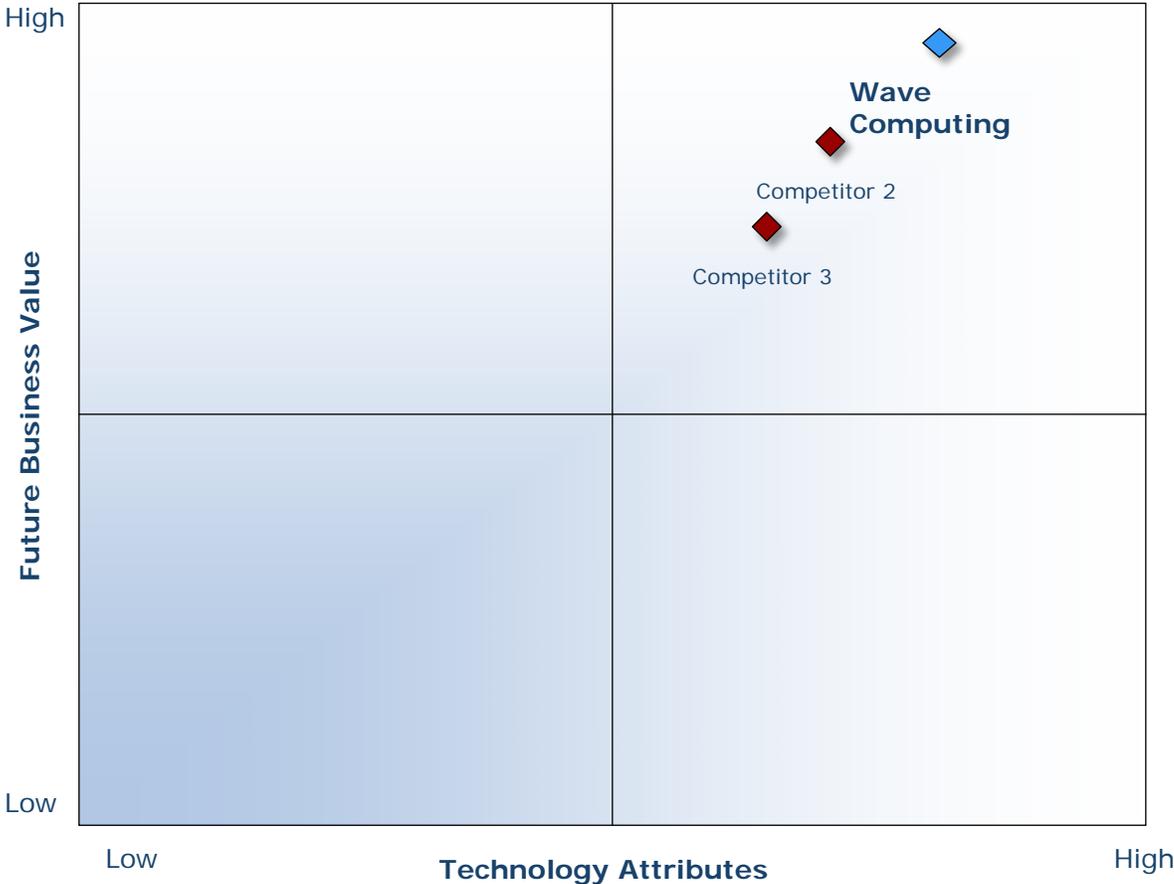
Requirement: New technology enhances the company’s brand, creating and/or nurturing brand loyalty.

Criterion 5: Human Capital

Requirement: Customer impact is enhanced through the leverage of specific technology, translating into positive impact on employee morale and retention.

Decision Support Matrix

Once all companies have been evaluated according to the Decision Support Scorecard, analysts then position the candidates on the matrix shown below, enabling them to visualize which companies are truly breakthrough and which ones are not yet operating at best-in-class levels.



Best Practices Recognition: 10 Steps to Researching, Identifying, and Recognizing Best Practices

Frost & Sullivan analysts follow a 10-step process to evaluate Award candidates and assess their fit with select best practice criteria. The reputation and integrity of the Awards are based on close adherence to this process.

STEP	OBJECTIVE	KEY ACTIVITIES	OUTPUT
1 Monitor, target, and screen	Identify Award recipient candidates from around the globe	<ul style="list-style-type: none"> • Conduct in-depth industry research • Identify emerging sectors • Scan multiple geographies 	Pipeline of candidates who potentially meet all best-practice criteria
2 Perform 360-degree research	Perform comprehensive, 360-degree research on all candidates in the pipeline	<ul style="list-style-type: none"> • Interview thought leaders and industry practitioners • Assess candidates' fit with best-practice criteria • Rank all candidates 	Matrix positioning of all candidates' performance relative to one another
3 Invite thought leadership in best practices	Perform in-depth examination of all candidates	<ul style="list-style-type: none"> • Confirm best-practice criteria • Examine eligibility of all candidates • Identify any information gaps 	Detailed profiles of all ranked candidates
4 Initiate research director review	Conduct an unbiased evaluation of all candidate profiles	<ul style="list-style-type: none"> • Brainstorm ranking options • Invite multiple perspectives on candidates' performance • Update candidate profiles 	Final prioritization of all eligible candidates and companion best-practice positioning paper
5 Assemble panel of industry experts	Present findings to an expert panel of industry thought leaders	<ul style="list-style-type: none"> • Share findings • Strengthen cases for candidate eligibility • Prioritize candidates 	Refined list of prioritized Award candidates
6 Conduct global industry review	Build consensus on Award candidates' eligibility	<ul style="list-style-type: none"> • Hold global team meeting to review all candidates • Pressure-test fit with criteria • Confirm inclusion of all eligible candidates 	Final list of eligible Award candidates, representing success stories worldwide
7 Perform quality check	Develop official Award consideration materials	<ul style="list-style-type: none"> • Perform final performance benchmarking activities • Write nominations • Perform quality review 	High-quality, accurate, and creative presentation of nominees' successes
8 Reconnect with panel of industry experts	Finalize the selection of the best-practice Award recipient	<ul style="list-style-type: none"> • Review analysis with panel • Build consensus • Select recipient 	Decision on which company performs best against all best-practice criteria
9 Communicate recognition	Inform Award recipient of Award recognition	<ul style="list-style-type: none"> • Present Award to the CEO • Inspire the organization for continued success • Celebrate the recipient's performance 	Announcement of Award and plan for how recipient can use the Award to enhance the brand
10 Take strategic action	Upon licensing, company is able to share Award news with stakeholders and customers	<ul style="list-style-type: none"> • Coordinate media outreach • Design a marketing plan • Assess Award's role in future strategic planning 	Widespread awareness of recipient's Award status among investors, media personnel, and employees

The Intersection between 360-Degree Research and Best Practices Awards

Research Methodology

Frost & Sullivan's 360-degree research methodology represents the analytical rigor of our research process. It offers a 360-degree-view of industry challenges, trends, and issues by integrating all 7 of Frost & Sullivan's research methodologies. Too often companies make important growth decisions based on a narrow understanding of their environment, leading to errors of both omission and commission. Successful growth strategies are founded on a thorough understanding of market, technical, economic, financial, customer, best practices, and demographic analyses. The integration of these research disciplines into the 360-degree research methodology provides an evaluation platform for benchmarking industry participants and for identifying those performing at best-in-class levels.

360-DEGREE RESEARCH: SEEING ORDER IN THE CHAOS



About Frost & Sullivan

Frost & Sullivan, the Growth Partnership Company, enables clients to accelerate growth and achieve best-in-class positions in growth, innovation and leadership. The company's Growth Partnership Service provides the CEO and the CEO's Growth Team with disciplined research and best practice models to drive the generation, evaluation and implementation of powerful growth strategies. Frost & Sullivan leverages more than 50 years of experience in partnering with Global 1000 companies, emerging businesses, and the investment community from 45 offices on six continents. To join our Growth Partnership, please visit <http://www.frost.com>.