



Accelerate the future by offloading computation

Higher performance for faster data insights

Data is coming at us faster than we could ever hope to process. But when your business model—or your customer's business model—depends on the real-time use of that flood of data, you need to increase efficiency immediately. It requires you to think beyond a traditional computing setup with hardware specifically designed for your purposes. You need to offload.

Break free from computational limitations

A traditional compute setup was designed for the at-rest (or pooled) data that has already been collected. Periodically combing through those data sets for useful information is among the earliest uses of computer systems and still provides very important functionality today. Newly available streams of real-time data drive

the need for instant decision making by combining this data with these at rest data sets. Real time data deserves a real time answer. Making real time decisions from multiple data sources exponentially compounds the compute power needed.

The right hardware for the job

Even with the versatility and reliability of a CPU, its capabilities have been long augmented with Graphics Processing Units (GPUs). Just as a GPU freed up the CPU for silky smooth graphics, people quickly learned how to push more work to the GPU that wasn't graphics related. General-purpose GPU (GPGPU) performs highly parallel calculations to further increase performance and unburden the CPU, but with roots in graphics processing, it can be somewhat limited.

Intelligent Processing Units (IPUs) are made specifically for dealing with the math and data typically used in AI training and workloads. While these IPU's are more limited in scope than a GPGPU, they can yield incredible results when running the workload they were designed to serve.

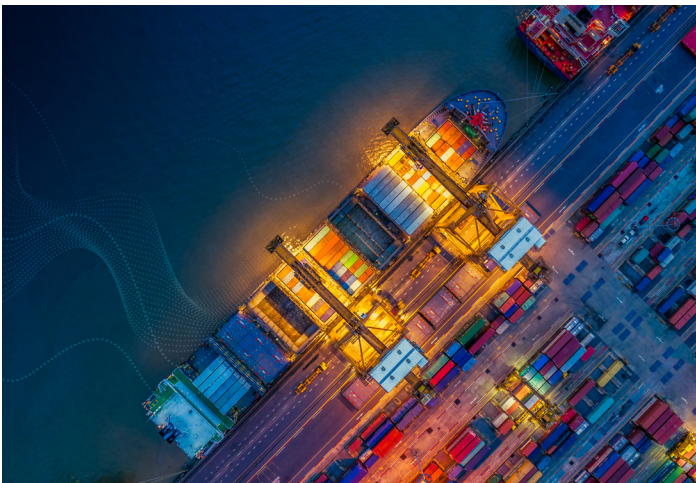
Smart NICs are another type of offload in our toolbox. Smart NICs are especially useful when it comes to in-flight data. As packets of data come in from the network, the NICs determine whether this data is important for consideration and routing to the proper processing engine. The ability to make decisions about incoming data and handle it properly inside the Smart NIC will dramatically decrease network latency. Almost all Smart NICs also bring along crypto and compression engines to increase security of this data and prevent the CPU from needing to do these compute intensive operations.



Increasingly complicated data calls for more complex solutions

Part of the value of data is the immediacy in which insights can be used. Efficiency is paramount, as is the customized approach to fit the exact needs of your business. At-rest data can be processed casually as resources permit, while in-flight data needs the ability to process quickly, often relying heavily on AI which might use some of the at-rest data in the decision-making process. The same equipment does not serve both roles equally well.

Processing in real-time at scale while maintaining profitable efficiency requires offloading. While scaling out to multiple traditional CPU based systems might work, right sized offload will yield savings in operational efficiency and increase performance. Just as one GPGPU, when applied to the right workload, can do the work of dozens of CPUs with a fraction of the power, the coming IPU's will be able to do the work of several GPUs while increasing power efficiency. To gain the largest savings and best performance, it's going to be critical to pair the right workload with the right offload.



An investment in higher performance

The flexibility and effectiveness in offloading is so impactful that resources allocated to those upgrades pay for themselves within weeks' time. It's a shift that impacts not just the bottom line and ability to adequately compete in the marketplace, but will help you achieve sustainability goals through immediate energy consumption improvements, overcome compute limitations and establish a new level of performance.

It's only the beginning

The complexities involved in adopting a strategic offload process are substantial. We're embarking on a future of possibilities that takes machine learning and applies it to day-to-day processes in businesses just like yours. We know it's an ambitious effort to design a solution that solves for managing data in real time at scale while maintaining profitable efficiency, and yet, we're already there. We will help you get the open system you need and embrace as many industry standards as possible. Dell Technologies OEM Solutions welcomes you as a partner in designing your intelligent future.



To start discovering the right technical solution for your workload, contact your account team representative.

To learn more, contact your account team representative or visit [DellTechnologies.com/OEM](https://www.delltechnologies.com/OEM)