Hitachi NEXT 2018

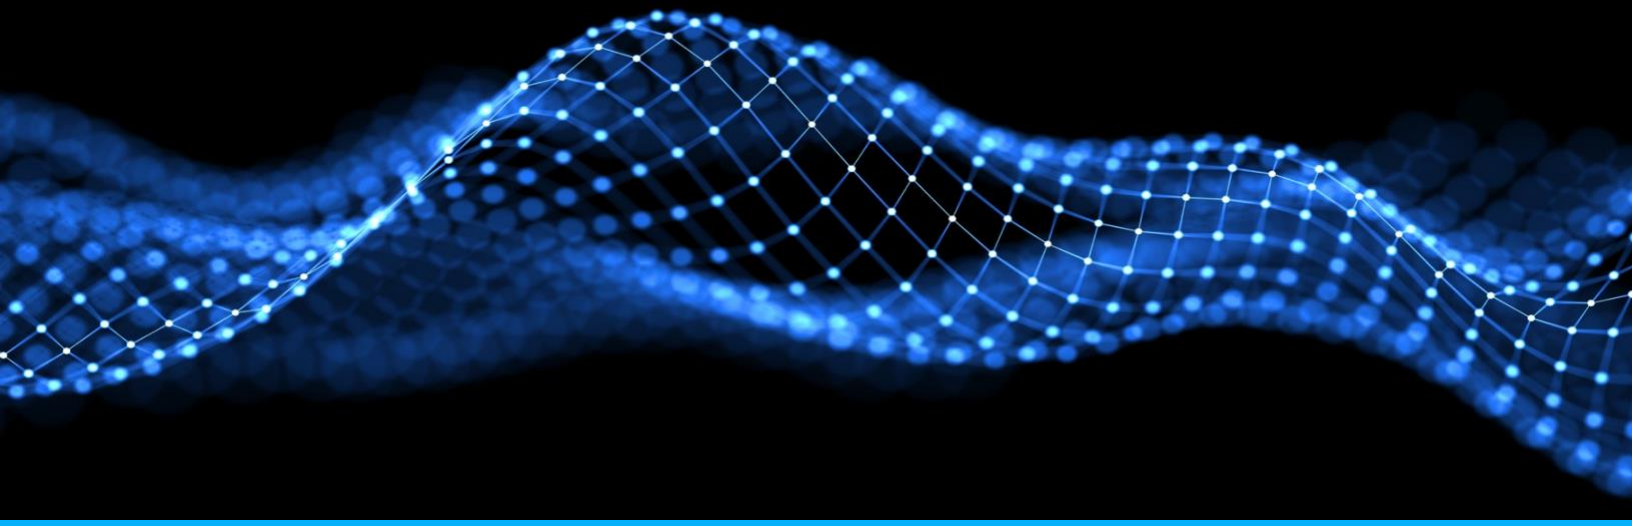

HITACHI VANTARA INNOVATION - TECHNOLOGY REVIEW

The Architect's View

Hitachi Vantara continues to build on a strategy of enabling customers to create, store, manage and exploit their data assets. At Hitachi NEXT 2018, the company announced new products and services for enterprise customers. Specifically, Hitachi Vantara discussed the following technologies in the keynote session:

- *IoT* - Video Analytics and LiDAR
- *Data/Analytics* - Machine Learning Model Management
- *Core Storage* - Hitachi Unified Compute Platform HC (UCP HC) – all-flash HCI
- *Infrastructure* – FPGA acceleration and intelligent offload
- *Infrastructure* – Digital Twins for Data Center modelling

These offerings meet the strategy of collecting data (increasingly at the Edge), storing and processing data on Hitachi core solutions and using Hitachi analytics products to derive value from the data. For customers, the focus is on using information and the insights they create to drive business agility and create competitive advantage. Hitachi Vantara has a portfolio of products and services that position the company as a leader in data management and analytics solutions for the enterprise. In this paper, we look at the announcements and how customers can use them to add value to their business.

Video Analytics & LIDAR

What is it? - LIDAR is a technology that performs real-time mapping of an area in 3D by dynamically scanning objects with laser beams and measuring the time it takes to receive the reflected signal back - the Time of Flight or ToF.

The technology is analogous to RADAR, which uses radio waves, but operates at much shorter ranges with higher degrees of accuracy. As an example, autonomous vehicles use expensive long-range LIDAR to detect objects in their travel path. Hitachi has developed a compact and affordable LIDAR device that can produce real-time data over short distances, such as in offices, airports and retail outlets, while delivering higher resolution than long distance LiDAR.

Video analytics use computer vision and machine learning (Artificial Intelligence, or AI) to distinguish and count types of objects in video data, such as people, vehicles, bags, etc. and deliver automated alerts and statistical insights about those objects. Hitachi Video Analytics (HVA) can analyse archived video data or real time video streams.

What are the applications? For customers, HVA and LIDAR can be used in conjunction with ML/AI analytics tools to create insights on customer behaviour. Examples include:

- Measuring footfall within and outside of retail outlets to determine how best to attract customers either with window displays or by improving the store layout. Existing LIDAR customers have been able to improve sales simply by identifying where best to place retail staff on the shop floor.
- Queue management – in large public areas such as airports, HVA can detect and highlight when queues build up, allowing the dynamic deployment of resources and staff to keep waiting times to a minimum in security, check-in and other areas.
- Video analytics and LIDAR have also been used to measure the attendance of classes within academic institutions and optimise the use of classrooms. This also resulted in the cost avoidance of extending existing facilities or building new classrooms.

Collecting and storing the data from LIDAR devices is only one aspect of using the technology. With Hitachi Video Analytics software, customers can already perform a range of tasks that include: Activity Visualiser, People Counter 3D, Traffic Analyser, Licence Plate Recogniser, Queue Detector, Intrusion Detector. The value for customers is in collecting new types of information like LIDAR, while finding new ways to exploit existing data sources like video. Hitachi's approach to these analytics minimizes false positives, improves accuracy and reliability, and offers privacy protection.

LIDAR and Hitachi Video Analytics both have wide applications across businesses internally and customer-facing scenarios. Businesses, airports, governments, and factories can use the technologies to improve the workplace, gain operational insights, and make their facilities and open areas safer. The technology also allows businesses to develop greater insights into customer activity in retail, transportation and campus verticals.

The following links provide more information on Hitachi Vantara LIDAR and Video Analytics solutions.

1. Hitachi Video Analytics datasheet (PDF, Hitachi Vantara Website)
2. Hitachi LIDAR and IoT Imaging Solutions examples (Hitachi America website)
3. Hitachi Smart Spaces and Video Intelligence (Hitachi Vantara Website)

Machine Learning Model Management

What is it? Historically, Machine Learning (ML) models (algorithms) have been developed in a linear fashion with data science projects collating data, training a model and iterating until the model is good enough to be used. This is a classic set of sequential steps:

Prepare Data->Engineer Features->Train, Test and Build Models->Deploy Best Performing Model

Once that model is in production, the challenge is how businesses deal with updating and maintaining the accuracy of those models as the underlying distribution of the data to what was predicted changes over time, for example, customer purchasing behavior. Hitachi Vantara Labs is working on tools and techniques that will effectively "close the loop" and make the process of training and implementing ML/AI models more streamlined and efficient. The result is an iterative approach of Monitor, Evaluate, Compare, Rebuild and repeat.

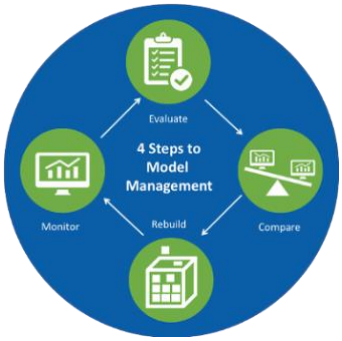

What are the benefits? For customers, developing ML./AI models represents a significant time and effort. There's a continuous need to ensure that model accuracy is improved and the benefits from using those models increases. As more businesses start to use data to increase competitive advantage, so the development and training of ML/AI models needs to be scalable. This is a three-step process:

A – get models into production faster
B – maximise the accuracy of models
C – implementing governance around the development of models.

Through Pentaho, Hitachi Vantara already has a rich set of tools to visualise data. The next step is to scale the ML/AI model management process efficiently by providing features in the Pentaho software that enables consistency when developing and testing models using disparate toolsets. Furthermore, productivity can be enhanced by allowing such development and management to be performed graphically, as opposed to having to write code. Hitachi Vantara Labs is actively developing new tools that will improve the ability to build, manage and maintain ML/AI models. These solutions are available to the open source community, already generating thousands of downloads via the Pentaho marketplace.

Customers that are already building and managing data models will be aware of the ongoing management scalability issues. Model management isn't a specific product, but more a statement of intention on understanding the way in which data models are built and helping customers make this process work more simply and efficiently. This will translate into features that could ultimately be integrated into the Pentaho platform.

The following links provide more information on Hitachi ML/AI solutions and model management.

1. 4-Steps to Machine Learning Model Management (Hitachi Vantara Community Blog)
2. Machine Intelligence Made Easy (Hitachi Vantara Community Blog)

Hitachi UCP HC

What is it? The new Hitachi UCP HC V124N is an all-NVMe enabled hyper-converged hardware platform using VMware vSphere and vSAN. The solution introduces new Intel 3D-XPoint technology (branded as Intel Optane) to further accelerate performance over what can be achieved purely with flash storage alone.

Why Optane? The Intel Optane architecture is a new form of persistent storage that operates at very low latency and high throughput. Optane is also designed for a greater endurance than flash storage, in that it can handle much greater volumes of write I/O without wearing out.

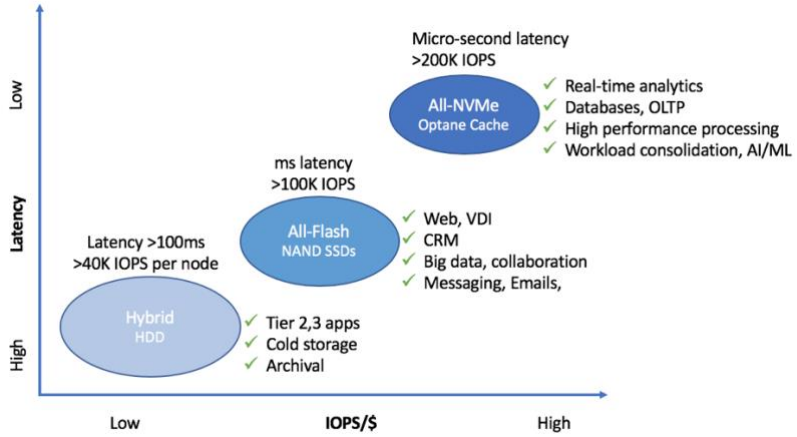

As the speed of processors is improved through adding more CPU cores, traditional SAS and SATA drives can't deliver the performance required to fully exploit the processor capabilities. NVMe and Optane are technologies that address this performance gap.

Why Optane and vSAN? VMware vSAN provides customers with a scale-out storage and compute platform that requires fewer technical skills to deploy and manage – hyper-converged infrastructure or HCI. The virtualisation administrator can handle the installation and management of storage, without needing a separate or dedicated team of storage specialists.

The architecture of vSAN splits data into two areas; active or hot I/O (both reads and writes) and cold or inactive data. Active I/O is served from cache, which in the past has been fast hard drives or resilient flash (e.g. SLC or MLC technology). In the V124N model, Optane is used as the cache tier. This allows the majority of I/O to be served from cache, while storing the remaining colder data on cheaper NAND flash.

What are the benefits? In performance tests with VMware HCIBench and HammerDB, the following improvements were measured, compared to a baseline system using SATA SSDs: *225% increase in IOPS*, with multiple virtual machines in use; *3x increase in IOPS* for write-intensive workloads; *70% reduction in disk latency* for write-intensive workloads; *4x greater IOPS/$*. For relational databases, where individual I/O latency is directly related to database performance, the use of Intel Optane in UCP HC resulted in the following improvements; *2x transactions per minute* throughput; *2x IOPS*; *1/6th the latency*. Another significant benefit, critical for enterprise applications, with UCP HC V124N is consistently low latency even at peak workload.

Flash and Optane are expensive resources, so getting the right balance of the two to obtain the best price/performance ratio requires engineering testing and development. Hitachi is one of the first companies to bring this kind of solution to the market, having significant previous experience with solid-state storage solutions.

The platform aligns with existing HCI offerings in adding increased levels of performance, while targeting applications in today's enterprise that need high throughput and low latency. This includes traditional database applications, workload consolidation and new ML/AI workloads. V124N is designed for write-intensive applications that also need low latency.

The following links provide more information on Hitachi UCP, Optane and VMware vSAN.

1. Hitachi Vantara UCP HC Page (Hitachi Vantara Website)
2. VMware vSAN Product Page (VMware website)
3. Hitachi UCP HC Family Specification Sheet (PDF, Hitachi Vantara Website)
4. Enterprise Storage Arrays and NVMe (Hu Yoshida, Hitachi Vantara Community Blog)
5. Hitachi +VMware vSAN + Intel Optane NVMe = Turbocharged HCI (Dinesh Singh, Hitachi Community Blog)
6. The Race towards End-to-End NVMe in the Data Centre (Architecting IT Blog)

FPGA Acceleration & Intelligent Offload

What is it? FPGAs (Field Programmable Gate Arrays) are a type of processor that can be programmed in place (in the field) to improve the processing of specific workloads and applications, compared to general CPUs or processors. An FPGA can be used to accelerate or improve the parallelisation of a set of tasks at a fine-grained level, improving the processing performance compared to standard processors alone.

What are the benefits? Hitachi Labs are working on using FPGA acceleration to improve the performance of applications by offloading tasks to FPGAs for common workloads such as traditional databases, machine learning and analytics. In the Hitachi Vantara Lab, FPGA acceleration has been demonstrated to improve the performance of a Pentaho analytics engine, delivering between 10 and 100 times performance improvement compared the same test without FPGAs. The benefit for customers is in achieving faster time to value and being able to process more data for quicker and more accurate results with less hardware and the associated cost.

Hitachi already has over 90 patents in this area and had dedicated thousands of man hours of work in developing solutions. Historically, Hitachi has been using FPGAs for many years in the NAS line of storage products. Today FPGA acceleration is being designed to work with Hadoop, SQL databases and MLP neural networks. To make the FPGAs practical to deploy, Hitachi is developing a framework that will allow FPGA technology to be applied to a range of workloads, rather than being specific to one type of application. This makes the solutions more "software-defined" and so more easily adaptable to differing workloads, reducing the need to use dedicated or custom hardware for each application use case.

FPGA acceleration and offload have significant benefits for real-time or time-sensitive analytics applications. The technology is also expected to provide significant benefits at the edge, by allowing improved processing without requiring large amounts of hardware to be deployed.

The following links provide some additional background on Hitachi Vantara's use of FPGAs and the importance of hardware.

1. Does Hardware Matter?
2. Digital Transformation Requires New Architectures for Network Attached Storage

Digital Twins

What is it? As IT infrastructure becomes more distributed and complex, companies need greater insight into how components of technology interact and behave in normal and failure scenarios. A digital twin is a virtual model of infrastructure that can be used to simulate, visualise and model the real world in order to manage the uptime and resiliency of infrastructure. The digital twin can mirror configuration information, track performance data and capacity utilisation of data centre infrastructure. This includes traditional servers, storage and networking as well as telemetry information coming directly from applications.

What are the benefits? In large distributed environments, it becomes impossible for any one individual or even teams to understand the complex interactions between components of infrastructure and evaluate the impact of failure to ongoing operations. Building a digital twin enables IT organisations to manage and de-risk issues that could impact resiliency and uptime. ML/AI capabilities extend this efficiency to identify issues that can lead to overspending or poor customer experiences. Hitachi Vantara is bringing digital twin technology into the data centre to monitor assets, identify bottlenecks and provider much better planning around budgets and spending on both hardware and software resources.

With a digital model, businesses can play "what if" scenarios much more accurately, modelling the results of device failures or other stresses on the system, such as high volumes of application traffic. With more accurate modelling data, business can build out contingency plans more effectively, rather than, for example, simply buying more hardware.

Hitachi already has many production implementations of digital twin technology for IoT applications. This allows Asset Avatars (Hitachi's name for the digital representation of a physical asset) to be developed from real-world sources such as sensors. The asset avatars work with historical data, Hitachi Lumada and Pentaho platforms to create analytics insights that can be developed without large amounts of coding. Within the data centre Hitachi Vantara already collects data from hardware infrastructure. The principles of developing asset avatars can now be applied to infrastructure telemetry, providing insights on the way IT organisations are deploying and using their internal infrastructure.

Any large organizations that have significant, complex and distributed data centres and IT infrastructure will be experiencing challenges in managing uptime, resiliency and costs. As Digital twins come to the data centre, these customers will be able to benefit from the advantage the technology delivers in improving the capabilities of Hitachi Vantara's existing data centre management tools.

The following links provide more information on understanding the concept and usage of digital twins.

1. IoT Minute Episode 69: What in the World is a Digital Twin?

 DocID: BRKWP0103. Last revised 13 December 2018. Revision 1.1.