HPE Edgeline and Microsoft Azure:

Designing cutting-edge video analytics solutions from the intelligent edge to the cloud

Abstract

Artificial Intelligence (AI)-enabled video analytics solutions are being built today to deliver disruptive business outcomes in different industry verticals. This whitepaper describes how the combination of HPE's Edgeline Converged Edge Systems portfolio of systems at the intelligent edge coupled with Microsoft-supported accelerators and the Microsoft Azure cloud is helping create the next generation of AI-enabled surveillance applications.

Contents

The evolution of video surveillance

Video today serves as an integral component of almost any company's security profile. From shoplifters to threat assessments, video helps companies capture and store a visual record that can assist law enforcement in apprehending individuals before, during, and after they commit crimes.

Up until now, the video and security narrative focused on video surveillance with human oversight. But as technology has evolved, so too has possibility. In place of rooms full of monitors scrutinized by security personnel, companies are now turning to smart technology with automated components and built-in analytics to help identify suspicious behavior.

With all these amazing capabilities, however, has come new challenges. One concern's storage. Companies are collecting massive amounts of video imagery for analysis. Not only that, but much of the content is high-definition video, which requires extra storage space. Due to cost and scalability issues, many companies can only store video for the previous thirty days. That's not very useful in issues that stretch beyond that timeframe.

A second challenge is extracting the most value from collected data. We've reached a point where the sheer amount of video imagery being generated has exceeded the human ability to accurately comprehend, quickly analyze, and effectively respond to them. The Internet of Things (IoT) and the rise of intelligent devices connected to the internet are creating massive amounts of data. In fact, every year brings with it enough data to exceed all data collected previously online.

A final challenge is the issue of ethics, privacy, and analytics. While the potential of artificial intelligence (AI) is boundless, there needs to be principles in place to ensure video systems designed and deployed with AI and machine learning are done in an ethical way that protects user and company privacy. Anything else becomes invasive and a danger to business viability.

IoT can convert images into metadata that can be centrally aggregated, analyzed, and stored, enabling analysts to view video streams from multiple cameras in multiple locations. With the use of machine learning to identify important bits, they can get faster and more accurate video analysis. This helps businesses be more efficient at a lower cost, cities be much safer, and citizens more secure. IoT also helps replace the manual responsibility of having people monitor multiple screens eight hours a day, a task no one really enjoys and one that could have health impacts later, such as with sight and mobility.

That's not to say there still aren't obstacles for IoT to overcome. Some companies view sending massive amounts of video data to a central location for analysis as an unacceptable risk. Take insight-delivery latency, which can threaten employee safety. Imagine a manufacturing facility where an employee's about to cross over the perimeter of a dangerous machine. A lack of real-time notification or a dependency on the cloud prevents immediate action. Or consider a train platform where there's an unattended piece of luggage. It could potentially threaten hundreds of lives if not quickly addressed. Such situations require real-time alerts to avoid potential disasters.

To address the challenges of notification and response, technology providers are working on ways to take in data and process it faster. One such way is to process information at the edge, where it occurs and analyze it locally, saving the time and dependency of data transport. Converged edge computing offers several benefits in general and specific benefits to businesses. These are outlined in Figure 1.

Video analytics on the edge – the intelligent edge

We previously reflected upon the evolution of video surveillance to analytics. It's important to remember that video analytics within the business world has its own history too. In the past, analytic processes, especially out in the field, often proved laborious. Information would get collected offsite and then sent to analysts for review. That information would then get sent to decision makers who would take actions, which then would get relayed back to the field.

Now take that same scenario and apply it to IoT and the intelligent edge. Information gets collected where it lives and stored locally. Machine learning and AI, deployed and executed at the edge, drive fast analytics that significantly reduce latency and bandwidth while increasing the immediacy and value of insights. With thirty million cameras already in use, generating more than four billion hours of video footage per week, the impact is significant.

Improvements in systems-design technology are working to streamline this process even more. The availability of uncompromised enterprise-grade compute and acceleration at the edge leads to an analytical continuum from edge to the cloud whereby enterprises of the future can be edge-centered, cloud-enabled, and data-driven.

Figure 1. Convergence at the edge

These are the seven key drivers for converged edge computing and their business impact.

It's important to remember that video data isn't the only example that illustrates the value of analysis at the edge. There's also operations technology (OT) data collected from networks and often used for industrial control and supervisory control and data acquisition systems (ICS/SCADA). This data relies on a collection of devices that work together to monitor and control how physical devices perform. When collected and analyzed together, these two data streams differ enough to warrant using them both for analysis and insight.

HPE not only recognizes the differences between these data sources but also their combined value. By enabling their physical convergence for an optimized and purpose-built product line, they can help deliver edge video analytics from not only multiple cameras in multiple locations, but also across multiple edges.

AI for video analytics

Solving today's business problems

With intelligent, converged edge video analytics, organizations today can deliver solutions that address pressing business problems. Below are three industry examples with the potential to significantly impact the bottom line of different industries.

Retail: Out-of-stock management

When customers go shopping, they expect the items they want to be in-stock. As we all know, that isn't always the case, which impacts customer experiences. Proper inventory management could prevent or significantly reduce such inventory issues. Traditional systems however, have proven somewhat limited. On the one hand, they can scan items as they come in and again as they leave a store. But they can't tell you where an item is located within the store or if it's been misplaced. Is it on the floor or back in the warehouse? Questions like these require sending employees running around the store in search of answers.

Video analytics powered by AI can focus on detecting product on the shelf (object detection). Retailers can extend the use of their vast camera streams to build an "eyes on the ground" approach to identifying, counting, and tracking inventory items in real-time. This saves employees the time required to locate items, freeing them for more productive tasks. And with more accurate inventory and item location, customers can find the items they want to purchase. Meanwhile employees can be more productive while focusing on helping customers, which improves customer satisfaction and helps build customer loyalty.

Travel: Airline safety

Passenger safety is always a concern in the airline industry. To that end, regulators set certain industry standards and practices that must always be complied with. Failure to do so can lead to unsafe passenger conditions, bad passenger experiences, and large overall fines.

By tapping into camera video streams and using AI that focuses on aircraft activity (object detection), ground crew personnel can monitor aircraft speed and paths in real-time during parking and pushback from busy airport gates. Non-compliance conditions can be quickly identified, driving faster reactions to correct conditions. When deployed across an entire airport, ground crew managers can gain an "eye in the sky" perspective across all gates, instead of a single gate at a time.

Manufacturing: Industrial quality assurance (QA)

Manufacturing companies rely on product lines in their facilities to streamline processes, reduce cost, ensure quality, and maximize profit. That's where quality assurance (QA) comes in. Within the QA process, inspectors rely on product images to determine if an item gets a "pass" or "fail." Their decisions come from their line of sight and visual perspectives. Mistakes can easily occur, which can prove costly. For example, if an item gets flagged when it shouldn't, the downtime to verify the situation slows down production. Even worse is if something that should be flagged slips through and reaches the customer. Defective products can hurt a brand, prove costly to recall, repair, or replace and even more important, put consumers at risk.

Video analytics and real-time AI can optimize factory processes and accuracy as well as increase throughput. The onus of QA shifts from people to computers running a trained AI model to detect anomalies and failures. More items can be quality tested too. This doesn't mean automation replaces the human element. In fact, it helps QA personnel focus on failures that require more in-depth scrutiny (optimized human capital). That leads to better results. It also enables manufacturers to more quickly address production issues.

How AI works at the intelligent edge for vision

What is an AI model?

The key to edge intelligence comes from AI. Trained to think the same way humans do, AI associates neurons with patterns, compiled to a model, which is then updated and pushed out to hardware at the edge. The diagram below illustrates how a computer processes a single image from a video stream to determine what is being identified.

How AI helps a computer processes an image

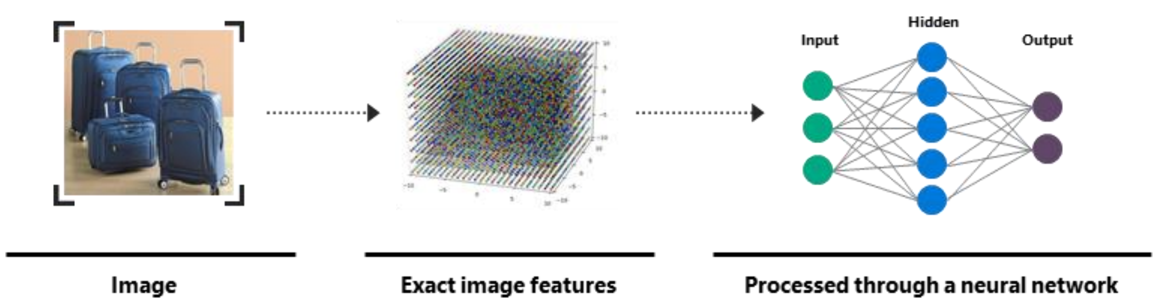

Figure 2. Images are processed through a "trained" neural network for decision-making (Inferencing).

This 'neural network' functions much like the human brain. It takes an image, intuitively notices features (color, size, curves etc.) and associates those features with past experiences to derive a decision about the image.

While the human brain does this amazingly quickly, computers aren't as fast and require large amounts of processing (compute power) to try to think as humans do. This is where 'hardware acceleration' or specialized hardware optimized for neural network processing helps.

Just like people, once the model's built and in use, it can only evaluate what it's been taught and trained on. It requires constant retraining based on new and changing factors. In the aforementioned case of unattended luggage at a train depot, retraining might include new styles or designs of luggage, changing patterns on how people pack luggage, or seasonal changes like snow skies during the winter.

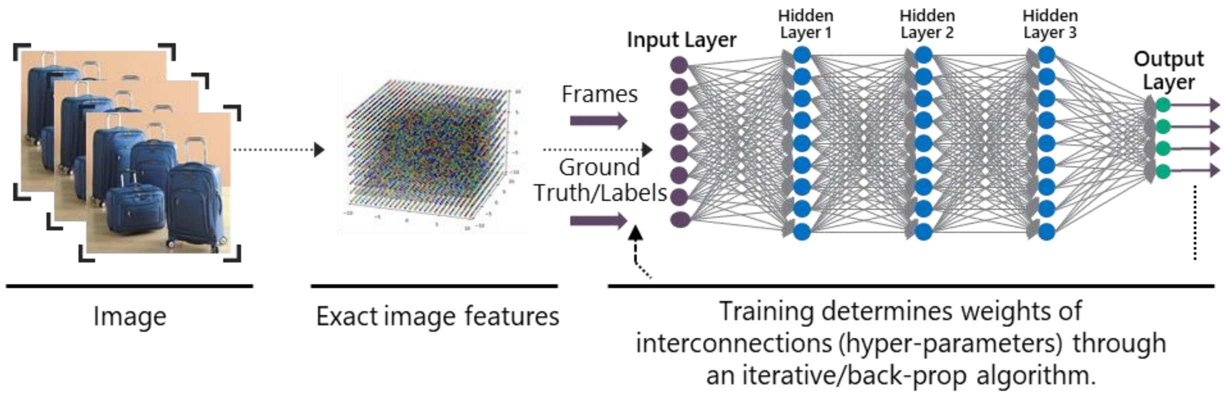

Figure 3. Creating a "trained" neural network using pre-trained images and ground-truth labels (Training).

AI at the edge architecture

When thinking about a video analytics at the edge solution, we commonly see three main components: Cameras/sensors, one or more edge systems, and one or more clouds (see Figure 4).

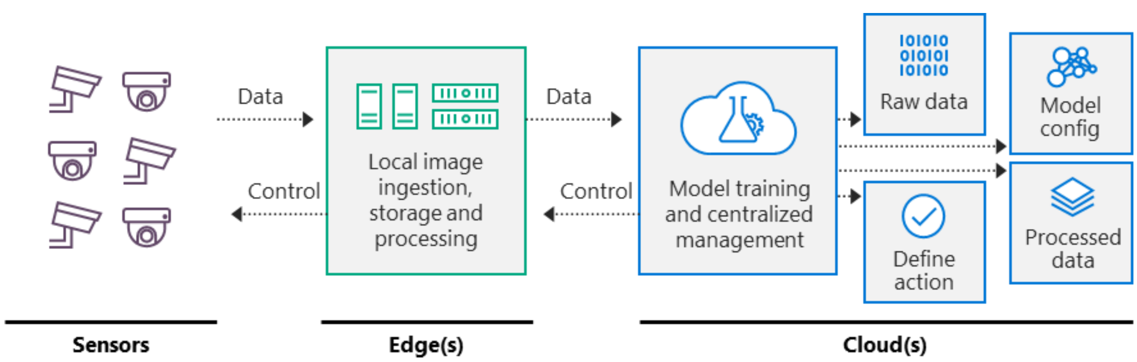

Figure 4. Components of a video analytics at the edge solution.

In this architecture, cameras are just one of many types of sensors generating data to be locally processed at the edge device. The edge device is responsible for short-term storage, data processing, local actions, controlling the cameras, and relaying data, as needed, to the cloud for later processing or storage. The cloud is responsible for long-term data storage, processing data and data relationships, creating, training, and retraining the AI models, and defining configuration for the edge device and cameras. While there's a one-to-many relationship with the single cloud to multiple edge systems, there's also a one-to-many relationship with the edge systems to multiple cameras, all working together as a single solution.

At any level, you can have more than one component, such as one edge device sending data to another edge device or an edge device sending data to one or more public and private clouds.

Each component is hardware fit for purpose, optimizing cost while addressing compute and acceleration needs that are different at each level. Each component is software fit for purpose, where more complex technologies exist in the cloud with optimized technologies run on the edge and even on smart cameras – which can combine some of the edge processing task at the camera.

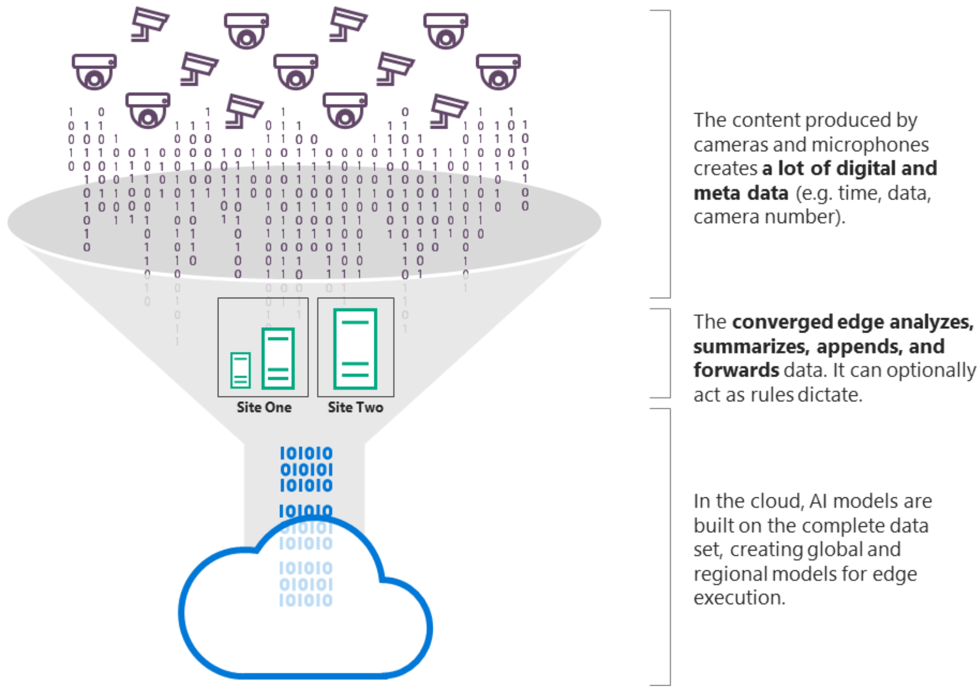

Figure 5. How video analytics at the edge works.

Lastly, there are at least two types of data flows: A control plane and a data plane. The control plane's about controlling or managing the devices which can include device or configuration updates. The data plane's for data, such as video streams, binary data or meta data relating to the items being analyzed.

Sensors and cameras

There are many camera and sensor vendors in different industries. When performing QA inspection, for example, manufacturers often need special high-speed or infrared cameras. Other industries need outdoor-certified, wireless cameras with night vision. Whether using a wired camera connected with USB or other media or using a network-based camera or camera concentrator accessed using real-time streaming protocol (RTSP), cameras serve as a data input and often can be shared with existing or legacy needs.

As mentioned above, some smart cameras can serve as a camera and an edge processor, talking directly to the cloud or to other edge systems. One example is the Microsoft Vision AI Developer kit: **http://aka.ms/visionaidevkit**, which enables developers to prototype video analytics solutions.

Edge systems

Edge systems are computer-based and have the processing power and storage to function autonomously for limited network connectivity solutions. When processing AI models [neural networks] the edge hardware needs to be sized for the appropriate need, including compute, storage, hardware accelerators, environment condition certification and secure enclosure. Selecting the right edge device is important for successful video analytics solutions.

Edge systems combine operating systems, hardware accelerators, accelerator runtime/drivers, edge runtimes, and AI runtimes to deliver the complete solution. They also need to have the appropriate I/O interfaces to acquire data from a variety of vision and industrial systems, and drive control actions based on the result of video analytics. Figure 6 shows the components at the edge working together:

For other examples of AI platform leadership, visit our Windows AI blog.

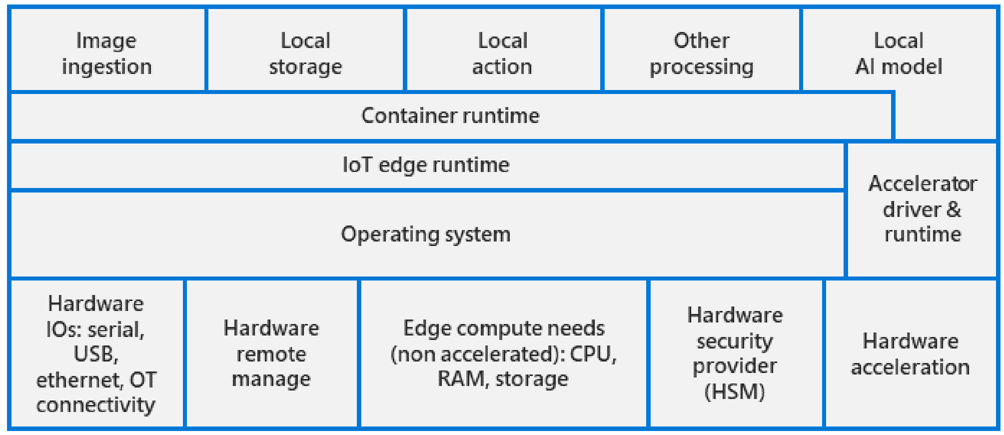

Figure 6. Sample edge components required for video analytics solutions

AI performance at the intelligent edge—hardware acceleration

To accomplish real-time analytics within the edge device, we often need hardware acceleration and the associated software acceleration frameworks to process multiple frames per second. Edge hardware varies in size, performance, and acceleration capabilities. The smallest of these edge devices may take multiple seconds to process a single image/frame (endpoint computing), whereas more performant devices, such as HPE Edgeline Converged Edge Systems, combine data-center grade CPUs and accelerators to enable tens to hundreds of frames per second.

There are four basic approaches to hardware acceleration embraced by smart cameras and edge systems (see Figure 7).

HPE Edgeline brings Enterprise computational accelerators to the edge

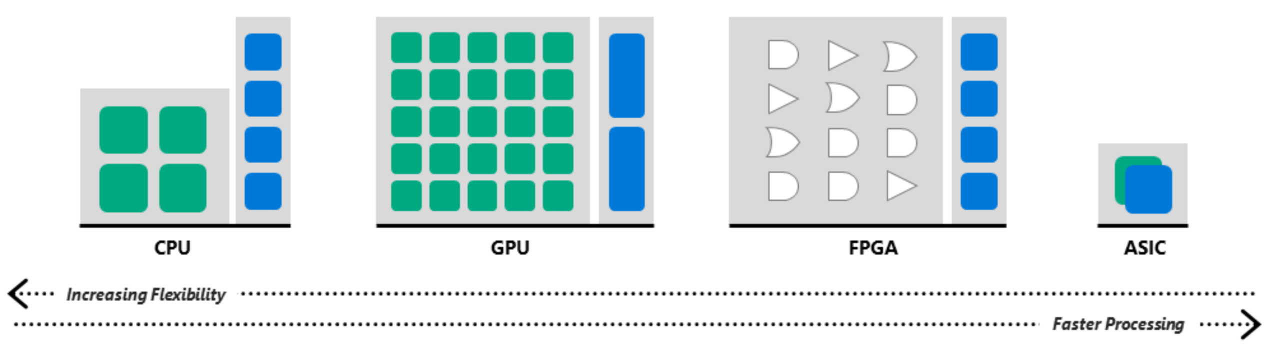

Figure 7. Hardware accelerators available for use at the edge.

Central Processing Units [CPUs]. CPUs offer the most flexibility, yet perform the slowest. Some CPUs alone without hardware acceleration may take seconds to process a single image.

Graphical Processing Units [GPUs]. Often included as part of special CPUs or as separate processors, GPUs offer flexibility and increased processing speed, but at a higher cost that's often prohibitive at the edge. Because of the flexibility of GPUs, they're often used at scale in the cloud or in edge systems for training models or for real-time inferencing of specific AI workloads and models.

Field Programable Gate Arrays [FPGAs]. Although they cost a fraction of what GPUs do, FPGAs are limited to inferencing only and can't be used to build models. FPGA's also offer much lower latencies for inferencing than with GPUs (which are throughput-efficient). FPGA's do offer some software programable flexibility, so they can be adapted to different or changing use cases at the edge.

Application Specific Integrated Circuit [ASICs]. ASICs are faster and often provide the least costly approach, but because the logic is placed in silicon, they are fixed-function and inflexible.

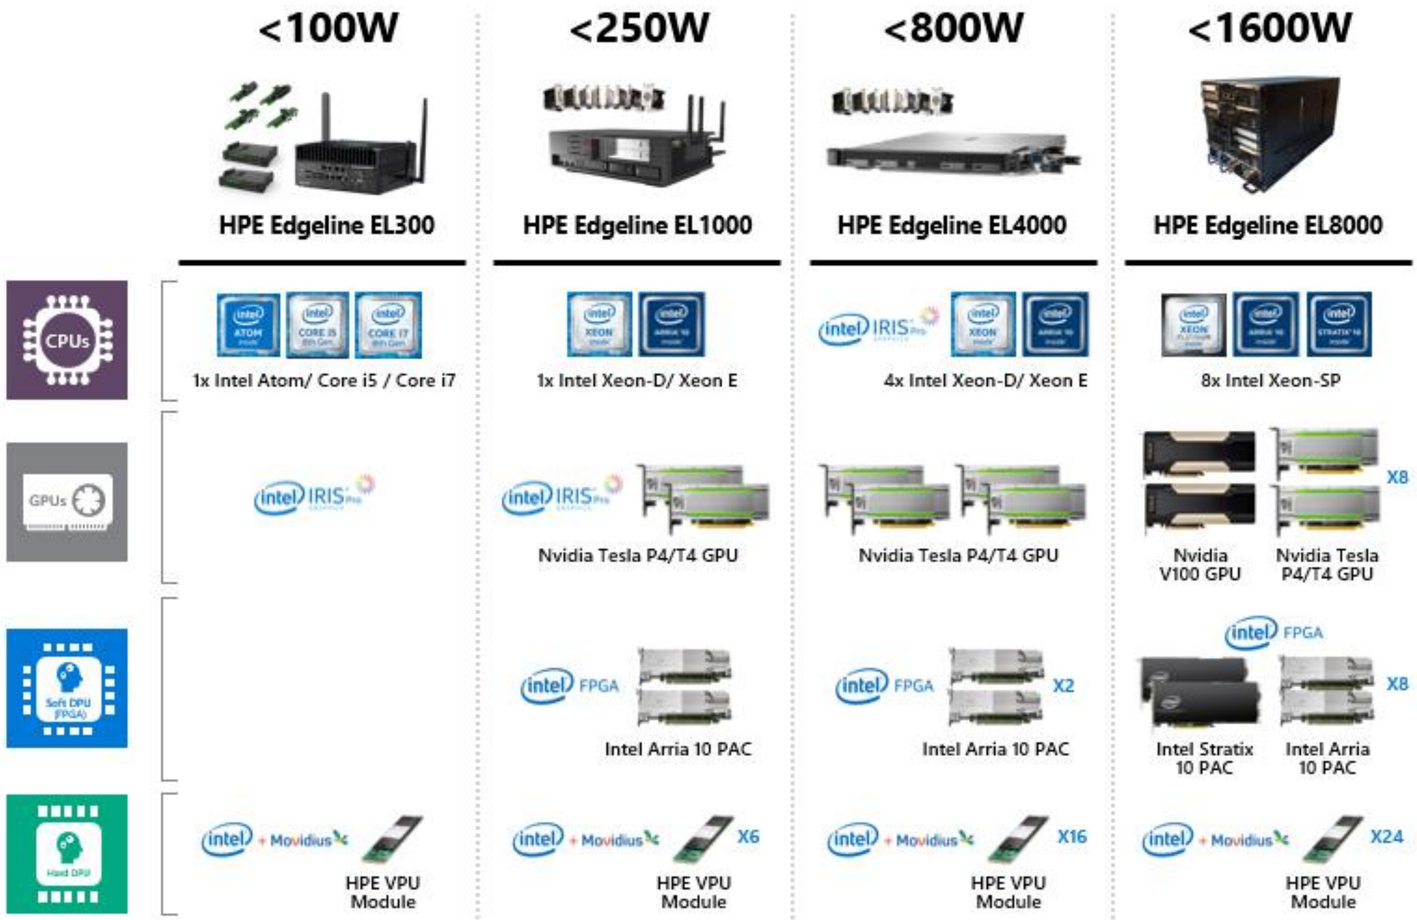

Figure 8. HPE Edgeline Converged Edge Systems provide CPU-, GPU-, FPGA- and ASIC-based AI Accelerator options

Enterprise-class converged edge systems, pioneered by HPE, support all classes of these accelerators as shown in Figure 8. When powered by Azure IoT Edge and AI by Microsoft, they can deliver video analytics at the intelligent edge as shown in Figure 10 below.

HPE Edgeline Converged Edge Systems and Operational Technology (OT) Link Technology

HPE Edgeline converges operational technology (OT), such as data acquisition, control systems, and industrial networks, with enterprise-class IT in a single, rugged system suited for harsh edge environments. Each compact HPE Edgeline system is outfitted with robust Intel compute, high-capacity fast storage, high-performance IT/OT wired and wireless connectivity, and industry-leading remote systems management.

The embedded data acquisition capability of HPE Edgeline and OT Link hardware and software provides a conduit to rich sources of data, such as standard IP cameras, industrial devices, robotics systems, and SCADA systems. This data can be analyzed using a library of hardware accelerators, and resultant control actions (switching, actuation, alarms etc.) can be directly driven out of the Edgeline system to influence the business outcome. Furthermore, the HPE OT Link platform has native connectivity to Azure IoT Hub for extension of these intelligent edge capabilities into the cloud.

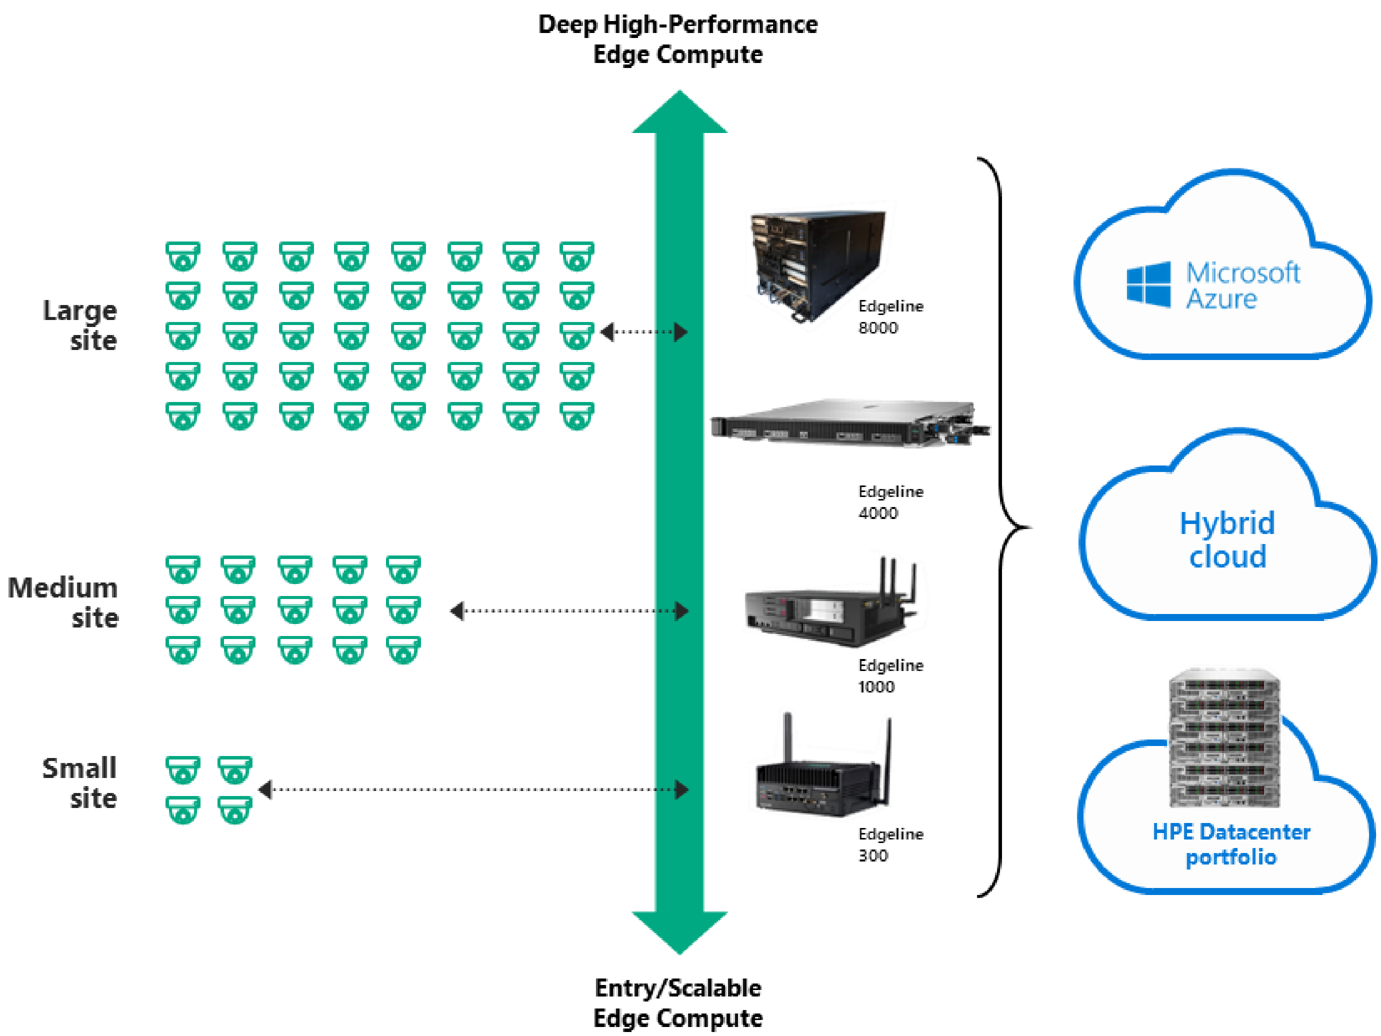

Figure 9. HPE Edgeline Converged Edge Systems offer right-sized data acquisition, analysis, and control for a variety of edge use cases.

HPE Edgeline currently supports a wide spectrum of computational accelerators across the portfolio (Figure 8), starting with high-performance Intel Xeon SP (Cascade Lake with specialized Intel DLBoost), Xeon E (with integrated GPUs), Xeon D, and Core i5 CPUs. This enables Edgeline to offer right-sized support for various use cases. Accelerated Edge AI inferencing is enabled by NVIDIA Tesla P4 and T4 (equipped with specialized Tensor Cores) discrete GPUs. The Edgeline EL8000 also supports top-tier GPUs such as the NVIDIA Tesla V100.

In addition, Edgeline enables FPGAs such as Intel Arria 10 and Intel Stratix 10 as discrete PCIe accelerator cards.

HPE Edgeline Converged Edge Systems also support a multitude of Intel Movidius Myriad X Vision Processing Units (VPUs) (from 4 to 64 chips) as product options. HPE and Microsoft are jointly looking at enabling all future AI inference silicon options from other vendors from the edge to the cloud.

HPE Edgeline Converged Edge Systems are built to provide enterprise-class embedded and on-premises consolidated management for systems at the edge for faster time to deployment, higher system and application availability, and significantly quicker time to resolution when issues arise. Embedded management begins with the core capabilities of Integrated Lights Out (iLO) built into the HPE Edgeline EL1000, EL4000, and EL8000 systems. It's extended with edge-specific capabilities enabled by the HPE Edgeline Integrated System Manager (iSM) embedded in the EL300 system. The HPE Edgeline Infrastructure Manager (EIM) provides the ability to manage multiple HPE Edgeline systems, with key capabilities including device discovery, health monitoring, firmware updates, and a summary dashboard view of all managed systems. HPE brings its expertise in datacenter management out to the edge.

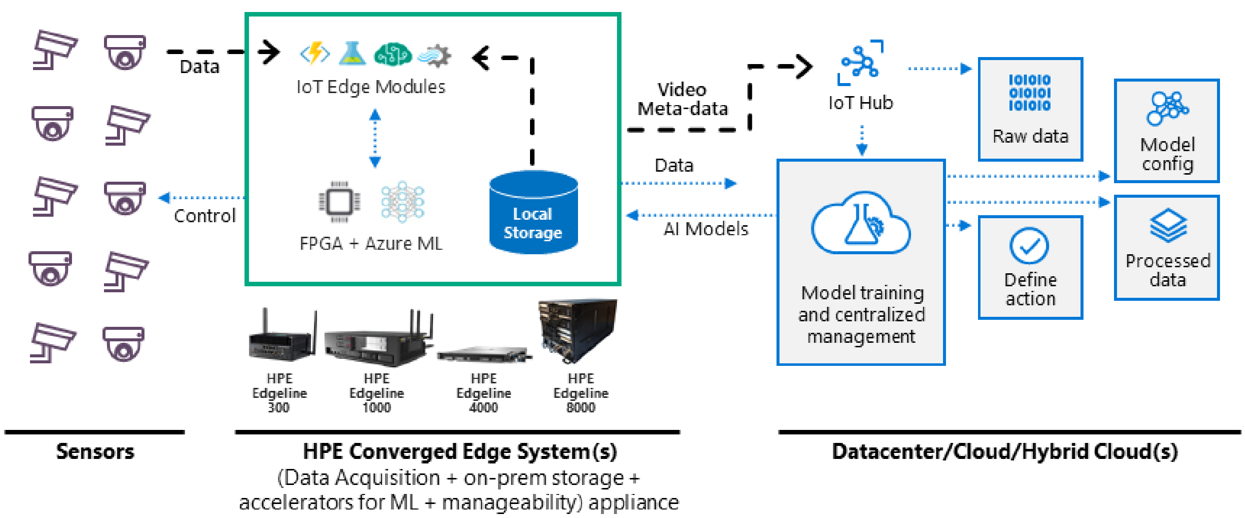

Figure 10. Microsoft Azure IoT Edge with Azure ML on HPE Edgeline Converged Edge Systems provide scalable edge to data-center/cloud/hybrid cloud solutions for next-generation IoT video analytics & management solutions

Data from connected cameras/sensors (typically H.264 compressed) is decoded and processed via IoT Edge Modules in real-time, raw frames can be passed on to an accelerator enabled Azure ML stack on the Converged Edge System to perform inferencing. The resulting metadata (object identifier, accuracy estimate, co-ordinates of bounding boxes) generated in real-time is passed on to the IoT Edge Hub (Figure 10) which in turn can pass it to a centralized Azure IoT Hub. Raw video (or other IoT sensor data) can also be retained on local storage on the HPE converged Edge system (up to 0.5 PB) and if necessary (based on inference results) be transmitted over time to Azure Block storage for further analysis. Azure ML running in the cloud/datacenter (as Azure stack) or in a hybrid cloud environment can re-create new AI models for inferencing which can be delivered to distributed Converged Edge Systems for more accurate/efficient inferencing.

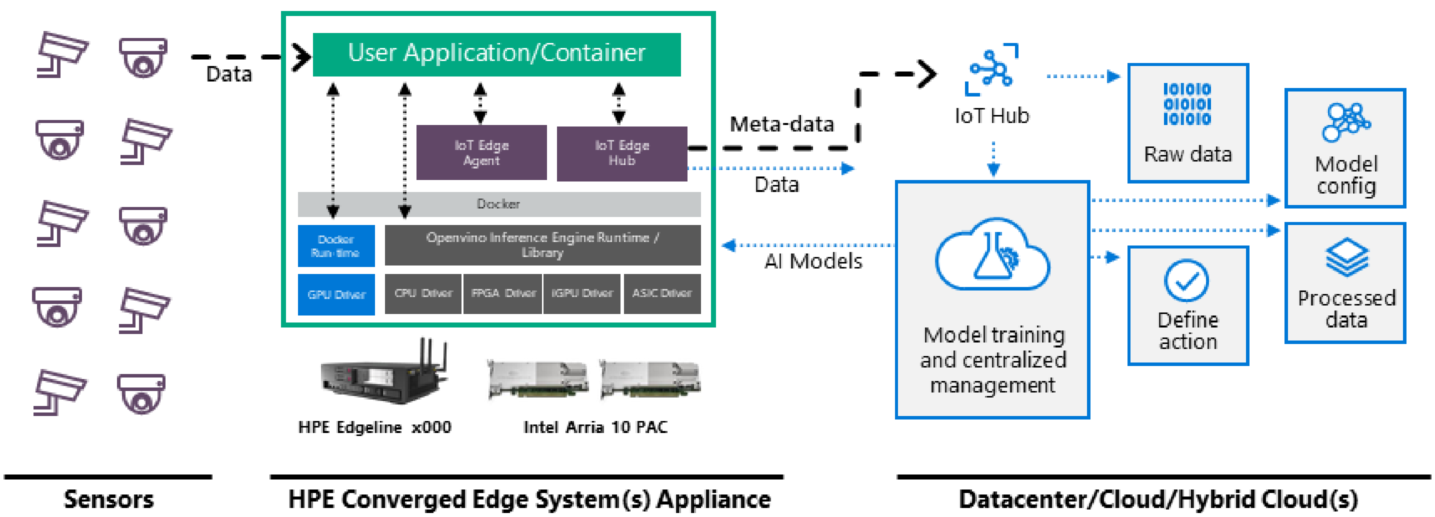

Figure 11. Software stack to enable hardware-accelerated GPU/FPGA/ASIC driven acceleration with Microsoft Azure IoT Edge and Azure ML on HPE Edgeline Converged Edge Systems for next-generation video analytics & management solutions

Figure 11 shows a software centric view of the same capability in an accelerator-centric view. Each class of accelerator requires its own device driver (under OS support) and may utilize optimized overlay libraries (or run-times) to execute Deep Neural Network inferencing (DNN libraries). These libraries and run-times are generally provided by the corresponding silicon vendors (Figure 11). Intel OpenVino (Open Visual Inferencing and Neural Network optimization toolkit) stack is unique in that it incorporates the necessary components for x86 CPUs, integrated GPUs (iGPU), FPGA's and ASICs. As described earlier, FPGA's require a custom implementation of hardware (at gate-level) to execute a neural network; several commonly used neural networks in the form of Systolic arrays are included as programmable bit-streams in the OpenVino toolkit (Figure 12).

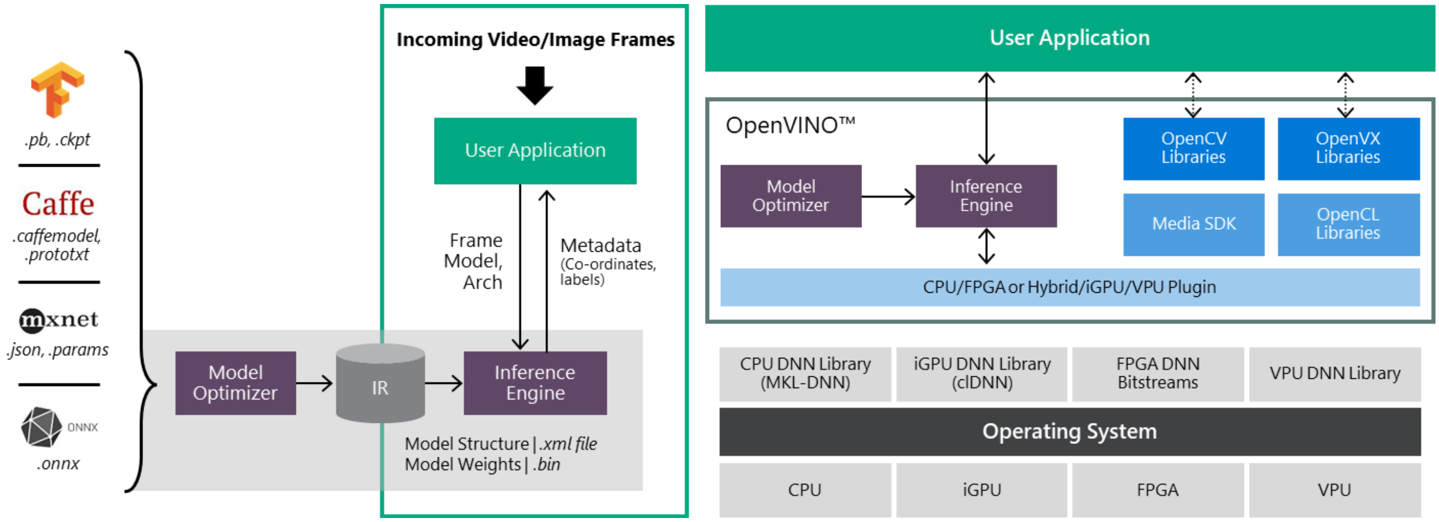

Figure 12. Intel OpenVino Toolkit for CPU-, iGPU-, FPGA-, VPU (ASIC)-based video analytics at the edge

For other instruction programmable accelerators, such as CPU, iGPU and VPU, an optimized library is included. A user-application performing AI-inferencing needs to link to an Inference Engine (provided as a library) and operate on intermediate representation (IR) versions of an AI model. An IR version of an AI model is necessary to maintain compatibility between model-representation, due to differences in various popular open-source frameworks (for performing AI training) such as Caffe, Tensorflow, MXNet, CNTK and ONNX (Open Neural Network Interchange Format). This step is performed by a Model Optimizer code. For all practical purposes, using this framework (irrespective of the kind of AI accelerator used), raw images/frames incident upon a user-application can be operated upon in real-time to generate meta-data (co-ordinates of bounding boxes, labels, estimate of accuracy) which can then be passed on from the edge to the Azure IoT Hub.

Figure 13 shows about 450 images/sec at a median latency of less than 10 ms using Intel Arria 10-based FPGAs using Resnet-10 equivalent model. This performs image inferencing on video frames from an assembly line to detect whether components on a PCB have been installed correctly or are in error. It must easy to recognize that dense PCB's may have hundreds of discrete components requiring very low-latency and fast turn-around of a component-by-component visual inspection. This is an application where an FPGA's low-power, low-latency AI inferencing capability is a great choice. The tables in Figures 14-16 compare this to the model running on a Xeon D 16-core CPU.

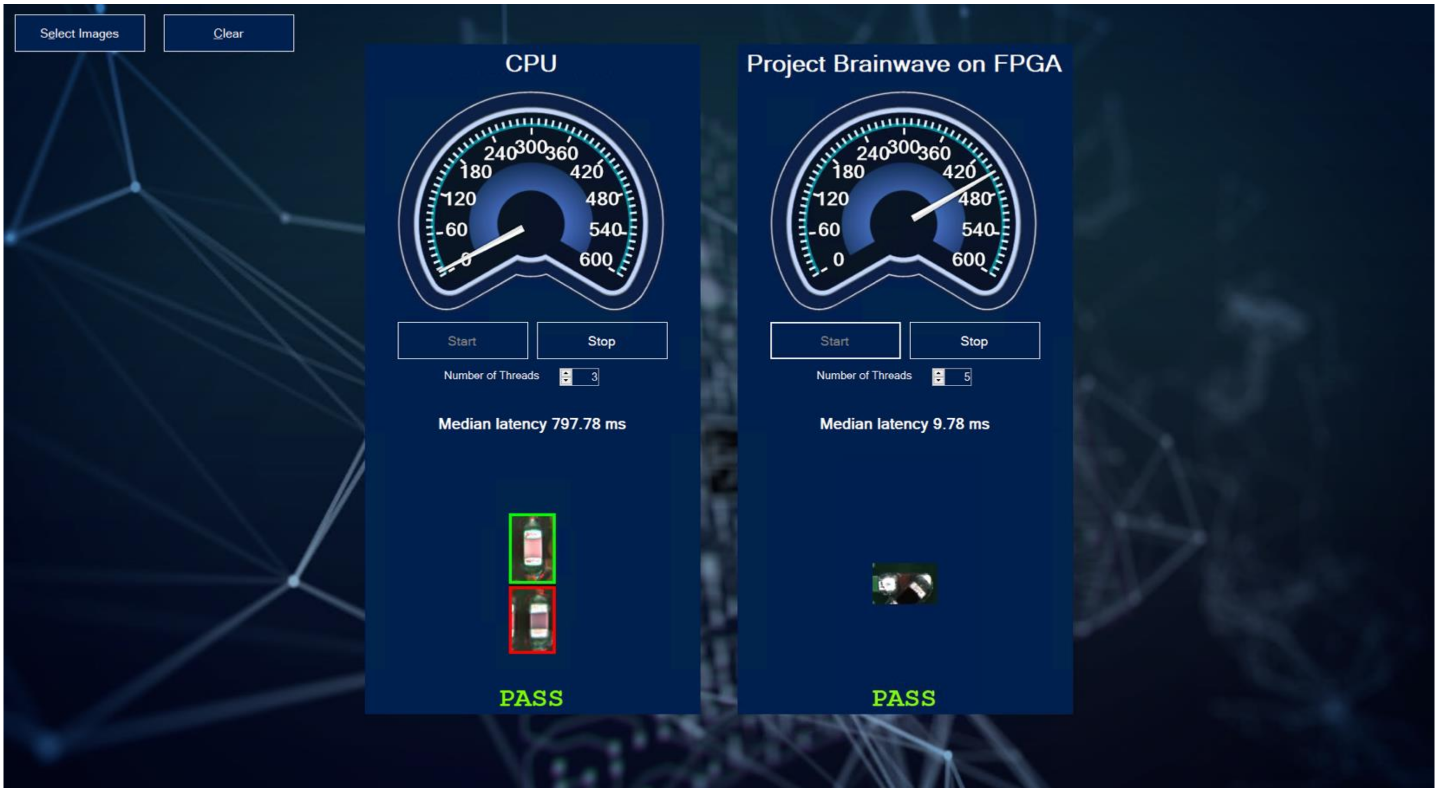

Figure 13. Lower Latency with FPGA-based image inferencing for a PCB components inspection (QA) use-case as compared to CPU's.

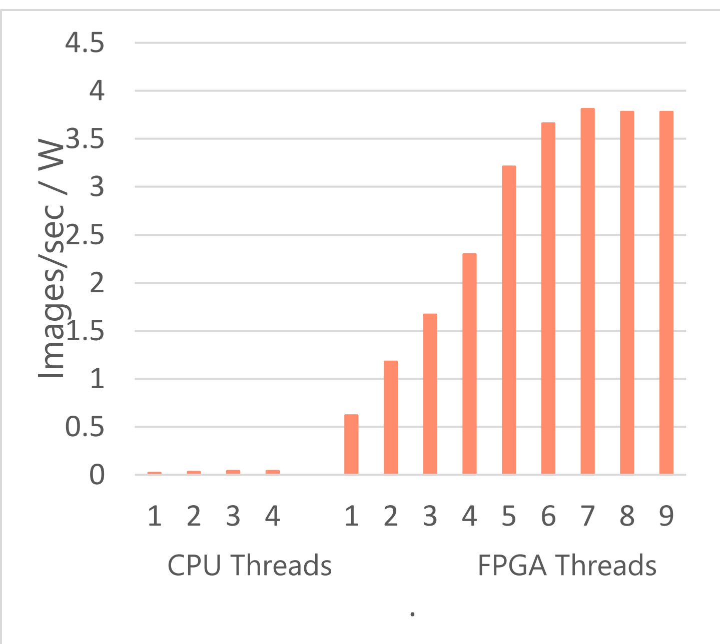

Figure 14. Higher inference rate per watt with FPGA-based image inferencing compared to CPU's.

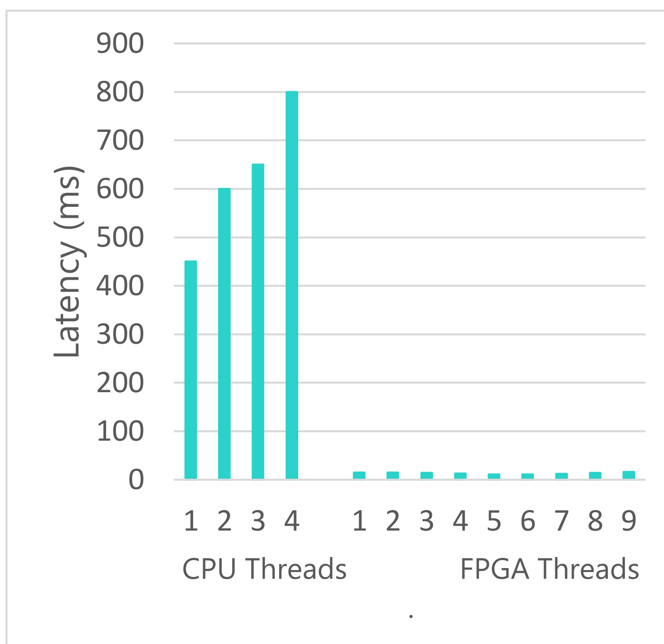

Figure 15. Lower latency with FPGA-based image inferencing compared to CPU's.

Threads	CPU Model			FPGA Model		
	latency (ms)	images per sec.	power (W)	latency (ms)	images per sec.	power (W)
1	450	<10	81	15	60	97
2	600	<10	89	15	120	102
3	650	<10	104	14	190	114
4*	800	<10	112	13	280	122
5	-	-	-	11	430	134
6	-	-	-	11	520	142
7	-	-	-	12	560	147
8	-	-	-	14	560	148
9	-	-	-	16	560	148

Figure 16. Data comparison between CPU-based and FPGA-based model

Edge to cloud

While ingesting data and performing AI at the edge are the primary needs of a video analytics solution, all the edge device, cameras/sensors, AI models, configuration, and data must be managed and kept working together as a complete solution. To address this need, Microsoft introduced Azure IoT Edge, which when combined with Azure IoT Hub and machine learning, delivers a complete solution.

Azure IoT Edge runs on Windows and Linux. It's responsible for managing all the AI components, routing data to and from edge modules, and sending data to and from the cloud as well as updating new AI models. This orchestration, at scale, enables large manageable video analytics solutions, enabling partners and customers to focus on the AI model development and actions instead of the platform. Delivered as a Platform as a Service (PaaS), it enables secure rapid deployment, data, and configuration management.

In the cloud

Driving all this capability and retaining the module are the responsivity of the cloud, either a public cloud like Azure, a hybrid cloud like Azure Stack, or a private cloud. Oftentimes, there are many clouds involved, for different needs such as data retention, data sovereignty, or business relationships and data ownership.

Azure IoT and machine learning work together for model authoring, building, and deploying (Figure 17). Whether leveraging an existing model from the machine learning gallery or building your own model using common frameworks like Tensorflow on multiple GPUs, Azure enables a complete machine learning model management delivered from the cloud.

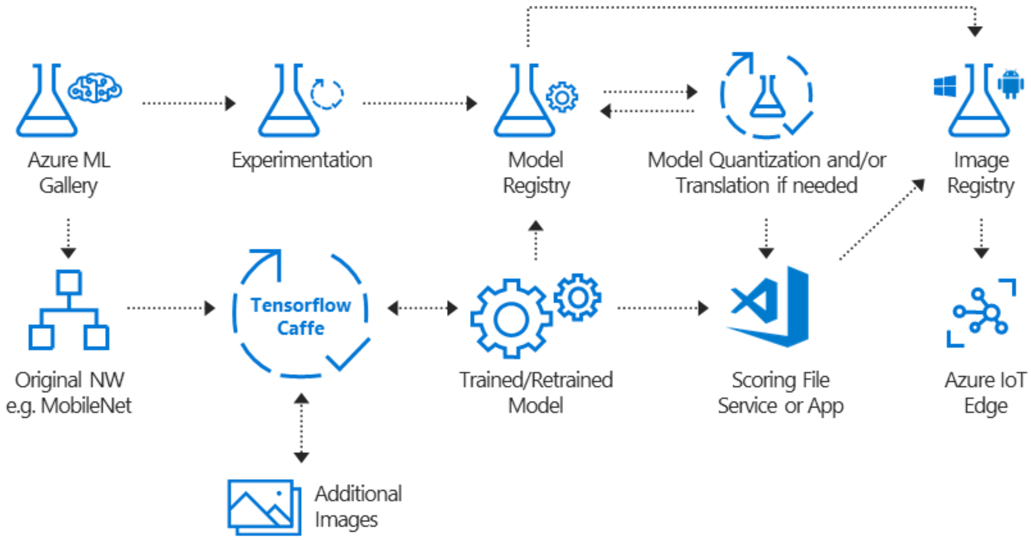

Figure 17. Azure enables AI on the edge with easy model conversion.

One of the main drivers for public clouds like Azure is scale and ability. Machine learning's a great example, where you can automatically build models from large datasets on multiple GPUs paying only for the GPU time used when building and storing the finalized model in a registry to be pushed down to edge systems.

For toolsets, Microsoft provides authoring tools like Visual Studio, Visual Studio Code, Azure Machine Learning Studio, and Azure Notebooks. For developers or data scientists who have a preferred integrated development environment [IDE], such as Jupyter, Microsoft supports open-source frameworks such as PyTorch, Tensorflow, MXNet, and scikit-lern.

For compute, Microsoft provides both Data Science Virtual Machines and PaaS, such as Azure Machine Learning Service, Databricks, and Azure Machine Learning Studio. For developers with their own proprietary framework and models, Microsoft also provides GPU-based virtual machines and Azure Batch AI for out-batch processing of large data sets.

Conclusion

Video analytics solves real-world challenges, both traditional and new opportunities. Applying AI to make decisions on what is seen isn't new, but with advancements in hardware acceleration and deep neural networks, we can now enable computers to do boring repetitive tasks, freeing people to do more productive work and not just stare at video screens.

With the advancements in hardware accelerators, neural networks, enterprise-class edge hardware — all powered by the cloud, video analytics is a core use case for the intelligent edge. Microsoft and HPE with the extensive partner network, can help deliver such video analytics and more.

Whether you're a skilled data scientist who just needs compute power and robust frameworks to get the job done, an operations professional who's tasked with building out the intelligent edge at each location or a business partner who needs help getting started, HPE's Edgeline Systems and OT Link technology connected to Azure, coupled with HPE Field Application Engineering (FAE) services is a great starting point to assess, prove, and rollout a complete edge AI video analytics solution. HPE Pointnext also provides a variety of advisory and professional services to help customers on their IoT transformation journey.

To experience this joint HPE and Microsoft solution in action, connect with us at HPE IoT Innovation Labs at iot-innovation-labs@hpe.com

Additional resources:

HPE Edgeline Converged Edge Systems	https://www.hpe.com/us/en/servers/edgeline-systems.html
Azure IoT Edge on Nvidia	https://azure.microsoft.com/en-us/blog/microsoft-and-nvidia-extend-video-analytics-to-the-intelligent-edge/
Azure IoT Edge on Qualcomm	https://azure.microsoft.com/en-us/blog/accelerating-ai-on-the-intelligent-edge-microsoft-and-qualcomm-create-vision-ai-developer-kit/
Artificial Intelligence on Intel FPGA	https://software.intel.com/en-us/blogs/microsoft-turbocharged-ai-with-fpgas,
Intel OpenVINO on Windows	https://software.intel.com/en-us/articles/OpenVINO-Install-Windows#inpage-nav-4-2
Intel Movidius on Windows ML	https://newsroom.intel.com/news/intel-microsoft-enable-ai-inference-edge-intel-movidius-vision-processing-units-windows-ml/#gs.2uambo
Ethical Artificial Intelligence	https://www.microsoft.com/en-us/ai/ai-for-good
Azure Machine Learning	https://docs.microsoft.com/en-us/azure/machine-learning/
Building AI apps for Windows Developers	https://blogs.windows.com/buildingapps/2018/03/07/ai-platform-windows-developers/
HPE Pointnext Services	https://www.hpe.com/Pointnext