

City-Scale Edge Inference, Built in 3 Weeks

Video surveillance is growing at a [breakneck pace](#), creating urgent demand for automated analytics. Cities and other large users are deploying systems with hundreds or even thousands of cameras—resulting in far too many video streams for humans to monitor.

To help cities make better use of this data, companies like Agent Video Intelligence ([Agent Vi](#)) already offer cloud-based AI that can analyze large numbers of video streams. But as the number of cameras grow—and resolutions and frame rates increase—cloud-based systems are running into a variety of limitations, including:

- Bandwidth costs associated with transmitting HD video streams to the cloud

- Latency of responses from cloud-based systems
- Security vulnerabilities from passing through various networks

To address these issues, Agent Vi partnered with Intel® to deploy deep learning at the edge (**Figure 1**). By performing functions like inferencing at the edge, video surveillance scenes with no activity can be discarded or tagged as lower priority while footage of interest can be flagged for immediate transmission to more powerful cloud-based analytics systems.

Because there is less data, transmission costs are also lowered and the possibility of loss or theft is reduced.

System Architecture

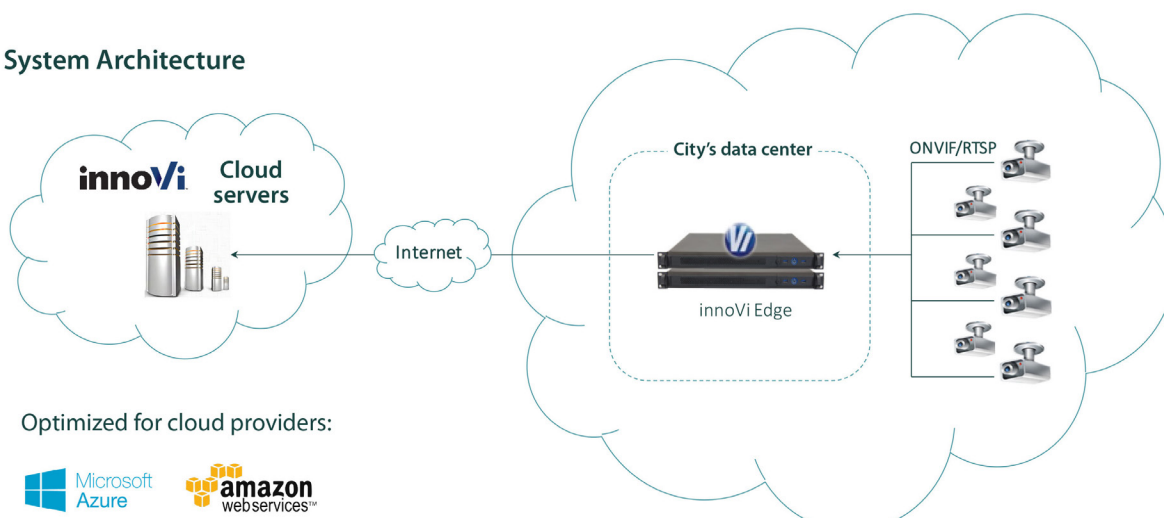


Figure 1. The Intel® Distribution of the OpenVINO™ toolkit enabled Agent Vi to scale its innovi deep learning video analytics capabilities to any installed IP or analog cameras. (Source: [Agent Vi](#))

Extending Deep Learning to Legacy Systems

But smart city surveillance systems have long lifecycle requirements and are deployed in a variety of environments. This meant that Agent Vi had to enable a wide variety of new and legacy equipment with deep learning technology, including cameras, encoders, and video management systems (VMS).

For example, the company wanted to upgrade its [innoVi Edge](#) appliance, which performs visual decoding for fixed IP and analog cameras. And the company wanted its software to work on a variety of third-party systems as well.

To make this vision a reality, Agent Vi's video analytics algorithms had to be tuned for a range of CPUs, GPUs, and other hardware accelerators. To enable this on a range of new and legacy smart city surveillance cameras, Agent Vi leveraged the [Intel® Distribution of the OpenVINO™ Toolkit](#).

Openvino Toolkit Accelerates Agent Vi Edge Surveillance

The OpenVINO™ Toolkit is an architecture-agnostic development suite that helps engineers deploy deep

learning across CPUs, GPUs, FPGAs, and vision processing units (VPUs). Plus, the toolkit integrates a suite of graphics and image libraries that optimize inferencing performance (**Figure 2**).

Pairing the OpenVINO Toolkit with innoVi delivered a number of benefits for Agent Vi and its customers. For one, the innoVi Edge appliance was suddenly deep learning-enabled, bringing distributed video analytics inferencing closer to the edge. This reduces network transmission costs and latency versus sending multiple, full-HD smart city video surveillance streams directly to the cloud.

In terms of performance, OpenVINO not only improved the accuracy of visual analysis on edge appliances, it also enhanced developer productivity. According to Zvika Ashani, CTO and Co-Founder of the company, the toolkit improved both development speed and capacity.

"With the OpenVINO toolkit, the results have been impressive, enabling us to move from supporting three cameras to 14 with one developer, in under three weeks," he said. "We will be able to fully scale our solutions to the edge with the right performance per dollar while leveraging Intel® Movidius™ Myriad™ VPU and Intel® FPGA solutions."

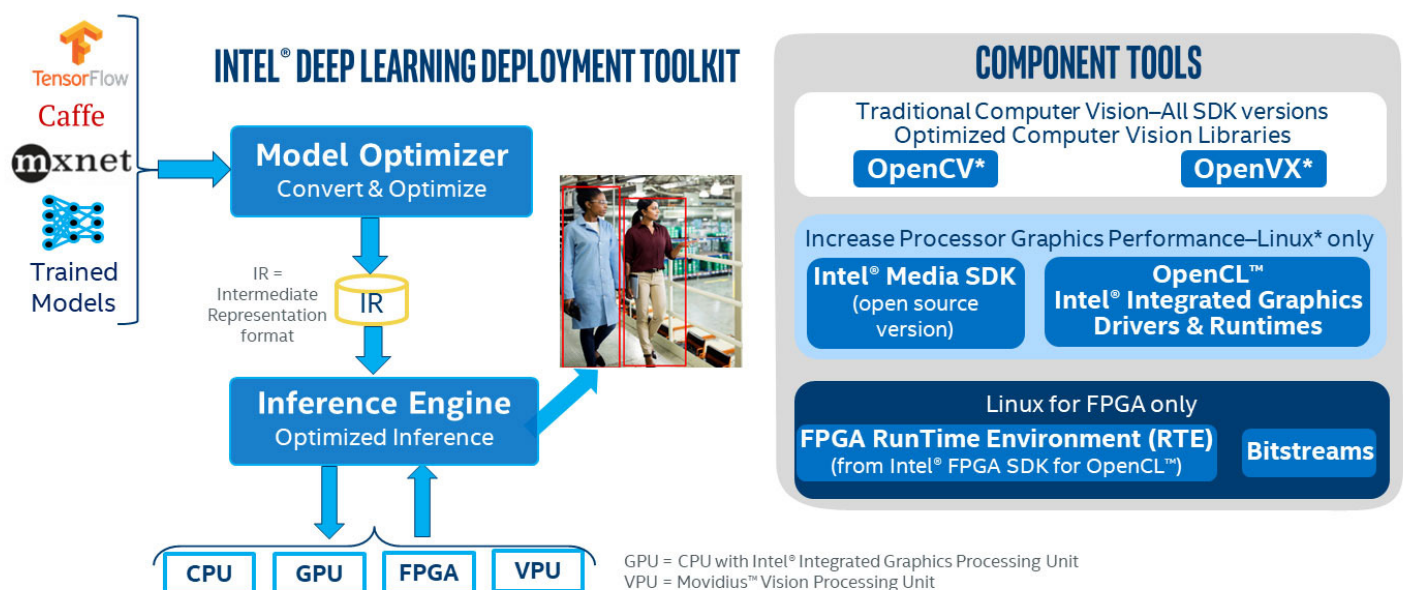


Figure 2. The Intel® Distribution of the OpenVINO™ toolkit supports a variety of development frameworks, vision/graphics libraries, and hardware architectures so that developers can deploy optimized algorithms almost anywhere. (Source: [Intel® Corp.](#))

Finally, the flexibility afforded by OpenVINO ensures that innoVi software can operate with a broader set of surveillance equipment, now and in the future. Not only are Agent Vi's algorithms compatible with a broader set of compute architectures, they're also future-proofed against infrastructure upgrades that have historically required re-engineered software.

Not only did this ease integration with existing on-premises environments and compute platforms, it also future-proofed the solution against infrastructure upgrades that often require new or modified software.

In short, innoVi was primed for massive scalability in smart city surveillance applications like public safety, critical infrastructure monitoring, and traffic management.

Deeper Learning for Smarter Cities

With the innoVi platform now extended to the edge, smart city engineers can apply the technology in ways that transform urban life.

For instance, traffic cameras can use deep learning to track the number and type of vehicles on a road, then recommend that intelligent transportation systems (ITS) reroute commercial trucks. Surveillance systems monitoring an area can alert officials of unexpected human activity in real time, allowing authorities to respond quickly and eliminating the need for dedicated human security guards.

Without deep learning at the edge, the data from these systems would likely be used for historical analysis or retroactive scanning rather than real-time intervention. Costs would also increase.

With solutions like innoVi and the Intel Distribution of the OpenVINO toolkit, security system operators, transit companies, and other smart city stakeholders can now benefit from true intelligence at the edge.