

IronYun Vaidio® AI Vision Platform 8.0 for MSP's

Vaidio Kubernetes Architecture and Analytics Virtualization
for Managed Service Providers – an IronYun Technical Whitepaper



Summary

Managed service providers are constantly looking for proven solutions that will improve operational efficiency and deliver greater value to their end users. Vaidio 8.0 is purpose built for Service Providers with a Kubernetes and Microservices-based architecture that virtualizes analytics functions and enables scalable, easy to manage, compute efficient, public-, private-, or hybrid-cloud- and subscription-based video analytics-as-a-service.

Key Features

Vaidio 8.0 has four key features that integrate and work together to enable service providers, three of which are new and all four of which are unique to the market: Kubernetes integration, floating analytics licenses, analytics virtualization, and the Vaidio Service Portal.

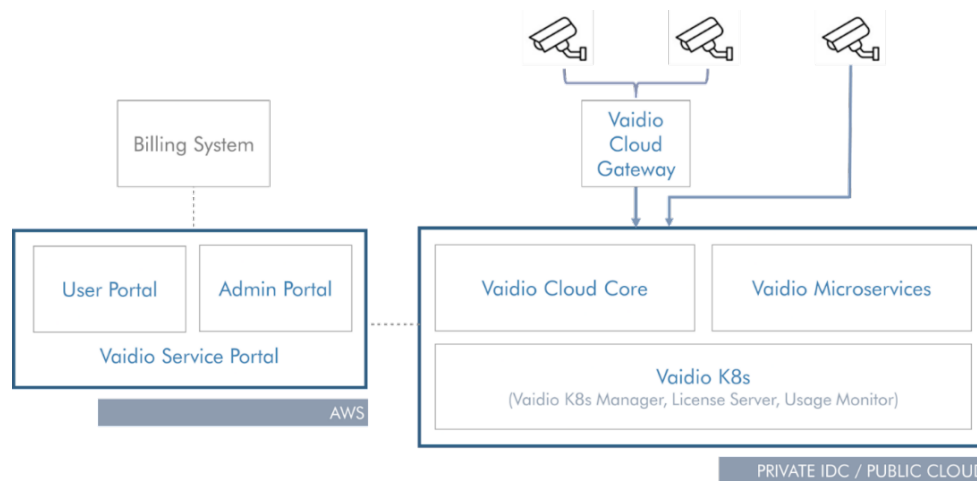


Figure 1: Vaidio 8.0 High Level Kubernetes Architecture

Kubernetes Integration

Vaidio 8.0 offers a new architecture based on integrating our Linux core with Kubernetes. Kubernetes (K8s) is a container orchestration system. It runs virtual or physical machines in clusters which in turn run workloads in containers. This makes managing and scaling workloads in clouds and data centers much easier.

Vaidio 8.0 offers a Kubernetes-based cloud architecture that can be deployed on-prem or on private, public or hybrid clouds. It supports on-demand, scalable auto-provisioning as part of a managed service. We go into detail below, but in short, this feature coordinates 6 components: our Linux core with deep Kubernetes integration, unique use of containers, an optimized node structure, analytics microservices, pre-processing and automated workload monitoring.

This new architecture maximizes compute efficiency, availability, and dynamic scalability. It serves as an optimized foundation to enable service providers to deliver advanced video analytics as a service – while minimizing resource costs, operating costs, and scaling costs.

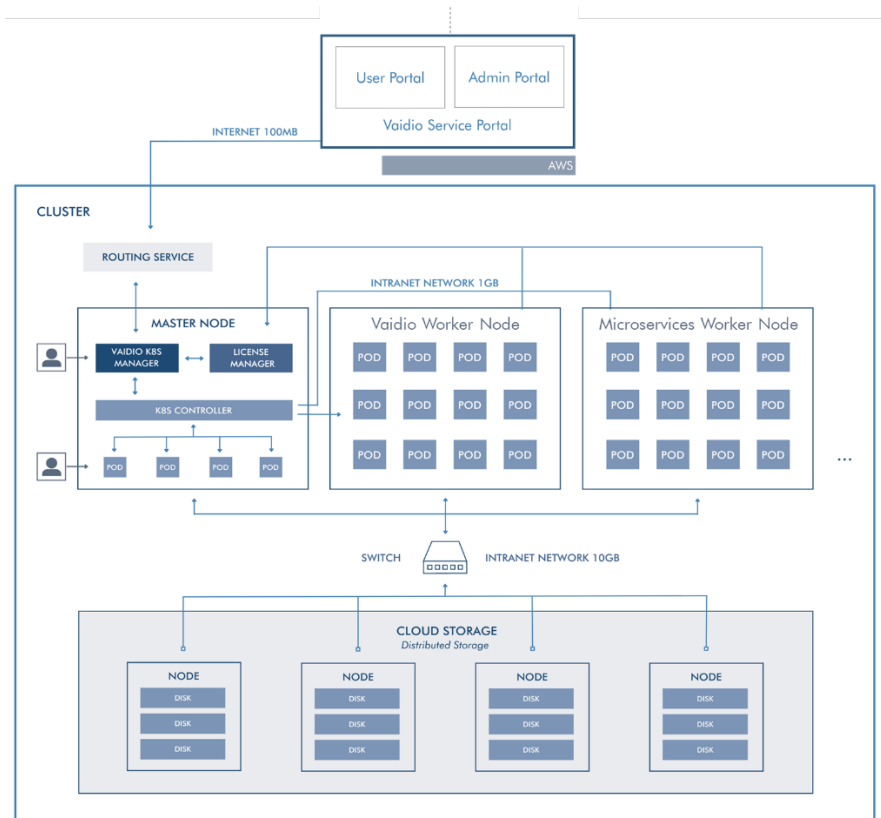


Figure 2: Vaidio 8.0 Kubernetes Architecture Detail

Technical detail on Kubernetes integration:

Vaidio 8.0 Kubernetes clusters contain three nodes: a master node for K8s, workload, and license management; a worker node that manages containers for Vaidio core instances, metadata, and camera connections; and a microservices node that orchestrates containers for analytics microservices. Vaidio 8.0 also puts Core and Analytics Microservices into independent containers.

Vaidio K8s integration also incorporates low-compute preprocess algorithms to filter irrelevant frames, sending only relevant frames to microservices for processing; as well as automated workload monitoring, which spins microservices down when they are not in use.

Vaidio orchestrates preprocessing and workload management, analytics virtualization, and architecturally independent containers to maximize compute efficiency, availability, and dynamic scalability – thereby minimizing ongoing compute and management costs for both end-user and service provider.

Feature 2: Floating Analytics Licenses

Vaidio has long featured the ability to ‘float’ analytics licenses across cameras. We highlight this feature here because it is a critical component of the new 8.0 architecture. Vaidio analytics licenses are not locked to a specific camera. They are re-deployable on schedule or on demand, and, they allow for multiple analytics per camera. So, for example, a facility could use the same camera to do LPR and vehicle counting during the day, and then switch out to do intrusion detection at night. This ability to dynamically allocate licenses from a license pool is a differentiating technology and a core enabling technology in a dynamic Kubernetes-based environment.

Feature 3: Analytics Virtualization

In the server world, hypervisors enable server virtualization. Previously, ten applications had to run on ten physical machines. Hypervisors add a virtualization layer that enables the ten applications to share the resources of a single physical machine – consolidating ten machines down to one.

Likewise: prior Vaidio releases had to reserve a fixed amount of physical resources for each camera. With Vaidio 8.0, microservices virtualize the analytics and enable the cameras to share the physical GPU resources.

We estimate that Vaidio analytics virtualization will enable a 3 to 12X increase in the number of cameras that one physical GPU can support.

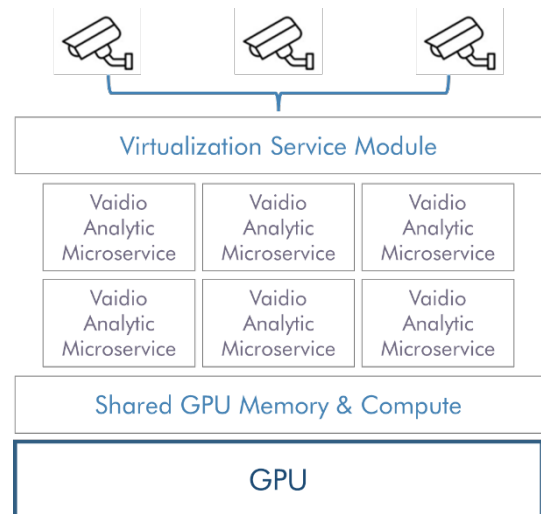


Figure 3: Vaidio 8.0 Analytics Virtualization

Vaidio analytics virtualization represents a groundbreaking advance in our space, a huge increase in efficiency and scalability, and the potential for enormous resource and hardware cost savings.

Feature 4: The Vaidio Service Portal

The Vaidio Service Portal is a new component of Vaidio 8.0 designed specifically for service providers to deliver and manage Vaidio V1aaS (vision intelligence as a service) at scale. The Vaidio Service Portal is a cloud-based portal with User and Service Admin components. It can be rebranded to the Service Provider’s brand. It can support thousands of end-users with secure multitenancy. It can also support a two-tier service provider model. Both User and Admin portals offer new-user registration, configuration, usage and subscription monitoring, integrated billing and more. The Admin portal also offers instance management, license management, aggregated user management and physical and cloud-based workload monitoring. The Vaidio Service Portal is a fundamental component that enables service providers to easily and efficiently deliver Vaidio-as-a-Service at scale.

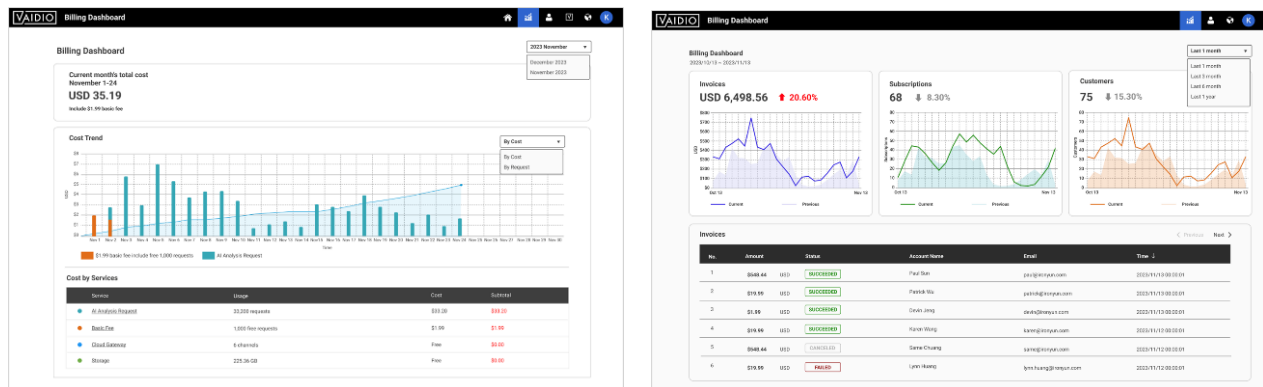


Figure 4: Vaidio Service Portal – user and admin billing dashboards

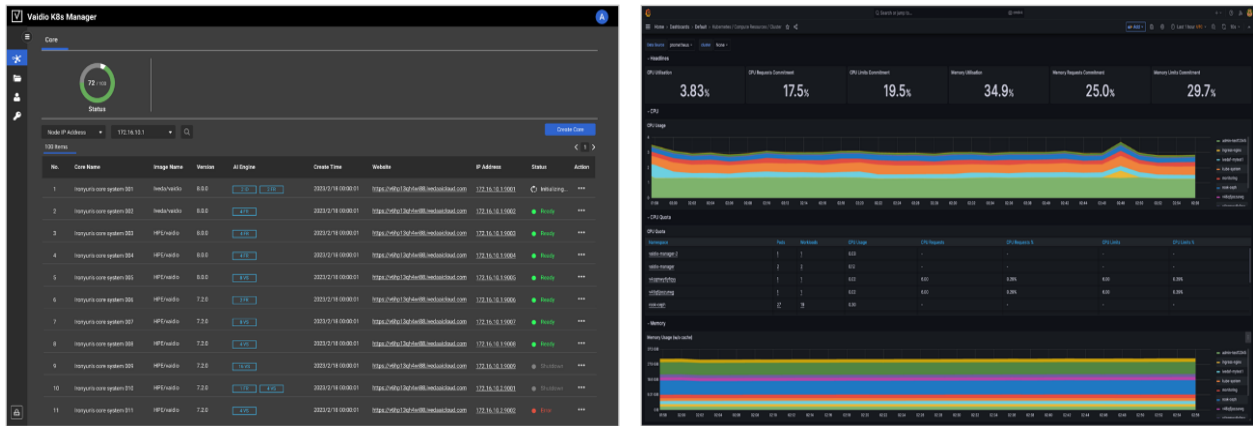


Figure 5: Vaidio Service Portal – instance and workload management

Conclusion

Vaidio 8.0 is designed to make it both easy and cost effective to offer a comprehensive set of the world’s most advanced video analytics applications for safety, security, operations and beyond as subscription-based software-as-a-service. Every component of the 8.0 release, from Kubernetes integration, to analytics virtualization, to the Vaidio Service Portal, is designed to maximize ease of use and ease of service delivery, to minimize resource and operating costs, and to enable a scalable, on-demand, and auto-provisioning framework for managed service providers.

The Vaidio Service Portal, Kubernetes integration, and analytics virtualization are parts of a massive expansion of the capabilities of the Vaidio AI Vision Platform with the 8.0 release. Vaidio 8.0 also includes the industry’s first concrete applications for generative AI, complementing our deep learning models with natural language video search, complex object and behavior recognition, massive object vocabulary expansion, and accelerated custom model development.

These Vaidio GenAI functions enhance platform security and safety capabilities, and open up huge potential for new applications that can drive significant operational and business value. In sum, the new architecture and capabilities built into Vaidio 8.0 are designed to drive significant cost savings and new value for end-users and managed service providers alike.

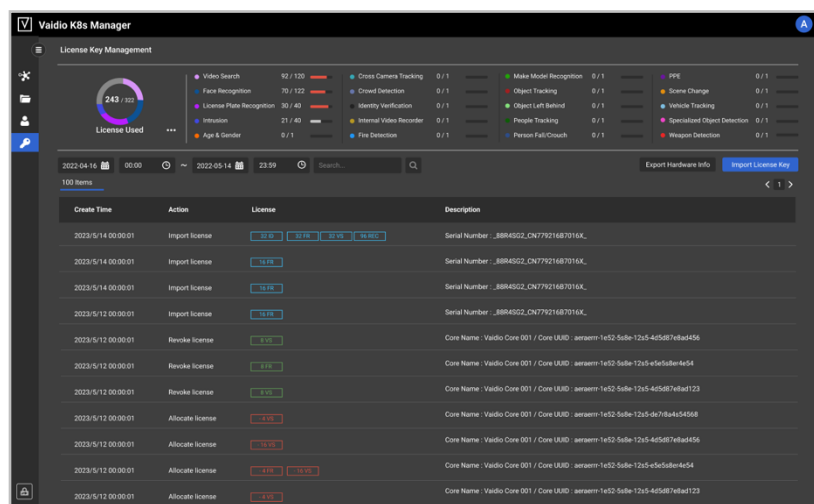


Figure 5: Vaidio 8.0: Vaidio Service Portal – license key management across the Vaidio analytics floating license pool