Hewlett Packard Enterprise

DEEPSIG

Technical white paper

Check if the document is available in the language of your choice.

WIRELESS SIGNAL IDENTIFICATION AND ANALYSIS

Enabling AI and Analytics at the Edge

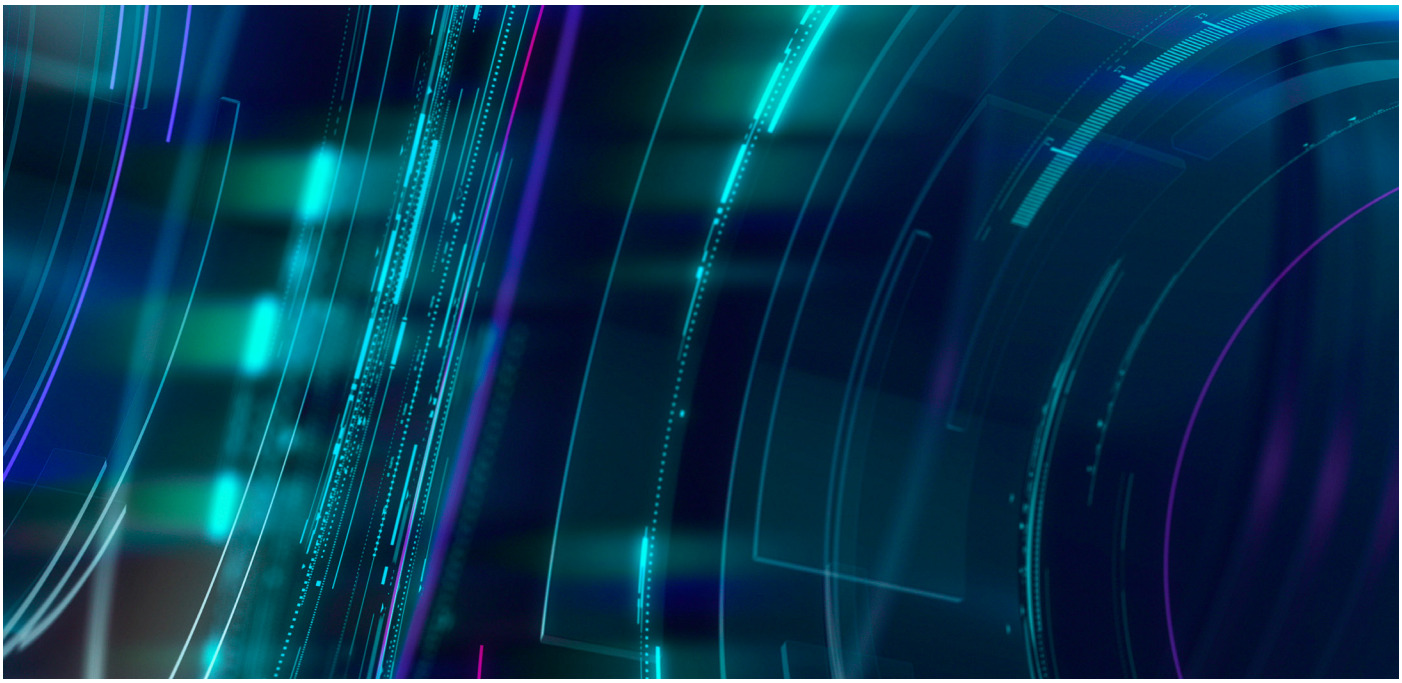

CONTENTS

EXECUTIVE SUMMARY

Leveraging Artificial Intelligence (AI) at the edge opens new frontiers in Radio Frequency (RF) signal analysis. Non-interrupted wireless connectivity will soon be mission-critical for many applications, resulting in RF signal analysis becoming essential for a host of new industries. 5G and 4G wireless, autonomous vehicles and national security are only the front line of near-term applications for AI-powered signal intelligence. Real-time electromagnetic signal identification and tracking will no doubt prove beneficial to future operations in many other industry verticals as well. Hewlett Packard Enterprise and DeepSig Inc. have partnered to provide an accelerated, deep-learning-at-the-edge signal identification solution for any field application. HPE's solution begins with the HPE Edgeline EL8000, a small toolbox-sized multi-node server platform (approx. 9" x 9" x 17"), running DeepSig's OmniSIG™ Sensor application. Together, Edgeline and OmniSIG create a next-generation platform for monitoring the wireless spectrum for a large number of signals. OmniSIG's Software Development Kit (SDK) enables straightforward, simple and customizable retraining of the AI model—even in the field. With OmniSIG's SDK, the user dictates the core sensor capabilities and capability updates, eliminating the need for a machine learning expert in the loop. HPE Edgeline EL8000/ OmniSIG solution is described, including configuration detail, applications and use cases.

INTRODUCTION

Wireless communications rely on a busy and complex portion of the microwave and radio-wave electromagnetic spectrum, with thousands of signals vying for prominence at any moment.

Companies whose business models demand up-to-the-instant spectrum usage intelligence are constantly working to improve their analysis and identification pipeline. The compute power needed to run a smart signal identification dashboard has traditionally lain in the data center. However, real-time applications often lack the compute power needed to gather data at the network edge, perform ingest functions, process/ perform data reduction and then present results to users.

The problem of RF interference results in lost capacity, upset users and potential loss of critical services. However, current RF interference detection solutions are limited to specific frequency bands, while RF interference, either non-intentional or intentional is not limited to any band. It is also typically difficult to add capability to current RF interference detection solutions once they deployed, and they often cannot scan fast enough to detect all forms of interference—especially adversarial signals that may be "frequency hopping."

The wireless spectrum is rapidly evolving as new services and eventually 5G technology is rolled out. Next-generation communication systems must be aware of many different non-contiguous frequency bands and many different signal types in order to ensure compliance and continuous operations without loss of capacity.

It takes a versatile platform to monitor and analyze the broader ranges of spectrum that constitute wireless communications in the years ahead. Responsive, near-instant signal intelligence requires smart, accelerated, AI-enabled signal analysis and identification in a portable server platform that is powerful enough to perform all the processing functions required to enable agile situational awareness and decision making.

ELECTROMAGNETIC SIGNAL ANALYSIS AND IDENTIFICATION AT A GLANCE

Commercial spectrum analysis, deconfliction and sharing applications

Wireless data streams are so commonplace today that it's easy to overlook *how* those data streams traverse to and from edge devices. Yet increasingly many companies rely on rapid-response, zero-lag versatility and always-on analytics. As 4G and 5G channels saturate and become overcrowded, companies cannot assume all the spectrum your applications need will always be there.

Many companies today stand up their own private 4G LTE networks for proprietary data communications and analysis. They, of course, want to monitor and protect their network's spectrum—ensuring it's operating at full potential, with minimal interference or security threats. In these applications, real-time signal intelligence is near essential.

Competing monitoring systems that track many different wireless signals simultaneously require sometimes tens of minutes to scan and identify all signals of interest. Such long lag times can add significant burdens on the network and those tasked with ensuring its security. These other systems also lack the capability to learn their own standard environment, resulting in generalized solutions that don't truly meet the needs of the user.

Security applications

National security applications demand an always-on and always-aware signal intelligence presence, even, and especially, when under hostile attack.

In the defense and national security space, there is a constant need for monitoring the wireless spectrum to confirm known good signals, to detect anomalous signals and rogue waveforms. Any wireless interface—from Wi-Fi access points, to a smartphone's Bluetooth® interface, or even to wireless security camera backhaul connections—is a potential attack vector for cyber adversaries.

Traditional wireless detection systems are often unable to monitor the entire width of spectrum that such high-demand and high-security applications demand. While looking for many different types of signals simultaneously, other solutions focus on common signal types such as Wi-Fi or Bluetooth. Of course, adversary or security applications are not constrained to known, pre-existing, common signal types.

Security threats in this space include:

- Cheap communications equipment that does not adhere to Federal regulations
- Jamming of commercial cellular and/or public safety wireless communications
- Wireless as a cybercriminal inroad into secured networks
- Man-in-the-middle attacks that use rogue Wi-Fi base stations
- Cellular tracking that use rogue cellular base stations.

BUILDING A SIGNAL ANALYSIS AND IDENTIFICATION SOLUTION

By using a data-driven Machine Learning (ML) approach to signal analysis and identification, OmniSIG ensures real world performance spectrum awareness. ML-driven spectrum sensing enables rapid signal awareness and identification of RF devices & activity not previously feasible. By using Artificial Intelligence (AI) techniques to train models OmniSIG can identify these signals types "out of the box": Cellular and infrastructure signals such as GSM, LTE, CDMA2000, and WCDMA; ISM-band signals such as WiFi and Bluetooth; Mobile radio services such as P25 and DMR; IoT signals such as LoRA; and Commercial aircraft RADAR signals. These signal models live in a neural network that can be customized to include other signals of interest via the OmniSIG Software Development Kit (SDK).

OmniSIG's SDK is a complementary tool to the OmniSIG sensor that provides a new class of RF sensing using DeepSig's pioneering application of Artificial Intelligence (AI) to radio systems. The SDK gives users of OmniSIG the tools to embed even more knowledge and understanding of known and unknown signals, without the need for a machine learning expert in the loop. The SDK gives end users the ability to create a custom, AI-enabled, RF detection and classification system optimized for their specific needs. Users with no experience in AI development are now able to add custom signals to DeepSig's signal detection and classification product (OmniSIG) or simply improve the performance of signals that are already supported. Using a combination of the latest web technologies and advanced server-side tools, users are able to load a handful of RF snapshots of signals of interest, use drag and drop methods to label and annotate the custom signals, and then use this new training data to evolve a neural network without writing any DSP code or having to capture thousands of snapshots

Hardware

The Edgeline (EL) 8000 is HPE's new converged edge compute system, providing data center computing in a compact form factor. HPE Edgeline EL8000 is designed for austere operations. It incorporates the latest in Intel® processors, NVIDIA® GPUs, NVMe SSD storage technology, networking options, HPE iLO security and remote management features.

The HPE Edgeline EL8000 is engineered to enable true data center performance and capability with a rugged feature set that can operate in challenging environments. HPE Edgeline 8000 also interfaces with a wide range of sensors, controllers, vehicle and system electronics, and industrial class systems.

The HPE Edgeline EL8000's small size, weight and power (SWaP) factors put it in a class of servers unto itself. The advanced applications it runs—such as computer vision, AI, machine learning, and analytics—mimic workloads found on data center-ready servers. However, its slight footprint and lean operating envelope enable enterprise computing with the HPE Edgeline EL8000 closer to the tactical and mobile edge.

HPE Edgeline 8000 is a modular 5U chassis and weighs 50 lbs. It's a half rack in width (8.7" wide), 17" deep, and 8.63" (5U) high. HPE Edgeline EL8000 consists of a controller/power/network base (bottom 1U) and space for up to four (1U) compute blades (or optional storage).

The modular chassis provides substantial computing and data storage capacity:

- Up to 4 independent server nodes with data center class Intel® Xeon® Scalable Processors
- Up to 112 cores per chassis
- Up to 28 cores per server
- Up to 1.5 TB of memory per server (using DDR4 DIMMs, greater capacities achievable using Intel Persistent Memory)
- Up to 6 TB of memory per chassis
- Up to 20 TB of m.2 NVMe storage per server
- Up to 80 TB of m.2 NVMe storage per chassis
- Up to 120 TB of SSD storage with optional 2U storage shelf
- Integrated GPU and FPGA options for Artificial Intelligence applications
- Up to 8 GPUs per chassis
- PCIe interface and expansion options
- Integrated 10 GbE ethernet switching, native 4 x 10 GbE on each blade
- Redundant power, AC or DC

The HPE Edgeline EL8000 Converged Edge System is designed to reside where data is created, which is sometimes in harsh, space-constrained and/or grimy, dirty environments. This compact, ruggedized system is designed to withstand increased shock and vibration. It also tolerates high ambient operating temperatures up to 55°C (131°F)—well beyond what traditional servers can typically endure.

The HPE Edgeline EL8000 chassis accommodates:

- Up to 4 x 1U HPE ProLiant e910 Server blades
- Up to 2 x 2U HPE ProLiant e910 Server blades
- A mix of 2x 1U and 1x2U HPE ProLiant e910 Server blades

The HPE Edgeline EL8000 chassis also provides redundant cooling, 1+1 redundant power supply and aggregated management of its HPE ProLiant e910 blades—which are themselves separate, independent, hot-pluggable compute nodes. TPM2.0 is supported by HPE Edgeline EL8000 on a cartridge and chassis level.

The integration of NVIDIA GPUs into HPE's powerful Edgeline platform provides the speed and agility to sense a wide frequency spectrum in near real-time. Each HPE ProLiant e910 1U blade can accommodate an NVIDIA T4 GPU via the low profile PCIe slot. The 2U HPE ProLiant e910 blade can accommodate up to four NVIDIA T4 GPUs if a higher GPU to CPU ratio is desired. The 2U HPE ProLiant e910 can alternatively accommodate full size GPUs (e.g. NVIDIA v100s, NVIDIA RTXs, etc).

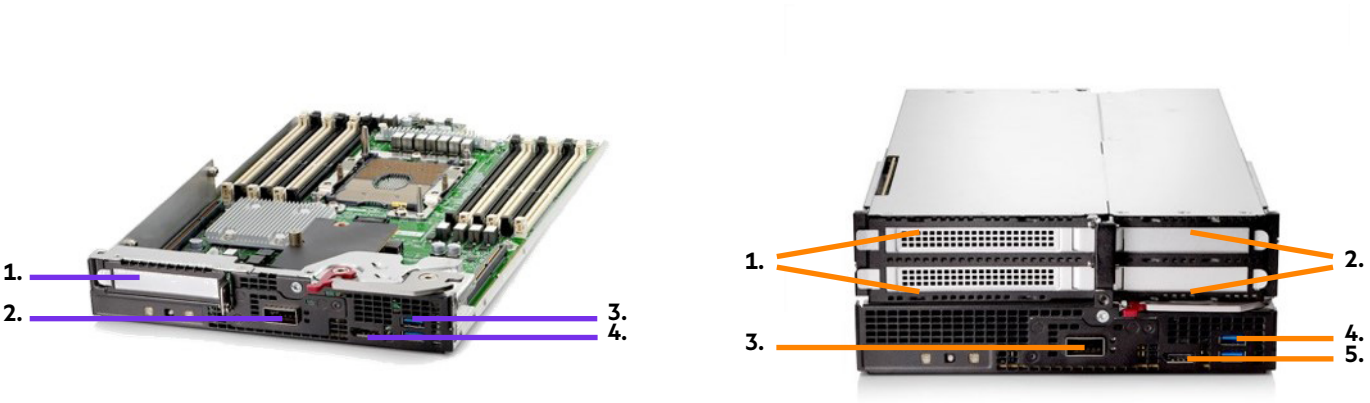

FIGURE 1. HPE ProLiant e910 compute cartridges, 1U and 2U versions

1. HALF HEIGHT, HALF LENGTH PCIE 3.0 SLOT

2. BUILT-IN NIC, UP TO 4X10 Gbps WITH QSFP+, SFP+ OPTIONS

3. USB 3.0 PORTS

4. DISPLAY PORT (DP)

1. PCIE SLOTS FOR FULL-HEIGHT FULL-LENGTH CARD

2. PCIE SLOTS FOR LOW PROFILE HALF-LENGTH CARD

3. BUILT-IN NIC, UP TO 4X10 Gbps WITH QSFP+, SFP+ OPTIONS

4. USB 3.0 PORTS

5. DISPLAY PORT (DP)

The following tables provide detailed technical specification for HPE Edgeline EL8000 chassis and HPE ProLiant e910 Server blades.

TABLE 1. HPE Edgeline EL8000 chassis technical specifications

HPE Edgeline EL8000 chassis	
Dimensions	8.63" (5U) height x 8.7" width x 16.91" depth (two chassis side-by-side fit into standard 19 or 23 inches rack)
Weight (single 5U chassis 2x2U cartridges)	50 lbs
Cooling	Front-to-back
Power supply	• 1+1 redundant PSU, 1500W each, -36-72VDC or 95-265VAC • 400W maximum power envelope per each 1U in the chassis
Environment specification*	Continuous operation 0 – 55C, Storage -40 – 70C, 8 to 90% operational humidity non-condensing; 5 to 95% non-operating humidity non-condensing.
Management and security	Chassis level: built-in dedicated silicon for out of band management for common chassis elements and optional aggregated management of individual cartridges, full RedFish compliance. Cartridge level: individual iLO5 chips on board of each cartridge, full RedFish compliance; chassis provides isolated network for aggregating out-of-band management from cartridges to single physical port off the chassis.

TABLE 2. Edgeline HPE ProLiant e910 compute cartridge technical specification

HPE ProLiant e910 compute cartridge	
CPU	Single Socket Intel® Xeon® Scalable, up to 205W TDP (supporting up to 8200 Platinum series). 8, 18, 24, and 28 Core Intel® Xeon® Scalable Processors are available and the processor stack will stay current with Intel's update cycle.
Memory	12 x DDR4 RDIMMS/LRDIMMS (2666/2933MHz 1DPC) (8 GB-128 GB) with ECC protection, Maximum Configuration of 1 TB (DRR4—subject to CPU capability, larger memory footprint option with Intel® Optane® Memory)
Storage	Two 2280 NVME M.2 flash storage slots. Additional M.2, both NVME and SATA are supported by optional mezzanine riser (adds up to four (4) x 3.84 TB 22110 m.2 NVME devices)
NIC	Integrated 4x10G, off cartridge faceplate as QSFP+ or integrated 2x10G off cartridge faceplate as 2xSFP+ Supporting wide range of PCIe-based NICs for support of up to 100 Gbps
I/O	1U version: 1xHHHL Gen3 PCIe (x16), 2x Type A USB 3.0 ports, Display port, optional NVMe riser (up to 4 additional m.2 NVMe devices) 2U version: 2xHHHL Gen3 PCIe (x16 lane dedicated to single card or 2x8 lanes in case of two cards) and 2xFHFL Gen3 PCIe (x16 lane dedicated to single card or 2x8 lanes in case of two cards), 2x Type A USB 3.0 ports, Display port. Optional internal NVMe riser (up to 4 additional m.2 NVMe devices) can be used on HHHL (i.e. right) side instead of external PCIe adapters.
Examples of PCIe accelerator support	Intel® Arria®10 FPGA, NVIDIA A100/V100 GPU, NVIDIA T4 GPU, NVIDIA RTX6000 GPU, and others.
Management and security	Individual iLO5 chips on board of each cartridge, full RedFish compliance.
Environment specification*	Continuous operation 0 – 55C, Storage -40 – 70C, 8 to 90% operational humidity non-condensing; 5 to 95% non-operating humidity non-condensing*.

- Additional detail on the HPE Edgeline EL8000 chassis
- Additional detail on the HPE ProLiant e910 server blades

*figures provided are valid for compute cartridges, could be subject to environmental limitations of components it carries (e.g. PCIe plug-in cards)

Software

The OmniSIG Sensor software package from DeepSig provides a new class of RF sensing and awareness. Going beyond the capabilities of existing spectrum monitoring solutions, OmniSIG is able to not only detect and classify signals but also to understand the spectrum environment to inform contextual analysis and decision making. Compared to traditional approaches to spectrum analysis, OmniSIG Sensor provides higher sensitivity and is more robust in harsh and dynamic spectrum environments, while requiring fewer computational resources and hence less power.

It provides a web-based UI as well as an open, low-latency streaming interface and control API for seamless integration into customer systems and applications. OmniSIG Sensor detects and classifies RF emissions across large bandwidths on the order of milliseconds, giving it the ability to report anomalies, changes or threats in near real-time. It works with both wide-band and narrow-band environments.

Detection and recognition has been validated across a wide range of signal types, including analog single carrier modulations, multi-carrier modulation schemes, cellular and infrastructure signals (e.g., GSM, WCDMA, CDMA2K, LTE), ISM-band signals (e.g., Wi-Fi, Bluetooth) and mobile radio services (e.g., P25, PTT). It can also be readily extended to include additional signals and protocols based on customer requirements and applications. It is robust to interference, both intentional and unintentional, and to other impairments like those caused by receiver hardware.

The OmniSIG Software Development Kit (SDK) is a complementary tool to OmniSIG Sensor that provides a new, powerful and accessible class of RF sensing. The SDK provides a tool suite that enables customers to rapidly curate their RF datasets and train state-of-the-art deep learning inference models for custom wireless sensing applications. The tools it embeds provide even more knowledge and understanding of known and unknown signals, without the need for a machine learning expert in the loop.

Using OmniSIG's SDK, users can generate custom, AI-enabled, RF detection and classification systems optimized for their specific needs. Using a combination of the latest web technologies and advanced server-side tools, users can load snapshots of RF signals of interest, label and annotate those custom signals, and then evolve a neural network sensitive to these new inputs without writing any DSP code or having to capture thousands of snapshots.

OmniSIG was developed using internal funding and is distributed as a binary release. Its design allows for integration with other technologies, since it is a containerized application that is of course fully compatible with all HPE Edgeline EL8000 operations.

Major OmniSIG updates are currently pushed as new Docker images and are tested prior to release. Minor neural network model updates are pushed as small packages that contain new model weights for the neural networks. This is a file that contains a new set of numbers to be applied to the nodes in the network. A major OmniSIG release occurs quarterly (every three months) with minor bug releases shipped approximately once per month.

USE CASES

Commercial Signal Identification

In the commercial space, as noted above, RF interference (whether on purpose or accidental) represents both productivity loss and revenue loss as well as an agitation (at best) to the user base. The HPE Edgeline EL8000 coupled with OmniSIG and its versatile SDK represents the ideal solution.

Spectrum deconflicting challenges is an emerging problem for commercial users—and is solved with HPE Edgeline EL8000/OmniSIG signal identification. HPE and OmniSIG's solution is a fast-sensing wireless signal detection system that operates over wide frequency bands with accessible protocols requiring no previous experience with deep learning algorithms.

Commercial users may also find the HPE/OmniSIG solution compelling for reasons of cybersecurity and cell jamming by non-compliant hardware (whether originating from innocent incompatibility issues or from rogue third party operators). In this scenario, OmniSIG's neural networks, bolstered by a robust interface for on-the-fly training via its SDK, rapidly learn standard patterns of the frequency space and thus can quickly isolate and identify anomalies that might otherwise go undetected.

Federal and national security signal identification

Defense and security customers require fine-grained spectrum analysis above and beyond what many commercial applications demand. However, other spectrum awareness solutions cannot provide millisecond-level timescales while also covering many GHz of spectrum. This persistent blind spot may represent a substantial shortfall in situational awareness and real-time threat assessment.

See Figure 2, "Operational View Diagram", for more details about the operational view of the present use case.

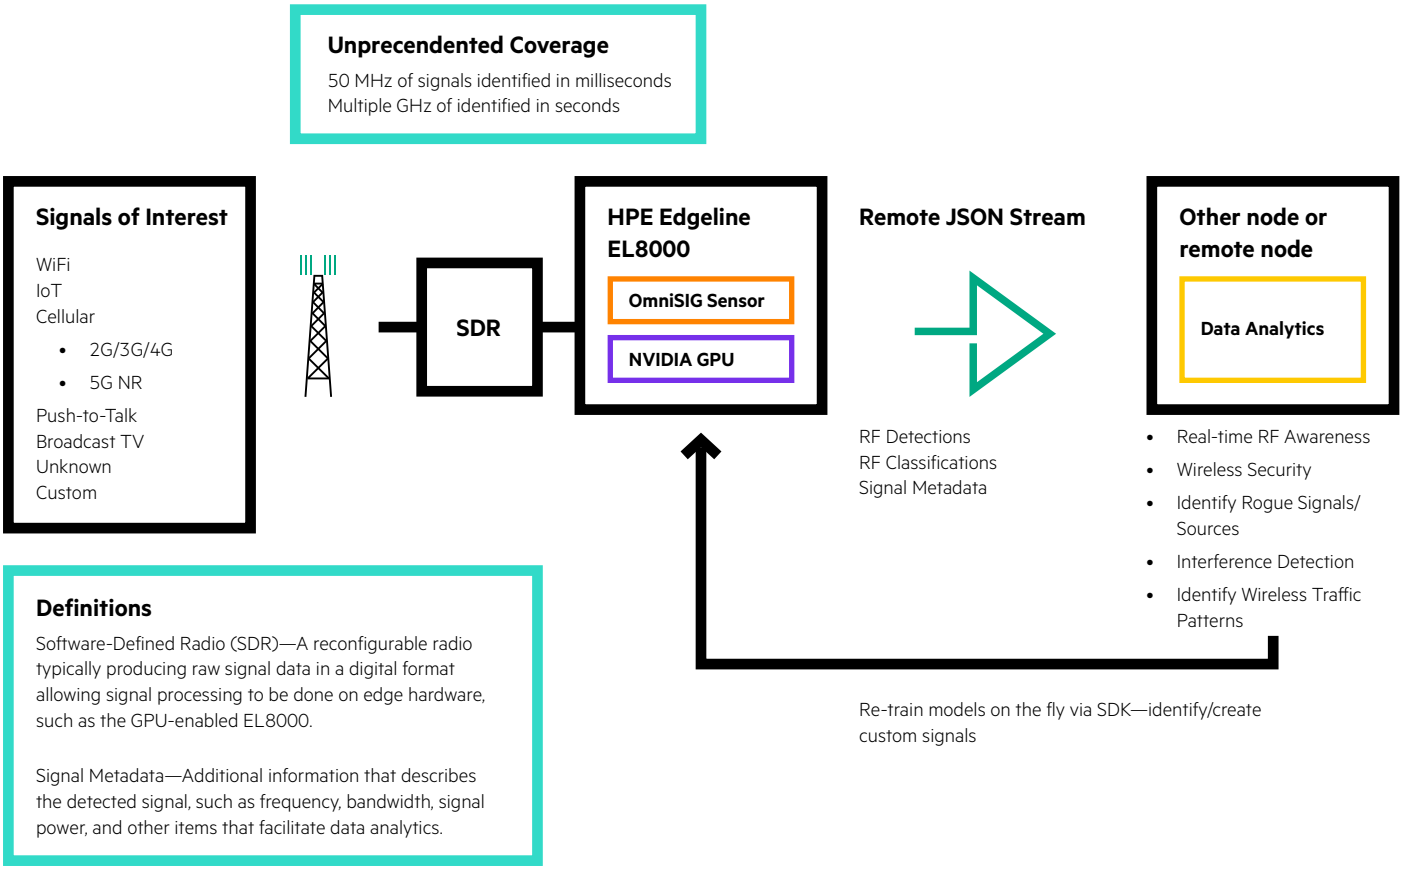

FIGURE 2. Operational View (OV) Diagram

As with the commercial use case, HPE Edgeline EL8000 runs OmniSIG with, in this case, GPU acceleration. Detecting wireless threats in milliseconds, OmniSIG is portable amongst many HPE hardware platforms and offers rapid notification of unknown signals. Its out-of-the-box capabilities include identification of standard commercial signals such as Bluetooth®, cellular, Wi-Fi and other radio signals.

Because the wireless battlefield is changing rapidly, new friendly and adversarial wireless signals are deployed nearly daily. This means that any signal identification system must be robust to such a dynamic threat space. Your electronic threat and surveillance systems must, in other words, adapt quicker than new systems can become operational.

OmniSIG's SDK means your signal intelligence platform can adapt in a matter of minutes. Your team will not be tasked with coding back at a data center. Instead, they can concentrate on capturing unknown signals, which the network can train itself on and then deploy for rapid discovery.

SERVICES

HPE Data Platform Design & Planning Services

Define, architect, and build an integrated IT strategy for AI and your data. Using an established methodology that brings together all the parts of your organization, HPE Pointnext Services consultants can devise the most efficient and effective means to capture, consolidate, manage, and protect business aligned information that suit your organization's use cases.

Consumption-based IT services

Gain the flexibility and control of the on-premises public cloud with HPE GreenLake—a set of consumption-based IT solutions. HPE Pointnext Services will implement and operate these solutions for you, enabling you to focus your own IT resources where they add the most business value. Or you can consume the technology of your choice, also using the pay-per-use model, in a manner suited to how you operate IT.

Operational Support Services

HPE offers a range of predefined, packaged support services from the following options:

- HPE Foundation Care is a cost-effective support to help customers when there is a problem
- HPE Proactive Care builds on HPE Foundation Care but provides an Enhanced Call Experience in case of a failure, while HPE Proactive Care Advanced includes a dedicated resource for more personalized support
- HPE Datacenter Care is a combination of all support services for all infrastructure components into a single support contract

SAMPLE CONFIGURATIONS

Since each HPE ProLiant e910 compute node can be independently configured to meet the workload, each node can be built with different CPUs, memory, NICs, optional GPUs, etc. Two representative examples are detailed below—the first containing four GPUs (one per node), the second containing a mix of GPU-accelerated and non-GPU accelerated nodes:

HPE Edgeline EL8000 AI Inference optimized (4 nodes) (4 x 1U nodes with NVIDIA T4)

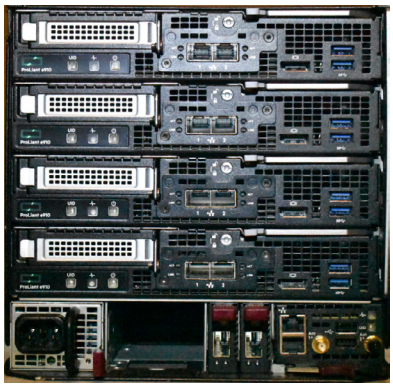

1 - 5U HPE Edgeline EL8000 Chassis

FIGURE 3. HPE Edgeline EL8000 AI Interference Optimized (4 nodes) (4x1U nodes with NVIDIA T4s)

BILL OF MATERIAL

- HPE Edgeline EL8000 AC Chassis w/ Switch
- 4 HPE ProLiant e910 Server Nodes

NODE DETAIL:

- 6212U Intel® Xeon® Scalable Processors 24 Core
- 576 GB Memory (6x64 GB + 6X32 GB)
- 2 x 240 GB m.2 boot
- 4 X 4 TB m.2 NVMe capacity tier
- NVIDIA T4 16 GB GPU

NOTES:

- 96 cores
- 2.3 TB Memory
- 64 TB raw NVMe storage
- 4 NVIDIA T4 GPUs
- 50lbs EL8000

HPE Edgeline EL8000 Accelerated Analytics/Training Optimized (3 nodes) (2 x 1U nodes, 1 x 2U node with NVIDIA v100)

1 - 5U HPE Edgeline EL8000 Chassis

FIGURE 4. HPE Edgeline EL8000 Accelerated Analytics/Training Optimized (3 nodes) (2x1U nodes, 1 x 2U node with NVIDIA v100)

BILL OF MATERIAL

- HPE Edgeline EL8000 AC Chassis w/ Switch
- 3 HPE ProLiant e910 Server Nodes

NODE DETAIL:

1U Nodes

- 6212U Intel® Xeon® Scalable Processors 24 Core
- 288 GB Memory (6x32 GB + 6X16 GB)
- 2 X 240 GB m.2 boot
- 4 x 4 TB m.2 NVMe

2U NODE

- 6212U Intel® Xeon® Scalable Processors 24 Core
- 576 GB Memory (6x64 GB + 6X32 GB)
- 2 X 240 GB m.2 boot
- 4 x 4 TB m.2 NVMe
- NVIDIA v100 32 GB GPU

NOTES:

- 72 cores
- 1 TB Memory
- 48 TB raw NVMe storage
- 50lbs EL8000

SUMMARY

New frontiers in wireless communication demand new wireless spectrum analysis solutions. Increasingly—for rapid allocation, analysis, security, threat identification and anomaly detection—this analysis must be performed at the network edge. Because of the stringent compute demands, the server must be both high-performance and ruggedized with portable size, weight and power (SWaP) factors. As the wireless spectrum broadens and expands in signal complexity, AI and machine learning solutions are increasingly essential to not only detect and classify signals, but also to understand the spectrum environment to inform contextual analysis and decision making.

AI-enabled applications at the edge present new challenges. For example, the computing resources needed for signal intelligence, computer vision and video analytics traditionally necessitated datacenters for CPUs, GPUs, memory and storage—usually requiring multiple systems working in unison. These components require data center-grade resources to accomplish their tasks. However, there simply is not enough space, storage, power, or cooling at the edge to handle today's AI and analytics demands using yesterday's data-center systems.

HPE Edgeline EL8000 coupled with DeepSig's OmniSIG Sensor software uniquely provide both inference (signal identification using a neural network) and training (the ability to re-train the neural network) at the tactical edge. Since the HPE EL8000 is a multi-node platform, a mix of workloads can be present in a single chassis. All of this processing power is contained within a 5U footprint that runs on standard 120V power and is hardened for use in the field. Coupled with the optional OmniSIG SDK, users can now re-train signal intelligence-based neural networks on the fly.

Ultimately, processing and analyzing spectrum data at the edge—pushing only the key data to the core that was surfaced in edge analysis—will help your teams make smarter and more informed decisions, reduce their time to decision and improve overall accuracy and value.

To explore the new directions that AI-powered wireless signal identification and analysis can take your team, contact your HPE representative for a demonstration today.

a00107131enw, October 2020

Hewlett Packard Enterprise