

ATLAS-1000

Fully Integrated, 1U OEM Application Acceleration Platform



Introduction

The Accolade Technology ATLAS-1000 is a fully integrated, 1U application acceleration platform exclusively designed for OEM customers developing network/cyber security and monitoring appliances. Like COTS appliances, the ATLAS-1000 provides a multi-core Intel CPU, memory and storage but that is where the similarities end. The platform integrates Accolade's APP (Advanced Packet Processor) in the form of an onboard FPGA with acceleration features such as lossless packet capture, nanosecond precision timestamping, packet merging, packet filtering, flow classification, packet steering and more. The API implemented for the ATLAS-1000 is also common to Accolade's ANIC packet capture adapters ensuring seamless migration across Accolade's comprehensive product offerings.

For a detailed description of ANIC features please visit:



<https://accoladetechnology.com/features/>

ATLAS-1000 has a very small form factor: 1U (1.75 in, 4.45 cm) height, 8.25 inches (20.96 cm) width, 14 inches (35.56 cm) depth. For higher density, two ATLAS- 1000 platforms can be mounted side-by-side in a standard 19-inch rack and multiple 19-inch racks can be populated with ATLAS-1000s to offer customers unprecedented scalability

In addition, ATLAS-1000 has a pluggable interface module (PIM) which can accommodate all the following network interface types:

- 4-port 1 Gigabit Ethernet, RJ-45 (Copper), with integrated passive copper TAP
- 4-port 1/10 Gigabit Ethernet, SFP+
- 4-port 1/10 Gigabit Ethernet, SFP+, with integrated passive optical TAP
- 4-port 1/10 Gigabit Ethernet, SFP+, with bypass in case of power loss
- 2-port 40 Gigabit Ethernet, QSFP+
- 2-port 40 Gigabit Ethernet, QSFP+, with integrated passive optical TAP
- 2-port 40 Gigabit Ethernet, QSFP+, with bypass in case of power loss

Customer Engagement

Accolade offers the ATLAS-1000 platform exclusively to OEM Customers and systems integrators in the security/network monitoring markets and **DOES NOT** develop or market any end user networking or security applications or appliances. Only when Accolade's OEM customers integrate their software applications with the ATLAS-1000 platform's acceleration features sets does the combined product turn into a network or security appliance.

For ANIC adapters, Accolade works closely with each OEM customer to successfully integrate the packet capture adapter into the OEM's network appliance. The customer relationship is identical with the ATLAS-1000 but of course no hardware integration is required because ATLAS-1000 is the hardware, packet capture card and acceleration features all bundled together in a small form factor. A major benefit of the ATLAS-1000 is that no hardware integration is required for the OEM customer since the ATLAS-1000 provides a comprehensive platform combining both a scalable CPU complex and the most advanced acceleration features in a compact package.

ATLAS-1000 strategically allows customers to concentrate on developing leading edge, value added security or networking software applications without having to deal with hardware design and acceleration feature development. All customers have to do is port their applications to run on ATLAS-1000 (utilizing Accolade's user friendly, comprehensive API) and the result is a unique, differentiated and complete solution for their end customers.

Architecture

Figure 1 shows the size and architectural layout of the ATLAS-1000.

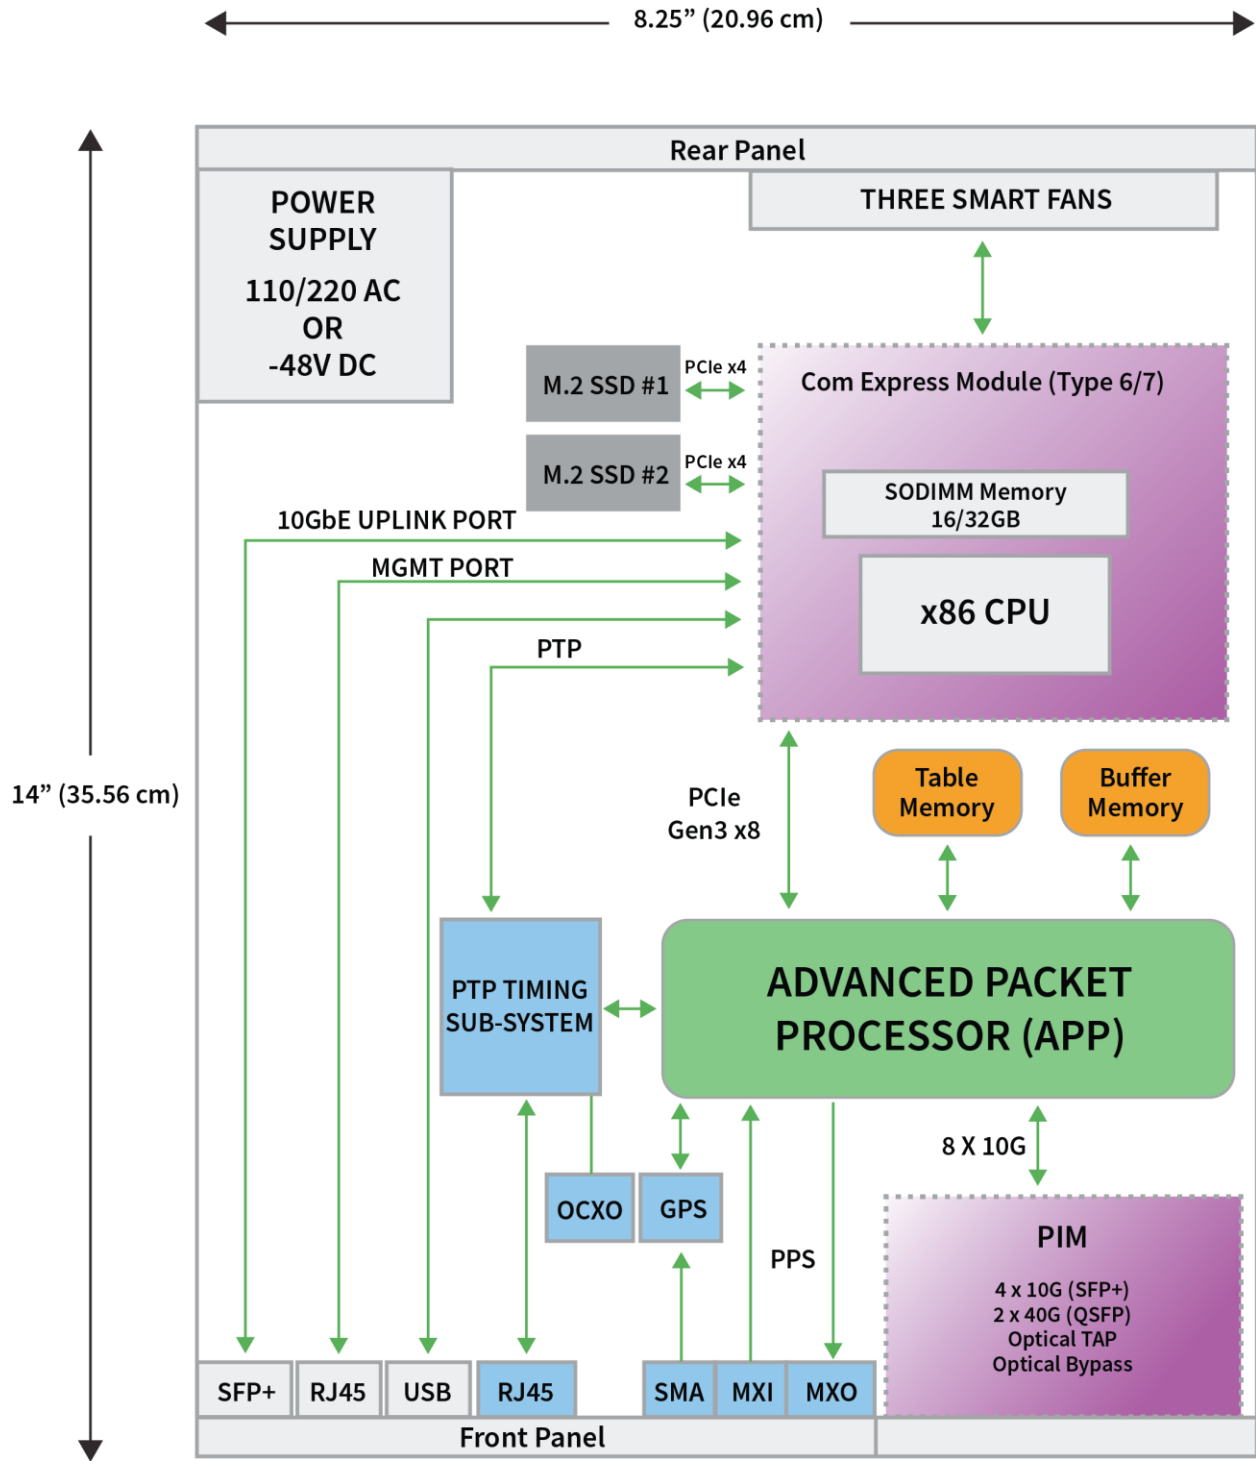


Figure 1: ATLAS-1000 Architecture

Advanced Packet Processor

The heart of ATLAS-1000 is the Advanced Packet Processor (APP) which is a Xilinx FPGA. This powerful FPGA along with the associated table and buffer memory provides packet processing and acceleration features such as flow classification, packet merging, deduplication, packet filtering, nanosecond precision timestamping, packet filtering, direct memory access (DMA) and more. Accolade engineers have designed all the intellectual property associated with the APP and provide support and custom firmware enhancements based upon specific customer requirements.

Details about all the features are available at: <https://accoladetechnology.com/features/>

Pluggable Interface Module (PIM)

The network ports on the platform are supplied via a pluggable interface module (PIM). As the name suggests, this PIM is replaceable and thus different interface platform configurations can be built and shipped with relative ease. PIMs are also field replaceable and can be ordered separately. Below is a list of available PIMs for customers to choose from:

- 4-port 1 Gigabit Ethernet, RJ-45 (Copper), with integrated passive copper TAP
- 4-port 1/10 Gigabit Ethernet, SFP+
- 4-port 1/10 Gigabit Ethernet, SFP+, with integrated passive optical TAP
- 4-port 1/10 Gigabit Ethernet, SFP+, with bypass in case of power loss
- 2-port 40 Gigabit Ethernet, QSFP+
- 2-port 40 Gigabit Ethernet, QSFP+, with integrated passive optical TAP
- 2-port 40 Gigabit Ethernet, QSFP+, with bypass in case of power loss

The PIMs with integrated passive optical or copper TAP are unique in that they eliminate the need for an external optical or copper TAP. With this combination, live optical fibers can be tapped and monitored very cost effectively in a single platform with minimal footprint. The split ratio for the optical TAP is 70/30 and they come in both MMF and SMF versions.

Figure 2 shows the front plate of a 10GbE PIM with integrated optical TAP (40GbE is very similar and can be used to monitor one 40G bi-directional link using Cisco BiDi compatible transceivers). There are four 1/10 GbE transceivers (SFP+) labeled P0, P1, P2 and P3 across the bottom. Each transceiver has a transmit (TX) and receive (RX) side. With four transceivers, this configuration allows two bi-directional links (4 fibers total) to be monitored. Above the transceivers are eight blue and four red optical fiber receptacles. A single optical fiber goes in or comes out each of these receptacles.




Figure 2: Integrated Optical TAP Front Plate

We will describe an example configuration that is monitoring both sides of a single, live bi-directional link. This example can easily be extended to accommodate an additional bi-directional link. We refer to “live” and “tapped” traffic to distinguish between traffic that is flowing in a live production network and traffic that has been replicated off the live production network. Clearly there can be no interruption to the live traffic or data loss will occur.

In Figure 2 the live traffic comes IN and OUT (as marked) the blue receptacles at the top and the tapped traffic out the red receptacles. There are two rows of receptacles marked “A” and “B”. The A row is for the fiber (of a pair) that is going in one direction and its corresponding fiber going in the other direction is on the B row.

By example, assume one live fiber is coming in to the blue receptacle labeled “IN 1” in row A. That tapped traffic then comes out of the red receptacle next door marked “TAP 1” and is jumped over to the RX (left) side of the transceiver marked P0. This tapped traffic goes into the ATLAS-1000 for analysis. The live traffic then comes out the next blue receptacle marked “OUT 1” and goes back to the live network. The other side of this bi-directional link follows almost the exact same pattern except it comes in the left most blue receptacle in row B. The tapped traffic comes out the red receptacle next door and is jumped over to the RX (left) side of the transceiver marked P1. And as before the live traffic comes out the blue receptacle and goes back to the live network.

In an identical manner, one more bi-directional link can be tapped using the right side receptacles that are marked “IN 2”, “TAP 2” and “OUT 2”.

Even if power is lost to the ATLAS-1000 the live optical fibers are not effected and traffic will continue to flow. The reason is because they are passive TAPs meaning they require no power to operate. In fact, the PIM could be physically removed from the platform



chassis and as long as the optical fibers remain intact, traffic will flow as expected without any interruption.

COM Express Module

The **computer-on-module** (COM) form factor is a highly integrated and compact module that can be plugged into any platform design. COM Express was developed and is maintained by PICMG (PCI Industrial Computer Manufacturers Group). COM Express was released in the summer of 2005 and is the most widely used COM standard. The standard defines the physical size, interconnect, and thermal interface for a COM. ATLAS-1000 accommodates a basic size, type 6/7 COM Express that measures 95mm x 125mm (3.7in x 4.9in) and contains a multi-core x86 CPU (e.g. Xeon, Core, or Celeron) along with up to 32GB of pluggable DRAM (SODIMM) memory.

The advantage of using COM Express is that based upon specific customer needs the CPU and memory can be modified during platform assembly by changing out the COM Express module. For example, some customers will require the most powerful Intel CPU available on the market with a large amount of memory while others require much less horsepower and can choose a more cost effective CPU with less memory. These seemingly contradictory requirements can both be satisfied with relative ease and efficiency.

Precise Timing

ATLAS-1000 is designed to be a very high precision timing device. To that end, there is an **optional** OCXO (Oven-Controlled Crystal Oscillator) on the motherboard. This type of oscillator is not susceptible to frequency change due to variations in ambient temperature and achieves the highest frequency stability possible with a crystal. These types of crystals are often used in military grade equipment and cellular base stations.

All timing related components are shown in Figure 1 in light blue color including the related RJ-45, SMA and MXC connectors.

ATLAS-1000 supports both PTP (Precision Time Protocol) and PPS (Pulse Per Second) timing. PPS timing is often supplied via GPS or a CDMA cellular network. For both PTP and GPS the related decode chips and circuitry are integrated on the motherboard.

There are three different ways to get timing input in to the platform and one way to get timing output.

1. For PTP (IEEE 1588) input there is an RJ-45 jack on the front panel. The uplink port can also support in-band PTP traffic.
2. For GPS timing, there is an SMA coaxial input on the front panel. The GPS signal is decoded (with the shown GPS block) and outputs a 1 PPS timing signal, time of day and location information to the APP.

3. There are two MCX (micro coaxial) connectors on the front panel labeled MCI (input) and MCO (output). A 1 PPS signal can be fed directly into the ATLAS- 1000 via MCI and a 1 PPS signal can be output by the ATLAS-1000 via MCO. A common use of these two ports is to daisy chain timing from one ATLAS-1000 to another. In other words, output a 1 PPS signal (MCO) from ATLAS-1000 #1 and send it as an input (MCI) to ATLAS-1000 #2.

Persistent Storage (SSD)

ATLAS-1000 is primarily designed to be a packet in, packet out platform; offering 1 terabyte (TB) of on-board, solid-state drive (SSD) memory. This memory can be utilized by the host application to persistently store data within the platform. The storage is broken up into two physically separate 512GB M.2 SSD. Each drive has a PCIe Gen3 x4 interface to the CPU complex for ultrafast read/write access.

Power

There are two power supply options: AC and DC. The AC option is typically required for enterprise use and the DC for telecom.

- AC – Universal AC Power Supply (100-240VAC – 47/63 Hz – 2.5A)
- DC – Telecom Grade (48VDC@2.8A)

Cooling (Smart Fans)


Due to the small form factor and versatility of the product, cooling is a major design consideration. Cooling is front-to-back with air entering the front panel and funneled out the back by 3 separate smart fans. We call them smart fans because their speed is regulated by onboard thermal sensors. When the platform is idle or receiving cold air flow, the fans may not run at all or just one runs at a relatively low speed. As the platform heats up more fans kick in and their speed also increases to keep the platform at an optimal temperature. The fans are also equipped with fan failure detection circuitry to alert the application in the unlikely event of a failure.

Uplink Port

On the front panel is a 10 Gigabit Ethernet (10GbE) port (SFP+ form factor) that the host application can use to send (or receive) data to the network. The 10GbE port is routed directly to the COM Express Module and therefore is controlled exclusively by the host application. The port can be used for a variety of functions such as in a security application to forward specific traffic to an upstream server that performs some targeted data analysis.

The uplink port can also support in-band PTP traffic.

Management Port



On the front panel is a Gigabit Ethernet (GbE) port (RJ-45) that the host application can use for management purposes. This port serves a standard management purpose and is typically used for configuration or alerting (e.g. SNMP trap) purposes.

USB Port

There is a USB port on the front which can be used for many purposes such as loading host application software into memory during testing and debug. Software can also be loaded from the SSD memory.

Serial Port (RS-232)

The platform also has a traditional front panel serial port (RS-232) that can be used as a control console for diagnostic and configuration purposes. One advantage of the serial port is it requires very little supporting software from the host system.

The serial port may also be used to communicate with external CDMA timing sources.

Application Programming Interface (API)

We've just described the architecture of ATLAS-1000, but how does a networking or security (host) application take advantage of the powerful packet processing and acceleration features available right on the motherboard?

Figure 3 shows a complete depiction of the various components involved to integrate an application into ATLAS-1000.

As with all Accolade products, the host application interacts with the APP (FPGA) via an application programming interface (API); a lightweight, C language shared library which is linked to the host networking or security application.

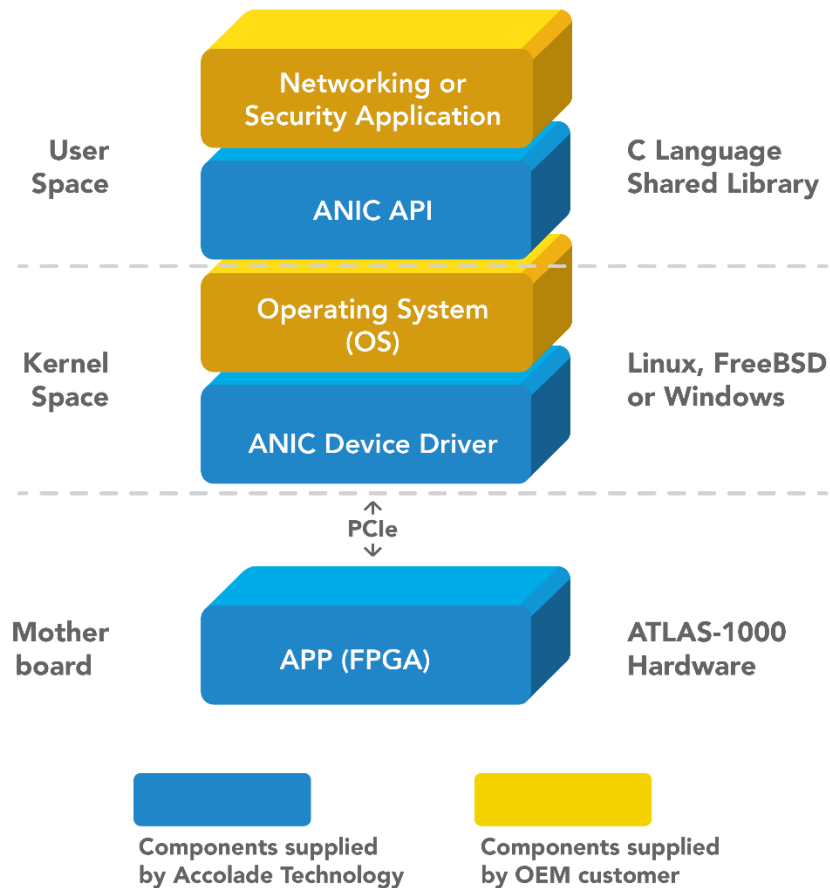


Figure 3: Host Application Integration

A comprehensive set of API calls are provided to perform functions including retrieving entries from a flow table, configuring policies, reading port status, etc. Based on specific customer requirements additional API calls can also be added.

Use Cases

ATLAS-1000 is a very powerful and flexible platform that can be used in a variety of different situations. Following are a few representative use cases to give readers some sense of how to utilize the product. This list is not exhaustive and the platform can be adapted to most any networking or security application. It is particularly effective when there is some specific packet processing or manipulation feature that must be performed at scale.

Flow Exporter

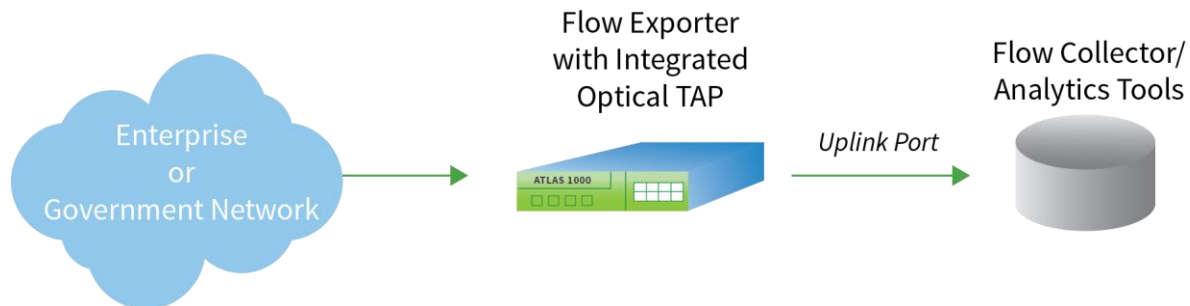


Figure 4: ATLAS-1000 as Flow Exporter Appliance with Integrated Optical TAP

Capturing flow records is an important monitoring function in computer networks and the de-facto standard format is NetFlow. In some cases, switches or other devices in the network can output NetFlow records. However, there are often points in the network where records are not available or perhaps must be output in another format. In this case, a standalone Flow Exporter appliance is very useful and with ATLAS-1000 the OEM vendor only has to supply the requisite software; leaving all the hardware and advanced packet processing responsibilities to the Accolade platform. The architecture shown in Figure 4 also subsumes the network TAP function directly into the appliance thus eliminating an element from the network.

Data Deduplicator




Figure 5: ATLAS-1000 as Data Deduplicator Appliance

Duplicated data poses a significant burden on network monitoring appliances. Sometimes as much as 50% of all traffic is duplicate. This is often because a SPAN port is configured to copy ingress and egress data from every switch port, which leads to duplicate packets for every packet that goes into and then out of a network switch. With ATLAS-1000, duplicate packets can be eliminated in hardware at 10 or 40G speeds before they ever reach the network monitoring appliance, thus saving a tremendous number of wasted processing cycles.

ATLAS-1000 can also be deployed as an inline data duplication appliance thus operating on live network traffic. In this scenario, for maximum data availability, the platform can be equipped with an optional PIM that includes optical bypass. With bypass, even if the data duplicator loses power or otherwise fails, traffic will continue to flow through the appliance (with no deduplication) uninterrupted.

Cluster Load Balancer

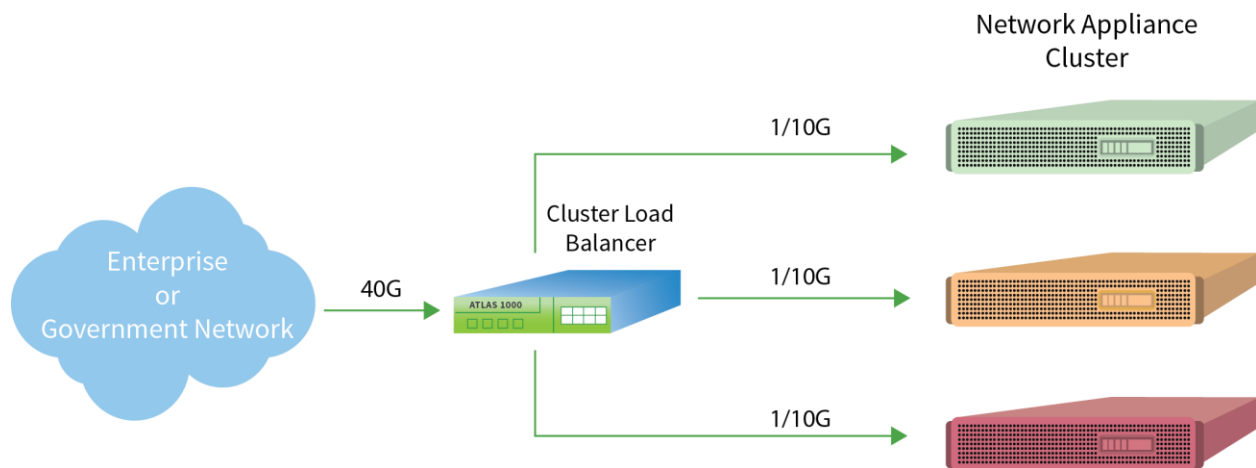


Figure 6: ATLAS-1000 as Cluster Load Balancer Appliance

For scale purposes a group or cluster of network monitoring appliances (NMAs) is often front-ended by a load balancer or packet broker. At a high level these devices take in packet data from the network, inspect the traffic and intelligently (based upon some rules or policy) distribute the traffic to some number of NMAs that form the cluster. The NMAs may be homogenous (different instances of the same function) or heterogeneous (cluster of NMAs with different purposes such as security and performance monitoring). Figure 6 shows a heterogeneous cluster of NMAs front-ended by an ATLAS-1000.

There are numerous load balancer and packet broker products available in the market from companies such as F5 Networks and Gigamon. These products are certainly appropriate in some situations but they are also very complicated and very expensive. Sometimes a simpler solution that has a targeted purpose with a much lower price point is more appropriate. This is where the ATLAS-1000 platform shines and delivers tremendous value to networking/security companies and their end users alike.

Inline DPI




Figure 7: ATLAS-1000 as Inline Deep Packet Inspection (DPI) Appliance

ATLAS-1000 can be used to limit or completely eliminate certain application traffic (e.g. Skype) within or between networks. The platform provides the hardware, optical bypass (in case of appliance power loss) and advanced flow classification. These capabilities combined with deep packet inspection (DPI) software from a leading vendor yields a robust, compact and cost effective solution.

Hardware-based flow classification is the key ATLAS-1000 feature that enables a robust inline DPI solution. Details on how flow classification can be used for this purpose can be found at: <https://accoladetechnology.com/hardware-based-flow-classification-of-up-to-32-million-ip-flows/>

Company Profile

Accolade is the technology leader in advanced, lossless packet capture and acceleration adapters and OEM acceleration platforms. Accolade's 1-100GE ANIC FPGA-based adapters and ATLAS series of acceleration platforms help accelerate network/cyber security and monitoring applications developed by the world's leading networking companies. ANIC adapters are fully PCIe compliant and seamlessly integrate into standard servers offered by companies such as Cisco, Dell, HP, Super Micro and others. Accolade's OEM customers offer products for network security and monitoring, flow classification, deep packet inspection, network test and measurement, video stream monitoring, high frequency trading (HFT), and more.

ID:170119

Corporate Headquarters:

124 Grove Street, Suite 315
Franklin, MA 02038

T: 877.653.1261

F: 208.275.4679

Silicon Valley:

980 Mission Court,
Fremont, CA 94539

T: 844.399.9904

South East U.S. Regional:

2997 Cobb Parkway,
Atlanta, GA 30339

T: 404.556.9393

www.accoladetechnology.com