Tech Brief

Next-Gen Brokering and Monitoring Architecture

Why Network Packet Brokers Should Not Be Designed Like Network Switches

Technology Benefits

- No Packet Drops

  A truly distributed smart-port based architecture scales with traffic without dropping packets

- Real-Time Monitoring

  FPGA-based probe design provides a 2-in-1 solution: monitoring and brokering – with advanced ultra-low-latency metrics

- All-in-One Solution

  No separate processing modules and no separate feature licenses make cVu an extremely simple and economical product to deploy and manage

Business Benefits

- De-Risking the Business

  By providing lossless data delivery to security and performance tools, you reduce the risk of security attacks and application down time

- Prolonged Investments

  Consolidating your data and delivery centrally and offloading tools for monitoring reduces costs, complexity, and tool sprawl

- Operational Efficiency

  Pervasive visibility allows you to optimize applications, security posture, and network capacity with timely upgrades

The Challenge

Today's mission critical networks and ever-increasing security threats to them require deep, fast, and high-performance network visibility and real-time monitoring. The network data plane has moved to 100Gbps, yet the monitoring architecture however is lagging behind. While the network I/O connectivity on the monitoring and network packet brokering (NPB) devices has been upgraded to match the data plane spine-leaf switches, the associated processing power lags behind. It's much like giving your decade old car a new paint job and tires but the engine remains the same.

The fundamental reason for this is the way those packet brokers are designed - originally inspired by the network switch designs and hence referred to as the "monitoring switches". However, today's packet brokering demands a very different architecture under the hood. Older designs suffer during peak network performance, drop packets, and cause serious blind spots for the security and performance monitoring tools - putting your business at risk. You can invest all you want in tools but if your packet broker drops packets, the tools are no good.

The Technology

Traditional network switches in the compact form factor are designed with one or two central packet processor (PP) which is a specialized application specific integrated circuit (ASIC) optimized for network packet lookup, manipulation, and forwarding in the data plane. If the switches support "stacking", a fabric chip (FC) may be added to support the interconnect, otherwise mostly it is "local switching" among the front ports. There are two mainstream switching architectures with their respective pros and cons: hash-based and cell-based, the scope of which is beyond the conversation here. Simply put, hash-based switching is based on cut-through and low-latency but with shallow buffers. Cell-based switching is not as fast but is good for deep buffer applications.

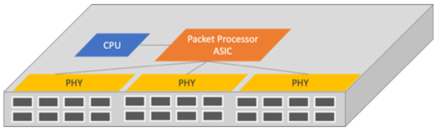

Compact switch

In case of modular switch chassis, the design differs. A rather older design is based on a backplane architecture. In this design, a backplane with copper traces runs at the rear of the chassis connecting different slots. A primary and standby management module (MM) or supervisor module (SM) hosts the CPU(s) for central control plane, packet lookup, and processing - and consumes a couple of those slots. In some cases, a fabric chip is also hosted on the MM or in some cases separate fabric modules (FM) are used to provide the data plane switching between different slots or I/O modules. With this design, bandwidth per slots is usually fixed and limited by the number of MM or FM (in load-sharing fashion). Moreover, the location of the FM can impact the signal integrity across different slots.

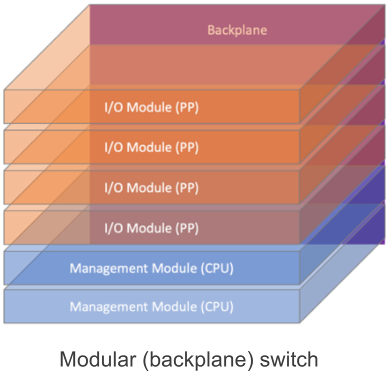

Modular (backplane) switch

Most of the switches today exercise multi-layer switching (MLS) or distributed switching where the very first packet in a traffic flow is always routed to the MM (CPU) for a lookup and forwarding decision. The decision is then downloaded and programed into the ASIC or PP on the I/O modules, so all the follow-up packets are locally switched through an I/O module or through the FM connecting different I/O modules. This results in "wire-speed" forwarding.

The newer "orthogonal" modular design eliminates the backplane mostly and hence many bottlenecks. In this case, more than one FM (usually in N+1 redundancy) connect to the I/O slots at 90-degrees and directly connect with each I/O and MM module. This results in faster and efficient switching suited for low-latency, high-performance applications. Many newer spine switches for data centers employ this architecture.

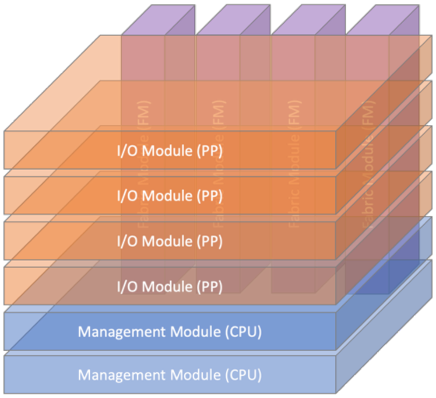

Modular (orthogonal) switch

Traditional Packet Broker Hardware Architecture

Many mainstream packet broker providers add specialized brokering features to fundamental switch architectures (such as filtering, packet slicing, deduplication, tunneling, replication) by adding "smart modules" and call it a packet broker. This solution works at low network speeds or when the number of packet manipulation features being turned on aligns with the total traffic. The problem starts when either the network speed at the I/O ports (also called port modules) approaches line rates or when more packet processing is required for every port, even at moderate speeds. As you add more ports (port modules), you need to supply more CPU power by adding more "smart modules" – which aren't cheap. Even then, there is no guarantee that you will not drop packets.

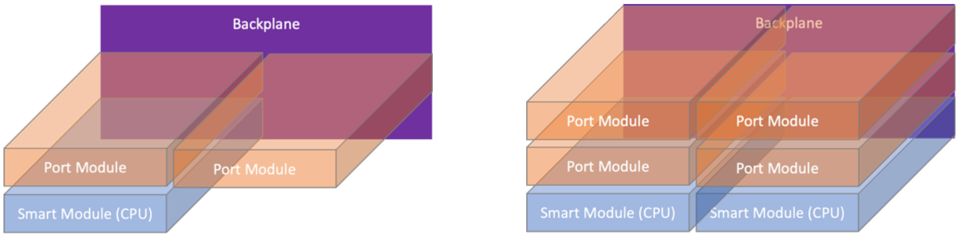

A generic Packet Broker needs port modules and matching Smart Modules

This "central architecture" simply does not scale at higher network speeds or intensive visibility applications, adds cost, adds complexity, and forces you to trade-off between port density vs. processing power since the number of slots does not change. Moreover, think of the backplane architecture issues discussed earlier, limiting the long-term bandwidth capacity per slot. Although the port modules may still have a localized ASIC, it doesn't get used to the full extent because unlike a switch, the packet broker aggregates traffic from multiple switches and the flows are short-lived and ever-changing, causing every other packet routed to the CPU. Now you are trapped as a Network Architect. You have invested too much money into a packet broker, it's all wired up, and you continue to spend more money to scale, but your security and application teams keep complaining about the missed packets.

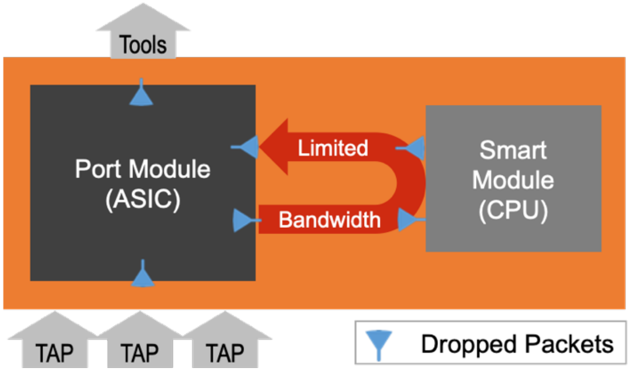

Central architecture causes severe choke points and dropped packets while adding cost and complexity.

Traditional Packet Broker Software Architecture

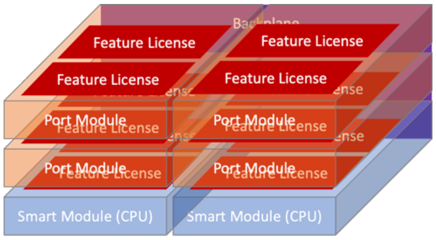

Complex licensing scheme adds cost and complexity

The pain does not end here. Many of the traditional packet broker vendors have come up with a distributed software licensing model on top of the centralized hardware model. The type of brokering features you can deploy depends on the smart modules and port modules you procure and the licenses you are entitled to. For example, you may need to buy a license for deduplication, another license for tunneling, and yet another license for advanced filtering. To make it more complicated, those licenses are restricted by the module type, not by the packet broker. So now you have a nightmare of tracking and managing licenses while you were already trying to justify the costs to your boss.

This is when many network engineers start looking for alternative packet brokering solutions. If this was not frustrating enough, then it may be a Saturday in the war room with the CISO since a data breach was detected too late because the security tools never saw the critical packets, they needed for detecting a threat.

The Solution

This is where a truly "distributed architecture" with clean thinking is required. If you put localized processing intelligence behind every I/O port, the packets do not need to be routed all the way to a central CPU, at least not often. Also, the packet manipulation happens as soon as it enters a "smart port" at the ingress, or right before it leaves a smart port at the egress, making it much faster, efficient, and real-time. This is ideal for ultra-low-latency applications and this is where the cPacket technology advantage begins.

cPacket Packet Broker Hardware Architecture

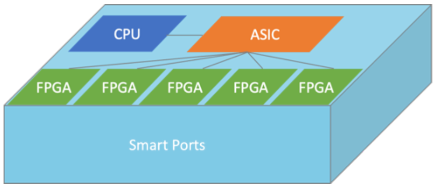

cPacket cVu Packet Broker+ uses a FPGA-based distributed processing

cPacket cVu® series packet broker+ hosts a field programmable gate array (FPGA) behind every I/O port – called a "smart port". The FPGA runs specialized code for advanced features at wire-speed operation. Think of the FPGA as a packet processor per port which is faster and more flexible than a centralized merchant ASIC and also upgradable. This allows cPacket to continuously innovate and keep the customer investments prolonged. Think of an F-16 with multiple "blocks" where a software upgrade can take the avionics to the next level. Or think of a Tesla vs. a traditional car.

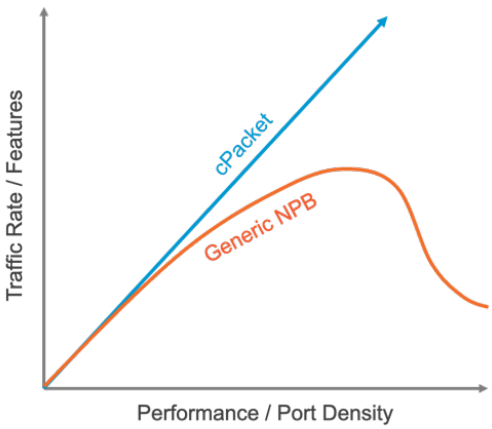

Behind the FPGAs, there is also an ASIC, but that is for basic operations such as packet switching and forwarding between the ports. The packet brokering features though are mostly pushed to the FPGA. So, comparing a next-gen packet broker+ to a legacy packet broker, the cVu has no packet drop issues and scales at lines rates up to 100Gbps speeds while you can turn on any of the supported advanced features at the same time. This is the architecture under the hood in the cPacket next-gen (NG) series of packet brokers – a key part of cPacket Intelligent Observability Platform: cVu 16100NG, 8100NG, 4100NG, 3240NG, 2440NG, 560NG, 400NG, 240NG, and 160NG. See the Product Quick Reference Guide for details.

This is one of the fundamental reasons why the cVu NG series packet broker+ continue to replace other packet brokers in mission critical environments such as financial services, high-frequency trading, high-performance computing, healthcare, government, and many other applications.

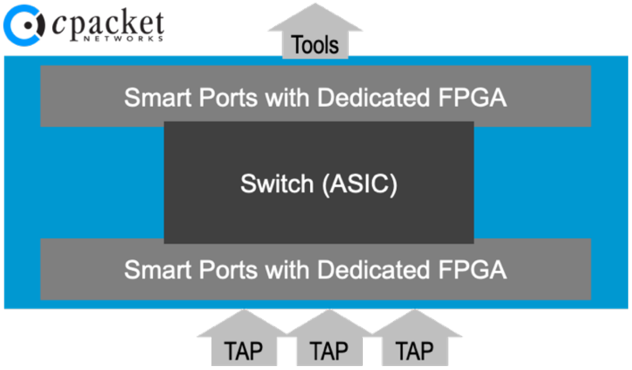
cpacket NETWORKS
Tools
Smart Ports with Dedicated FPGA
Switch (ASIC)
Smart Ports with Dedicated FPGA
TAP
TAP
TAP

Distributed Architecture scales with traffic and packet processing requirements and reduces cost and complexity.

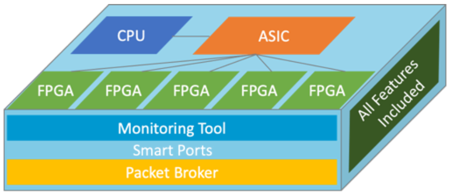

cPacket cVu acts as a Packet Broker as well as a monitoring tool

The technology and commercial advantages of cPacket do not end here. Where the return-on-investment really doubles is that FPGAs play an additional role of being multiple "monitoring probes" as well. So, while the network traffic is passing through those smart ports, why not analyze it and extract some key metrics? This turns cVu into a 2-in-1 platform: a packet broker as well as a monitoring device or a tool at the same time. Some of the key metrics reported include precision timestamping, micro-burst analysis, bandwidth utilization, and one-way latency. For this reason, one of the most popular products deployed by financial services is cPacket's cBurst.

cPacket Packet Broker Software Architecture

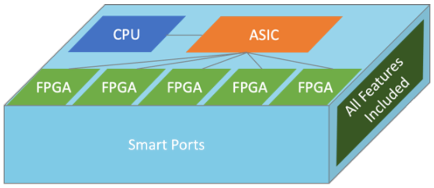

cPacket cVu is an all-inclusive product

Another major advantage of the cPacket cVu NG series architecture is that it does not apply a complicated software licensing model. What you buy is what you get and what you deploy. The value that the cVu NG series delivers as a platform far exceeds its total cost of ownership. That makes the cVu a really hard to beat platform from both a technical and an economical angle.

Overall, (1) cVu NG acts as a data consolidation and brokering platform for IT security and performance monitoring tools, reducing tool sprawl and costs, (2) cVu performs network monitoring offloading or even avoiding the need for such specialized tools, again saving costs, and (3) cVu measures the bandwidth utilization for every tool connected using features like cBurst, helping in planning tool upgrades and prolonging the return on investment.

Call to Action

Want to see the cVu NG series architecture in action? Request a product demo today.

About cPacket Networks

cPacket Networks enables IT through network-aware application performance and security assurance across the distributed hybrid environment. Our AIOps-ready single-pane-of-glass analytics provide the deep network visibility required for today's complex IT environments. With cPacket, you can efficiently manage, secure, and future-proof your network - enabling digital transformation. cPacket solutions are fully reliable, tightly integrated, and consistently simple. cPacket enables organizations around the world to keep their business running. Our cutting-edge technology enables network, application, and security teams to proactively identify issues before negatively impacting the business. The result: increased security, reduced complexity, and increased operational efficiency. Learn more at www.cpacket.com