

KEY CHALLENGES:

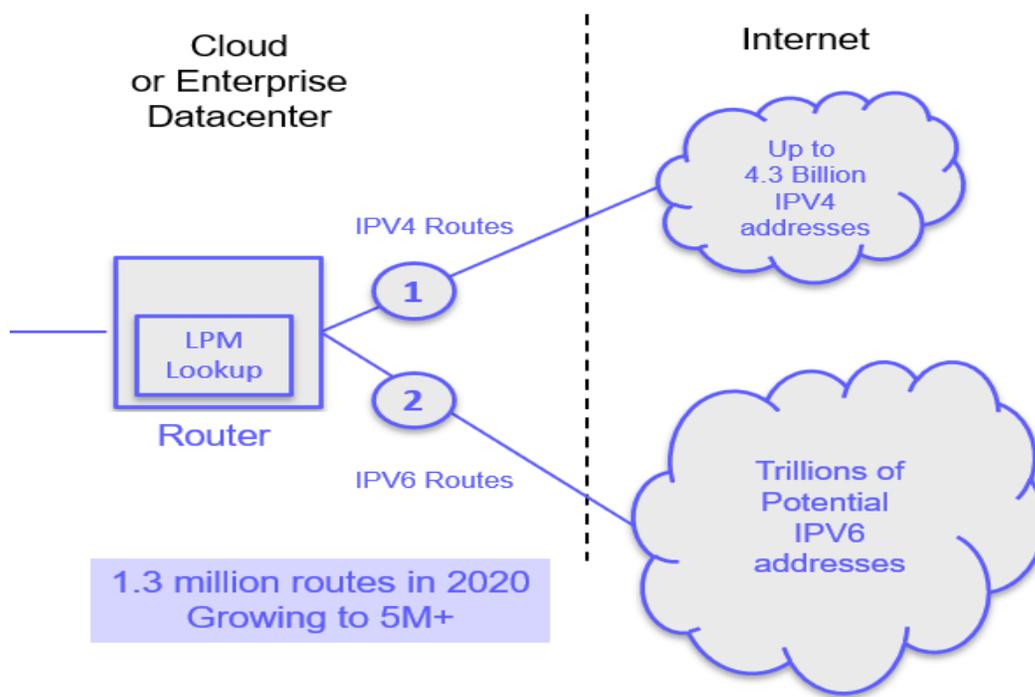
The latest information publicly available puts the number of IPV4 routes at close to a million routes with potentially more than 4 billion connected devices using IPV4. IPV6 routes are still relatively small – less than 300 thousand. However, with the advent of 5G cellular wireless and IoT (Internet of things) devices, that number is expected to increase dramatically over the next few years. This is because IPV6 has 128 address bits versus only 32 bits for IPV4 and it can potentially connect to 3.4×10^{38} devices. Yes, that is 38 zeros.

Therefore, being able to successfully route packets efficiently to those devices is going to be very important.

We are already seeing new designs for Cloud and Enterprise Datacenter Routers demanding this sort of capacity with any mix of IPV4 or IPV6

Routing requires that the headers of each packet is examined in real-time. This is called Deep Header Inspection (DHI) – and searches can use a combination of Exact Match, Longest Prefix Match (LPM) or even Access Control List match (ACL).

Depending upon the system, the total number of rules, the number of rules that match, the complexity of the rules and the speed of the searches (Millions of searches per sec) determine the overall performance level.



KEY SYSTEM CONSIDERATIONS:

- Cloud and Enterprise Routers typically sit at the edge of the datacenter or the edge of an Enterprise network
- Deep Packet Header Inspection requirements – general purpose routes must be able to support millions of rules and do lookups in real time
- Typical 2 – 4+ Million rules – mix of IPV4 and IPV6
- Typical 50 – 400+ Million searches per second
- Typical 100Gbps to Terabits/s with low latency a must

Meeting the search requirements can be accomplished in several ways:

- Execution in Software – but this can be too slow
- Using a standalone TCAM chip – but this can be high cost and high power
- Using a hardware accelerated Algorithmic TCAM in ASIC or FPGA
- Combining a multi-terabit SmartSwitch with any of the above

MOSYS SOLUTION

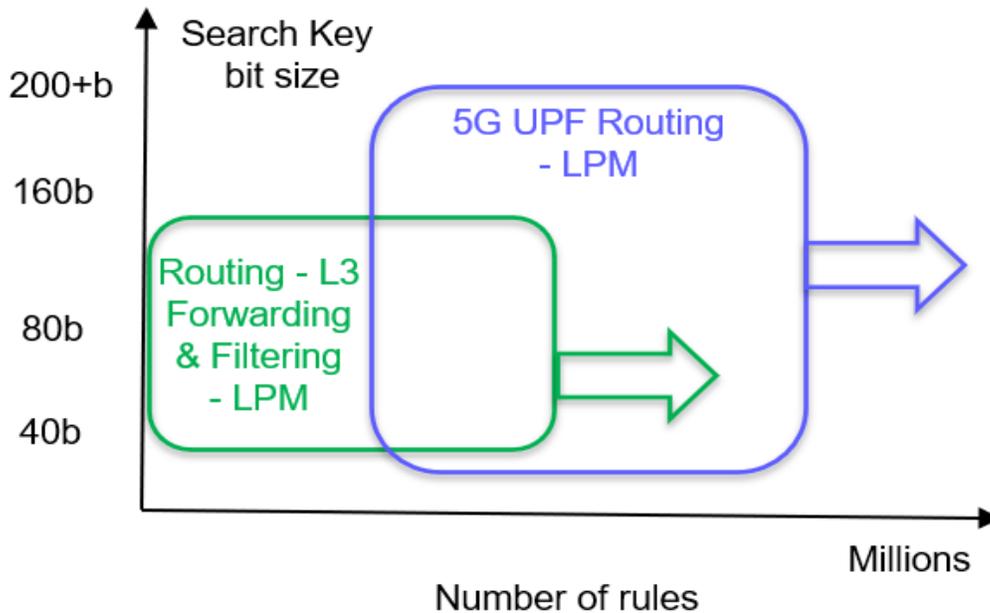
The MoSys Stellar Packet Classification Platform – High Performance LPM is provided as Intellectual Property (IP) that uses a hardware accelerated Algorithmic TCAM-like approach to help ensure that a router can keep up with the huge volume of routing decisions per second that it has to process.

High Performance LPM

- Ultra-High-Speed Search Engine IP
- Deep Header Inspection (DHI) solution
- Available for ASIC or FPGA
- Optimized for high performance routing
 - Can add other functions
 - Security, load balancing...
- Tuned for Longest Prefix Match (LPM)
 - Any mix of IPV4 and IPV6 lookups
 - Optimized for 1 or 2 tuple matches
 - Supports virtual routes
- Provides scalable performance
 - Uses Graph Memory Engine (GME)
 - 100s of Million lookups per second
 - Low latency solution
 - Very efficient memory usage
 - Extremely efficient use of logic gates
 - Very fast rule updates
 - No need to recompile rules
 - No need to preconfigure table sizes
 - On the fly updates – no need to stop traffic
 - Ability to receive updates from AI/ML logic
- Supports large number of bits for next hop data
- Capacities and key sizes beyond normal routing
- Up to multigigabit TCAM equivalence
- Supports broad range of devices
 - Can utilize hybrid mix of memories
 - Internal SRAM and/or external DDR, HBM
 - Can also use MoSys memories, but can operate without any MoSys silicon present
 - Supports RTL for Intel Stratix 10, Xilinx UltraScale+ FPGAs, or ASIC/SoC/DPU...
 - Replaces multiple expensive and power hungry TCAM chips
 - Common API for software interface – easier to port applications
- Applicable to designs based on NIC, SmartNIC, DPU, Standalone SoC, SmartSwitch...

KEY POINTS SUMMARY:

- Very High-Performance solution designed to accelerate one of the main Cloud and Enterprise Datacenter bottlenecks – IPV4 and IPV6 Routing
- Very flexible design – MoSys IP easily integrated
- Takes advantage of available gates and memory in FPGA or ASIC
- Helps future proof designs by supporting wide range of key sizes, n+ tuple looks ups, very large number of rules at a very high performance in very efficient logic



ADDITIONAL RESOURCES:

- [Stellar Virtual Acceleration Engines](#)
- [Stellar Virtual Acceleration Platform](#)
- [Virtual Acceleration: The MoSys Approach](#)
- [Cheetah Development Kit](#)