

The ACOS Advantage: **Advanced Core Operating System**

Technical Brief

AX Series New Generation Application Delivery & Server Load Balancing

Superior Performance By Design

Today, business agility requires flexible application delivery platforms optimized for high-performance. However, high-performance deep packet inspection, application acceleration, and application off-load combined with more connection intensive technologies are driving legacy OS designs to their limits. A10's ACOS with true Scalable Symmetrical Multi-Processing (SSMP) is the solution.



SSMP: Scalable SMP OS Design

ACOS was designed from the ground up to meet the requirements for scalability and business agility. Key industry differentiators are:

- Designed from inception for multi-core CPUs
- Decoupled CPU architecture
- Shared memory/No IPC (Inter-process Communication)
- 32-bit or 64-bit OS
- ASIC/FPGA hardware design

ACOS has zero IPC, zero interrupt, zero scheduling, zero locking and



zero buffer copy, in a linear non-blocking architecture that ensures the lowest latency in packet processing.

Decoupled Multi-core CPU Architecture: Faster, Truly Independent

Designed for modern multi-core CPUs from day one, ACOS removes single CPU and single core restrictions. With no wasted cycles or unneeded processes, ACOS ensures each processor is truly independent from each other. No single processor or core is reliant on another processor to complete a cycle.

Shared Memory: Accurate Data, Faster Processing

Shared memory ensures all needed information is instantly accessible, enabling accurate, real-time, single sourced data for global decision making criteria. IPC is eliminated. Multiple copies of the same data are no longer required per processor, using less memory space, which can be significant when multiplied by a large number of cores.



No Inter-process Communication (IPC)

As ACOS has removed the requirement for IPC, no updates to other processors and their dedicated memory are required. Whether via traditional IPC or "messaging sub-systems", accuracy, reliability and overhead goals are not met without shared memory. This is most notable for internal efficiency and accuracy at high traffic volume. Updates are reduced to nanoseconds with instruction level waits versus multiple milli-seconds for IPC overhead.

Technical Brief

IPC Examples

Centralized Method: Each data core updates its own data/ state information via IPC message to the control core. Upon receiving the message, the control core decodes the message, updates the relevant data, and then sends the corresponding data/state information to all the other cores.



 Distributed Method: Each data core updates its own data/ state information via IPC multicast message to all other data cores.



IPC Message Overhead and Latency

In a connection rate example, using the distributed method, each core sends an IPC message to update the state based on a periodic timer. Periodic updates reduce the number of IPC messages sent, but at the expense of accuracy of the state; because notification of the state change is delayed.

In a system delivering 1 million connections per second, where the maximum connection rate is updated via IPC once every 100 ms, it can cause a 100,000+ maximum rate connection inaccuracy. This would be the same for any feature requiring accurate and real-time accessible data.

With ACOS there is no IPC, so all global state information is always accurate.

32-bit or 64-bit OS: Faster, More System Resources, More Scalable

ACOS was designed for 32-bit and 64-bit processors. With 64-bit ACOS, integers, memory addresses or other data units that are 64-bits (8 octets) wide can be processed in a single cycle. Secondly, this frees individual CPUs from the 4 GB memory limitation per CPU, resulting in more addressable memory available for the OS's system resources. In practical terms this allows room for increased maximum session limits and any other feature that relies on memory.

When porting legacy software to 64-bit, systems could use memory and processing cycles more inefficiently, actually resulting in a 64-bit penalty. 64-bit design as demonstrated by ACOS from conception is critical for future performance and scalability, resulting in optimal performance.

ASIC/FPGA Hardware Design

ACOS leverages the best combination of hardware and software for maximizing performance. Additional hardware streamlines the packet flow and off-loads CPU intensive applications to amplify performance.

ACOS can leverage switching/routing ASICs for layer 2/3 network processing, SSL ASICs for decryption/encryption, flexible traffic ASICs or FTA (based on FPGA technology to allow reprogramming) and Compression ASICs for HTTP compression.

The FTA technology not only optimizes flow distribution, it also allows DDoS SYN flood protection, ensuring an attack will add 0% additional CPU utilization required for legitimate traffic processing. Today, ACOS can take advantage of the latest processors, which outperform all leading vendors' current choices.

However, ACOS can capitalize on any new processor architecture as required, due to the fact ACOS virtualizes the processing logic independently of other logic for maximum flexibility.

Highest Performance Today

ACOS represents a new generation of Application Delivery Controller (ADC). While internal architectures differ, a proof point can be found in performance. For example, the AX Series 5200 platform utilizes the same number of cores or less for over twice the throughput and 3x the transactions per second, in a form factor a third the size of competitors' latest platforms.

This results in greater performance, and the greatest performance per Watt, making the AX Series the most flexible, scalable, energyefficient solution on the market today.