



White Paper

Array SpeedCore™ Architecture

Application Delivery Networking

Multi-Processing Technology

Introduction	3
CPU vs. ASIC	3
SpeedCore Overview	4
SpeedCore Logical Architecture Diagram	4
Lock-Free Multi-Core Design	5
Content-Aware Traffic Engines	5
Fault-Tolerant Architecture	5
Combining Intelligence & Performance	6
SpeedCore Detailed Architecture Diagram	6
SpeedCore Benefits	7
Breakthrough Performance	7
Application Intelligent Processing	7
Dynamic Scalability	7
Reduced Packet Drops	7
High Reliability	7
Extensibility	7
Adaptable Software Architecture	8
System Management	8
About Array Networks	9

Introduction

Growth in Web and application traffic – as well as trends in mobile and cloud computing – are placing increased demands on application delivery networking solutions. As a consequence, today’s application delivery solutions must be rooted in an architecture that can deliver both performance and scalability and the agility to operate in any environment and respond quickly to an evolving technology landscape.

This paper provides an overview of Array’s SpeedCore multi-processing architecture, detailing the advantages of CPU-based multi-processing as compared to ASICs and providing an under-the-hood look at Array’s core technology and its associated benefits for application delivery networking.

CPU vs. ASIC

A decade ago, CPU load balancing lost the battle to ASIC load balancing because general purpose CPUs and operating systems could not deliver sufficient performance. However, performance improvements enabled by ASICs did come with downsides including higher engineering costs, longer time to market and greater complexity in fixing bugs.

With the evolution of general-purpose processors and multi-core technology, the pendulum is swinging away from ASICs because optimized software can take advantage of off-the-shelf hardware to deliver equivalent or better performance with increased agility and greater affordability.

With optimized software running on cutting edge off-the-shelf processors, vendors can introduce new features and enhancements without requiring customers to rip and replace existing products or take products and services offline to perform a hardware upgrade.

This is particularly important for cloud service providers and emerging services where application delivery requirements are rapidly evolving. Vendors with a CPU-centric approach have the ability to rapidly meet these needs in a timely manner with features and customizations that can be deployed with a simple non-disruptive software upgrade.

Packaging optimized software in purpose-built hardware appliances minimizes complexity and maintains a controlled environment that guarantees high performance and reliability and keeps support costs low.

To get the best of all worlds – cutting edge performance and scalability with agility and affordability – it is essential to run optimized software (such as Array SpeedCore) on purpose-built appliances that leverage the latest in multi-processing technology. This is because standard software, operating systems and processors share resources with tasks unrelated to networking and pay a performance penalty because underlying components are not optimized for the demands of application delivery networking.

SpeedCore Overview

SpeedCore is a new-generation software architecture that allows Array products to scale beyond 10Gbps to meet the requirements of complex and demanding application delivery networking environments.

SpeedCore leverages multi-core technology to increase performance by up to ten times vs. standard processors and operating systems. Performance gains enabled by SpeedCore significantly improve HTTP, TCP and SSL processing and take full advantage of today's 10Gbps network interface cards. Combined with efficient packet processing, a 64-bit operating system and larger memory size, SpeedCore allows Array to introduce reliable, high-performance and fully-featured products with rapid time to market. In essence, Array has developed an efficient, optimized, scalable and flexible architecture that works well under any load and in any use case without compromising on price or performance.

Linear Scalability Across Multiple Processors & Cores

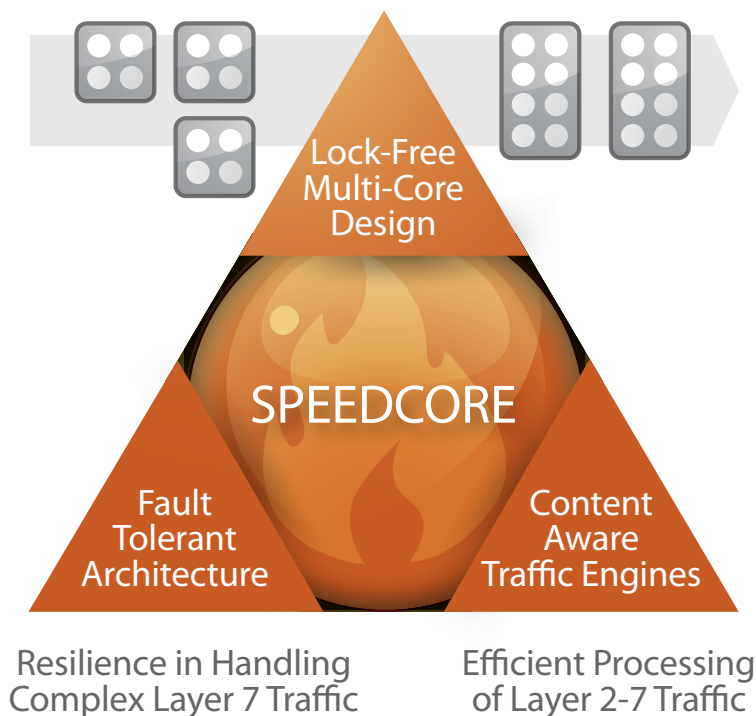


Figure 1: SpeedCore Logical Architecture

Lock-Free Multi-Core Design

- A key attribute of the SpeedCore architecture is a lock-free design that enables linear scaling with the addition of processors and cores
- The lock-free architecture supports multiple applications simultaneously without negatively impacting user experience or application uptime, enabling more applications and services to be deployed in parallel and in real-time with greater efficiency

Content-Aware Traffic Engines

- Fast and intelligent application and content processing engines accelerate and scale the performance of enterprise and office applications such as SharePoint, Exchange, CRM, ERP, SAP, Oracle and multimedia SIP/RTSP applications
- Allows for selective acceleration and optimization based on the nature of the content and type of application. Adaptive engine usage is based on Layer 2, Layer 4 and Layer 7 content. Content-aware traffic engines efficiently process Layer 4 connections, Layer 7 application requests and SSL traffic
- Enables network managers to deploy more applications and services in real-time, while running the processor-intensive tasks in the backgrounds

Fault-Tolerant Architecture

- The fault-tolerant design provides application and service flow isolation and separation, improving scalability and stability
- Processing L7 traffic can be complex and error-prone, as application connections, packets and content must be “glued” together to preserve the context of the flow. This gluing enables SpeedCore to apply application and content-specific policies to increase system stability and availability
- SpeedCore has built-in health and memory usage monitoring that enables the system to be self-healing and adaptable under abnormal and surge traffic conditions. The architecture isolates abnormal traffic patterns to a specific process on the system to prevent system crashes and allow continued processing of normal user traffic with zero service interruption

Combining Intelligence & Performance

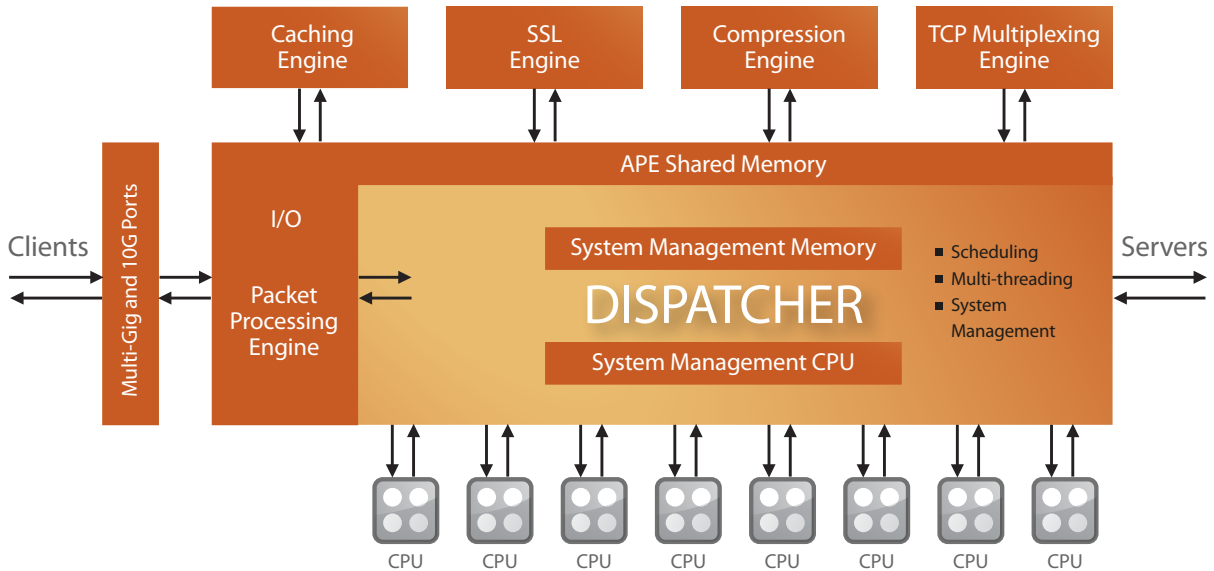


Figure 2: SpeedCore Detailed Architecture

SpeedCore's dispatcher provides ultra-high performance by taking advantage of multi-core CPUs for efficient handling of Layer 4-7 traffic while maintaining 10Gbps+ application traffic rates. Dedicated application processing engines (APEs) are highly scalable; all run concurrently without performance degradation.

With today's increasing application complexity, it takes an intelligent design to maximize multi-core performance. The system must be able to pool application traffic flows, understand and inspect packets, divide flows among multiple cores in parallel and asynchronously process application traffic flows in real-time. In addition, all of these tasks must be executed without latency. Resolving data bottlenecks, thread synchronization issues and process contention issues between multiple cores are the most common hurdles.

The SpeedCore dispatcher resolves synchronization, process contention issues and data bottlenecks with the help of application processing engines (APEs). The dispatcher is always aware of the overall systems computing capacity and availability, and therefore can intelligently distribute application flows across multiple cores. The dispatcher receives application traffic flows from the I/O subsystem, forwards them to the appropriate APE and allocates available CPU cores on-demand – thereby removing

data bottlenecks, synchronization overhead and parallel processing complexity associated with multi-core processing of application traffic. These capabilities create an architecture that drives ultra-high performance in the most demanding environments and enables Array's application delivery products to run all features concurrently.

SpeedCore Benefits

Breakthrough Performance SpeedCore reduces response times for a broad range of legacy, new and evolving applications, including rich Internet technologies. The architecture builds upon 64-bit addressing to provide more I/O bus bandwidth (to deliver higher aggregate throughput) and larger memory addressability (to handle more connections and transactions); the benefits are more throughput and a larger system cache. In addition, multi-core processors are well-positioned to bring performance benefits to software models such as virtualization, service-oriented architecture (SOA) and cloud and Web services which require concurrent execution of multiple processes.

Application Intelligent Processing Array has extensively explored ways to best utilize the power of multi-core processors with Layer 4-7 application traffic flows. For example, compare two applications - one throughput intensive (file shares) and the other interactive (AJAX). Both have different performance characteristics. One is known for its moderate memory load and scales well with clock speed and cache size. The other is more demanding on memory bandwidth, but is still sensitive to cache size. This is a very common challenge today and SpeedCore automatically tunes itself to applications to achieve the best possible performance.

Dynamic Scalability SpeedCore was built to allow processing units to spin up as necessary to handle variable loads (we use two pairs of ring queues instead of one in our messaging layer to allow for mismatches between the number of user space processing units and the number of network stacks in kernel). This allows SpeedCore to dynamically alter the number of instances that need to be deployed.

Reduced Packet Drops SpeedCore supports (as an infrastructure element) flow control, and every application running over SpeedCore obeys the flow control. As a result, any mismatch between the rate of outgoing packets and the rate of incoming packets is relayed to upstream and downstream entities (clients or servers) to minimize packet loss and buffering.

High Reliability SpeedCore splits up data path processing into modules (e.g. IP, TCP, messaging layer, SSL, compression, HTTP proxy, L4 proxy, caching) that do one thing and do it well (UNIX philosophy). As a result, each module is simplified. Inter-module communication uses explicit message passing (instead of shared memory to improve reliability) except where it interferes with the zero copy property.

Extensibility SpeedCore's extensible architecture offers major performance improvements and easy integration with third-party applications and hardware. For ultra-high performance and low application latency, third-party application and protocol specific hardware (e.g. add-on NIC cards, ASICs, SSL accelerators, TCP accelerators, compression) can be integrated to upgrade performance in very short timeframes.

Adaptable Software Architecture As new applications and Web technologies emerge, SpeedCore's framework can be easily extended to accommodate them. This preserves the adaptability, stability and safety of SpeedCore for easy addition of new features, faster development of new software and faster turn-around time when fixing field problems. As an example, Array handles IPv6 traffic exactly the same as it handles IPv4 – a feat that was achieved rapidly in software without the need to re-spin proprietary hardware or ASICs.

System Management The system's dispatcher always reserves computing capacity on the cores in a private space so that it can always perform management tasks - even under heavy load. This enables the dispatcher to always debug issues when problems occur. The system management module has detailed built-in usage counters to facilitate faster debugging. SpeedCore's enhanced logging and statistics collection helps resolve application and content processing errors for timely field resolution of problems.

White Paper

SpeedCore | Multi-Processing Technology

About Array Networks

Array Networks is a global leader in application delivery networking with over 5000 worldwide customer deployments. Powered by award-winning SpeedCore™ software, Array application delivery, WAN optimization and secure access solutions are recognized by leading enterprise, service provider and public sector organizations for unmatched performance and total value of ownership. Array is headquartered in Silicon Valley, is backed by over 300 employees worldwide and is a profitable company with strong investors, management and revenue growth. Poised to capitalize on explosive growth in the areas of mobile and cloud computing, analysts and thought leaders including Deloitte, IDC and Frost & Sullivan have recognized Array Networks for its technical innovation, operational excellence and market opportunity.

Corporate Headquarters
info@arraynetworks.com
408-240-8700
1 866 MY-ARRAY
www.arraynetworks.com

Belgium
+32 2 6336382

China
support@arraynetworks.com.cn
+010-84446688

France
infosfrance@arraynetworks.com
+33 (0) 180 886 086

India
isales@arraynetworks.com
+91-080-41329296

Japan
sales-japan@arraynetworks.com
+81-45-664-6116

Korea
array-sales@arraynetwork.co.kr
+82(2)3461-8124

Taiwan
support.taiwan@arraynetworks.com
886-2-7718-2750

UK
infoeurope@arraynetworks.com
+44 (0) 7717 153 159



To purchase Array Networks Solutions, please contact your Array Networks representative at 1-866 MY-ARRAY (692-7729) or authorized reseller.

Aug-2013 rev. a

© 2013 Array Networks, Inc. All rights reserved. Array Networks, the Array Networks logo, AppVelocity, NetVelocity, DesktopDirect and SpeedCore are all trademarks of Array Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Array Networks assumes no responsibility for any inaccuracies in this document. Array Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.