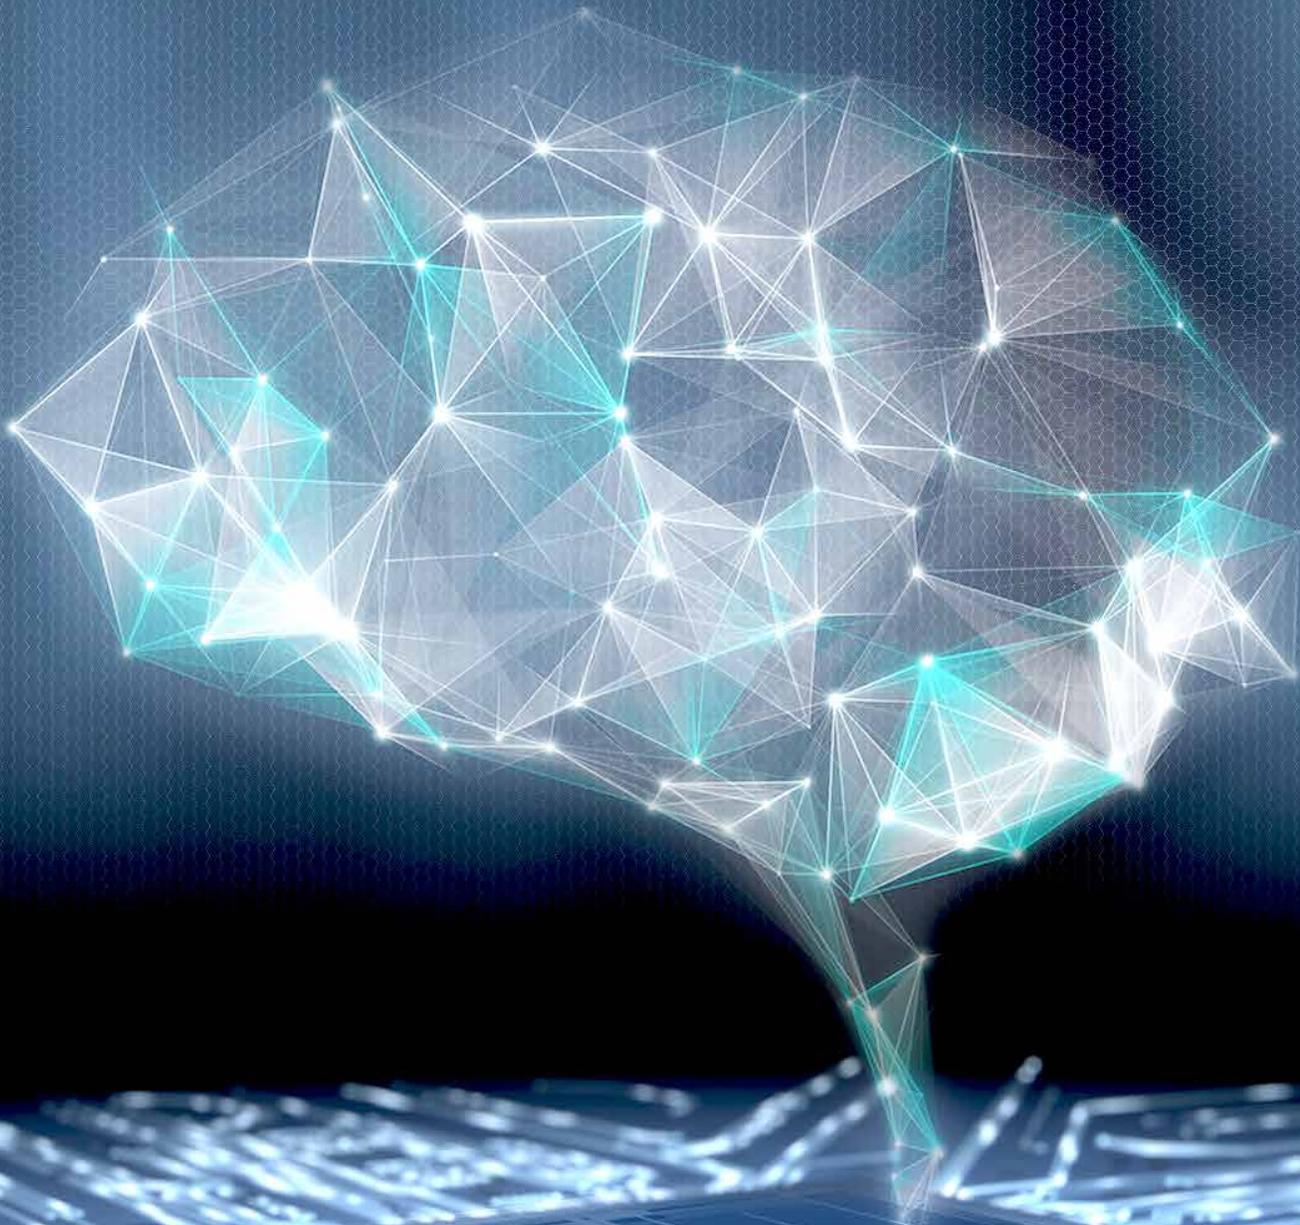
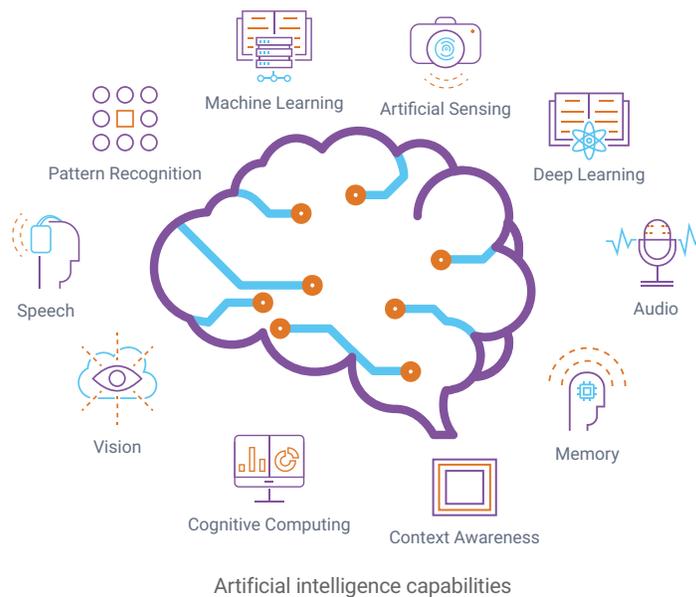


# DesignWare IP for Amazing Artificial Intelligence SoCs



## Amazing AI SoCs Start with DesignWare IP

Recent innovations in deep learning algorithms and neural network processing is driving new technology requirements for artificial intelligence (AI) SoCs. Deep learning capabilities for vision, speech, context awareness, general data pattern recognition and more are being added to SoCs across all markets. Synopsys DesignWare IP enables the specialized processing capabilities, high-bandwidth memory throughput, and reliable high-performance connectivity demands in AI chips for mobile, IoT, data center, automotive, digital home, and other markets. Synopsys' silicon-proven DesignWare® IP portfolio addresses the diverse processing, memory, connectivity, and security requirements of each market.



## Deep Learning SoC Design Challenges

When incorporating deep learning capabilities, SoC designers encounter specific requirements for specialized processing, memory architectures, data connectivity, and security. Specialized processing is required to manage massive and changing compute intensities for machine and deep learning tasks. Memory performance becomes a critical design consideration to support new complex artificial intelligence models. Capacity and bandwidth are primary concerns for training, while inference optimizations create irregular memory access challenges. Real-time data connectivity between sensors, such as CMOS image sensors for vision, and deep learning accelerators become key components. Power consumption is reduced by minimizing data movement between the processor and memory, using key power management features, and designing in advanced FinFET technologies.

## Benefits of DesignWare IP for Deep Learning

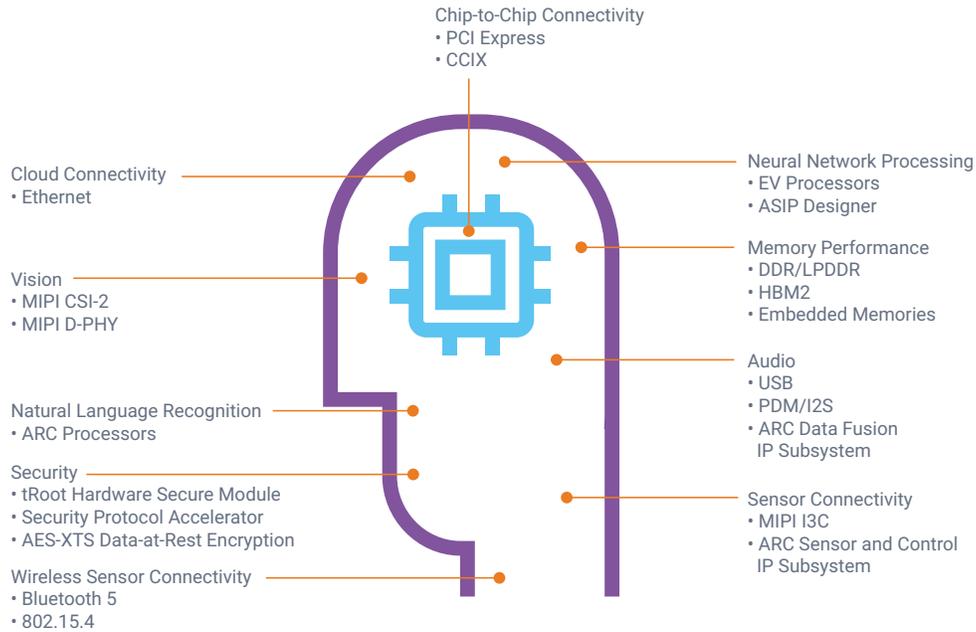
### Specialized Processing

Synopsys provides a portfolio of embedded processors to efficiently execute the varied workloads of AI applications. This includes IP and tools for scalar, vector, and neural network processing. The ARC® EV Processors integrate heterogeneous computing elements optimized for embedded vision applications, including convolutional

neural networks (CNNs). The ARC HS and EM Processors combine RISC and DSP processing capabilities to deliver the best balance of performance, power, and area. ARC's extensible instruction set architecture enables users to add their own instructions or hardware to accelerate AI algorithms and tightly couple memories and peripherals to the processor to reduce system bottlenecks. For custom AI workloads that benefit from a high degree of parallelism and specialized datapath elements, Synopsys' ASIP Designer tool automates the development of custom processors and hardware accelerators.

### Memory Performance

Synopsys provides memory IP solutions to support efficient architectures for different AI memory constraints: bandwidth, capacity, cache coherency. The latest DDR IP addresses capacity needs for data center AI SoCs. HBM2 IP addresses the bandwidth bottleneck while providing an optimized off-chip picojoules (pJ) per bit memory access. CCIX IP enables cache coherency with virtualized memory capabilities for AI heterogeneous compute and reduced latency in AI applications. A wide array of embedded memory compilers enable high density, low leakage, and high performance on-chip SRAMs options including, TCAMs, and multi-port memories.



DesignWare IP for artificial intelligence designs

## Real-Time Data Connectivity

Synopsys provides reliable connectivity to CMOS image sensors, microphones, and motion sensors for AI applications including vision, natural language understanding, and context awareness. The interface IP portfolio in advanced FinFET process technologies reduce power and support a range of widely used standard specifications like MIPI, USB/DisplayPort, HDMI, PCI Express, Cache Coherent Interconnect for Accelerators (CCIX), Ethernet, and more.

## Security

As AI becomes pervasive in computing applications, so too does the need for high-grade security in all levels of the system. Security needs to be integral in the AI process. The protection of AI systems, their data, and their communications is critical for users' safety and privacy, as well as for protecting businesses' investments.

The DesignWare Security IP portfolio includes tRoot Hardware Secure Modules with Root of Trust, which offer diverse security functions in a trusted execution environment (TEE) as a companion to one or more host processors, including secure identification and authentication, secure boot, secure updates, secure debug and key management. DesignWare Security Protocol Accelerators are highly integrated embedded security solutions with efficient encryption and authentication capabilities and provide increased performance, ease-of-use, and advanced security features such as quality-of-service, virtualization, and secure command processing. They are used to protect AI models and enable secure communications to the cloud or other devices. High performance AES-XTS cores protect data-at-rest for runtime memories.

Trained Model	User Private Data	Data Integrity
<p><b>Protect model confidentiality in use and during updates</b></p> <ul style="list-style-type: none"> <li>• Safeguard the large investment to create (money and time)</li> </ul>	<p><b>Ensure user data privacy</b></p> <ul style="list-style-type: none"> <li>• Systems operate on sensitive user data, e.g., facial, biometric</li> <li>• Meet government regulations</li> </ul>	<p><b>Protect input and output data</b></p> <ul style="list-style-type: none"> <li>• Prevent manipulation of coefficients in embedded or external memories</li> <li>• Maintain integrity of data sent from the cloud</li> </ul>

Protecting assets in a neural network is critical to safeguarding company and user information

## AI Training in Data Centers

Specialized Processing	ASIP Designer	Industry-leading tool to design your own scalar, vector, and deep learning processors; deploys hardware parallelism and custom datapaths while retaining programmability.
	Foundation Cores	Flexible primitive math operations including 'dot product' form the building blocks of specialized AI processor designs for several leading AI SoC designs.
Memory Data Throughput	Embedded Memories and Logic Libraries	Efficient on-chip memory configurations via a wide array of embedded memories, TCAMs, logic libraries, and multi-port memory solutions optimized for AI applications; design analysis enables large SRAM arrays for maximizing local memory. Design expertise for optimized cells for density and leakage enable competitive designs.
	HBM2/2e	High bandwidth capabilities with the lowest pJ/bit memory access address the critical bottleneck and cooling costs of data center AI SoCs.
	DDR5	Maximum available capacity for AI training; supports the latest DDR specification.
	CXL and CCIX	Cache coherency and virtualized memory capabilities enable AI heterogeneous compute and reduced latency; Flexible interface support for easy integration with 3rd party, or customer's own coherent fabrics.
	High-speed interfaces	For proprietary accelerator interconnect, high-speed interfaces provide the necessary throughput at 56G and 112G.
Real-Time Data Connectivity	PCI Express 5.0	The best connectivity for Host to Accelerator PCIe offers high speed capabilities that enable AI accelerators to be ready for plug and play interconnect; supporting all the latest PCIe features including 32GT/s for maximum performance.
Security	tRoot HSMS	Manage the overall security in an SoC by providing a high-security grade TEE in which to process sensitive data and operations. Features include secure boot, key management, secure updates, secure debug/JTAG access. Permissions and policies in the hardware root of trust enforce that application layer clients can manage the keys only indirectly through well-defined APIs.
	Security Protocol Accelerators	Provide efficient encryption and authentication for updating AI models, for secure communication, and optionally for protecting data to and from peripherals within the SoC.
	AES-XTS Core	Provides high performance AES-XTS encryption/decryption to protect data-at-rest (memory protection).

## AI Inference in Data Centers, Edge Servers, and Small Cells

Specialized Processing	ASIP Designer	Industry-leading tool to design your own scalar, vector, and deep learning processors; deploys hardware parallelism and custom datapaths while retaining programmability. Includes RISC-V pre-configured setup.
Memory Data Throughput	Embedded Memories and Logic Libraries	Efficient on-chip memory configurations via a wide array of embedded memories, logic libraries, and optimized for AI applications including near threshold voltage options reducing power, and customized cell and array development for the most advanced designs.
	HBM2/2e	High bandwidth capabilities with the lowest pJ/bit memory access address the critical bottleneck of AI SoCs.
	LPDDR5	Supports low voltages and high performance to reduce the power consumption while satisfying the bandwidth needs for AI SoCs.
	CXL and CCIX	Cache coherency and virtualized memory capabilities enable AI heterogeneous compute and reduced latency; Flexible interface support for easy integration with 3rd party, or customer's own coherent fabrics.
Real-Time Data Connectivity	PCI Express 5.0	High speed capabilities enable AI accelerators that are configurable, optimized, and ready for plug and play interconnect; supports all the latest PCIe features including 32GT/s for maximum performance.
	Ethernet	Comprehensive portfolio for 10M through 100G Ethernet applications; supports the latest IEEE specifications.
Security	tRoot HSMS	Manage the overall security in an SoC by providing a high-security grade TEE in which to process sensitive data and operations. Features include secure boot, key management, secure updates, secure debug/JTAG access. Permissions and policies in the hardware root of trust enforce that application layer clients can manage the keys only indirectly through well-defined APIs.
	Security Protocol Accelerators	Provide efficient encryption and authentication for updating AI models, for secure communication, and optionally for protecting data to and from peripherals within the SoC.
	AES-XTS Core	Provides high performance AES-XTS encryption/decryption to protect data-at-rest (memory protection).

Inference at the Edge		
Specialized Processing	Embedded Vision Processors	Verified embedded vision solution including the scalar, vector DSP, and high-performance CNN engine and a suite of software tools supporting frameworks and mapping tools for AI model optimizations.
	ARC Processors	Scalar and vector processing capabilities with APEX extensions, combined with tightly coupled memories for AI acceleration and addressing the key deep learning bottlenecks.
	ASIP Designer	Industry-leading tool to design your own scalar, vector, and deep learning processors; deploys hardware parallelism and custom datapaths while retaining programmability.
Memory Data Throughput	Embedded Memories and Logic Libraries	Efficient on-chip memory configurations via a wide array of embedded memories, logic libraries, near threshold voltage options reduce power, and design analysis enables large SRAM arrays for maximizing local memory.
	LPDDR4/4x and LPDDR5	Supports the latest LPDDR4/4x and LPDDR5 standards, enabling the lowest interface voltage for low power AI inference.
Real-Time Data Connectivity	MIPI CSI-2	Complete camera solution for machine vision; support for multiple image pixel interfaces to merge multiple streams; 1 to 8 RX data lanes with D-PHY Protocol Interface (PPI).
	MIPI D-PHY	Direct CMOS image sensor connectivity; supports v1.2 specification at 2.5Gbps/lane; available in wide range of processes including FinFET.
	MIPI I3C	Multiple sensor integration for context awareness or other deep learning applications; backward compatible with the I2C slave devices at data rates up to 33.4 Mbps; supports master or slave configurations.
Security	tRoot HSMs	Manage the overall security in an SoC by providing a high-security grade TEE in which to process sensitive data and operations. Features include secure boot, key management, secure updates, secure debug/JTAG access. Permissions and policies in the hardware root of trust enforce that application layer clients can manage the keys only indirectly through well-defined APIs.
	Security Protocol Accelerators	Provide efficient encryption and authentication for updating AI models, for secure communication, and optionally for protecting data to and from peripherals within the SoC.
	AES-XTS Core	Provides high performance AES-XTS encryption/decryption to protect data-at-rest (memory protection).

DesignWare ARC Subsystems for Specialized Processing, Memory Performance, and Real-Time Data Connectivity	
ARC Sensor and Control IP Subsystem	Pre-validated, tightly integrated (memories and peripherals) subsystem for significant power savings with AI-specific hardware accelerators (i.e. sensor fusion, motor control).
ARC Data Fusion IP Subsystem	Efficient, low-power solution for AI triggering functions including voice and gesture recognition, face detection; tightly coupled PDM and I2S peripherals simplify integration of external audio devices; audio processing libraries provide building blocks for AI audio applications.
ARC Secure IP Subsystem	Solution for protecting high value targets against wide variety of malicious attacks; high performance cryptography protects AI communications with the cloud; anti-tamper features protect AI edge devices from local attacks.

## Nurturing AI SoCs Goes Beyond the Silicon Design

Adding AI capabilities into SoCs has highlighted weaknesses with today's SoC architectures for AI. Vision, voice recognition, and other deep learning/machine learning algorithms are resource-starved when implemented on SoCs built for non-AI applications. Selecting and integrating blocks of IP clearly determines the baseline effectiveness of an AI SoC, which makes up the "DNA," or nature, of the AI SoC. For example, introducing custom processors, or arrays of processors, can accelerate the massive matrix multiplications needed in AI applications.

However, the element of nurturing the design affects how the pieces function together in hardware or how IP can be optimized for more effective and optimized AI SoCs. The design process to optimize, test, and benchmark the performance of the SoC requires tools and expertise to optimize the AI system. Nurturing the design during the design process with customizations and optimizations can ultimately determine the SoC's success in the market.

Using tools, design expertise, and benchmarking expertise to enhance power consumption, performance, and cost is becoming required to architect a winning SoC architecture. Designers need a wide array of nurturing methods to accelerate their design process and silicon success.

## Tools for Software Development, Verification, and Benchmarking

To nurture the design implementation, designers require advanced simulation and prototyping solutions, such as those provided by Synopsys, with support for early software development, performance validation, and most importantly architectural optimization. These tools are being adopted much more regularly for AI, again due to the immaturity and complexity of the designs.

Benchmarking different AI graphs easily and quickly comes with expertise and established tool chains. Hand writing these graphs for benchmarking activities can be an arduous task but a necessary one to understand if the SoC design can provide the needed value. Relying on processors, such as the DesignWare ARC EV processors, that have tools to benchmark these graphs effectively and quickly can expedite the system design, ensuring it meets your requirements.

## Expertise and Customization

A targeted set of embedded memories can help designers address the challenges of high density and low leakage with customized solutions from Synopsys. Support for near-threshold logic libraries can bring energy harvesting capabilities to significantly lower power for AI accelerators.

Synopsys' Foundation IP portfolio includes an HPC Design Kit, a collection of logic library cells and memories, that has been co-optimized with EDA tools on advanced nodes to push the PPA envelope of any design and has been optimized for AI-enabled designs.

In addition to a rich silicon-validated portfolio that achieves superior PPA, Synopsys' support for customization to all IP titles to meet individual design needs makes the offering more flexible than any other.

Front end design expertise can leverage pre-built AI SoC verification environments by experienced designers. Therefore, design services and companies designing the second and subsequent generation chipsets have inherent advantages over first comers in time-to-market. Designers can rely on the knowledge of Synopsys' experienced designers as an effective way to accelerate time-to-market, freeing up internal design teams to focus on differentiating features of the SoC.

Interface IP hardening offers an additional optimization path for lower power and lower area implementations. Hardened IP makes room on the SoC for valuable on-chip SRAM and processor components needed for better AI performance as well as improving PPA for IP titles such as PCIe, USB, and DDR.

## About DesignWare IP

Synopsys is a leading provider of high-quality, silicon-proven IP solutions for SoC designs. The broad DesignWare IP portfolio includes [logic libraries](#), [embedded memories](#), [embedded test](#), [analog IP](#), [wired and wireless interface IP](#), [security IP](#), [embedded processors](#), and [subsystems](#). To accelerate prototyping, software development and integration of IP into SoCs, Synopsys' [IP Accelerated initiative](#) offers [IP prototyping kits](#), IP software development kits, and [IP subsystems](#). Synopsys' extensive investment in IP quality, comprehensive technical support and robust IP development methodology enable designers to reduce integration risk and accelerate time-to-market.

For more information on DesignWare IP, visit [synopsys.com/designware](https://synopsys.com/designware).