

Agilex™ FPGAs Deliver a Game-Changing Combination of Flexibility and Agility for the Data-Centric World

Leveraging the full breadth of Intel innovation to redefine the FPGA

Author

Martin S. Won

Senior Member of Technical Staff
Intel Corp.

"Intel is broadly enabling heterogeneous integration of computing, memory, and communications. We are connecting and stacking diverse technologies in tiny footprints tuned for specific power envelopes, providing unique cost and performance characteristics with much greater flexibility. Imagine the possibilities as we combine into the most efficient of packages the more diverse capabilities—even data center class technologies—once miles apart in computing terms or based on incompatible processes."

—Dr. Venkata (Murthy) M. Renduchintala, group president, Technology, Systems Architecture & Client Group, and chief engineering officer, Intel Corporation

Table of Contents

Executive summary.....	1
Introduction: the challenge of data proliferation	2
Architecting the ideal solution	3
Agilex FPGA elements	4
Putting it all together	8
Accelerating application development for Agilex FPGAs	9
Sample use cases	9
Conclusion.....	9

Executive summary

Rapid transformation across diverse markets—from edge to core to data center—requires solution providers, application developers, industries, and businesses to deliver massive innovation at unprecedented speed. This is evident at the edge as embedded devices are expected to deliver near-real-time, actionable intelligence; in the network core, with network function virtualization (NFV) to aggregate and process massive amounts of data; and in data centers grappling with increasing analytics, memory, and storage requirements. Essentially, data is inundating infrastructure at each critical point—from edge to core to data center. The rate of change reinforces the need for flexibility—as all market sectors seek to structure and process the data. The new Agilex™ FPGA is more than the latest programmable logic offering—it brings together revolutionary innovation in multiple areas of Intel technology leadership to create new opportunities to derive value and meaning from this transformation from edge to data center.




Figure 1. Intel is enabling diverse markets to optimize and accelerate data processing from the edge to the network to the cloud

Introduction: the challenge of data proliferation

Data is now defining the future of computing technology—as massive data generation combines with a competitive imperative to increase analysis and produce actionable insights. Most, if not all, areas of infrastructure that handle and process data are experiencing significant market transformation. Demands for increased agility and flexibility are increasing, along with greater diversity in data types and new methods and algorithms for deriving value from data. New ways to handle data and demands for new data services are cropping up with increasing frequency, compounding

the complexity for decision-makers, solution providers, and application developers.

For example, practical applications of artificial intelligence (AI) continue to evolve, and it is still unclear exactly how AI can and will be deployed in many areas. This uncertainty means that flexibility will be needed to handle these evolving AI workloads. It is also unclear what level of precision may be needed for many AI applications. To handle different levels of precision, a flexible approach is required.




Figure 2: Massive data generation is rapidly increasing demands across the infrastructure

Edge and embedded

With a wide range of businesses and industries tied to the Internet of Things (IoT), rapidly increasing amounts of data are being created, processed, and transmitted at the edge.¹ This brings an expectation that embedded devices at the network edge will be acting on structured and unstructured data and generating insight in near-real time. Processing the raw, unfiltered data, structuring it, and conducting deep learning inference at the edge demands new levels of performance and capacity.

Network core

As data is ingested and moves from edge through the network core, ever-increasing amounts of traffic must be processed. When network capacity relies on centralized hardware, value-add features can only be deployed where the specific hardware is in place to support them. But today, more intelligence and flexibility are necessary at several key areas at the network level to both transport and process the massive data loads more quickly, and to decentralize functions that were implemented at the equipment level so they can be deployed where needed.

As a result, NFV has emerged as a means of providing the necessary resource flexibility to ensure high utilization and high levels of programmability to cover diverse networking workloads. NFV can be built on commercial off-the-shelf (COTS) infrastructure to provide cost-effective deployment. But virtualized network functions can have limitations in maintaining quality of service (QoS) due to the increasing amount of data and the rise of low-latency, high-bandwidth applications. FPGA acceleration optimizes the utilization and cost-effective scaling of Intel® architecture-based NFV and extends the capabilities of IPU and SmartNICs to accelerate networking as users, data, and applications increase over time.

Also, 5G places new demands on networks—they must be highly available and scale efficiently to accommodate a massive number of connected devices with the exponential rise of data traffic and different quality of service levels generated by different use cases. Current solutions that use specialized appliances will not be sufficient to support future 5G needs. Proprietary hardware may be prohibitively expensive for broad deployments and may not be flexible enough to respond to evolving customer needs, as well as leverage an available developer ecosystem. Here, too, flexibility is needed to accommodate a changing customer landscape and a variety of implementation and deployment options.

Data center

Many market sectors now conduct critical analysis based on data center processing utilizing data sets of ever-increasing size, with use cases ranging from financial analytics, database acceleration, and genomics to video transcoding, network and storage acceleration, and security.

Moreover, the kinds and types of data processing happening within the data center are in constant flux, in part because data center customers are increasing demands to perform different types of data analysis in real time. In addition, data center operators may not know what kinds of data services their customers will want in advance. If data center operators rely on a solution that only accelerates a specific operation or algorithm, they won't have the flexibility to respond to future acceleration demands. Advanced planning is required to handle processing tasks that have not yet been defined—another reason that flexibility at the hardware level (the type of flexibility FPGAs provide) is essential.

Utilizing data center processing resources most efficiently is crucial and acceleration can be instrumental in supporting data-centric compute. Because Intel® Xeon® Scalable processors power many of the world's data centers, there

is a need to accelerate functions that Intel Xeon Scalable processors perform with as low delay and as little latency as possible to deliver peak performance for the most valuable operations, in the most resource efficient way possible.

The need: Flexibility combined with market-specific requirements

As end customers expand their service offerings and solutions, there is a need for highly customizable processing of data wherever it is generated, processed, transported, or stored. Processing at this level requires specific functionality to achieve the optimal power and performance. Each application class across the various markets has specific, unique requirements for data handling. These can include functions that are power- and cost-efficient for small form factors or complex designs where power, space, and cost are at a premium, such as edge and embedded devices or

vehicles; handling the highest data traffic and Ethernet speeds in the network; or providing high-bandwidth, low-latency compute acceleration in data centers.

To meet their specific needs, product developers have the option to use standard off-the-shelf products, or to design and deploy custom ASICs. However, standard products may not precisely fit specific industry requirements or allow sufficient market differentiation. Custom ASICs can perform these specific functions, but are typically time and cost intensive to develop and have high ROI hurdles to be economically viable.

The ideal solution combines the best of two worlds: FPGA flexibility and ASIC-level performance and power efficiency for specific functions.

Architecting the ideal solution

A new architectural approach is needed to address these challenges. For the most advanced members of the Intel Agilex FPGA portfolio, Intel is leveraging a new and disruptive approach to FPGA architecture that creates tailored FPGA products designed to address the unique challenges in each application class.

The ideal solution combines flexibility with maximum power and performance efficiency. To make this possible and build the next generation of programmable logic, Intel's transformative approach to FPGA architecture enables the integration of a wide range of semiconductor elements into a single system-in-package (SiP).

This approach combines a high-performance FPGA core die built using advanced Intel semiconductor process technologies with function-specific chiplets, all integrated heterogeneously into a single product with advanced packaging techniques. This enables Intel to address a broad array of acceleration and other applications with tailored, yet flexible, solutions. The chiplets provide functionality such as PCIe Gen 5, 116 G transceivers, and coherent interfaces (Compute Express Link (CXL)) to Intel Xeon Scalable processors. Other chiplets are also possible, including other transceiver types, custom I/O, and custom compute functions. Leveraging Embedded Multi-Die Interconnect Bridge (EMIB) and other leading-edge proprietary integration and packaging technologies, the new architectural approach allows the combination of traditional FPGA die with purpose-built semiconductor die to create devices that are uniquely optimized for target applications.

The Intel Agilex FPGAs and SoCs enable next-generation, high-performance applications via higher fabric performance, lower power, gains in digital signal processing (DSP) functionality, and higher designer productivity compared to previous-generation FPGAs. The Intel Agilex FPGA meets the myriad challenges of data-centric compute while opening up new possibilities for business and industry. It brings together a general-purpose fabric for flexibility with highly efficient processing at the silicon level for the specific, customized functions demanded by each market.

Custom functions can also be rapidly integrated into Intel Agilex FPGAs through the proprietary and unique Intel® eASIC™ device technology. With Intel eASIC technology, customer FPGA designs can be converted into function-specific die that provide ASIC-level performance and power efficiency, integrated into a single component package along with other functions for advanced customization.

Flexibility combined with agility to meet target application requirements

FPGAs have been highly valued for the flexibility to meet evolving market requirements. However, there are some functions that have stabilized and no longer require as much flexibility. It is desirable to “harden” these functions as much as possible to get the most power efficiency and performance. Hardening these functions produces these benefits because added flexibility always comes with a necessary trade-off in lower power efficiency and higher power.

In the past, FPGAs were designed mainly using a single monolithic die or multiple instances of the same monolithic die type. Today, Intel is using new advanced packaging technologies to supplement the monolithic approach to combine multiple, disparate silicon die within a single package. By integrating die from different process types and functions, Intel offers unprecedented flexibility and customization. Examples of purpose-built die include:

- Interfaces for low-latency, cache-coherent processor acceleration
- Advanced analog functions like 116 G transceivers (XCVRs) and data converters
- Custom compute engines for application-specific functions
- Memories of different types and configurations that can be closely coupled to the logic fabric




Figure 3. Intel brings together unique architectural innovation across key areas in the most advanced members of the Intel Agilex™ FPGA portfolio

TARGETED OPTIMIZATIONS TO MEET MARKET NEEDS ACROSS THE COMPUTE SPECTRUM

Edge and embedded	Network core	Data center
Power area-efficient, AI inference	High-speed transceivers up to 116 Gbps	Low latency Intel® Xeon® Scalable processors acceleration
Custom data preprocessing and ingest	Up to 400 G Ethernet blocks	Power-efficient DSP blocks for AI and other algorithm acceleration

Intel Agilex FPGA elements

Advanced FPGA fabric

The FPGA fabric at the heart of every Intel Agilex FPGA is built using advanced Intel manufacturing process technologies, including Intel 10nm SuperFin and Intel 7.

The fabric leverages the second generation of Intel® Hyperflex™ FPGA Architecture, which uses registers, called Hyper-Registers, throughout the FPGA, optimized for leading performance. The second generation of Intel Hyperflex FPGA Architecture, combined with Intel® Quartus® Prime Software, delivers the optimized performance and productivity required for next-generation systems. In fact, when compared to our competition's 7 nm FPGA portfolio, Intel Agilex FPGAs deliver ~ 2X better fabric performance per watt.¹

The FPGA fabric also features architecture optimizations for accelerating AI functions and DSP operations through dedicated structures for half-precision floating point (FP16) and BFLOAT16, as well as increased DSP density compared to prior generation FPGAs.

Intel Agilex FPGAs can implement fixed-point and floating-point DSP operations with high efficiency. The DSP blocks provide 2X the number of 9x9 multipliers compared to the prior generation. This also doubles the amount of INT8 operations that Intel Agilex FPGAs can deliver per DSP block. The addition of new modes for FP8 and FP16 supports highly efficient implementations for specific AI workloads, such as convolutional neural networks (CNNs) for image and object detection with a lower device utilization and lower power compared to implementation with FP32.

In the Intel Agilex FPGAs targeted for intelligent applications at the edge, the fabric features Enhanced DSP with AI Tensor Blocks, an architectural element that increases the

INT8 compute density by more than an order of magnitude compared to prior generations, which is valuable for AI inferencing applications.

2nd Generation Intel Hyperflex FPGA Architecture

The innovative second generation of Intel Hyperflex FPGA Architecture supports levels of performance not possible with conventional architectures.

Like the first generation, the 2nd generation Intel Hyperflex FPGA Architecture employs additional registers, called Hyper-Registers, everywhere throughout the core fabric. These registers are available across the routing structures and at the inputs of all functional blocks. The Hyper-Registers provide a fine-grained solution to the problem of how to increase bandwidth and improve area and power efficiency. When Hyper-Registers are used to implement these techniques, all other FPGA logic resources are available for logic functions instead of being sacrificed as feed-through cells to reach conventional LUT registers.

In the second generation Intel Hyperflex FPGA architecture, several advances have been made to improve overall fabric performance while minimizing power consumption. One of the most significant improvements is the addition of a high-speed bypass to the Hyper-Registers, as shown in Figure 4.

On the left of Figure 4 is a representation of an Intel® Stratix® 10 FPGA Hyper-Register. You can see that there is a signal path that goes through the register and another signal path that bypasses it. Both signal paths go through a mux, which is controlled by configuration RAM. One of the ways Intel has improved Intel Hyperflex FPGA Architecture in the second generation is by accelerating the speed of the Hyper-Register bypass path. This improvement increases performance for both Intel Hyperflex FPGA Architecture-optimized designs and designs that are not optimized for Intel Hyperflex FPGA Architecture.




Figure 4. A high-speed bypass accelerates Hyper-Registers to improve fabric performance

In Figure 5, we see two design examples. The one on top is optimized for Intel Hyperflex FPGA Architecture; the one on the bottom is not. The adaptive logic modules (ALMs) in the Intel Agilex FPGA are shown in the large, light-blue boxes, and the Hyper-Registers are shown in two colors: orange for the unused Hyper-Registers, and blue with a gray outline for the used ones.

As you can see, the top design has used Hyper-Registers, indicating that it has been optimized for the Intel Hyperflex FPGA Architecture, whereas the bottom one has no used Hyper-Registers, indicating that it is not optimized for the Intel Hyperflex FPGA Architecture.




Figure 5. The top design is optimized for Intel® Hyperflex FPGA Architecture; the bottom one is not

In Figure 6, we see the signal delays from register to register in each design, which determine the critical path and, ultimately, the fmax of each design.




Figure 6. Signal delays between registers determine the fmax of each design

You can see in Figure 7 that adding the Hyper-Register high-speed bypasses and ALM performance improvements accelerates any potential critical path in either of the two design types. The signal travels faster through each of the unused Hyper-Register (the orange boxes) and also faster through ALMs where there are no synchronous clears and clock enables. In this way, the 2nd Generation Intel Hyperflex FPGA Architecture delivers higher performance for all designs, whether they are optimized for Intel Hyperflex FPGA Architecture or not; although those that are optimized will receive the greatest performance benefit.




Figure 7. Adding Intel® Hyperflex™ FPGA Architecture accelerates the critical path for both design types

Overall, the improvements to the 2nd generation of Intel Hyperflex FPGA Architecture contribute significantly to the performance benefits delivered by Intel Agilex FPGAs.

DSP acceleration functions

To meet demands for higher-precision signal processing, Intel developed the industry's first variable-precision DSP block. In prior generations of Intel® FPGAs, variable-precision DSP blocks were enhanced to include single-precision floating point (FP32) support. In Intel Agilex FPGAs, variable-precision DSP blocks have been enhanced to include support for half-precision floating point (FP16) and BFLOAT16 (BF16). With thousands of floating-point operators available via these hardened DSP blocks, Intel Agilex FPGAs provide up to 38 TFLOPs FP16 or BF16, or up to 19 TFLOPs FP32 DSP performance.

Development advantages of chiplet-based architecture

Historically, FPGAs have integrated functions monolithically, adding individual functions like XCVRs, memory, and I/O into the same die as the programmable logic. However, this approach has three distinct limitations. First, it requires all the individual functions on the FPGA to be developed on the same process. Second, it requires a new tape-out whenever any of the individual functions are updated or improved. And third, it requires a new tape-out to get different resource mixes and ratios of individual functions (i.e., more or fewer transceivers with the same amount of logic). Tape-outs are time consuming and expensive, so they need to be minimized to accelerate the pace of product innovation.




Figure 8. Traditional FPGA update schedule vs. chiplet-based updates

A better, more versatile approach can be achieved with Intel chiplets. A chiplet or tile is a physical intellectual property (IP) block designed to integrate with other chiplets through package-level integration and standardized interfaces. Like LEGO blocks, chiplets enable mix-and-match infrastructure. Chiplets allow new products to be created with greater functionality, improved agility, and shorter time to market, as shown by Intel FPGA early support for 116 G XCVR, CXL, and PCIe Gen5 technology. This advantage is shown in Figure 8.

The chiplets are built on the optimal semiconductor process for each function. The chiplets interface with the Intel Agilex FPGA fabric and can be mixed and matched to produce the exact combinations needed by a specific application class.

The manufacturing process utilizes advanced packaging technology that acts as a substrate, allowing the multiple chiplets to be assembled into a single package. For example, multiple cores, I/O, FPGAs, analog, network and communication, and memory accelerators can be integrated via chiplets. Third-party ecosystem technologies could be combined with Intel silicon.

Transceivers with high-speed I/O allow data to get on and off the chip very quickly. The high-speed connectivity allows

processing of massive amounts of data. With the Intel Agilex FPGA generation, Intel has increased speed from 58 gigabits/sec (58 G) to 116 G with a single transceiver channel. With chiplets, the number of transceivers is not limited by the amount of channels that can be floorplanned into the FPGA fabric die. The number of transceiver channels can be increased or decreased simply by using the desired amount of transceiver chiplets, rather than re-floorplanning the die to integrate a different number of channels.

There is a general industry perception is that smaller is better—hence the race to 14 nm, 10 nm, and so on. Sometimes, there are reasons *not* to use the latest semiconductor process for some functions. For example, some analog and memory functions do not scale as well as general-purpose logic with process geometry. Also, some existing silicon IP functions may have been developed on older process nodes, where the benefit to move to the new node does not outweigh the cost and time to do the migration. With Intel's chiplet methodology, each function can be built with the specific semiconductor process ideal for its implementation. Chiplets can then be mixed and matched to achieve the right solution for a specific market need.


Advanced packaging

To put everything together in a single package, you need technology that can cost-effectively integrate different die from different functions. Traditionally, one silicon die is used in a package. Now with advanced Intel technology, it is possible to put multiple die together from different process nodes into the same package. For example, one of the advanced packaging techniques used in Intel Agilex FPGAs is Embedded Multi-Die Interconnect Bridge (EMIB), which provides a high-performance, cost-effective means of integrating chiplets together with the FPGA die in the same package, as shown in Figure 9.

These latest innovations in Intel packaging technology enable the integration of heterogeneous elements together into a single package. Purpose-built silicon addresses specific market needs and a flexible core fabric joins the die together—resulting in a massively integrated system.

With this kind of integration, Intel is coupling the FPGA fabric along with purpose-built silicon to build targeted products for specific application spaces. The resulting architecture meets the needs of many different use cases and market sectors and can work for specific functions, purpose-built for a particular space.

Figure 9.
Flexible,
heterogeneous
packaging of multiple
function-specific die



High-performance processor interfaces

With additional cores per processor and the much higher memory and I/O bandwidth in Intel Xeon Scalable processors, the increased demands on the intra-chip interconnect can become a performance limiter.

To facilitate high-performance in-line and offload acceleration of processor functions, Intel Agilex FPGAs support the latest generation of high-performance processor interfaces, including PCIe Gen5 and CXL.

Cache- and memory-coherent interface to Intel Xeon Scalable processors via CXL

Intel Agilex FPGAs are the first FPGAs to integrate hard IP that supports CXL, which is a high-bandwidth, low-latency cache- and memory-coherent interface between CPUs and workload accelerators. CXL technology maintains memory coherency between the CPU memory space and memory on attached devices, which allows resource sharing for higher performance, reduced software stack complexity, and lower overall system cost. This feature enables users to simply focus on target workloads as opposed to the redundant memory management hardware in their accelerators. More information about CXL is available at computeexpresslink.org.

PCIe Gen5 interface

Intel Agilex FPGAs are the first to support PCI Express interfaces up to Gen5 x16. PCIe Gen5 brings a leap in transfer speeds, enabling greater performance capabilities to developers of PC interconnect, graphics adapters, and chip-level communications, among other applications using this ubiquitous technology.

Advanced memory hierarchy

Memory capacity and bandwidth can be critical bottlenecks for next-generation platforms. Computing elements, such as FPGAs or CPUs, operate with data types of varying sizes and access requirements. This means that various memory types and sizes are required to facilitate the most efficient and high-performance processing.

Intel Agilex FPGAs support a broad hierarchy of memory resources, including embedded memory resources, in-package memory, and off-chip memory via dedicated interfaces. The hierarchy is shown in Figure 10.

The hierarchy begins with embedded, on-chip memory, including MLABs, block RAM, and eSRAM, each of which offer different capacities to accommodate different processing requirements.

High Bandwidth Memory (HBM) is the next generation of high-speed memory built into Intel Agilex FPGAs using SIP technology. HBM enables the highest memory bandwidth, not feasible with other solutions. Multiple DRAM layers are connected to a base I/O layer to form 3D, high-speed memory connected to and controlled directly by hard memory controllers built into the Intel Agilex FPGA. Integrating the HBM die directly into the FPGA package reduces board size and cost, and simplifies and reduces power requirements.

Some Intel Agilex FPGAs integrate HBM2E memory in-package next to the core fabric. The interconnect between

the core fabric and memory is significantly shorter, compared to discrete memories like DDR, which reduces the amount of power traditionally spent driving long printed circuit board (PCB) traces. The traces are unterminated and there is reduced capacitive loading, which results in lower I/O current consumption. The net result is lower system power and optimum performance per watt. Finally, Intel Agilex FPGAs include interfaces to memory components external to the device, including DDR4/5, QDR, and RLDRAM, which provides a new, breakthrough option for FPGA users.




Figure 10. Intel Agilex™ FPGAs support an advanced memory hierarchy—supporting on-chip, in-package, and onboard memory

Putting it all together

Intel FPGA and SoC innovations combine to enable unprecedented levels of customization and flexibility in Intel Agilex FPGAs. The Intel Agilex FPGA fabric and innovative chiplet architecture—enabled by Intel's EMIB packaging technology for integration of a variety of chiplets within a single system in package—delivers an extensible FPGA platform that scales across a wide range of device densities, and brings these key features and benefits:

- The latest generation of high-bandwidth processor interface interconnect including PCIe Gen5 x16 support with data rates of 32 GT/s per lane for next-generation data centers.
- High-performance processor interface speeds for acceleration of cache-coherent operations via dedicated support for CXL.
- Architecture optimizations for accelerating AI functions and DSP operations through dedicated structures for half-precision FP16, BFLOAT16, as well as increased DSP block density, delivering up to 38 TFLOPs FP16/BFLOAT16 or up to 19 TFLOPS FP32 DSP performance.
- 116 Gbps serial transceiver links to support the most demanding bandwidth requirements in next-generation data center, enterprise, and networking environments.
- Rapid integration of customer IP with Intel eASIC device options for better cost and power efficiency.
- 2nd generation Intel® Hyperflex™ FPGA Architecture delivering **50% geometric mean (GM) higher performance or up to 40% lower power** (vs. Intel® Stratix® 10 FPGAs)¹
- When compared to our competitors 7 nm FPGA portfolio, Intel Agilex FPGAs delivers ~ 2X better fabric performance per watt¹
- **Chiplet integration capability** leveraging purpose-built semiconductor chiplets ranging from XCVR tiles to processor interfaces to deliver targeted capabilities for specific applications and market segments
- Increased memory capacity and **support for advanced memory types like DDR5, LPDDR4/5 SDRAM and HBM2e**
- Integrated hard processor options including: dual-core Arm A76, dual-core Arm A55, and quad-core Arm A53 enabling **system intelligence and flexible, embedded functionality** for a broad range of applications

Accelerating application development for Intel Agilex FPGAs

Design development support for Intel Agilex FPGAs and SoCs is provided via multiple options. The first is with Intel Quartus Prime Software tools. The Intel Quartus Prime Software includes everything you need to design for Intel FPGAs and SoCs, from design entry and synthesis to optimization, verification, and simulation. Dramatically increased capabilities on devices with multimillion logic elements are providing designers with the ideal platform to meet next-generation design opportunities.

In addition, Intel is opening up FPGA development access to software developers starting with the Intel® Distribution of OpenVINO™ toolkit, and continuing with oneAPI toolkit. Intel's oneAPI development environment simplifies the programming of diverse computing engines across CPU, graphics processing unit (GPU), FPGA, AI, and other accelerators. The oneAPI platform includes a comprehensive, unified portfolio of developer tools for mapping software

to the hardware that can best accelerate the optimized applications, middleware, and frameworks.


Intel® FPGA Add-on for oneAPI Base Toolkit will allow more designers to take advantage of Intel Agilex FPGAs. With oneAPI toolkit, developers can explore different implementations of a design across diverse hardware choices, so they can understand the particular power and performance capabilities of each implementation and ultimately choose the best fit for their specific needs.

Solution Components

Intel Agilex™ FPGA
Intel® Xeon® Scalable processors
Intel® Distribution of OpenVINO™ toolkit
Intel® Hyperflex™ FPGA Architecture
Intel® Quartus® Prime Software

Sample use cases

Intel Agilex FPGAs target the acceleration and compute needs for a wide range of applications from the edge to the core to the cloud—customized for specific market segments and applications—wherever data is generated, transported, stored, or processed.



Conclusion

The design and architectural innovation in the Intel Agilex FPGAs and SoCs—bringing together key areas of deep Intel expertise—are poised to create powerful opportunities from edge to network to cloud, delivering critical and previously unattainable levels of performance, accuracy, flexibility, and customization.

Learn more

To learn more about Intel Agilex FPGAs, go to intel.com/agilex.

For more information about Intel FPGAs, visit intel.com/fpga.



¹ Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.