



ALVEO™ PORTFOLIO PRODUCT SELECTION GUIDE

Accelerator Card	Network Interface	Optical Interface	Form Factor	Logic Density	DSP	DDR4	HBM	AI Engines	PCIe®	Max Power ¹
Alveo U30	-	-	HHHL	N/A	N/A	8 GB	-	-	Gen3x8, 2x Gen3x4	75W
Alveo MA35D	-	-	HHHL	N/A	-	-	-	-	Gen4x8	50W
Alveo U45N	2x 100G	2x QSFP28	FHHL	1M LUTs	1,320	8 GB	-	-	Gen4x8, Gen3x16	150W
Alveo U50	1x 100G	1x SFP28	HHHL	872K LUTs	5,952	-	8 GB	-	2x Gen4x8, Gen3x16	75W
Alveo U55C	2x 100G	2x QSFP28	FHHL	1.3M LUTs	9,024	-	16 GB	-	Gen3x16, 2x Gen4x8	150W
Alveo V80	4x 200G	4x QSFP56	FH ³ / ₄ L	2,574K LUTs	10,848	32 GB	32 GB	-	Gen4x16, 2x Gen5x8	300W
Alveo X3522PV	2x 10/25G	2x DSFP	HHHL	1M LUTs	1,320	8 GB	-	-	Gen4x8, Gen3x8	75W
Alveo UL3524	32x 10/25G	4x QSFPDD	FH ³ / ₄ L	787K LUTs	1,680	16 GB	-	-	Gen4x8	125W
Alveo V70	-	-	HHHL	N/A	N/A	16 GB	-	404 TOPs ²	Gen4x8, Gen5x8	75W

1. The maximum power represents the total electrical limit of the card. The thermal design power (TDP) of the Alveo card may vary depending on factors such as workload, inlet temperature, CFM/LFM airflow conditions, and the card's placement within the server. For specific TDP values, please refer to the associated Alveo accelerator card data sheet.

2. Alveo V70: INT8 Peak Performance with 50% Sparsity; <https://www.xilinx.com/applications/data-center/v70.html>


Our compute accelerators are among the most versatile cards in the portfolio, targeting big data workloads with large data sets and those demanding massive parallelism, including data analytics, HPC, network switching, computational storage, blockchain, and various applications in FinTech including algorithmic trading and backtesting. The [Alveo U50](#) is our smallest form factor HBM-based card for memory bound compute, while the [Alveo U55C](#) offers a balanced mix of HBM memory, logic density, and DSP resources in an FHHL form factor. The [Alveo V80](#) is powered by the Versal HBM adaptive SoC and delivers the highest logic density, HBM bandwidth, network bandwidth, and DSP resources in the Alveo portfolio—ideal for the most compute-intensive big data workloads.



Alveo™ U50



Alveo U55C



Alveo V80

Network Interface	1x100 Gb/s	2x 100 Gb/s	4x 200 Gb/s
Form Factor	HHHL	FHHL	FH¾L
Logic Density	872K LUTs	1,304K LUTs	2,574K LUTs
DSP Resources	5,952 Slices	9,024 Slices	10, 848 Slices
DDR4 Memory	-	-	32 GB
HBM Memory	8 GB	16 GB	32 GB
PCIe®	Gen4x8 or Gen3x16	Gen4x8 or Gen3x16	2x Gen5x8 or Gen4x16

Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo

Our portfolio of network accelerators offload server CPUs from infrastructure tasks in the data center, accelerating workloads such as virtual switching, load balancing, virtual firewall, and computational storage. The [Alveo U45N](#) is versatile for diverse networking applications, supported by the [AMD OpenNIC](#) reference design for ease of development of custom solutions. The [Alveo U55C](#) compute accelerator card can serve as a load balancer or Smart Top-of-Rack switch with the ability to rapidly process large flow tables in HBM.



Alveo U45N




Alveo U55C

Network Interface	2x100G	2x100G
Form Factor	FHHL	FHHL
Logic Density	1,030K LUTs	1,304K LUTs
Processor	Discrete 16-Core Arm Cortex-A72	–
DDR4 Memory	12 GB	–
HBM Memory	–	16 GB
PCIe®	Gen4x8 or Gen3x16	Gen4x8, Gen3x16
Target Workloads	<ul style="list-style-type: none"> • OVS • IPsec • 5G UPF 	<ul style="list-style-type: none"> • Load Balancer • Smart Top-of-Rack (ToR) Switch • Transport Layer Security (TLS)

Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo

Our FinTech accelerators address the unique requirements of the financial market, which often demands a convergence of low latency networking and hardware acceleration for algorithmic trading, pre-trade risk analysis, and market data delivery services at nanosecond speeds. The [Alveo X3522PV](#) and [Alveo UL3524](#) target low-latency and ultra-low latency trading, respectively, for algorithms that prioritize trade execution performance vs. algorithmic complexity. The [Alveo U55C](#) compute accelerators offer more resources for complex trading algorithms, risk & price modeling, as well as data analytics, while the [Alveo V80](#) offers the most compute and bandwidth capability in the portfolio, ideal for financial modeling, analytics, and strategy backtesting.




Alveo™ X3522PV



Alveo UL3524



Alveo U55C



Alveo V80

Network Interface	4x 10/25G	32x 10/25G	2x 100 Gb/s	4x 200 Gb/s
Form Factor	HHHL	FH ³ / ₄ L	FHHL	FH ³ / ₄ L
Logic Density	1,030K LUTs	787K LUTs	1,304K LUTs	2,574K LUTs
DSP Resources	1,320 Slices	1,680 Slices	9,024 Slices	10, 848 Slices
AI Engines	–	–	–	-
DDR4 Memory	8 GB	16 GB	–	32 GB
HBM Memory	–	–	16 GB	32 GB
PCIe®	Gen4x8 or Gen3x8	Gen4x8	Gen4 x8 or Gen3 x16	2x Gen5x8 or Gen4x16
Expansion	–	32x 10/25G ARF6	–	2x Gen5x4, 1x Gen5x8 MCIO
Target Workloads	Low Latency Trading	Ultra-Low Latency Trading	Compute & Analytics	Algo Trading

Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo

XMP451 (v2.0)

Our portfolio of Alveo™ media accelerators is enabling a new era of interactive media and live streaming applications that demand a high volume of concurrent streams—including video collaboration, watch parties, live shopping, social streaming, sports betting, online auctions, and more. The portfolio delivers high streaming density at low latency by accelerating the full video pipeline, from decoding, to adaptive-bitrate (ABR) scaling, to encoding.

The [Alveo U30](#) media accelerator provides low-latency live streaming (as low as 32 ms 4Kp60 transcode) and supports the mainstream H.264 and H.265 codec standards. In addition to on-premise deployment, the Alveo U30 card is available on the AWS cloud ([Amazon EC2 VT1 instances](#)) and through [Avnet co-location services](#) for evaluation or production deployment. The accelerator card also supports the [Wowza Streaming Engine](#) for ease of bring-up and deployment without the need for video frameworks.

The [Alveo MA35D](#) media accelerator provides ultra-low latency streaming (as low as 8 ms 4Kp60 transcode) and supports the AV1 codec standard for state-of-the-art compression, as well as the H.264 and H.264 standards. The 5 nm ASIC-based technology offers breakthrough preprocessing and postprocessing engines for next-generation use cases, including cloud gaming services.

Both accelerators are powered by a full video SDK software stack, supporting the familiar FFmpeg and GStreamer frameworks, as well as a C-API for further customization. The [Alveo U30 video SDK](#) and Alveo MA35D [AMD Media Acceleration SDK](#) are both available on GitHub.



Alveo™ U30



Alveo MA35D

PLATFORM		
Max Resolution	4K	8K
Transcode Latency (4K)	32 ms	8 ms
Video Codec	H.264, H.265	H.264, H.265, AV1
PCIe®	Gen3	Gen4
Typical Power	25W	30–45W
HHHL, Single Slot	✓	✓
CHANNEL DENSITY		
1080p60 (H.26x / AV1)	8 / -	16 / 32
4Kp60 (H.26x / AV1)	2 / -	4 / 8
8Kp30 (H.26x AV1)	-	4
HARDWARE ENGINES		
ABR Scaler	✓	✓
Look-Ahead Engine	✓	✓
AI Processor		✓
Compositor Engine		✓
VQ QoE Engine		✓
DEPLOYMENT		
On-Premise	✓	✓
Cloud Instances and Co-Location Availability	✓	
Wowza Streaming Engine Support	✓	

The Alveo™ V70 production card feature [AI Engines](#) for breakthrough AI compute as high as 404 INT8 Peak TOPS.¹ The card is powered by [Versal™ adaptive SoCs](#).

The [Alveo V70](#) card comes in a slim HHHL form factor and is tuned for video analytics and natural language processing applications. With industry-standard framework support, AI developers can directly compile models trained in TensorFlow and PyTorch.



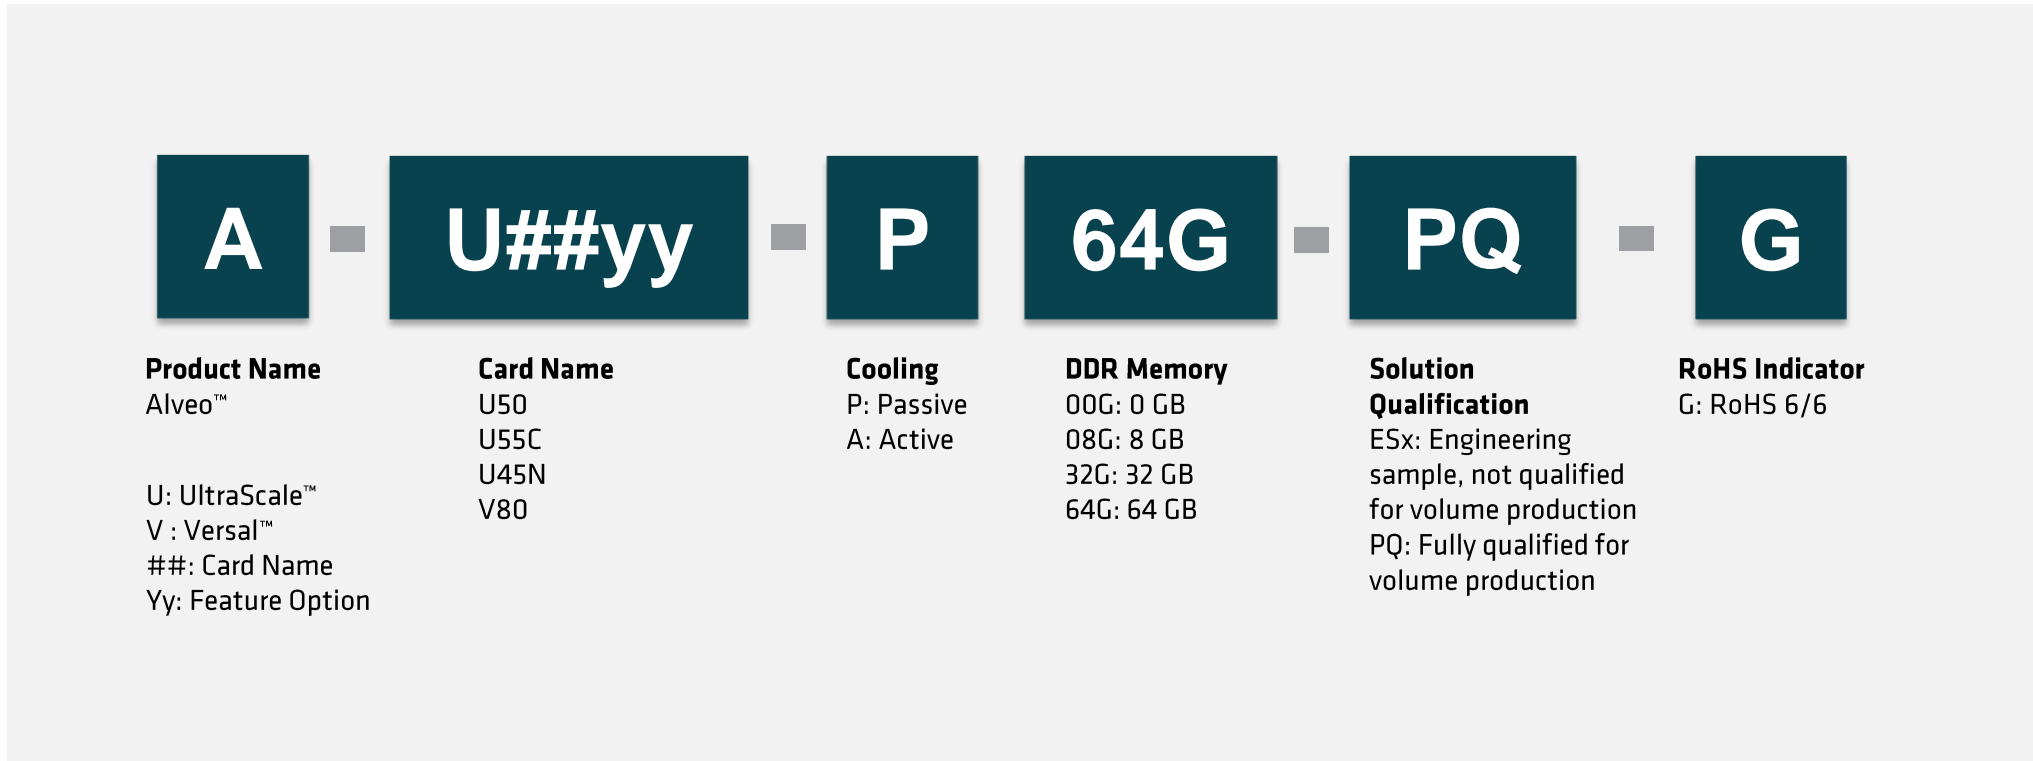
Alveo V70

AI Compute Technology	AI Engines-ML
INT8 TOPS	404 TOPS ¹
Network Interface	–
Form Factor	HHHL
Logic Density	N/A
DSP Resources	N/A
Video Decoder (H.264, H.265)	96 channels, 1920x1080p ²
DDR4 Memory	16 GB
PCIe®	Gen4x8, Gen5x8
Vivado™ Tool Support	–
Vitis™ Tool Support	–
Vitis™ AI Tool Support	✓

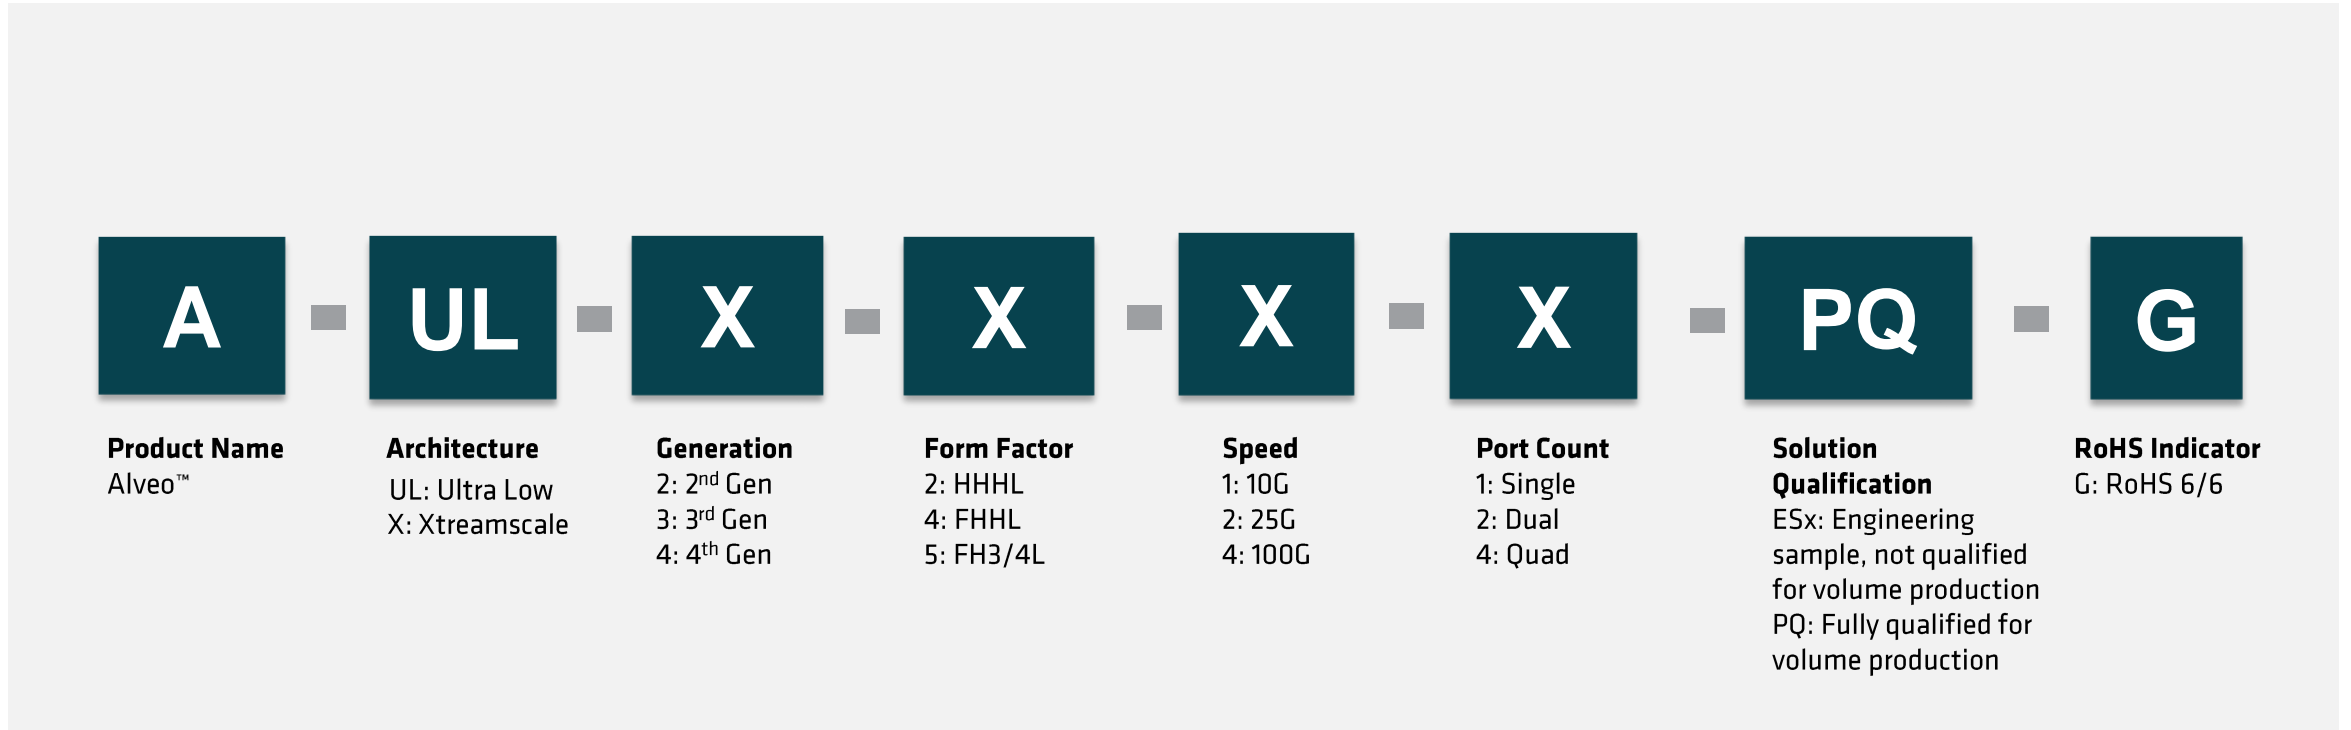
1: INT8 peak performance with 50% Sparsity: <https://www.xilinx.com/applications/data-center/v70.html>

2: 10fps, H.264/H.265

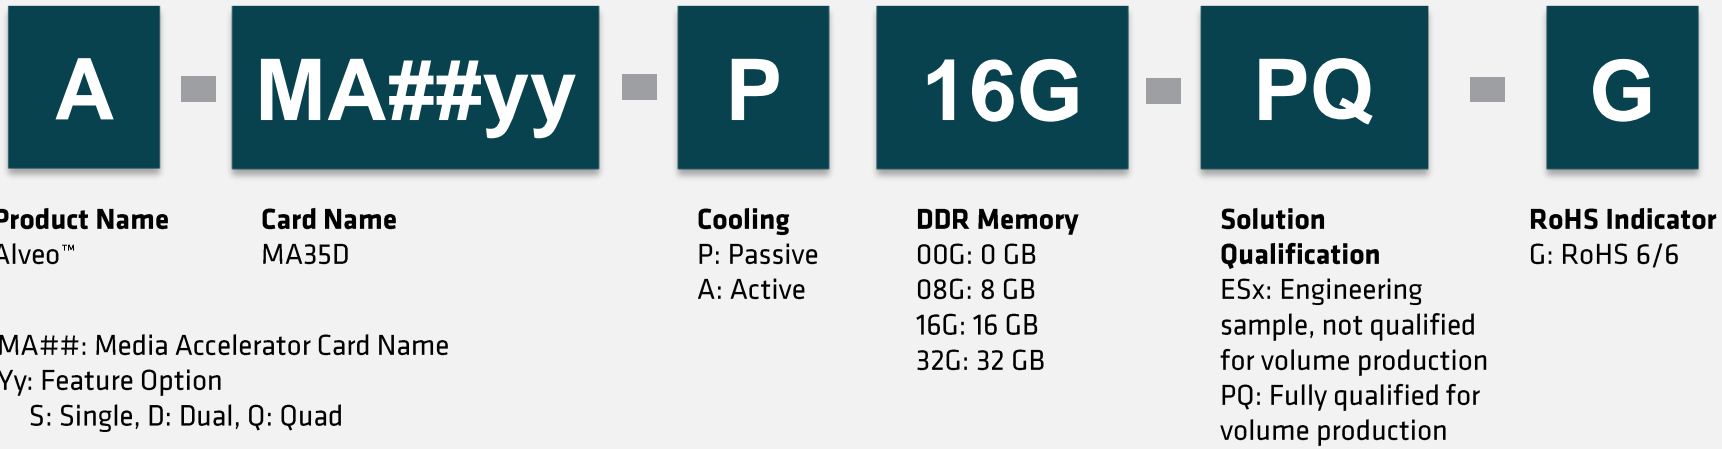
Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo



Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo



Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo



Verify all data in this document with the device data sheets or product guides found at www.amd.com/alveo

TAKE THE NEXT STEP

- To buy online or learn more, go to www.amd.com/alveo
- To enquire about volume purchase and lead times, contact your local sales representative or complete the [Alveo Inquiry Form](#).
- To start designing for Alveo™ accelerator cards using the Vivado™ Design Suite, visit www.amd.com/vivado-for-alveo
- Search through our partner network of OEM qualified servers at <https://www.xilinx.com/products/boards-and-kits/alveo/qualified-servers.html>

**DISCLAIMERS**

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

COPYRIGHT NOTICE

© Copyright 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Alveo, Versal, Vivado, Vitis, Vitis AI and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. PCIe, and PCI Express are trademarks of PCI-SIG and used under license. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. PID1978700

XMP451 (v2.0)