

WOLF Advanced Technology

COMPARATIVE ANALYSIS OF RUGGED WOLF XMC MODULES FOR EMBEDDED SYSTEMS

– WHITEPAPER

Written by: Tahshina Afridi, Technical Writer, Engineering and Marc Henney, Product Manager, Product Management



INTRODUCTION

This paper evaluates and draws a comparison between four of WOLF Advanced Technology's ruggedized XMC-GPU modules, including the [WOLF-3476 \(XMC-A2000E\)](#), [WOLF-3576 \(XMC-AD2000E\)](#), [WOLF-3696 \(XMC-BW500E\)](#), and [WOLF-3676 \(XMC-BW2000E\)](#).

All modules are ruggedized and defense-ready, but each serves a different class of compute mission. The analysis explores feature comparisons, ruggedization, compute capability, target applications, and integration considerations, with special emphasis on the evolution between the two Blackwell products as mission processing becomes increasingly AI-centric.

While the core XMC form factor and ruggedized design remain consistent across these modules, the primary differentiator lies in the updated GPU architectures and the integration of the latest software and firmware packages. These advancements include support for cutting-edge technologies, such as 9th-generation NVENC and NVDEC, updated DisplayPort compatibility, PCIe Gen 5 support, up to GDDR7 memory with higher bandwidth and ECC (module dependent), and the latest API compliance (Vulkan 1.4, OpenGL, and DirectX), ensuring enhanced performance, improved efficiency, and compatibility with the most current industry standards. This approach allows the boards to maintain their proven hardware reliability while significantly expanding their compute capabilities and enabling seamless integration into modern, AI-driven mission environments.

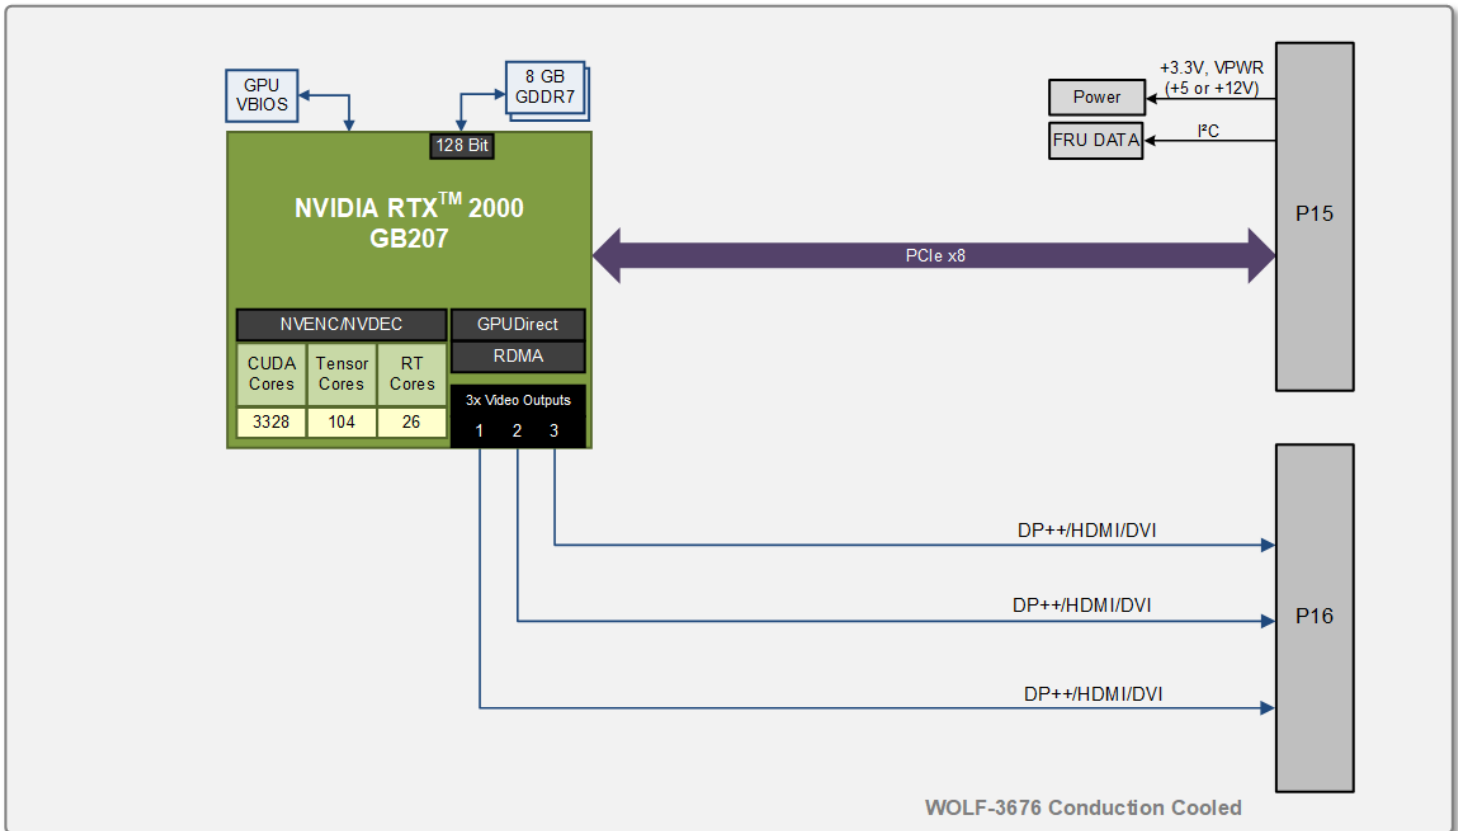


Figure 1: Block Diagram of WOLF-3676 XMC-BW2000 Module

AT-A-GLANCE COMPARISON

Each module features distinct GPU architectures and performance characteristics tailored to a variety of embedded and rugged computing applications.

WOLF-3476 is built on the Ampere architecture, while WOLF-3576 leverages the more advanced Ada Lovelace architecture. WOLF-3696 and WOLF-3676 introduce the latest Blackwell architecture, offering enhanced efficiency and computational capabilities for next-generation workloads.

Core specifications show progressive scaling across modules. CUDA core counts range from 1,792 on the WOLF-3696 to 3,328 on the WOLF-3676, with corresponding increases in Tensor and RT cores to support advanced AI inference and real-time ray tracing performance. Memory configurations evolve from 8GB of GDDR6 in the Ampere and Ada modules to 6GB and 8GB of GDDR7 in the Blackwell-based modules, delivering faster bandwidth and improved power efficiency.

Power consumption varies across modules, ranging from 26–50W in the lower-power configurations to 35–100W in the high-performance WOLF-3676, providing flexibility to meet diverse thermal and performance requirements.

Note: Performance at 35 Watts has yet to be evaluated for the WOLF-3676 module.

The table below presents a comparative overview of the four WOLF XMC modules.

| Module | WOLF-3476 (XMC – A2000E) | WOLF-3576 (XMC – AD2000E) | WOLF-3696 (XMC – BW500E) | WOLF-3676 (XMC – BW2000E) |
|------------------|-----------------------------|------------------------------|-----------------------------|------------------------------|
| GPU Architecture | Ampere | Ada | Blackwell | Blackwell |
| CUDA Cores | 2560 | 3072 | 1792 | 3328 |
| Tensor Cores | 80 | 96 | 56 | 104 |
| RT Cores | 20 | 24 | 14 | 26 |
| Memory | 8GB GDDR6 | 8GB GDDR6 | 6GB GDDR7 | 8GB GDDR7 |
| ECC Support | No | Yes | No | Yes |
| Power Range | 26 – 50W | 30 – 50W | 30 – 60W | 35 – 100W |

Note: For our XMC boards using the Ampere and Ada GPUs, we intentionally cap the maximum power at 50 W, even though the GPU is capable of operating at higher levels. This lets us remove one power stage and apply additional optimizations specifically suited for lower-power XMC designs.

COMPUTE AND AI PERFORMANCE

The following table presents the FP32 performance of the WOLF-XMC's GPUs operating at various power levels.

| Metric | WOLF-3476 (XMC - A2000E) | WOLF-3576 (XMC - AD2000E) | WOLF-3696 (XMC - BW500E) | WOLF-3676 (XMC - BW2000E) |
|--------------------------------------|-----------------------------|------------------------------|-----------------------------|------------------------------|
| FP32 Performance @25W (Boost Clock) | To be measured | To be measured | To be measured | Power too low |
| FP32 Performance @35W (Boost Clock) | 6 | 8.9 | 6.4 | To be measured |
| FP32 Performance @50W (Boost Clock) | 7.9 | 12 | 8.4 | 11.9 |
| FP32 Performance @80W (Boost Clock) | 9.3 | 14.4 | To be measured | 15.7 |
| FP32 Performance @100W (Boost Clock) | To be measured | 14.5 | Power too high | 17.4 |

The table below summarizes the relative AI Performance and PCIe interface capabilities of WOLF XMC modules.

| Metric | WOLF-3476 (XMC - A2000E) | WOLF-3576 (XMC - AD2000E) | WOLF-3696 (XMC - BW500E) | WOLF-3676 (XMC - BW2000E) |
|----------------|-----------------------------|------------------------------|-----------------------------|------------------------------|
| AI Performance | Mid | Improved | High per watt | Maximum |
| PCIe Interface | Gen4 x8 | Gen4 x8 | Gen4 x8 | Gen4 x8 (Gen5 capable) |

KEY TAKEAWAY

WOLF-3676 nearly doubles the compute performance of WOLF-3476 (Ampere 2000) and, at a higher power level, delivers more TFLOPS than WOLF-3576 (ADA 2000). Its BW2000 should also perform better on AI workloads because it uses 100% of its cores for FP32 and INT32 calculations, whereas Ada uses only 50%. Meanwhile, the BW500 in WOLF-3696 is optimized for compute efficiency at lower power envelopes, making it a good upgrade from the 3476 when a tighter power limit is required.

AI performance improves across modules, all of which use PCIe Gen4 x8, with WOLF-3676 additionally supporting Gen5 for higher data bandwidth.

MEMORY AND DATA BANDWIDTH

| Memory Attribute | WOLF-3476 (XMC – A2000E) | WOLF-3576 (XMC – AD2000E) | WOLF-3696 (XMC – BW500E) | WOLF-3676 (XMC – BW2000E) |
|------------------|-----------------------------|------------------------------|-----------------------------|------------------------------|
| Type | GDDR6 | GDDR6 | GDDR7 | GDDR7 |
| Bandwidth | 192 GB/s | 192 GB/s | 288 GB/s | 384 GB/s |
| ECC | No | Yes | No | Yes |

KEY TAKEAWAY

WOLF XMC modules have evolved from GDDR6 to faster GDDR7 memory with significantly higher bandwidth and optional ECC support for improved reliability. The GDDR6-based WOLF-3476 and WOLF-3576 offer 192 GB/s bandwidth, while the GDDR7-based WOLF-3696 and WOLF-3676 increase performance to 288 GB/s and 384 GB/s, respectively. This bandwidth jump from GDDR6 to GDDR7 significantly reduces latency, enabling faster throughput for demanding applications such as AI inference, signal processing, and embedded graphics workloads.

SYSTEM INTEGRATION AND SWAP IMPACT

| Integration Factor | WOLF-3476 (XMC – A2000E) | WOLF-3576 (XMC – AD2000E) | WOLF-3696 (XMC – BW500E) | WOLF-3676 (XMC – BW2000E) |
|---------------------------|-----------------------------|------------------------------|-----------------------------|------------------------------|
| Power Envelope | 26-50W | 30-50W | 30-60W | 35-100W |
| VPX/SOSA Migration | Yes | Yes | Yes | Yes |
| Display outputs (DP/HDMI) | 4 | 3 | 3 | 3 |
| Best Fit | Legacy upgrades | General purpose | SWaP-limited edge AI | High-density AI compute |

KEY TAKEAWAY

WOLF-3696 is uniquely positioned for SWaP-constrained ISR pods and UAV mission processors, while WOLF-3676 is a perfect central compute module for mission computers, high-mix sensor fusion, and onboard AI model hosting.

CONCLUSION

The WOLF XMC family provides mission-ready compute architectures for ISR, autonomy, and embedded AI.

- [WOLF-3476 \(XMC-A2000E\)](#) supports legacy ISR workloads with up to 7.9 TFLOPS at 50W, 4 DP/HDMI outputs, and a lower power envelope.
- [WOLF-3576 \(XMC-AD2000E\)](#) balances compute and graphics with up to 12 TFLOPS at 50W but only 3 DP/HDMI outputs.
- [WOLF-3696 \(XMC-BW500E\)](#) enables SWaP-focused AI edge processing, with up to 8.4 TFLOPS at 50W and up to 9.2 TFLOPS at 60W. It also has the lowest power envelope for the Blackwell chips. It's a great replacement for the 3476 when a tighter power limit is required with better performance on AI workloads, because it uses 100% of its cores for FP32 and INT32 calculations compared to 50% on ADA.
- [WOLF-3676 \(XMC-BW2000E\)](#) delivers centralized mission GPU capability for sensor fusion and tactical AI workloads. With 11.9 TFLOPS at 50W, up to 15.7 TFLOPS at 80W, and all the way to 17.4 TFLOPS at 100W, depending on module configuration and cooling. It has 8GB of GDDR7 memory and can support up to 3 DP/HDMI outputs. It has the highest power envelope in the product line but delivers better compute and AI performance at these higher power requirements.

This product line enables scalable, rugged compute modernization across defense and aerospace platforms.

While both WOLF-3676 and WOLF-3576 deliver similar overall performance, the former offers enhanced capabilities. Its increased power consumption supports advanced features and improved efficiency. A primary differentiator is the memory subsystem. It utilizes GDDR7 memory, which provides nearly double the bandwidth compared to the GDDR6 memory used in WOLF-3576, despite maintaining the same bus width. This increased memory bandwidth results in significantly faster data transfer rates and improved efficiency, particularly in memory-intensive applications.

Additionally, the WOLF-3676 features a more optimized core utilization, with 100% of its cores capable of handling both FP32 and INT32 operations, whereas WOLF-3576 supports this functionality on only 50% of its cores. This architectural enhancement contributes to the superior AI computational efficiency and throughput of WOLF-3676.

In terms of developing future-ready AI platforms, the Blackwell architecture modules (WOLF-3696 and WOLF-3676) provide the most flexible path, thanks to improvements in memory throughput, thermal scalability, and tensor acceleration efficiency, all while maintaining rugged deployment readiness.

When considering overall performance, WOLF-3676 outperforms the other XMC modules, followed by WOLF-3576, WOLF-3476, and WOLF-3696.

To learn more about WOLF's cutting-edge solutions, visit wolf-at.com or contact our team at sales@wolf-at.com.

