

Extending the Benefits of GDDR Beyond Graphics

Micron® High-Performance Graphics Memory Addresses the High-Bandwidth Demands of Networking Today and Tomorrow

Key Benefits of High-Performance Graphics Memory

- High Performance
- Increased Longevity
- Low Total Cost of Ownership
- Low Power Consumption
- Easy to Design in
- High Flexibility and Scalability



By 2020, the average bandwidth used each day is predicted to be about 1.5GB per person. Will networks be able to cope?

Overview

As networking and data center performance demands increase, higher-performance memory solutions are needed to address a variety of networking applications, including data processing and buffering. For system development beyond 400 Gb/s, efficiency and performance gains are becoming increasingly difficult to achieve in a cost-effective way.

Enter graphics double data rate (GDDR) memory as one solution to meet the relentless bandwidth needs in networking. Originally developed to address the high-bandwidth performance needs of game consoles and PC graphics, graphics memory is today being adopted into data center networking applications. GDDR offers a proven, low-risk, high-performance memory solution for networking that is both flexible and cost-effective.

This technical brief describes the trends in the industry driving these high-bandwidth and high-performance requirements. It also discusses how graphics memory, such as the GDDR5 and GDDR6 solutions offered by Micron, can satisfy those needs in a way that's lower-risk and more cost-effective than any other high-bandwidth memory solution available today.

Now Trending in Networking

Cisco's Global Cloud Index 2015 reported that Internet Protocol (IP) traffic was 4.7 zettabytes (ZB) in 2015 and is expected to move to 15.3ZB in 2020 (Figure 1), representing a 27% compound annual growth rate (CAGR). The bulk of this increase is related to the rise of big data and ever-increasing video traffic.

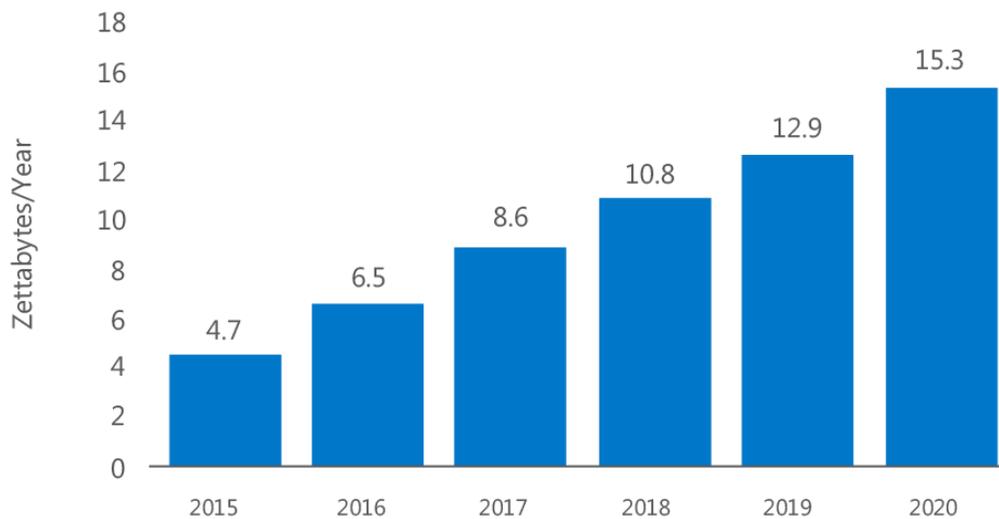


Figure 1: Global Data Center IP Traffic Growth From Cisco Global Cloud Index 2015

The same study predicts that the number of access devices, such as smartphones and tablets, will grow from 16 billion in 2015 to more than 26 billion devices in 2020, as shown in Figure 2.

What's Driving Networking Demands?

Growth in several significant areas is driving these unprecedented demands in the networking industry:

- **Mobile Data and Internet Video:** The need for accessing data on-the-go and streaming videos on demand over the Internet is ever-increasing.
- **Internet of Things (IoT):** The IoT is increasing the number of devices — like wearables, smart home appliances and cars — that have to access networks.
- **Cloud Services:** Countless businesses are moving their services to the cloud, along with new businesses being enabled by the cloud's capabilities opening up every day.
- **Intelligence/Analytics:** To make all the pieces of a complex networking puzzle work, edge devices like routers, switches and multiplexers within the network must have better and quicker insight into the data they carry.

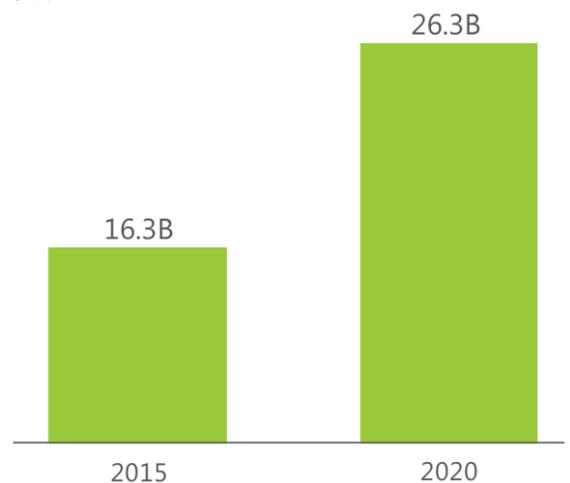


Figure 2: Networking Device Requiring Connectivity

In summary, the culmination of more people, devices, larger screens and the push to the cloud are driving an exponential increase in IP traffic. With no signs of this growth slowing, how can you address these demands starting today?

Why Graphics Memory for Networking?

Today, Micron graphics memory devices are available in densities up to 8Gb, soon to be available in 16Gb, on a par with the densest DDR4 memory.

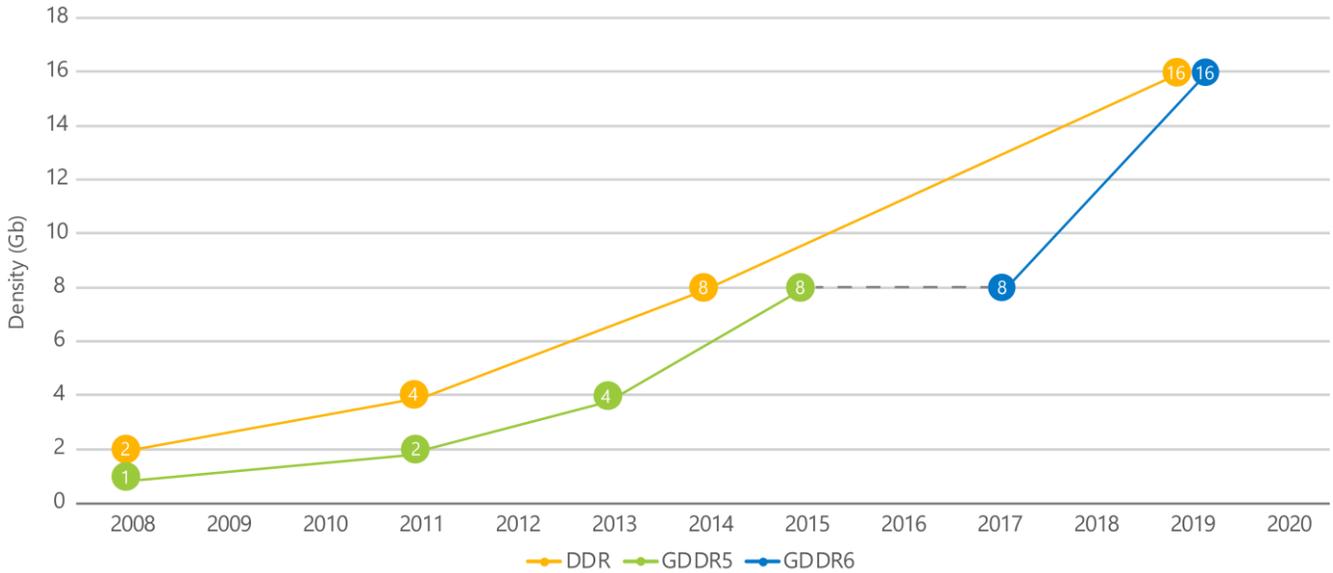


Figure 3: DDR vs. GDDR Family Density Comparison

The benefits of these high-performance graphics memory devices, offering a huge step in DRAM performance, make them ideal for addressing the demands of modern networks.

Performance Is Key

GDDR devices are in a league of their own, driving and exponential increase in bandwidth per packaged device.

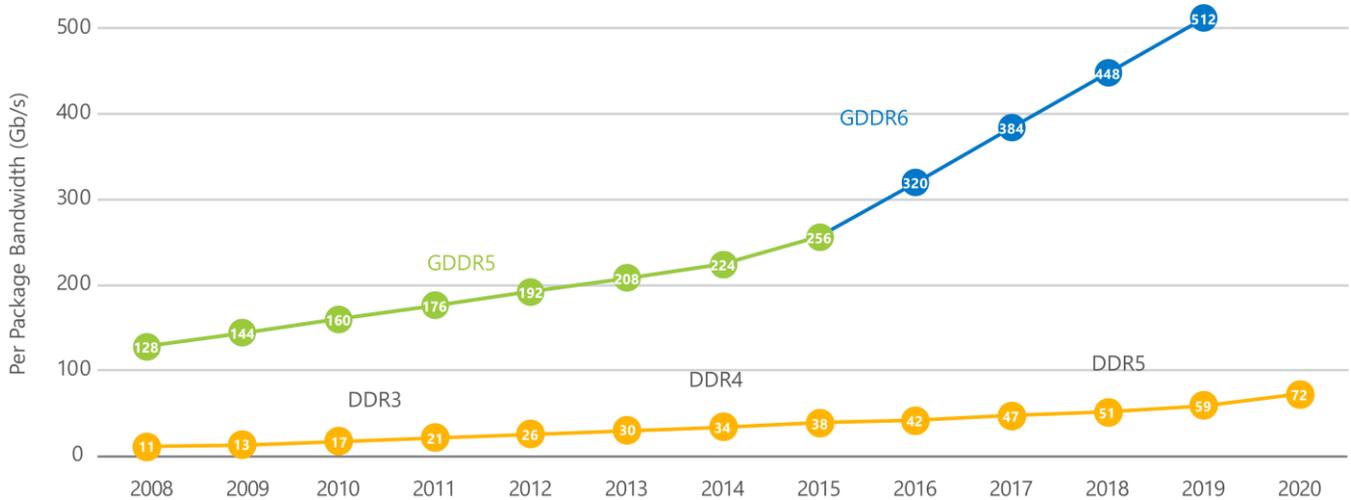


Figure 4: DDR vs. GDDR Family Bandwidth Comparison

While a maximum interface speed of 8 Gb/s can be achieved with GDDR5, next-generation GDDR6 devices will blow past this limitation, doubling the interface speed to 16 Gb/s and doubling the number of channels. This new dual-channel architecture enables a massive increase in performance while still providing backward compatibility to GDDR5.

With DDR4 today topping out at 3.2 Gb/s, GDDR5 at up to 8Gb/s and GDDR6 targeted at as much as 16 Gb/s, GDDR devices will provide more than 10X the bandwidth¹ of today’s leading-edge DDR4.

While DDR remains relevant to a myriad of applications, the latest high-performance GDDR products can address applications requiring the absolute high end of today’s performance curve.

Stability for the Long Haul



GDDR6 is the latest incarnation of graphics DRAM, which substantially extends the bandwidth capabilities of GDDR memory. This next innovation in performance continues the relevance and reach that graphics memory devices must offer. Coupled with the long-term reliability, design and system expertise provided by Micron, customers can design Micron graphics memory devices with confidence.

Lower Total Cost of Ownership

When considering overall total cost of ownership (TCO), it is important to look at all aspects of a design. Figure 5 compares three different approaches to address the needs of a 1Tb switching application. As shown, GDDR6 not only reduces design complexity by 80%, it reduces memory footprint by 82% and improves energy efficiency by 44%.

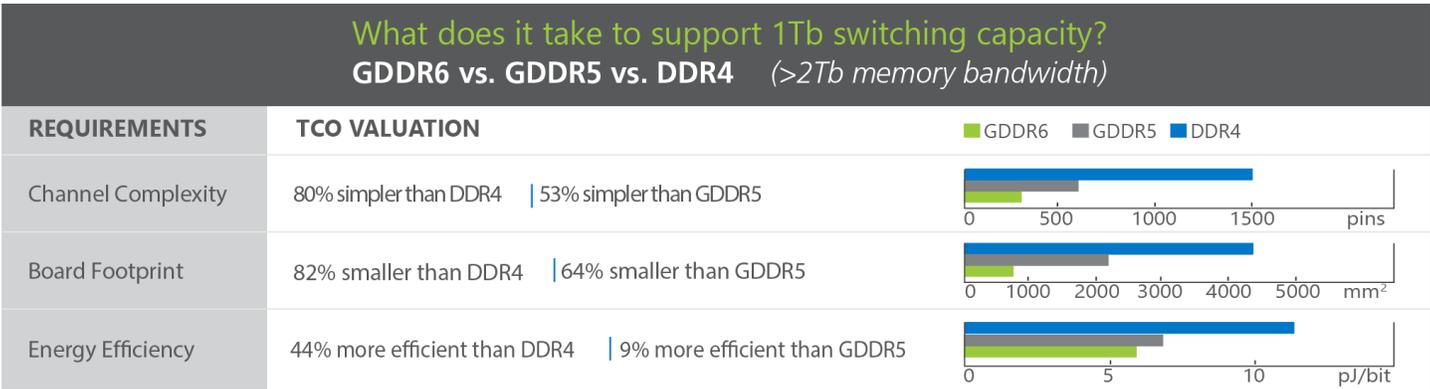


Figure 5: TCO Comparison of Mainstream High-Performance Memory^{2,3,4}



Industry-Leading Energy Efficiency

Power constraints and cooling are an increasingly important topic for any system design.

1. Per package bandwidth with the following assumptions: DDR4 = 3.2 Gb/s x 16, GDDR6 = 16 Gb/s x 32.
 2. Channel complexity equates to # of Pins – Includes only High speed address and data pins – does not include Power and GND.
 3. MM2 - assumes available (Micron) package outlines as dimension on the board – no routing breakout area.
 4. pJ/bit – is the device power over the BW, assuming peak power with 1:1 RD:WR and 30% precharge state.



Figure 6 compares the energy efficiency of GDDR5 and GDDR6 measured in a typical graphics application. While GDDR5 offered best-in-class power efficiency when it was introduced, the latest GDDR6 technology not only delivers higher system performance but does it even more efficiently than its GDDR5 predecessor.

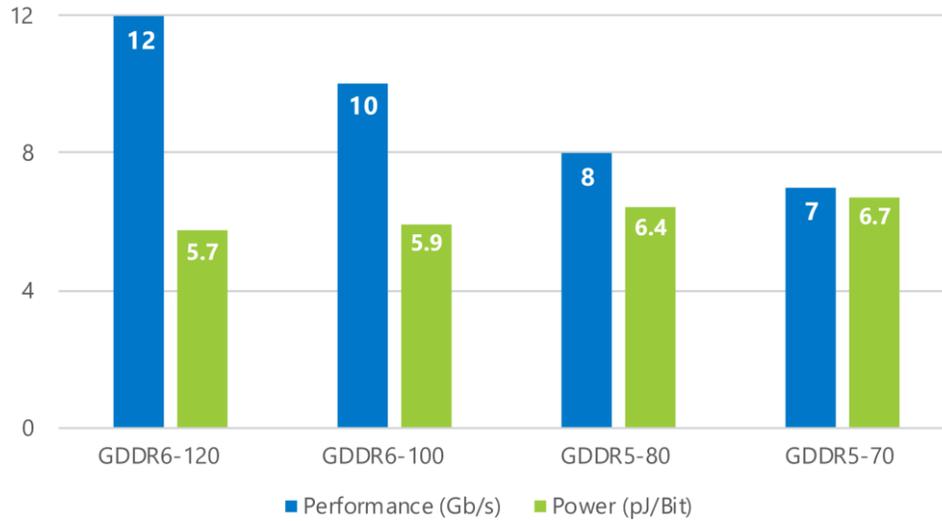


Figure 6: Comparing performance and power of GDDR6 and GDDR5⁴

Design With Ease

If you're already familiar with designing DDR devices, designing with GDDR memory will be a familiar, low-risk experience. GDDR5 and GDDR6 build upon conventional parallel memory techniques used by multiple generations of DDR technology.

GDDR6 adds a new ultra-high-speed mode and bandwidth increase, while leveraging a familiar controller design. The main difference between GDDR6 and GDDR5 is related to package and pinouts — so while the transition is not quite a drop-in replacement, GDDR6 follows the same design practice as GDDR5.



Flexible, Scalable Solutions

Micron's GDDR family of discrete packages provides compact, scalable and cost-effective designs. Their superb power efficiency enables the use of industry-standard thermal solutions, providing the designer versatile and flexible design and deployment options.

Using conventional DRAM package technology, we can leverage decades of manufacturing and testing experience to ensure the quality customers have come to expect and the ease of system design that a fully tested and packaged device provides.

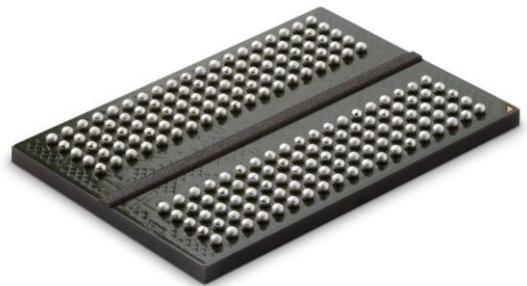


Figure 8: Discrete GDDR5 Memory Package

N is for Networking

Micron® GDDR5N and GDDR6N are JEDEC-compliant versions of GDDR5 and GDDR6, specifically designed and tested for the quality, reliability and availability demands of networking- and enterprise-class systems.

Available Now

There's no need to wait. Micron's GDDR5N devices are available today at micron.com, with GDDR6N devices coming soon.



Networking systems are more critical now than ever before. They need memory that enables high bandwidth and lower power.

Graphics Memory Solving Networking Problems Today

The rise of virtualization and its broad adoption in the data center means an increase in the demands placed on deep buffer switches. Oversubscribed, virtualized servers require smarter switches with features designed to mitigate network congestion and latency problems that affect performance. Merchant silicon providers are therefore creating silicon solutions for deep buffer switches focused on delivering the highest level of performance with low latency and high reliability.

One example of a deep buffer switch available today with Micron's GDDR5N: is Interface Masters® Edge Aggregation Switch. This flexible switch system uses Broadcom's Qumran chipset and 8GB of Micron GDDR5N. It is designed for optical transport, carrier Ethernet, edge switching, data center cloud and enterprise campus market segments.



Figure 9: Example of a deep buffer switch from Interface Masters fueled by GDDR5N

Conclusion

High-performance graphics memory, such as Micron's GDDR5 and GDDR6 technologies, has been adopted in market applications beyond graphics. Micron's GDDR5N and GDDR6N devices are an ideal fit for the bandwidth, power and scalability demands of networking applications. Micron GDDR technologies provide a flexible, low-risk and cost-effective solution to the memory demands of your next-generation designs.

Discover more at micron.com.

micron.com

This technical brief is published by Micron and has not been authorized, sponsored, or otherwise approved by Cisco or Broadcom. Products are warranted only to meet Micron's production data sheet specifications. Products and specifications are subject to change without notice.

©2017 Micron Technology, Inc. Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners. All rights reserved. Rev. B 9/17, CCMMD-676576390-10744