


AMD



PURE.
DATACENTER.
GRAPHICS.

MICHAEL DENEFFE

Director, Cloud Graphics

VISUALIZATION WORKLOADS MOVING TO THE CLOUD

MOBILE
WORKFORCE

72%

OF THE U.S.
WORKFORCE WILL BE
MOBILE BY 2020

NIMBLE
WORKFORCE

34%

OF TODAY'S
WORKFORCE ARE
CONTRACTORS

DATA
EXPLOSION

2X

THE SIZE OF THE DIGITAL
UNIVERSE WILL DOUBLE
EVERY TWO YEARS AT
LEAST

DEMAND FOR
SECURITY

#1

PRIORITY OF
IT MANAGERS

TECHNOLOGY
SAVINGS

93%

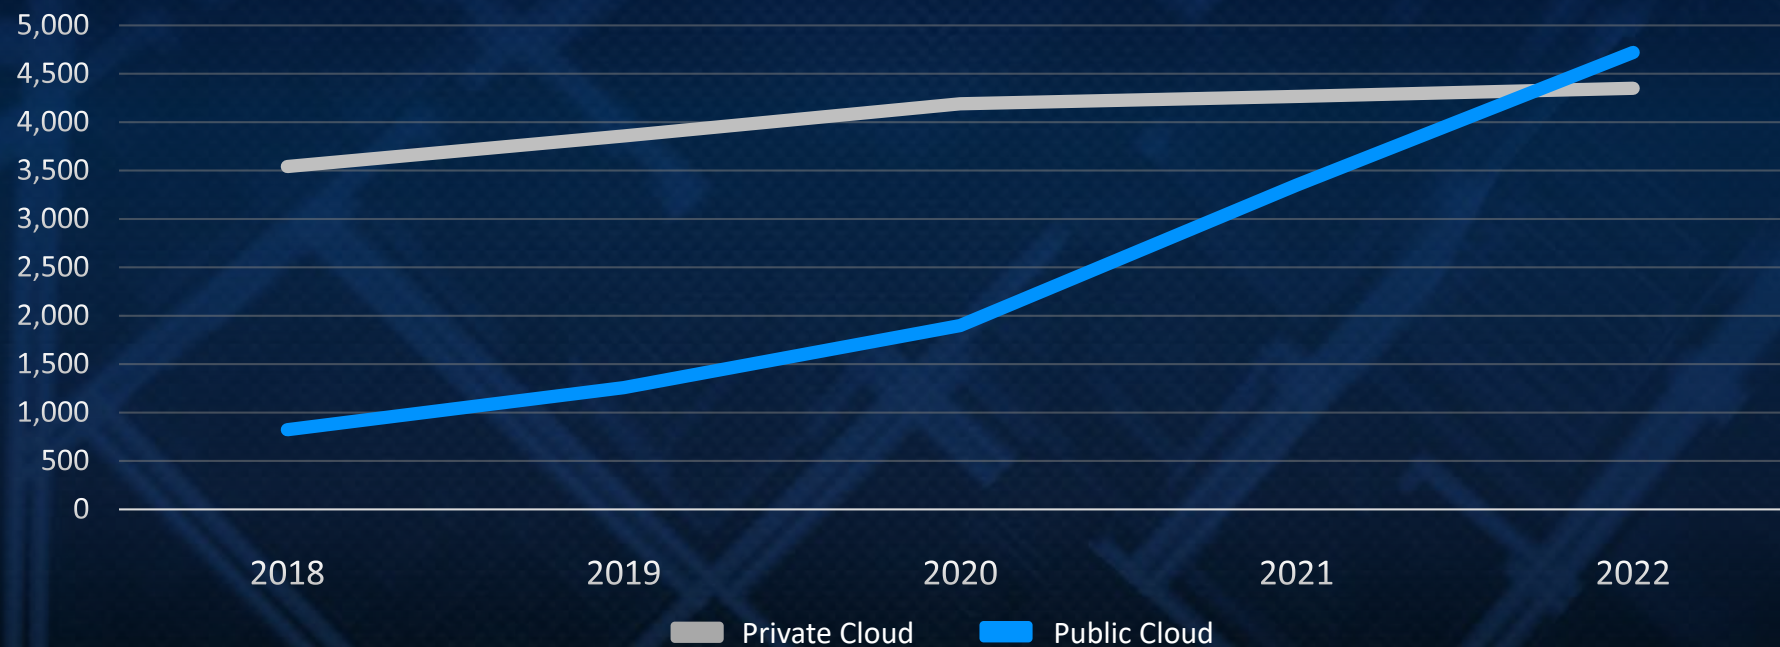
OF CITRIX USERS
REPORTED REDUCED
OVERALL IT COSTS

*SEE ENDNOTES

ACCELERATED ADOPTION OF PUBLIC CLOUD

CLOUD-BASED VISUALIZATION WORKLOADS

(K USERS)



SOURCE: INFOHOLIC RESEARCH LLC.

MORE USERS REQUIRE GRAPHICS ACCELERATION

GRAPHICS PERFORMANCE REQUIREMENTS

POWER USERS

CAD | CAE

20%

KNOWLEDGE WORKERS

MICROSOFT® WINDOWS® 10 AND OFFICE 365®

70%

TASK WORKERS

CALL CENTERS


10%

S7100X & S7150 x2

Existing product line addresses this area

*INTERNAL AMD ESTIMATES

RADEON VEGA IS EVERYWHERE



VEGA FOR
GAMING




VEGA FOR
PROFESSIONALS



VEGA FOR
COMPUTE

INTRODUCING
AMD RADEON™ PRO V340



REDEFINING HIGH PERFORMANCE DATACENTER GRAPHICS

AMD'S VISION FOR WORKSTATION-CLASS WORKLOADS

FROM DESKSIDE TO ON-PREMISES VDI TO CLOUD



DESKSIDE WORKSTATION



PRIVATE CLOUD



PUBLIC CLOUD



AMD RADEON™ PRO SOFTWARE ---- ONE DRIVER FOR ALL



AMD RADEON™ PRO
V340 GRAPHICS

2X "VEGA"
GRAPHICS PROCESSORS

56 x2
COMPUTE UNITS

32 GB
HBM2 MEMORY

UP TO **512 GB/s**
MEMORY BANDWIDTH

UP TO **32**
USERS PER CARD

BUILT-IN
SECURITY PROCESSOR

*SEE ENDNOTES

MORE USERS REQUIRE GRAPHICS ACCELERATION

GRAPHICS PERFORMANCE REQUIREMENTS

POWER USERS

CAD | CAE

20%

KNOWLEDGE WORKERS


MICROSOFT® WINDOWS® 10 AND OFFICE 365®

70%

TASK WORKERS

CALL CENTERS

10%



RADEON PRO
V340

1.9X PERFORMANCE

2X MEMORY

2X COMPUTE CORES

S7100X & S7150 x2

Existing product line addresses this area

*INTERNAL AMD ESTIMATES
SEE ENDNOTES

R A D E O N P R O V 3 4 0

INDUSTRY LEADING INNOVATION

2X

ENCODE ENGINES

2 Encode Engines
Delivering H.264 and H.265 4K
streams

UP TO **33%**

MORE DENSITY

33% more users per card than the
competition, reducing TCO

UP TO **8X**

FASTER MEMORY

Better power efficiency and
8X faster than the GDDR5

*SEE ENDNOTES

AMD
RADEON PRO
Software

SOFTWARE FEATURES

HYPERVERSORS

VMware® ESXi® and vSphere®

REMOTING PROTOCOLS

Citrix® XenDesktop® and XenApp®,
VMware® Horizon® View™

GRAPHICS DRIVERS


Radeon™ Pro Software

GUEST OS

Microsoft® Windows® 10,
Microsoft® Windows® Server

API

AMD RapidFire Encode



DELIVERING MODERN
VISUALIZATION WORKLOADS
FROM THE CLOUD



HARDWARE FEATURES

VIRTUALIZATION

SR-IOV Hardware Virtualization

ARCHITECTURE

AMD Radeon “Vega”

ENCODE ENGINES

H.264 & H.265


SECURITY ENGINE

Built-In

MEMORY

HBM2

*SEE ENDNOTES



RELIABILITY



PREDICTABILITY



PERFORMANCE



SECURITY



VALUE

RELIABILITY

OPTIMIZE YOUR GRAPHICS



AMD
RADEON PRO
Software

93%

DRIVER PASS RATE*


STRESS TESTED FOR
24/7 ENVIRONMENTS

EXTENSIVE OEM
PLATFORM TESTING

EXTENSIVE ISV
CERTIFICATION TESTING


**IN MAY 2018, AMD COMMISSIONED QA CONSULTANTS TO EVALUATE THE STABILITY OF SEVERAL OF THE LATEST GRAPHICS DRIVERS FROM BOTH AMD AND NVIDIA. SEE PAGE 1 OF THE "GRAPHICS DRIVER QUALITY - DETERMINATION OF STABILITY FROM LEADING MARKET VENDORS" REPORT AT [HTTPS://WWW.AMD.COM/SYSTEM/FILES/DOCUMENTS/GRAPHICS-DRIVER-QUALITY.PDF](https://www.amd.com/system/files/documents/graphics-driver-quality.pdf) FOR MORE DETAILS.

SR-IOV: DEMOCRATISED GPU



Each user is given a dedicated slice of the GPU and memory

The guest OS recognizes the GPU as a pass through device

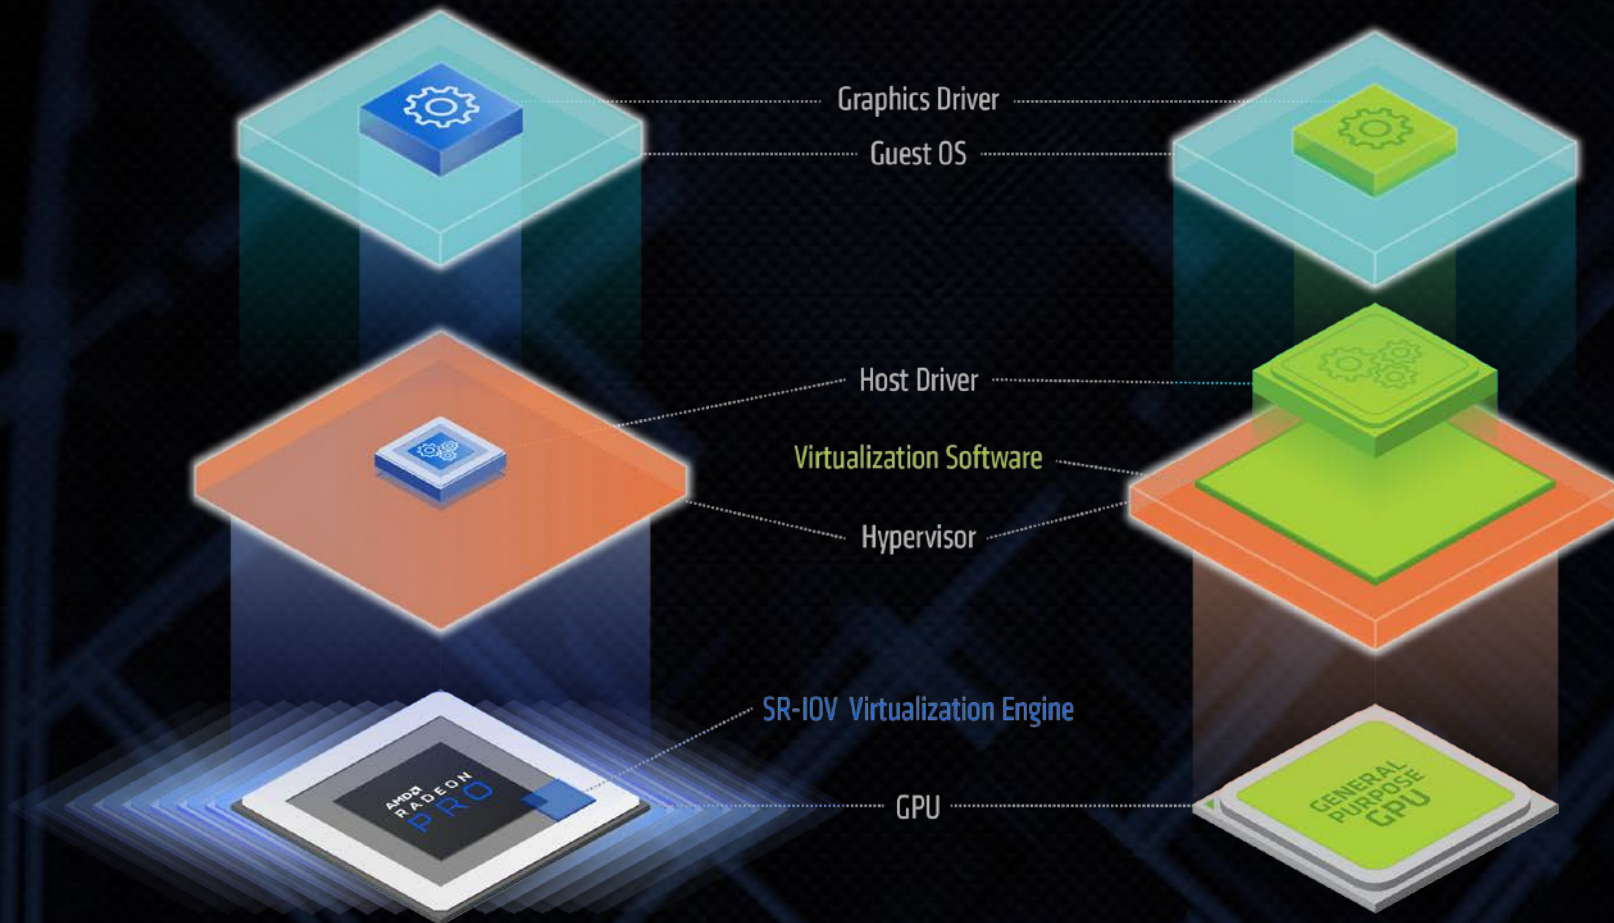


V I R T U A L I Z A T I O N

HARDWARE VS. SOFTWARE

HARDWARE VIRTUALIZATION

SOFTWARE VIRTUALIZATION





2X "VEGA" GRAPHICS PROCESSORS

56 x2 COMPUTE UNITS

32 GB HBM2 MEMORY

UP TO 512 GB/S MEMORY BANDWIDTH

UP TO 32 USERS PER CARD

BUILT-IN SECURITY PROCESSOR

AMD RADEON™ PRO V340 GRAPHICS


Redefining High Performance Datacenter Graphics

up to **33%**
MORE DENSITY

up to **8X**
FASTER MEMORY

2X
ENCODE ENGINES

*SEE ENDNOTES



AVAILABLE FROM LEADING
OEMs IN Q4, 2018



THANK YOU!

ENDNOTES

SLIDE 5: Visualization Workloads Moving to Cloud

- IDC, the mobile workforce is expected to be 105.4 million strong by 2020 and account for 72.3% of the workforce in the US
<https://www.businesswire.com/news/home/20150623005073/en/IDC-Forecasts-U.S.-Mobile-Worker-Population-Surpass>
- "The gig economy...is now estimated to be about 34% of the workforce and expected to be 43% by the year 2020. The gig economy is a term that refers to the increased tendency for businesses to hire independent contractors and short-term workers
<https://money.cnn.com/2017/05/24/news/economy/gig-economy-intuit/index.html>
<https://www.forbes.com/sites/larryalton/2018/01/24/why-the-gig-economy-is-the-best-and-worst-development-for-workers-under-30/#514be5246d76>
- The size of the digital universe will double every two years at least.
<https://insidebigdata.com/2017/02/16/the-exponential-growth-of-data/>
- 95 percent of CIOs surveyed said they expect cybersecurity threats to increase and impact their organization.
<https://www.gartner.com/en/newsroom/press-releases/2017-10-02-gartner-survey-of-more-than-3000-cios-confirms-the-changing-role-of-the-chief-information-officer>
- 93% of XenDesktop customers report reduced costs.
<https://www.citrix.com/blogs/2015/10/01/93-of-xendesktop-customers-report-reduced-costs-how-do-they-do-it/>

SLIDE 8: AMD Radeon Pro V340 Specifications

“Vega” is a codename for AMD architectures, and is not a product name. GD-122

Discrete AMD Radeon™ and Radeon™ Pro GPUs based on the Graphics Core Next architecture consist of multiple discrete execution engines known as a Compute Unit (“CU”). Each CU contains 64 shaders (“Stream Processors”) working together. GD-78

ECC support is limited to the HBM2 memory and ECC protection is not provided for internal GPU structures.

ENDNOTES

SLIDE 9: More Users Require Graphics Acceleration

Testing conducted by AMD Performance Labs as of August 20th, 2018, on a test system comprising of Intel Xeon Gold 6142, 8GB DDR4 system memory, Windows® 10 Professional 64-bit, AMD Radeon™ Pro V340 graphics, AMD FirePro™ S7150 x2 graphics, AMD professional driver 18.10 internal build. Benchmark Application: SPECviewperf® 12.1.1 geomean results. AMD Radeon™ Pro V340 geomean score = 71.31. AMD FirePro™ S7150 x2 score: 40.26. Performance Differential: $71.31/40.26 = 1.8x$ greater performance on AMD Radeon™ Pro V340 than on AMD FirePro™ S7150 x2. For more information about SPECviewperf, see <https://www.spec.org/benchmarks.html#gwpg>. PC manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers.

AMD Radeon™ Pro V340 graphics contains 56 x2 compute units, AMD FirePro™ S7150 x2 contains 28 x2 Compute units.

AMD Radeon™ Pro V340 graphics is equipped with 32 GB HBM2 memory, AMD FirePro™ S7150 x2 is equipped with 16 GB GDDR5 memory.

SLIDE 10: Radeon Pro V340 Graphics – Industry Leading Innovation

The AMD Radeon™ Pro V340 supports up to 32 virtual users with 1GB memory profiles per card. The NVIDIA Tesla P40 supports up to 24 virtual users with 1GB memory profiles per card. AMD user density advantage $32/24 =$ up to 33%.

HBM2: 8x Faster Memory vs. GDDR5: <https://www.samsung.com/semiconductor/dram/hbm2/>

SLIDE 11: MxGPU Technology

“Vega” is a codename for AMD architectures, and is not a product name. GD-122

SLIDE 16: AMD Radeon Pro V340 Summary

The AMD Radeon™ Pro V340 supports up to 32 virtual users with 1GB memory profiles per card. The NVIDIA Tesla P40 supports up to 24 virtual users with 1GB memory profiles per card. AMD user density advantage $32/24 =$ up to 33%.

HBM2: 8x Faster Memory vs. GDDR5: <https://www.samsung.com/semiconductor/dram/hbm2/>

“Vega” is a codename for AMD architectures, and is not a product name. GD-122

Discrete AMD Radeon™ and Radeon™ Pro GPUs based on the Graphics Core Next architecture consist of multiple discrete execution engines known as a Compute Unit (“CU”). Each CU contains 64 shaders (“Stream Processors”) working together. GD-78

ECC support is limited to the HBM2 memory and ECC protection is not provided for internal GPU structures.



DISCLAIMER

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 particular 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. Timelines, roadmaps, and/or product release dates shown in these slides are plans only and subject to change.

"Vega" is a codename for an AMD architecture, and is not a product name. GD-122

ATTRIBUTION

© 2018 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Radeon, LiquidVR, and combinations thereof are trademarks of Advanced Micro Devices, Inc. DirectX is a registered trademark of Microsoft Corporation in the US and other jurisdictions. OpenCL is a trademark of Apple Inc. used by permission by Khronos. OpenGL is a registered trademark of Silicon Graphics, Inc. used by permission by Khronos. Vulkan and the Vulkan logo are registered trademarks of Khronos Group, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.