



AMD FIREPRO

*Heterogeneous system architectures from
APUs to discrete GPUs*

April 2013

Confidential – NDA Required



Most parallel code runs on CPUs designed for scalar workloads



WASTES POWER



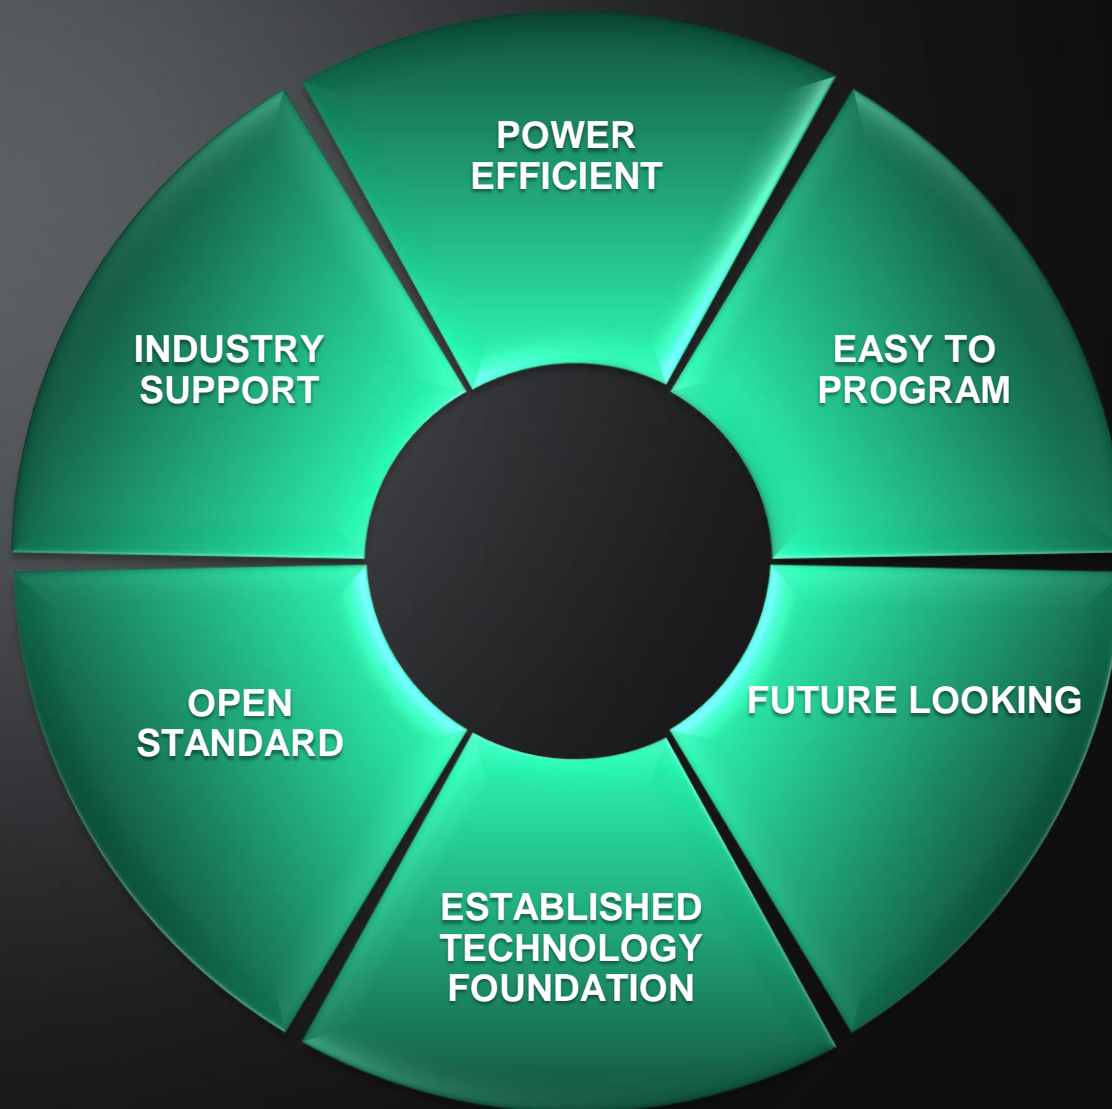
Typically platform builders create innovative new hardware and offer an API for software to access it
That tired thinking has only ever had niche success!

HSA BRINGS THE PLATFORM TO THE PROGRAMMER

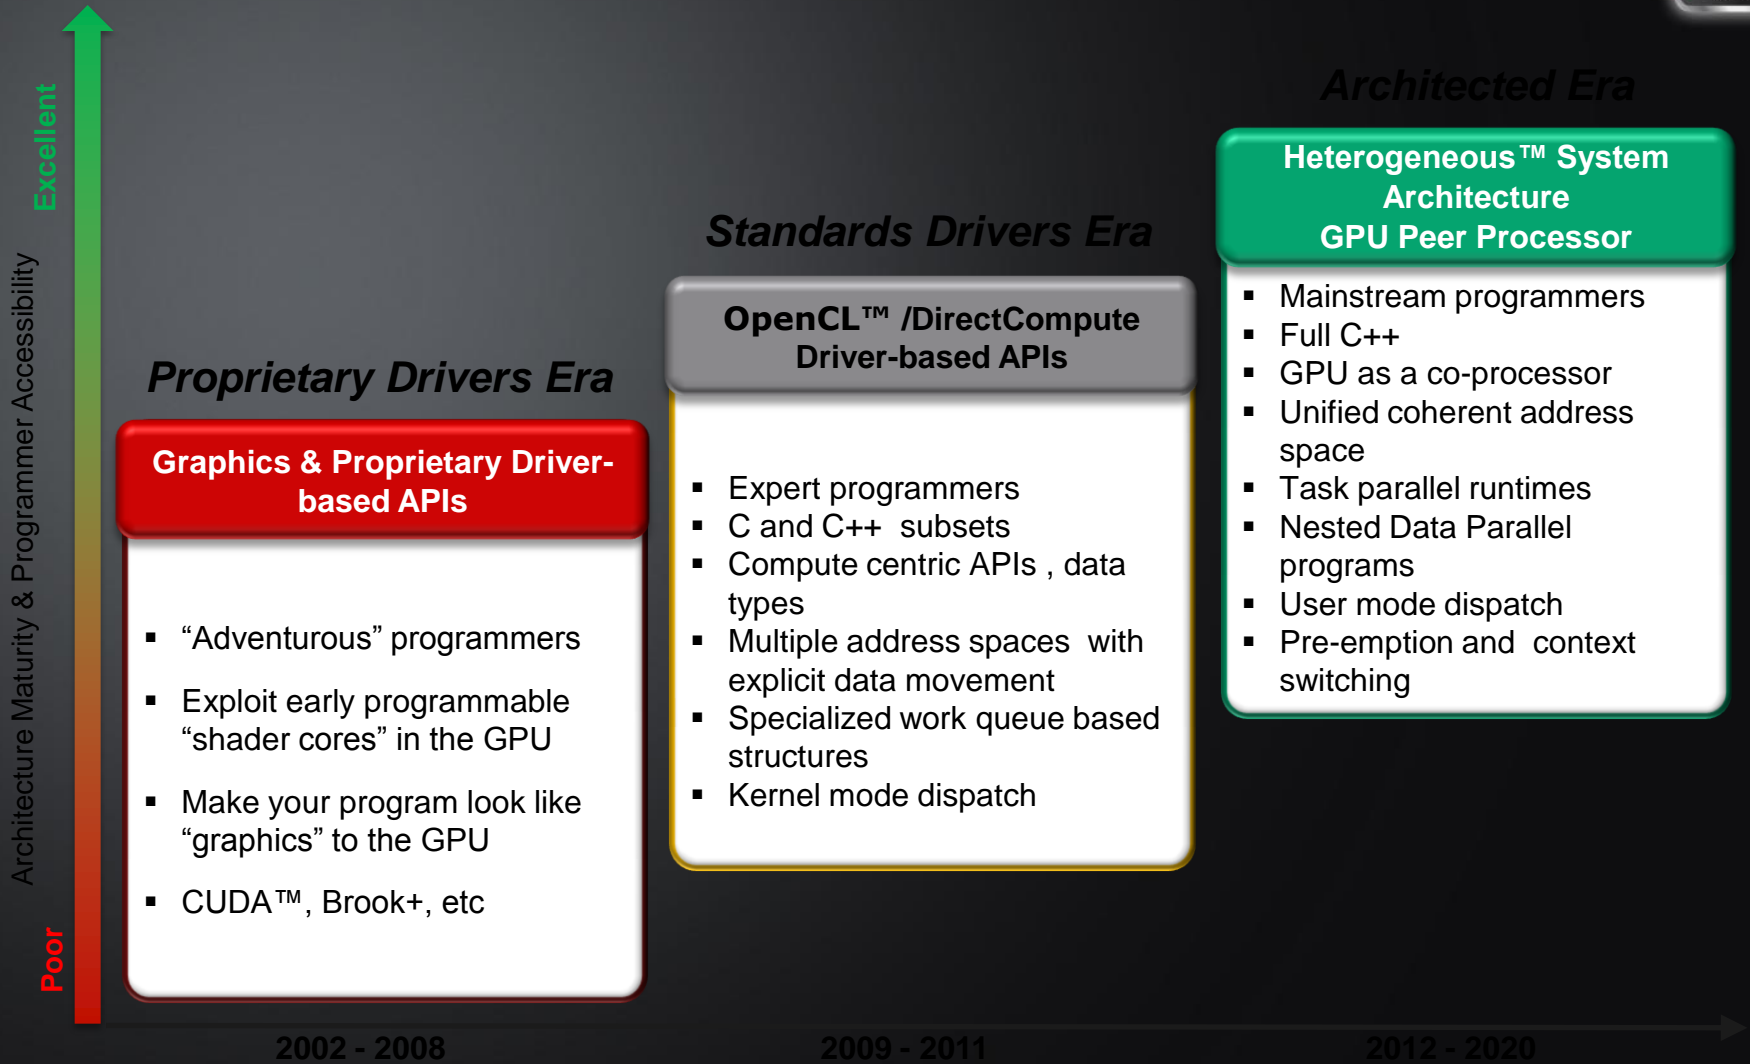


HETEROGENEOUS SYSTEM ARCHITECTURE

Brings All the Processors in a System into Unified Coherent Memory



EVOLUTION OF HETEROGENEOUS COMPUTING



HSA FEATURE ROADMAP



2011

2012

2013

2014

Physical Integration

Integrate CPU & GPU in silicon

Unified Memory Controller

Common Manufacturing Technology

Optimized Platforms

GPU Compute C++ support

User mode scheduling

Bi-Directional Power Mgmt between CPU and GPU

Architectural Integration

Unified Address Space for CPU and GPU

GPU uses pageable system memory via CPU pointers

Fully coherent memory between CPU & GPU

System Integration

GPU compute context switch

GPU graphics pre-emption

Quality of Service

Extend to Discrete GPU



Make the unprecedented processing capability of the APU as accessible to programmers as the CPU is today.




APU: ACCELERATED PROCESSING UNIT

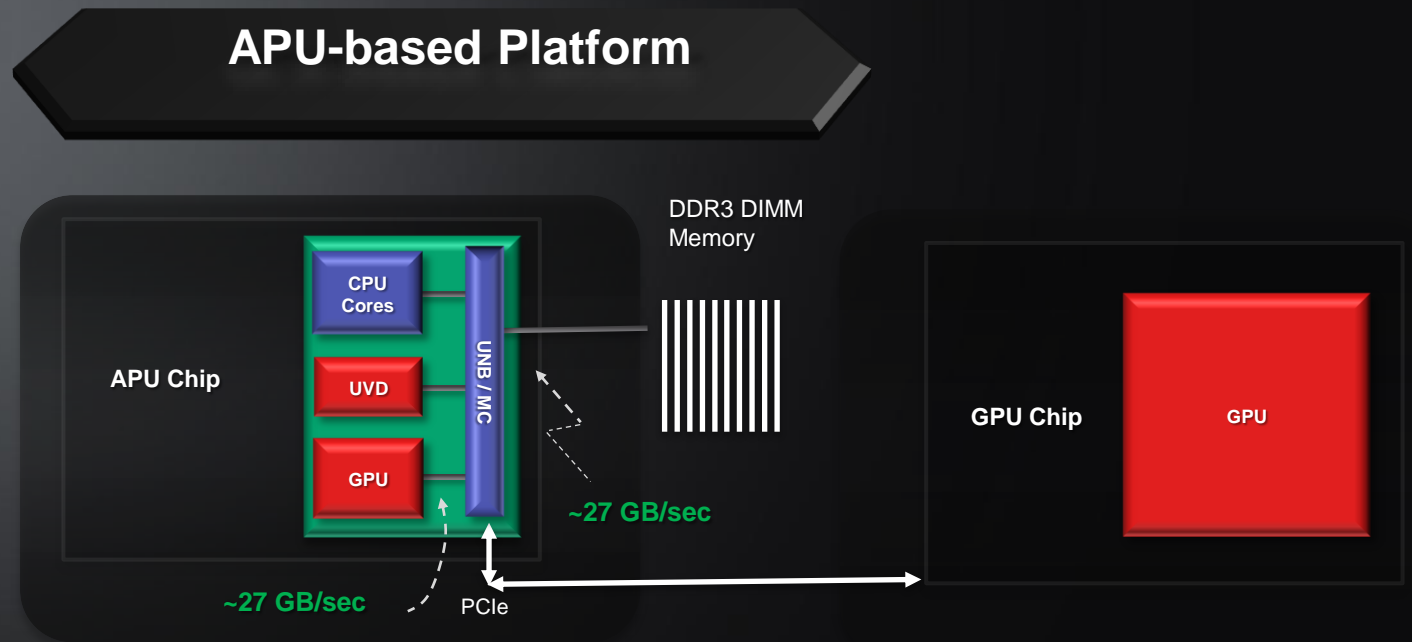


The APU has arrived and it is a great advance over previous platforms

- Combines scalar processing on CPU with parallel processing on the GPU and high bandwidth access to memory
- How do we make it even better going forward?
 - Easier to program
 - Easier to optimize
 - Easier to load balance
 - Higher performance
 - Lower power



- Heterogenous architecture APU + discrete GPU



- 3X bandwidth between GPU and memory
- Eliminate latency and power associated with the extra chip crossing

AMD FirePro S9000

High Performance

**3.23 TFlops SP / 806
GFlops DP**

High Bandwidth

**264 Gb/s
6Gb GDDR5**

Ambidextrous Multitasking

**Execute simultaneous
compute and 3D
graphics tasks**

Self execution Control

**GPU can execute tasks
by itself**

ECC

Memory protection


COMMITTED TO OPEN STANDARDS



- AMD drives open and de-facto standards
 - Compete on the best implementation
- Open standards are the basis for large ecosystems
- Open standards always win over time
 - SW developers want their applications to run on multiple platforms from multiple hardware vendors



OpenCL



DirectX®

JEDEC

dvi
digital visual interface


CERTIFIED
D™



OPENCL™ FOR NOW AND FOR LONG



- OPENCL is the STANDARD in the INDUSTRY today



...



Many popular applications benefiting from OpenCL™


More in the pipeline...




AccelerEyes Jacket (for Matlab)	Adobe Photoshop CS6	Adobe Premiere Pro CS6	ArcSoft Link+	Arcsoft MediaConverter	Arcsoft Media Impression	ArcSoft Panorama Maker Pro
ArcSoft ShowBiz	Arcsoft TotalMedia Theater	Arcsoft Webcam Companion	ANSYS Fluent	Assimilate Scratch	Autodesk Maya 2012 (Bullet Physics)	Autodesk Moldflow
BaoFeng Storm Player	Blender Blender	Corel VideoStudio Pro	Corel WinZIP	CyberLink PowerDirector	Dassault Systemes SIMULIA Abaqus/Standard	Dem Solutions EDEM
eyeon Fusion	GIMP	Handbrake	MainConcept H.264/AVC Pro Encoder	MotionDSP Ikena	MotionDSP vReveal	Nuvixa Stage Presence
Open Cascade	OPTIS THEIA RT	Paraken Musemage	Rovi TotalCode	Side Effects Software Houdini	Sony Software Movie Studio HD	Sony Software Vegas Pro
			Viewdle Photo Uploader			Wolfram Mathematica



PERFORMANCE GAINS WITH OPENCL™ AND AMD GRAPHICS (VS. CPU ONLY)



Abaqus/Standard 6.11 Physics Simulation Benchmarking (lower time = faster performance)



Testing conducted on a Dell T7400 w/ Intel Xeon E5405 CPU, 64GB RAM, Red Hat® OS v5.5 by Dassault Systemes, ATI FirePro™ V9800, AMD Catalyst Pro 10.12, Dassault Systemes SIMULIA® Abaqus/Standard 6.11 OpenCL™ is a trademark of Apple Inc., and is used with permission from Khronos.





AMD FirePro™ S10000 Server GPGPU

High Density / High Performance GPU Computing

8 GPU computing cards - 16 GPUs

11.84 TFLOPS peak double precision performance

47.28 TFLOPS peak single precision performance

1 Exxact GPU computing server



AMD FirePro™ S10000

- Dual Ultra High-End GPUs
- 6GB GDDR5
- 480 GB/s memory bandwidth
- 5.91 TFLOPS peak SPFP
- 1.48 TFLOPS peak DFPF
- 375W max power
- ECC memory support

World's Most Powerful GPU Computing Card

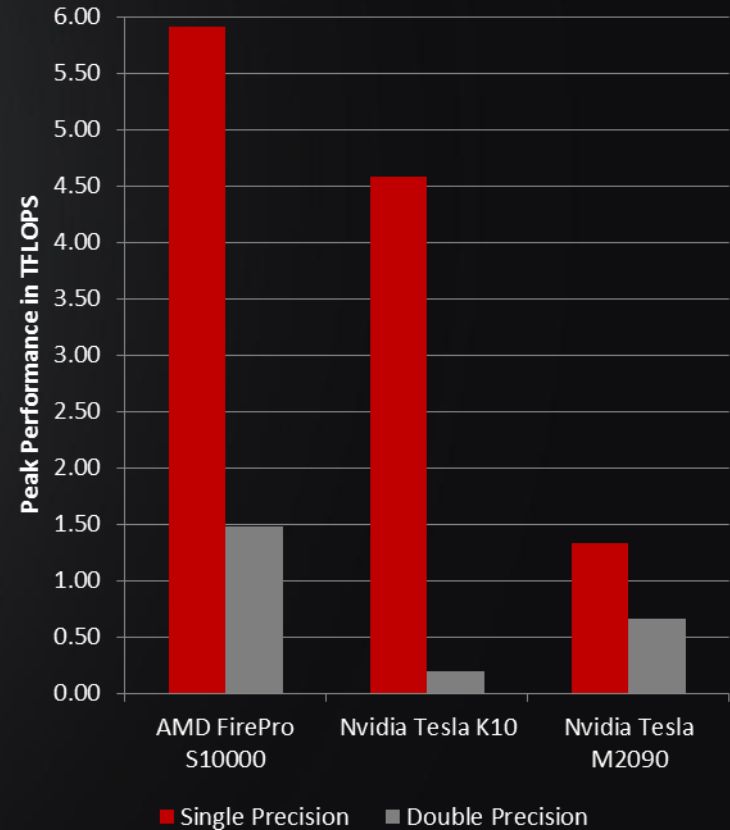


AMD FIREPRO™ S10000 SERVER GRAPHICS

High Performance and High Efficiency



- Highest single precision and dual precision floating point performance of any professional graphics solution³
 - 5.91 TFLOPS peak SPFP
Vs. Nvidia K10 at a max of 4.58 TFLOPS peak SPFP
 - 1.48 TFLOPS peak DFPF
Vs. Nvidia M2090 at a max of 0.665 TFLOPS peak DFPF
- Highest double precision / watt performance⁴:
 - 3.94 GFLOPS/Watt
Vs. Nvidia M2090 at 2.96 GFLOPS/Watt



AMD FIREPRO™ S10000 SERVER GRAPHICS

Industry Leading Compute Performance



The most powerful server graphics card ever created

- Up to 5.91 TFLOPS peak single precision performance and 1.48 TFLOPS peak double precision performance
- Up to 2.2X as fast as the M2090
- Up to 7.8X as fast as the K10 for peak double precision and 1.3X as fast for peak single precision³
- Delivers up to 1.5X the memory bandwidth as K10⁵

	AMD FirePro™ S10000	Nvidia Tesla K10	Nvidia Tesla M2090
ASIC Model	Tahiti (2x)	GK104(2x)	GF110(1x)
Form Factor	FH/FL Dual-slot	FH/FL Dual-slot	FH/FL Dual-slot
Core Clock Speed	825 Mhz	745 Mhz	650 Mhz
Cores (total)	3584	3072	512
Max Power	375W	225W	225W
Peak Single Precision	5.91 TFLOPS	4.58 TFLOPS	1.331 TFLOPS
Peak Double Precision	1.48 TFLOPS	0.19 TFLOPS	0.67 TFLOPS
Memory Size (total) / Type	6GB GDDR5	8GB GDDR5	6GB GDDR5
Memory Interface	384-bit	256-bit	384-bit
Memory Bandwidth (total, ECC off)	480 GB/s	320 GB/s	177 GB/s
ECC Memory Support	Yes	Yes	Yes
PCI Express®	X16, 3.0	X16, 3.0	X16, 2.0



AMD FIREPRO™ SERVER GRAPHICS


Product Family Comparison



	AMD FirePro™ S10000	AMD FirePro™ S9000	AMD FirePro™ S7000
ASIC Model	Tahiti (2x)	Tahiti	Pitcarin
Form Factor	FH/FL Dual-slot	FH/FL Dual-slot	FH/FL Single-slot
Core Clock Speed	825 Mhz	900 Mhz	950 Mhz
Total Cores	3584	1792	1280
Max Power	375W	225W	150W
Peak Single Precision	5.91 TFLOPS	3.23 TFLOPS	2.4 TFLOPs
Peak Dual Precision	1.48 TFLOPS	806 GFLOPS	152 GFLOPS
Memory Size / Type	6GB GDDR5	6GB GDDR5	4GB GDDR5
Memory Interface	384-bit	384-bit	256-bit
Memory Bandwidth	480 GB/S	264 GB/s	154GB/s
PCI Express®	3.0	3.0	3.0
Display Output	4x DisplayPort 1x DVI	1x DisplayPort	1x DisplayPort
Part Numbers	100-505779 (Retail) 100-505772 (OEM)	100-505748	100-505749



FirePro S7000



3x Single Precision

2x Memory Bandwidth

FirePro S9000 S1000



3x Double Precision

High Density

BEST GPU DENSITY



- Most efficient HW architecture and performance
- High Density system supported

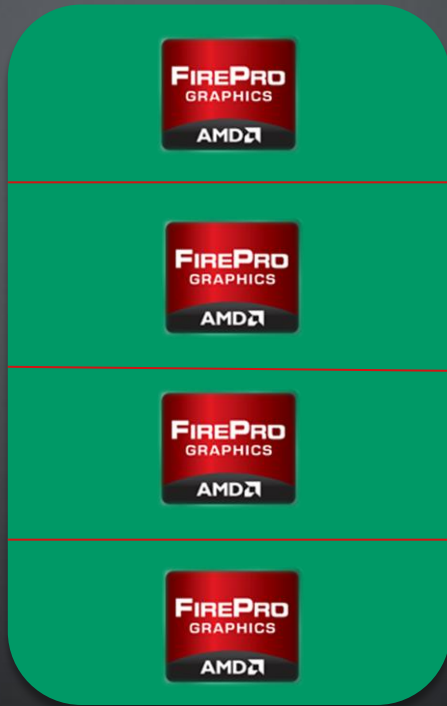
System	GPU Cards	Total SP <u>Tera-FLOPS</u>
Tyan FT72B7015	8 x S10000	47.28
Asus ESC400 G2	2 x S10000	11.82
Dell R720	4 x S7000	9.6
Dell T620	4 x W7000	9.73
Dell PEC 8220	16 x S9000	51.68
SuperMicro	4 x S9000	12.92
Tyan	8 x S10000	47.28



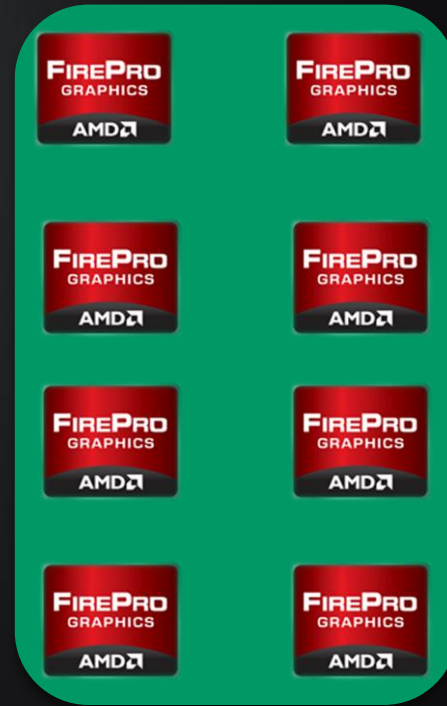
HIGH DENSITY



Reduce power and cost and increase performance



4 x S9000 GPU /
server
13 TFlops SP
3.2TFlops DP



4 x S10000 GPU /
server
23.6 TFlops SP
5.92TFlops DP



AMD AMBIDEXTROUS WITH HSA



Different solutions for each workflows



APU

Low Power
High Density

- High Performance/Watt/\$ ratio
- Balanced CPU/GPU
- Flexible workflows



GPU

High End GPU
Dual GPU card
High GPU density

Ultra High Performance/Watt/\$ ratio
High Performance parallel workflows

- OpenCL
- OpenMP with PGI, gcc
- Java, Python
- Scientific Libraries

Developer Tools

