




Penguin Computing Delivers AI

Sid Mair, SVP Federal Business

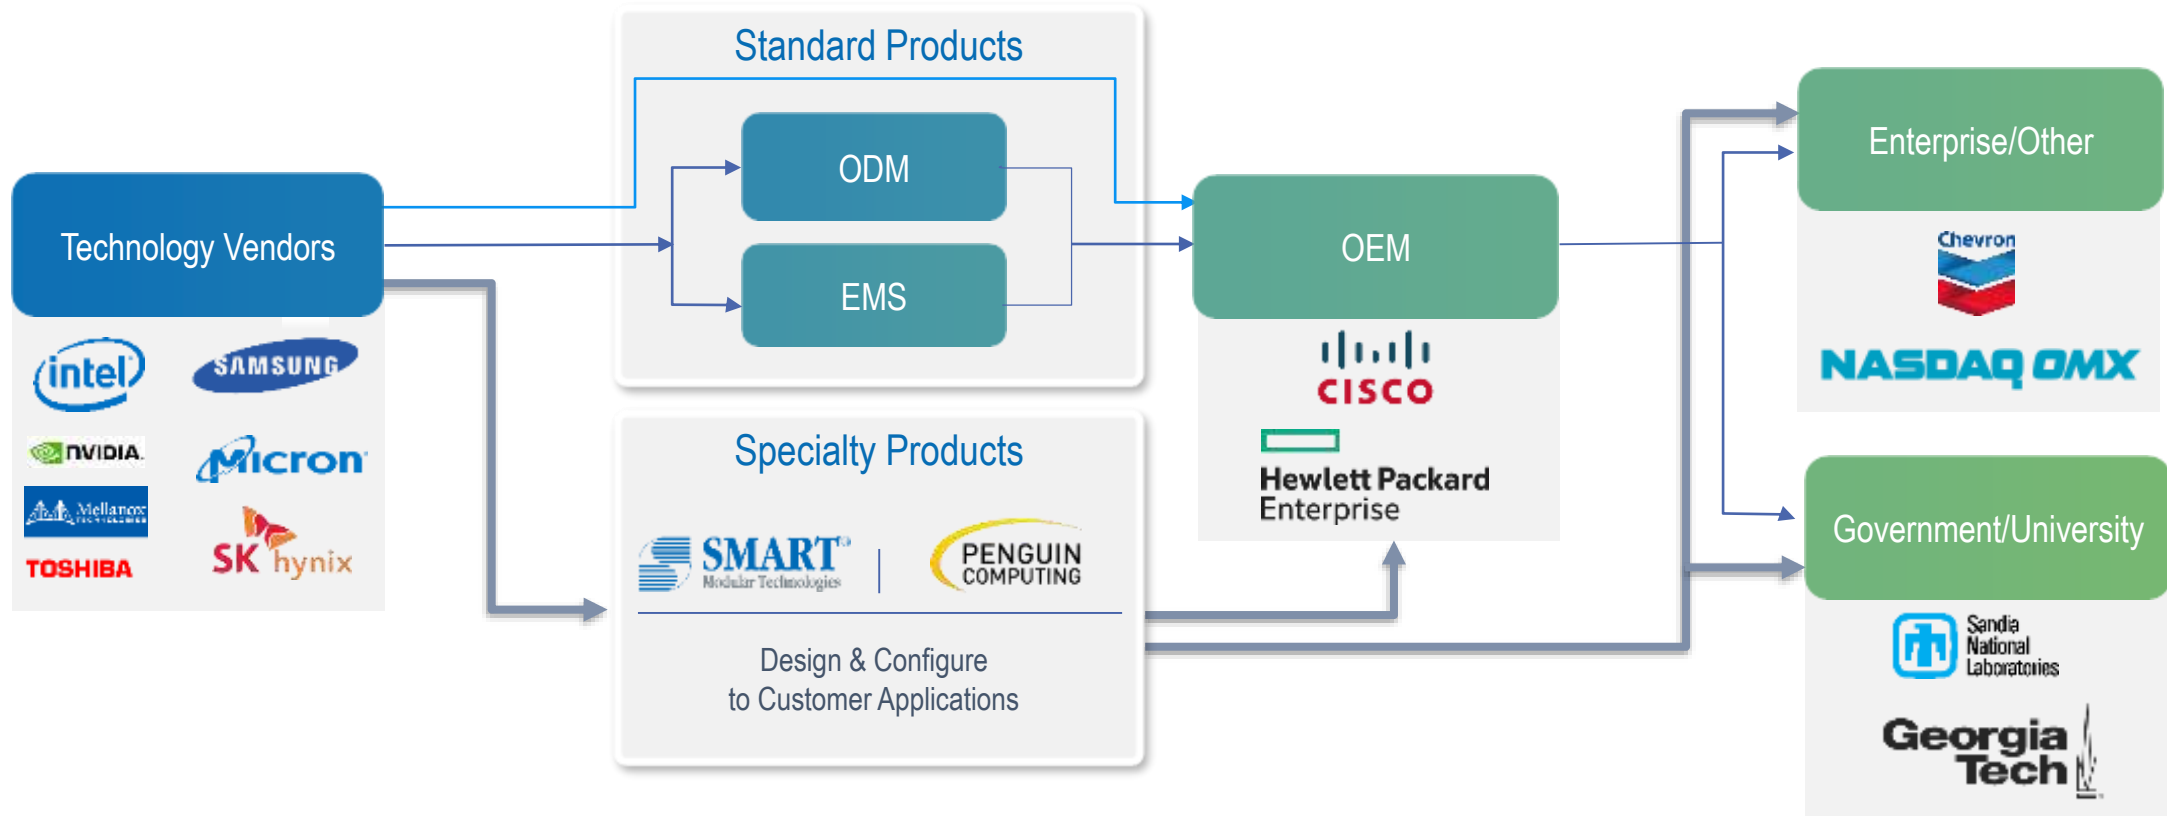
SMART Global Holdings

SMART's ability to rapidly integrate the latest, best of breed technologies with our IP provides differentiated, class-leading solutions to our customers in a diverse range of markets.



We integrate technology from multiple sources with our own technologies / IP to enable solutions for an application or geography.

SMART is Uniquely Positioned in the Technology Ecosystem



We work in partnership with key technology vendors in memory, computing and storage to define product and market solutions for application specific or market/geographic specific needs

Delivering End-to-End AI Capabilities

Designing & Building

- GPU-Accelerated Computing Systems
- Balanced AI Systems
- NVIDIA Partner of the Year Americas 2017
 - Driven by 778% increase in NVIDIA business over last four years

Professional Services

- Managed Services
- Hosting Systems



Full Set of GPU-Accelerated Computing Options



1U	Processor	PCIe Slots	GPU(s) Supported
Relion XE1114GT	Intel® Xeon® Scalable Processors	4x PCIe Gen3 x16 (GPU), 2x PCIe Gen3 x16 (LP-MD2)	Tesla P100-PCIe, Tesla V100-PCIe
Relion 1903GT	Intel Xeon E5-2600 v3 / v4 Series	3x PCIe Gen3 x16 (GPU), 1x PCIe Gen3 x8 (LP)	Tesla P40, Tesla M40-24GB, Tesla K80
Relion 1904GT	Intel Xeon E5-2600 v3 / v4 Series	4x PCIe Gen3 (GPU), 2x PCIe Gen3 x8 (LP)	Tesla P100-16GB, Tesla P100-12GB, Tesla P40, Tesla M40-24GB, Tesla K80
2U	Processor	PCIe Slots	GPU(s) Supported
Relion XE2112GT	Intel® Xeon® Scalable Processors	2x PCIe Gen3 x16 (GPU), 2x PCIe Gen3 x8 (LP), 2x OCP Mezz	Tesla P100, Tesla V100
Relion XE2118GT	Intel® Xeon® Scalable Processors	8x PCIe Gen3 x16 (GPU), 1x PCIe Gen3 x16 (LP), 1x PCIe Gen3 (Proprietary Mezz)	Tesla P100, Tesla V100
Relion 2903GT	Intel Xeon E5-2600 v3 / v4 Series	3x PCIe Gen3 x16 (GPU), 2x PCIe x8 (FHHL), 1x PCIe Gen3 x8 (Proprietary Mezz)	Tesla P100-16GB, Tesla P100-12GB, Tesla P40, Tesla M40-24GB, Tesla K80
Relion 2904GT	Intel Xeon E5-2600 v3 / v4 Series	4x PCIe Gen3 x16 (GPU), 1x PCIe Gen3 x8 (LP), 1x PCIe Gen3 x8 (Proprietary Mezz)	Tesla P100-16GB, Tesla P100-12GB, Tesla P40, Tesla M40-24GB, Tesla K80
Relion 2908GT	Intel Xeon E5-2600 v3 / v4 Series	8x PCIe Gen3 x16 (GPU), 1x PCIe Gen3 x8 (LP), 1x PCIe Gen3 x8 (Proprietary Mezz)	Tesla P100-16GB, Tesla P100-12GB, Tesla P40, Tesla M40-24GB, Tesla K80
4U	Processor	PCIe Slots	GPU(s) Supported
Relion XE4118GT	Intel® Xeon® Scalable Processors	8x PCIe Gen3 x16 (GPU), 2x PCIe Gen3 x16 (LP)	Tesla P100, Tesla V100
Relion XE4118GTS	Intel® Xeon® Scalable Processors	8x NVIDIA SXM2 (GPU), 2x PCIe Gen3 x16 (LP)	Tesla P100-SXM2, Tesla V100-SXM2
NVIDIA DGX-1	Intel Xeon Scalable Processors		Tesla V100-SXM2
NVIDIA DGX-2	Intel Xeon Scalable Processors		Tesla V100-SXM2
IBM Power AI	IBM Power Processors		Tesla V100-SXM2

Building Balanced AI Systems

Featuring

- Penguin Computing proprietary computing and storage technology
- IBM Power AI Enterprise and Spectrum Scale
- Plus NVIDIA DGX-1, DGX-2, and software stacks



Penguin on Demand (POD)

- Hybrid Cloud offerings virtually identical to on-prem work flows.
- Container and bursting support
- Bare metal - 25,000 Cores, OPA/Infiniband, 6 Que's
- High speed Parallel File System
- Core hour volume pricing, 3s billing resolution
- Generally \$0 for data movement
- Specialized PhD subject matter expertise
- Intel's 2017 Cloud Vendor of the Year



www.penguincomputing.com

1-888-PENGUIN



Platinum Sponsor

WASHINGTON, D.C.

OCTOBER 22-24, 2018

[The Ronald Reagan Building and International Trade Center](#)
[1300 Pennsylvania Ave NW, Washington, D.C. 20004](#)