

Outlook for HPC, AI, and the Cloud

June 2018



Intersect360
R E S E A R C H

Addison Snell

Christopher Willard, Ph.D.

Laura Segervall



Intersect360
R E S E A R C H

Intersect360 Research in 2018




More than 10 year history of HPC analyst business




Covering “high performance data center” markets, including traditional HPC and supercomputing, high-performance enterprise, cloud, big data, and hyperscale




New hyperscale advisory service



Weekly podcast, “This Week in HPC”



Market research partners of TOP500.org, contributing editors to The Next Platform, media partnerships with HPCwire, insideHPC




Sponsors of HPC Advisory Council




Just in! HPC Market for 2017

- Total worldwide HPC market (servers, storage, software, etc.) reached \$35.4 billion in 2017, up 1.6% from 2016.
- Demand-side indicators were up; supply-side indicators were down. What happened?
 - Commercial users moved many facilities costs out-of-budget, masking a large increase in spending on power and cooling for those that did not.
 - Cloud had a breakout year.
 - Storage was weak relative to server spending.
 - The software and services categories both declined.

HPC 2017 vs. 2016



HPC Cloud 2017 vs. 2016

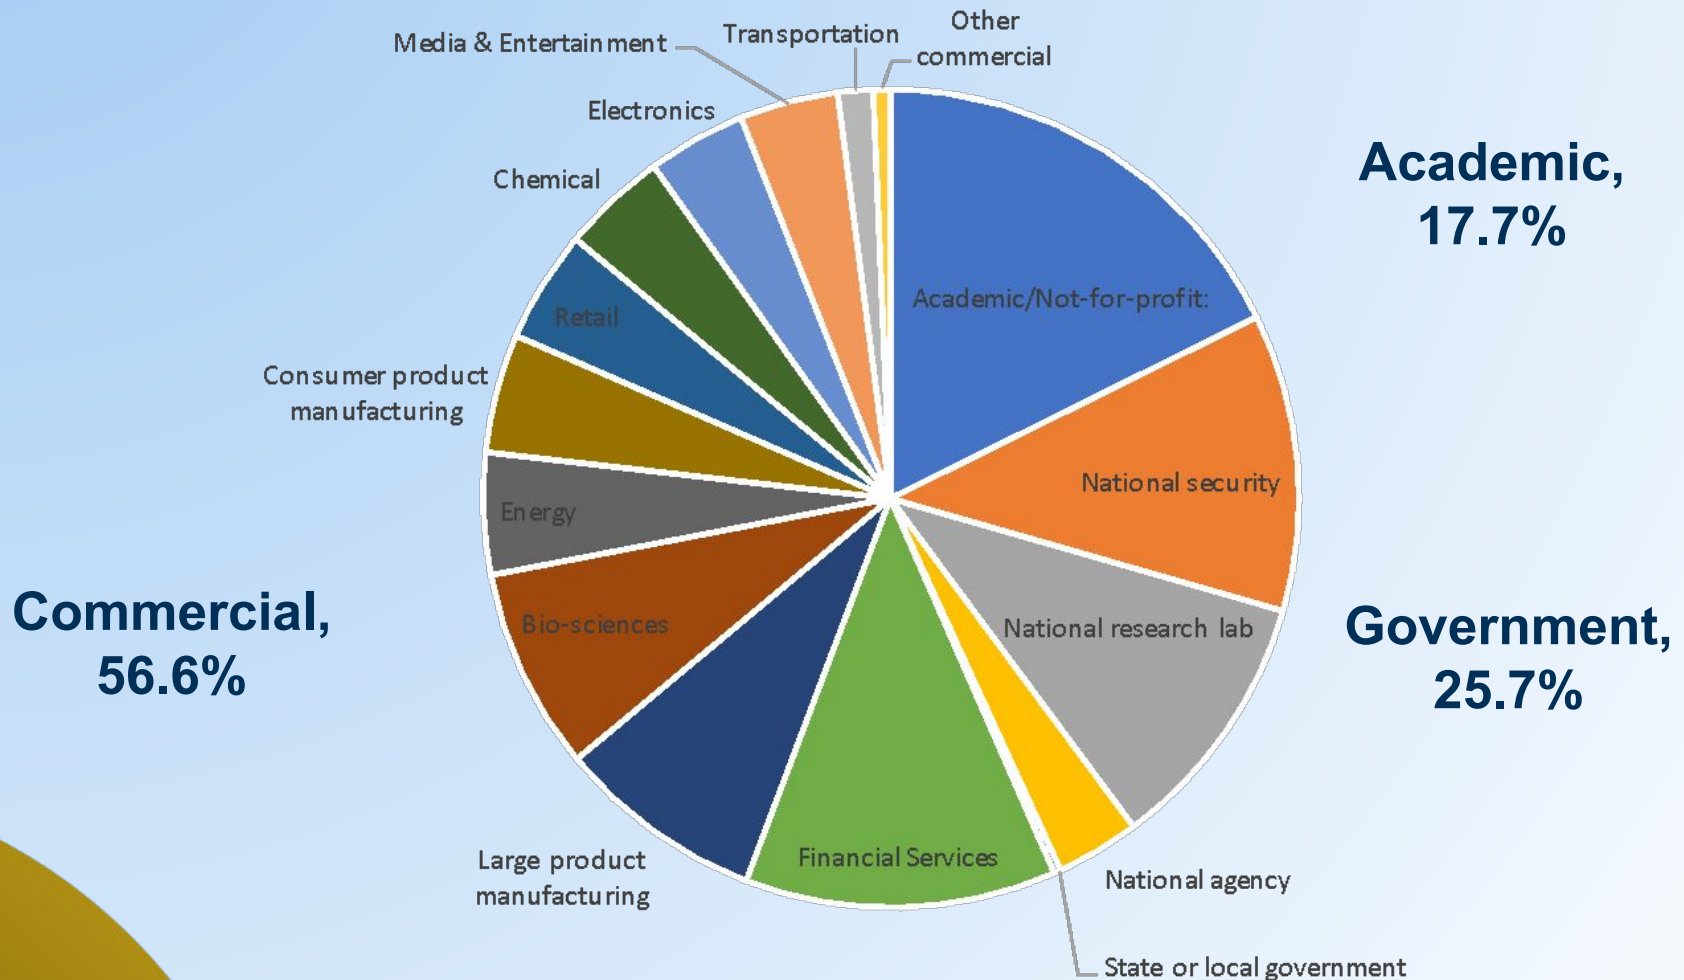


**Raw cycles grew only
along with overall
market for compute**

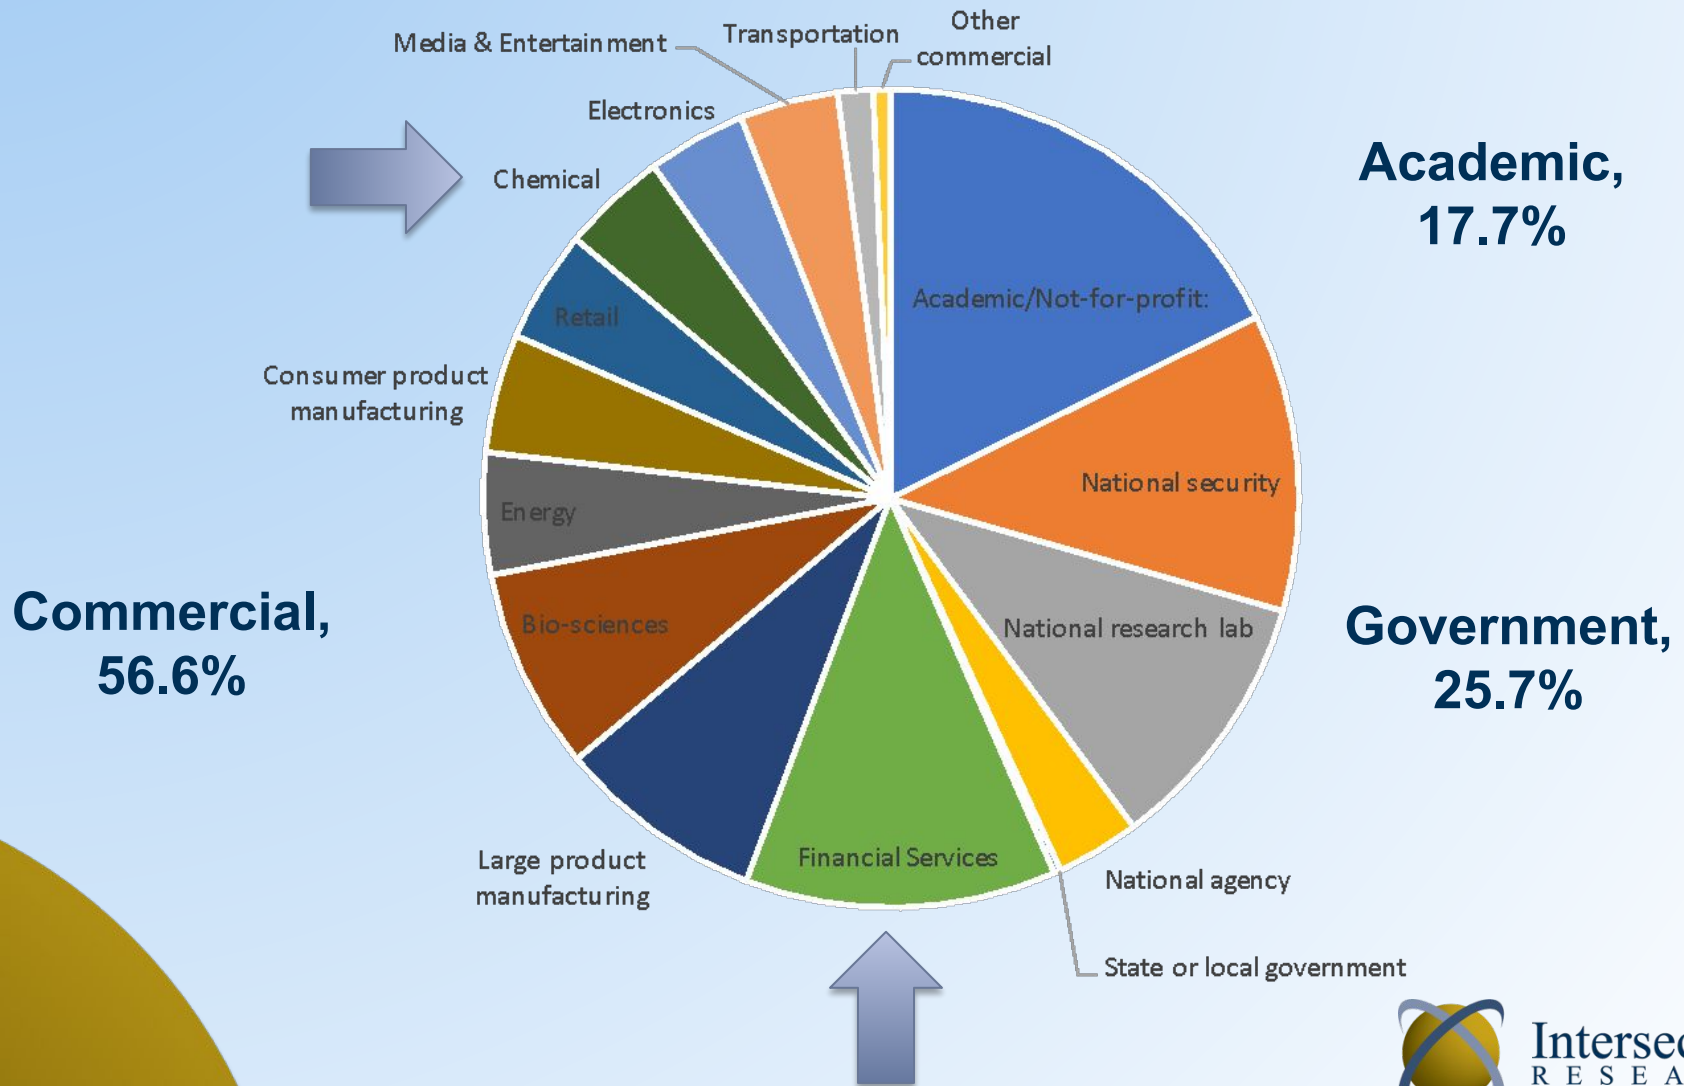
Fundamental Drivers

- Drivers for HPC:
 - There is always more science to do
 - There are always harder problems to solve
 - Most HPC usage is commercial
 - Plus: AI!
- Drivers for HPC in Cloud:
 - Suitable to variable workloads
 - Availability of high-performance components
 - Maturation of licensing models
 - Plus: AI!

HPC 2017 Revenue by Vertical



HPC 2017 Revenue by Vertical



AI / Deep Learning

- In 2016, Intersect360 Research placed AI in the hyperscale market
 - Similar to but distinct from HPC
 - Low precision, intensely parallel, strong affinity to public cloud
- DL is not a vertical market. It is an algorithmic or computational approach
- We are monitoring *spending patterns*
 - *Are budgets increasing?*
 - *What is the effect on configurations?*
 - *On premise versus cloud*

Sizing Deep Learning

- The great majority of dedicated spending is by hyperscale companies, not HPC
- About half of HPC organizations are running deep learning applications
 - Usually with same hardware as HPC, perhaps with consideration to configuration
 - Little effect on HPC budgets compared to what was already planned
 - To the extent that it is done as part of an HPC environment, it is already counted in our numbers

(For further reading, see “Big Data,” 2013)

The Big Challenge

- High-performance workloads are diversifying
 - Traditional HPC
 - Big Data / Analytics
 - AI / Machine Learning / Deep Learning
- More options in architectures
 - Processor options (x86, GPU, FPGA, ARM, ...)
 - Storage tiers (NVM, burst buffer, disk, archive ...)
 - Networking diversity (Ethernet, IB, OPA, ...)
- What to buy, with limited resources?

Future Outlook

- Cloud continues fast growth in HPC
 - Especially this year and next
 - ~\$2B in 2019, ~\$3B in 2022
- Almost all HPC Cloud usage will be hybrid
- AI remains a critical workload, becomes part of HPC (and other areas), not replacing
- Users need vehicles for reducing cost and risk in matching technologies to new workloads (cloud, managed services, ...)

Outlook for HPC, AI, and the Cloud

June 2018



Intersect360
R E S E A R C H

Addison Snell

Christopher Willard, Ph.D.

Laura Segervall



Intersect360
R E S E A R C H