



# Software defined infrastructure for hyperscale Data processing and HPC

Areg Melik-Adamyman, Engineering Manager



# Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Copyright © 2015, Intel Corporation. All rights reserved. Intel, Pentium, Xeon, Xeon Phi, Core, VTune, Cilk, and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

## Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804


# Cloud Infrastructure Enables New Usages

## Cloud 2015 Digital Services Economy



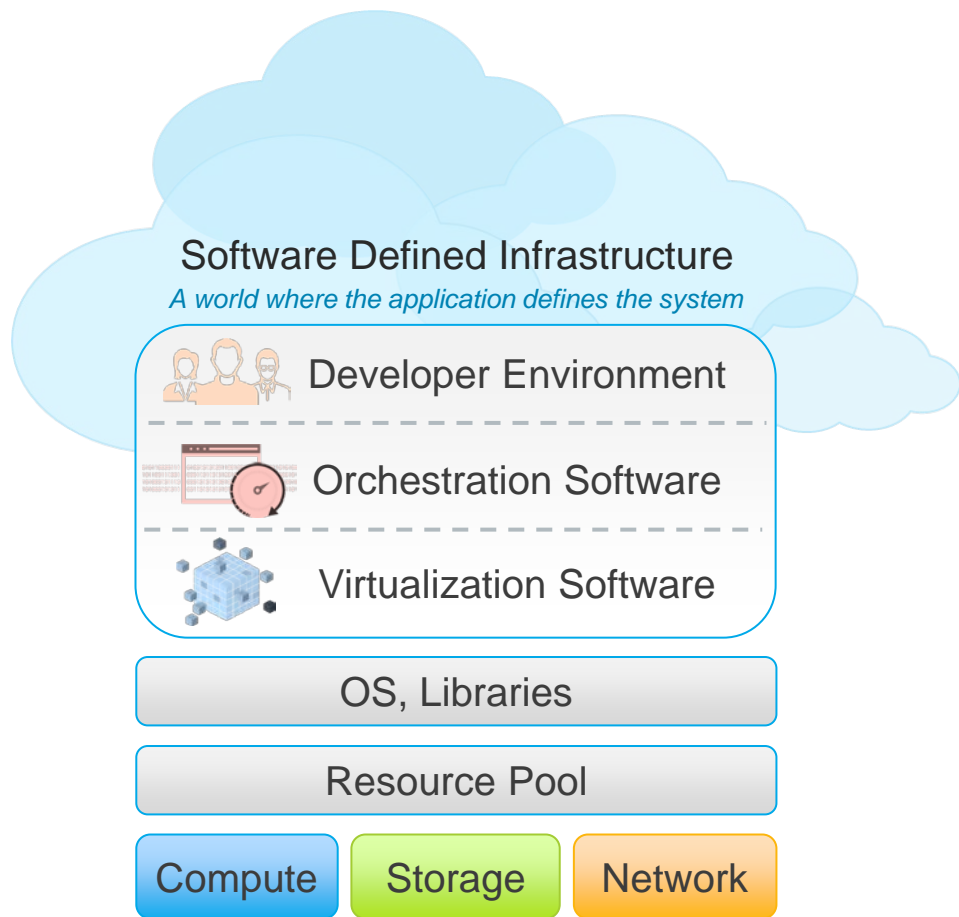
75% of current cloud demand comes<sup>1</sup> from consumer services

## Cloud 2020 IoT, Big Data, and Enterprise



By 2020, 65-85%<sup>2</sup> of apps will be delivered via cloud infrastructure

# But... Cloud Technology is Too Complex



## Merging Technology Paths

- Enterprises seeking Cloud Native Apps
- Hyperscale Providers expanding into traditional workloads

## Fragmented Cloud Ecosystem

- Growing Cloud ecosystem
- Limited out of the box solutions

# Industry Is Not Moving Fast Enough

## Hyperscale



Deliver the best performance / TCO

Continuous workload optimization

Volatility of demand

## Fast Followers




Deliver differentiated services at hyperscale economics

Difficult to achieve high efficiency and scale

Fragmented cloud stacks with significant feature gaps

## Broad Enterprise



Enable business innovation through the cloud

Proprietary solutions can be costly to deploy

Open source stacks are complex and lack enterprise features

GOAL

CHALLENGE

# Ultimate Need for Simple and Effective Cloud Infrastructure

## Hyperscale




Deliver the best performance / TCO

## Fast Followers



Deliver differentiated services at hyperscale economics

## Broad Enterprise



Enable business innovation through the cloud

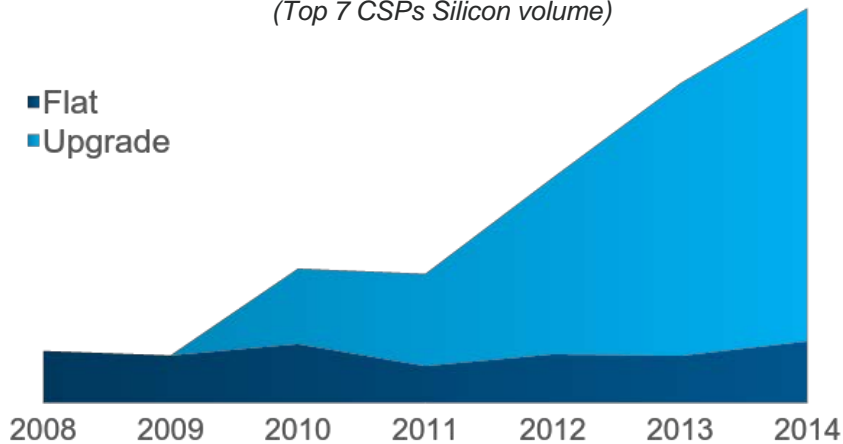
### INTEL STRATEGY:

- 1 **Optimize cloud infrastructure across a full range of workloads**
- 2 **Improve Software Defined Infrastructure efficiency and ease deployment**
- 3 **Collaborate with industry leaders**

# 1 Optimize Cloud Infrastructure across a range of workloads

## Optimizing to take full advantage of hardware capabilities

Intel Xeon processor trends  
(Top 7 CSPs Silicon volume)



Shipping custom Si to all hyperscale CSPs



2x+ performance on programmable algorithms



Early adoption of high performance components


Up to 50% cost reduction in service delivery



# Make SDI efficient and easier to deploy

2

So far...



What's next

## Accelerating technology maturity

- Enterprise ready: reliability features, installation, and operations
- Scale to thousands of nodes
- Intelligent automation and orchestration

## Exposing underlying hardware technology

- Support for Xeon, Xeon Phi, SoCs, ...
- Workload specific accelerators (eg: FPGA)
- Expose platform telemetry
- Hardware-enabled security

Reduce infrastructure deployment time from weeks to hours



# Align the industry

3

## Reference Architectures

vmware®

REDAPT



TECTONIC  
by CoreOS

## Open Solutions and Industry Standards



Open Container Project



Clear Linux Project

Rackscale Architecture

## Routes to market

Cloud SW

OEMs

VARs

Enable thousands of new cloud deployments

# Making Cloud Simple & Efficient WILL Accelerate USE

## Intel Strategy

1

Optimize cloud infrastructure across a full range of workloads

2


Make SDI efficient and easier to deploy

3

Align the industry

- Up to 50% savings in relative cost per VM instance
- Reduce infrastructure setup time from weeks to hours
- Enable thousands of new cloud deployments

## Emerging Cloud Usages 4



# Different Workloads - Two Large-Scale Systems

## Modern supercomputer



- Run programs in hours or days that would require decades or centuries on normal machine
- Designed for numerically-intensive applications

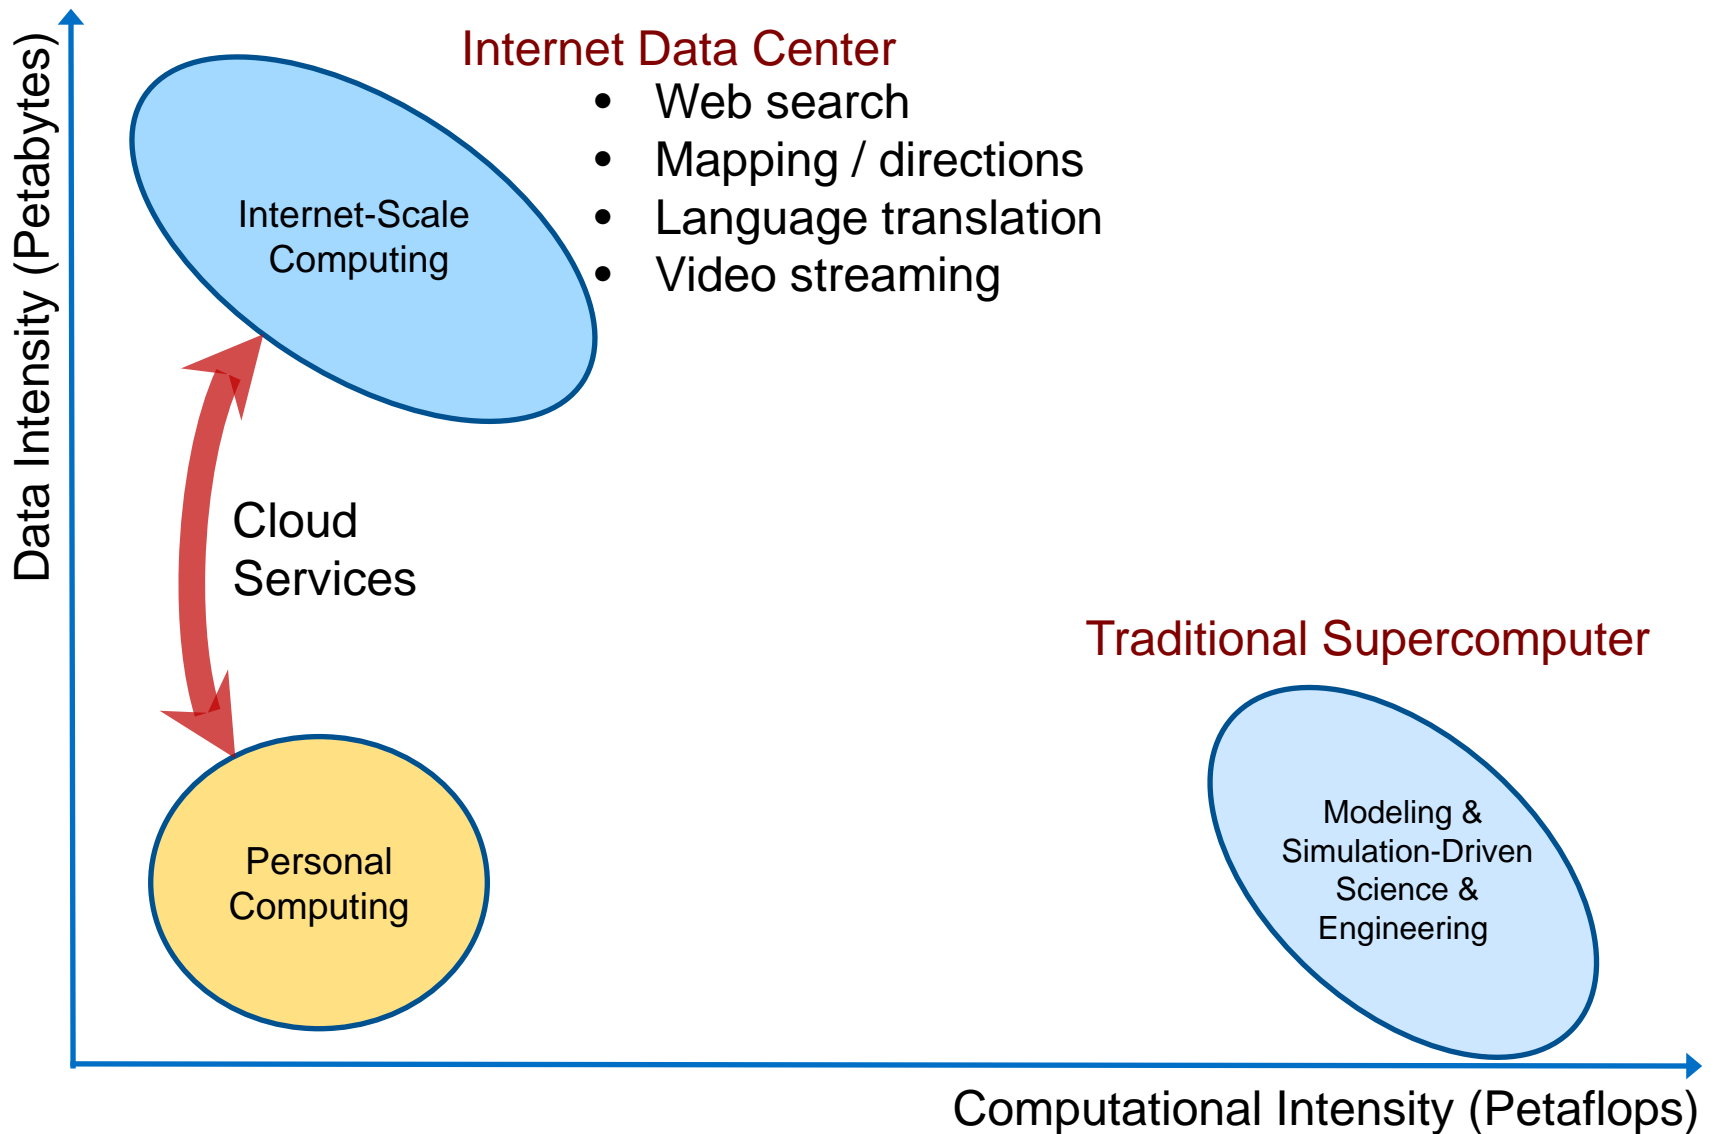
## Internet Data Center



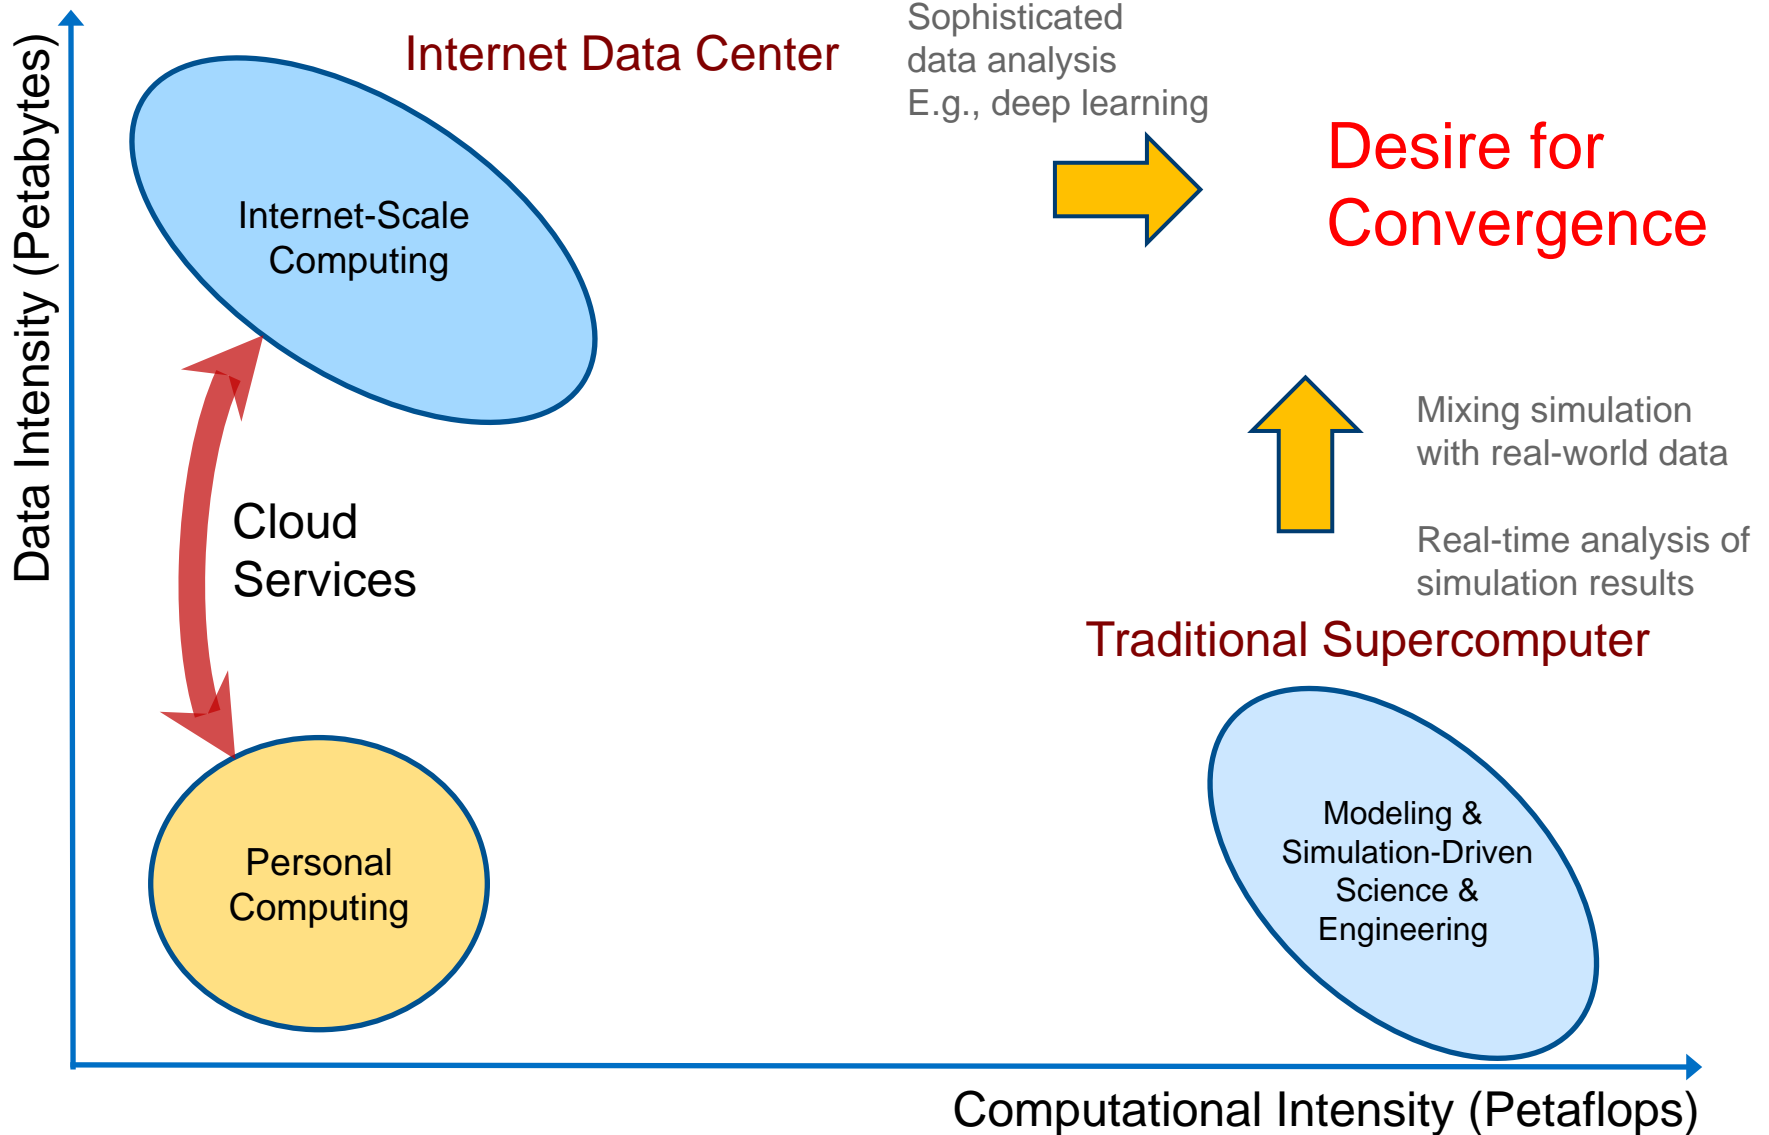
- Support millions of customers
  - Mostly small transactions
  - ...and large-scale analytics
- Designed for data collection, storage, and analysis

NCSI August, 2015

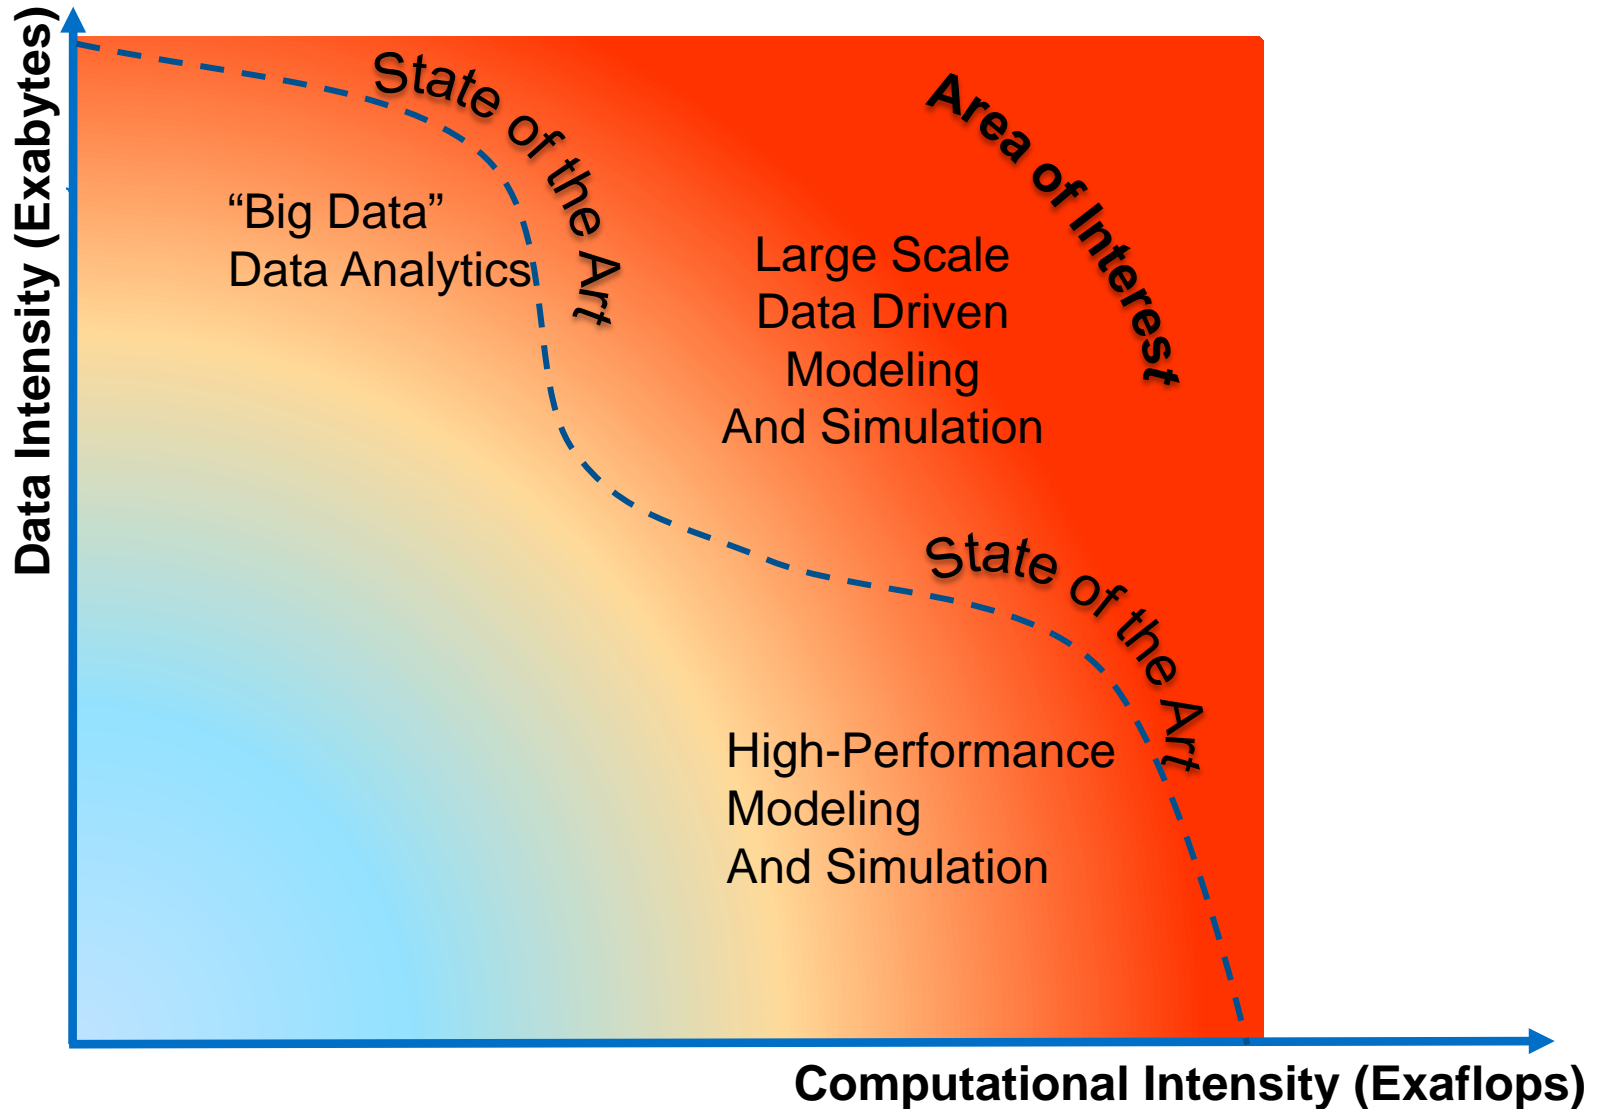
# Computing Landscape



# Computing Landscape



# Aspirations for Convergence



NCSI, 2015

# Observations

## Cannot Measure a Machine by Single Benchmark

- Top500 benchmark: Perform Gaussian elimination of dense matrix
- Not representative of real-world applications

## Cannot Measure a Country's HPC Capabilities by Its Biggest Machine

- Want diversity (size and design) of machines
  - Encourages use and innovation
- Want them widely deployed
- Want well-developed software infrastructure
- Want machines to be running useful applications across many disciplines
- Want research program to ensure continued progress

NCSI, 2015



# Modern Supercomputer Programming

## System Level

- Message-Passing Interface (MPI) supports node computation, synchronization and communication

## Node Level

- OpenMP supports thread-level operation of node CPU
- Exploiting large vectors (SIMD) Xeon Phi and CUDA programming environment for GPUs
  - Performance degrades quickly if don't have perfect balance among memories and processors

## Result

- Single program is complex combination of multiple programming paradigms
- Tend to optimize for specific hardware configuration
  - “As the performance of HPC machines approaches infinity, the number of people who program them is a approaching zero.”
    - Dan Reed

NCSI, 2015

# Supercomputer Programming Model

- Program on top of bare hardware

## Performance

- Low-level programming to maximize node performance
- Keep everything globally synchronized and balanced


## Reliability

- Single failure causes major delay
- Engineer hardware to minimize failures

## Drawbacks

- Application development requires resources and expertise
- Hard to port applications between systems
- Low degree of software reuse

Machine-Dependent  
Programming Model



# Example Cluster Programming Systems


## Hadoop



- Parallel file system aggregating disk drives across cluster
- Map/Reduce programming model providing data parallel programming abstraction

## Spark Project

- at U.C., Berkeley
- Grown to have large open source community



## GraphLab

- Environment for describing machine-learning algorithms
  - Sparse matrix structure described by graph
  - Computation proceeds by updating node values asynchronously



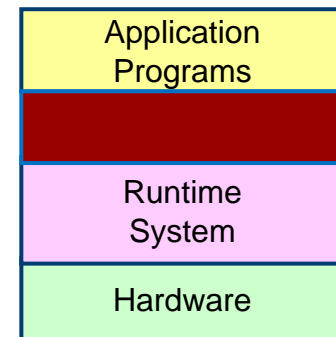
# Cluster Programming Model

- Application programs written in terms of high-level operations on data
- Runtime system controls scheduling, load balancing, fault tolerance

## Performance Challenges

- Centralized scheduler forms bottleneck
- Copying to/from disk very costly
  - vs. memory-resident computation
- Hard to limit data movement
  - Significant performance factor

Machine-Independent  
Programming Model



# Application Developer's Dream

Develop application based on well-defined abstractions

- Computation
- Data organization
- Resource allocation
- Fault tolerance

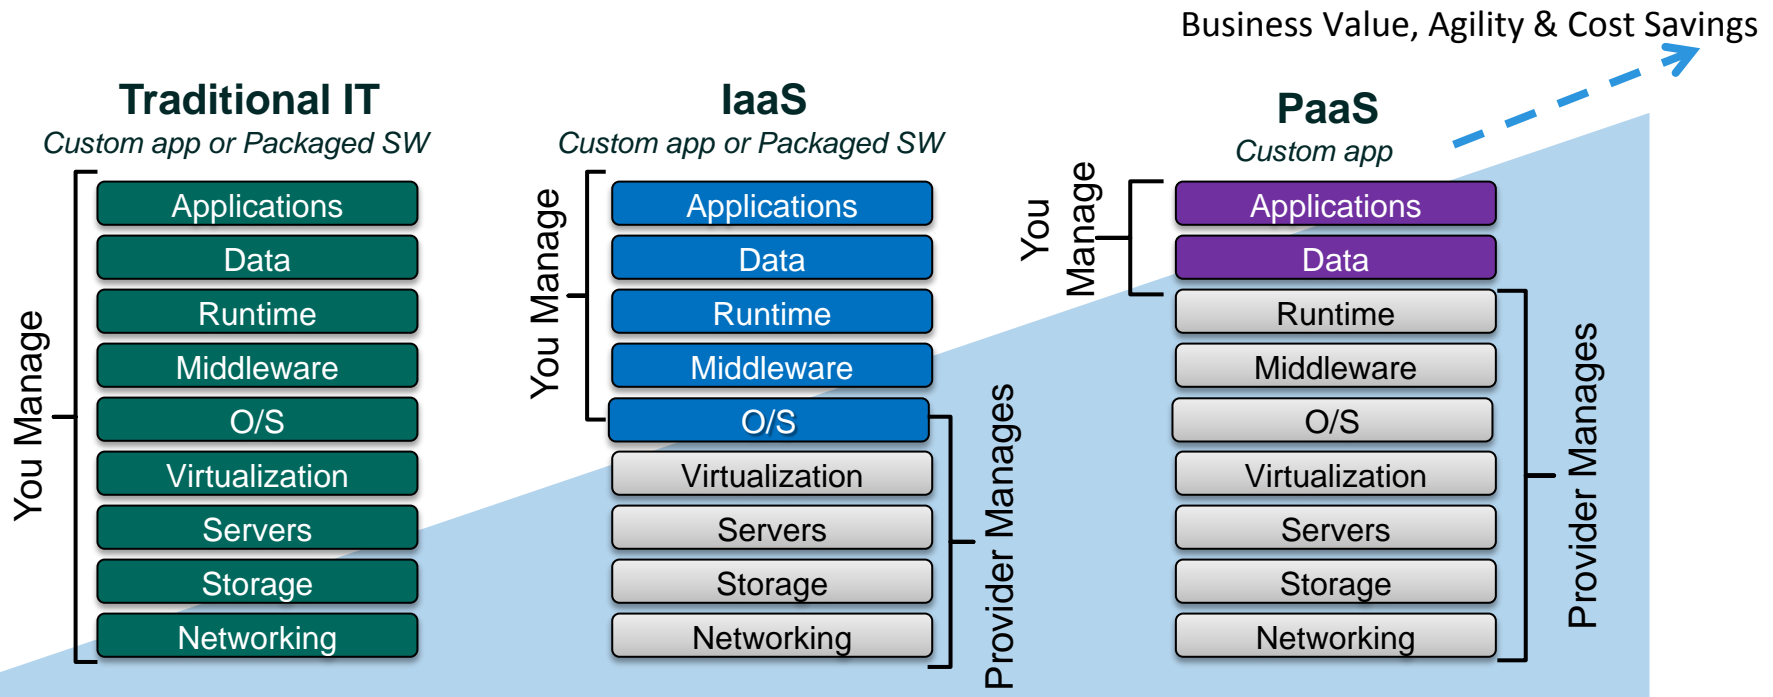
Libraries, compilers, autotuners take care of machine-specific mappings

- Performance comparable to what can be obtained by hand tuning

Rich collection of domain-independent and domain-specific software modules

- Enable sharing and reuse of software

# Let Some Dreams Come True - PaaS

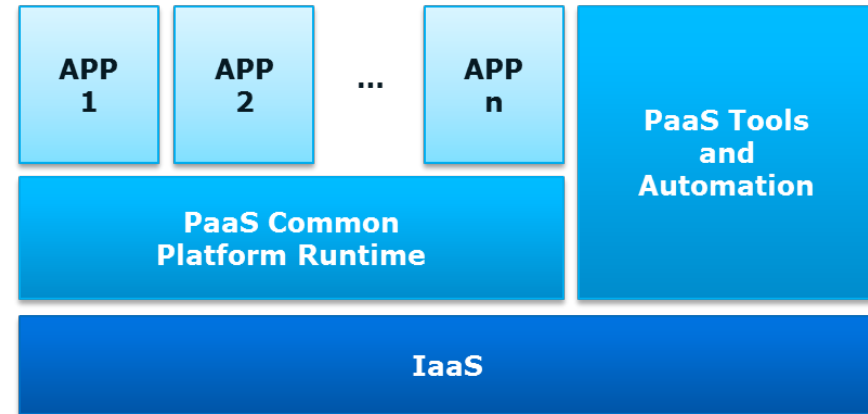


**PaaS enables developers to use self-service to build, deploy and manage their custom apps**

# Why PaaS?

## Developers code their app & deploy into production without IT assistance

- Cloud tooling: self-service, on-demand, multi-tenant, metered
- Pre-provisioned common platform of abstracted middleware & infrastructure



## Facilitates creation of cloud-ready applications

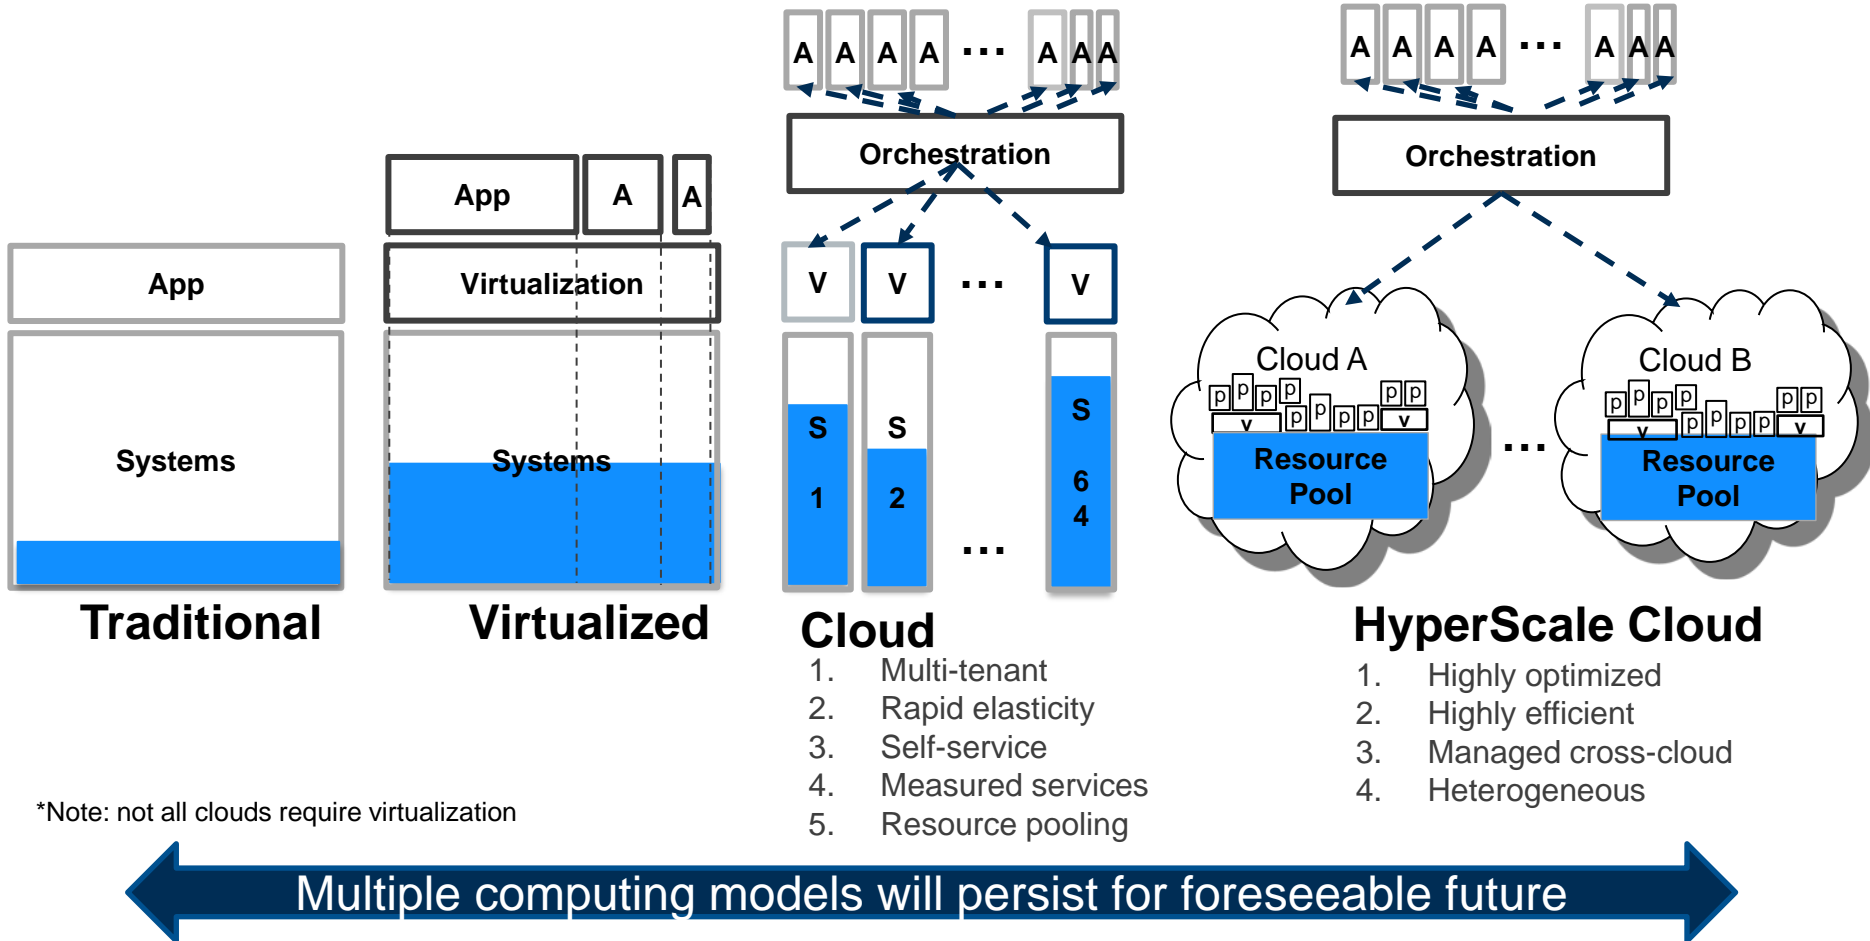
- Platform provides runtime container, elastic scaling and high availability
- Maximize resource sharing via multi-tenancy and reusable web services

## Speed to market and reduced SW delivery risk

- Reduced time to deploy a new app from 6 months to minutes
- PaaS helps accelerate innovation improving margin vs. reducing cost (IaaS)



# Computing Progression for SDI



# Go Application Development Platform

Build applications for Distributed, Concurrent & Cloud Computing Applications

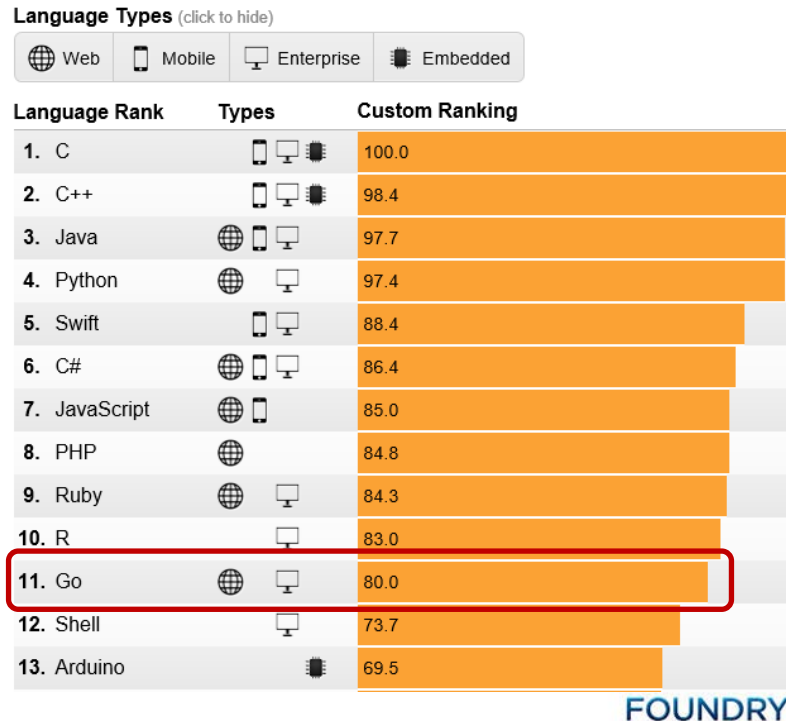
## Key features

- Statically typed, g support at language
- Support for concu and faster compil
- Productivity of a d
- Multiplatform, Por

## Appeals to..

- SDI domain
- Python & Ruby de performance
- C programmers th productivity

## Challenges



community support



# Call to action!

- Enhance local contribution to Go and PaaS (e.g. CloudFoundry)
- Research the different cloud models – fun and profitable (maybe)

