Hewlett Packard Enterprise

April 3, 2017.

# HPC Innovation in the FSI Market

## HPE Keynote Session – HPC on Wall St.

Lacee McGee, Sr. FSI Vertical Manager, HPE

# Table of Contents

# HPC in the Enterprise is driving change

## HPC in the Enterprise

| | Traditional HPC | New applications | Tiered Storage | Delivery Flexibility |
|---|---|---|---|---|
| **Challenges** | • Managing complex relationships<br>• Models increasing in size and have inefficiencies when scaling over clusters<br>• Microprocessor improvements limited by slowing up scaling | • Need for readily available data access & real time analytics<br>• Increasing data demands due to IoT & Artificial Intelligence<br>• Divide between machine learning and HPDA algorithms<br>• Energy cost associated with moving data | • Data overload<br>• Compliance Regulations<br>• Technical Computing Apps require petabyte level storage<br>• Traditional storage methods are no longer economical | • Capacity planning is hard<br>• 59% of enterprises face +3 month delay<br>• 50% of enterprises have suffered downtime as a result of poor capacity planning<br>• Overprovisioning is the norm |
| **Trends** | • Performance improvements require code modernization<br>• Shift from processor centric computing to pooled resources<br>• Increasing number of connections per node | • HPDA is inevitable<br>• In memory solutions dominant by 2019<br>• New applications with more parallelism<br>• Companies turning to artificial intelligence for deeper insights | • Storage tiers collapsing<br>• Data management Software becoming the norm<br>• Shift from compute centric to data friendly configurations | • Cloud (Public, Private, Hybrid-IT)<br>**FSI Cloud Adoption**<br>  <br>• Flexible consumption models<br>• HPC as a Service |
| | intel Hewlett Packard Enterprise |  Hewlett Packard Enterprise | Hewlett Packard Enterprise SUSE We adapt. You succeed. |  |

Hewlett Packard Enterprise

# How is HPE addressing these needs?

## Attacking the HPC Market

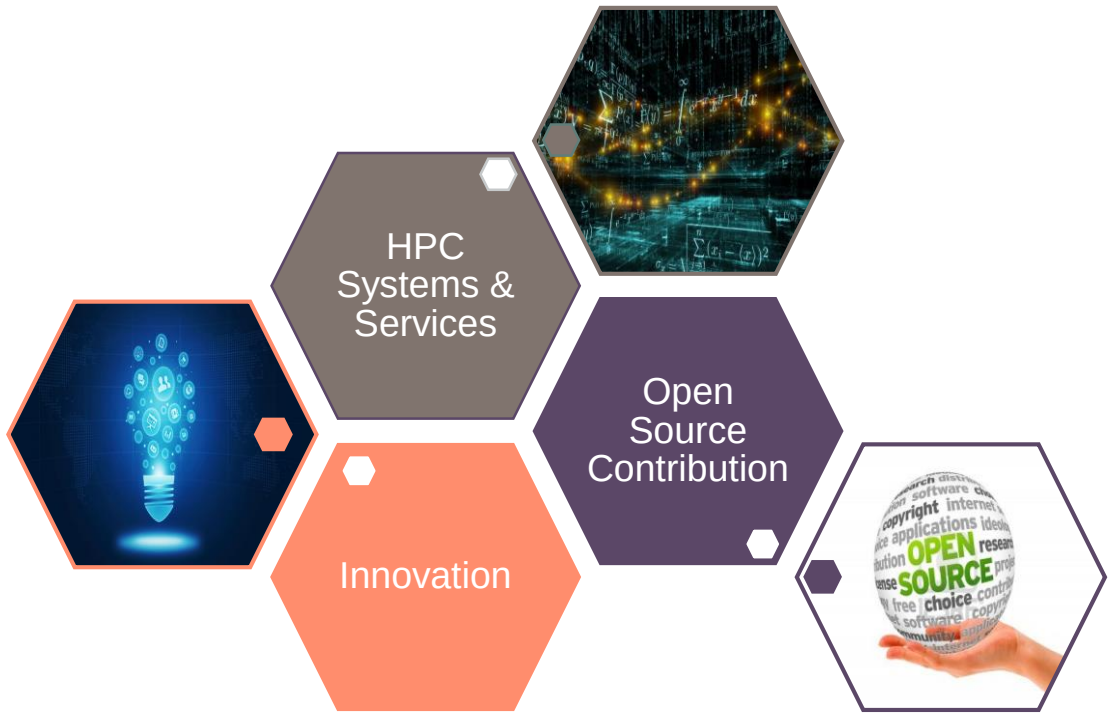


## Industry Leading Partner Ecosystem

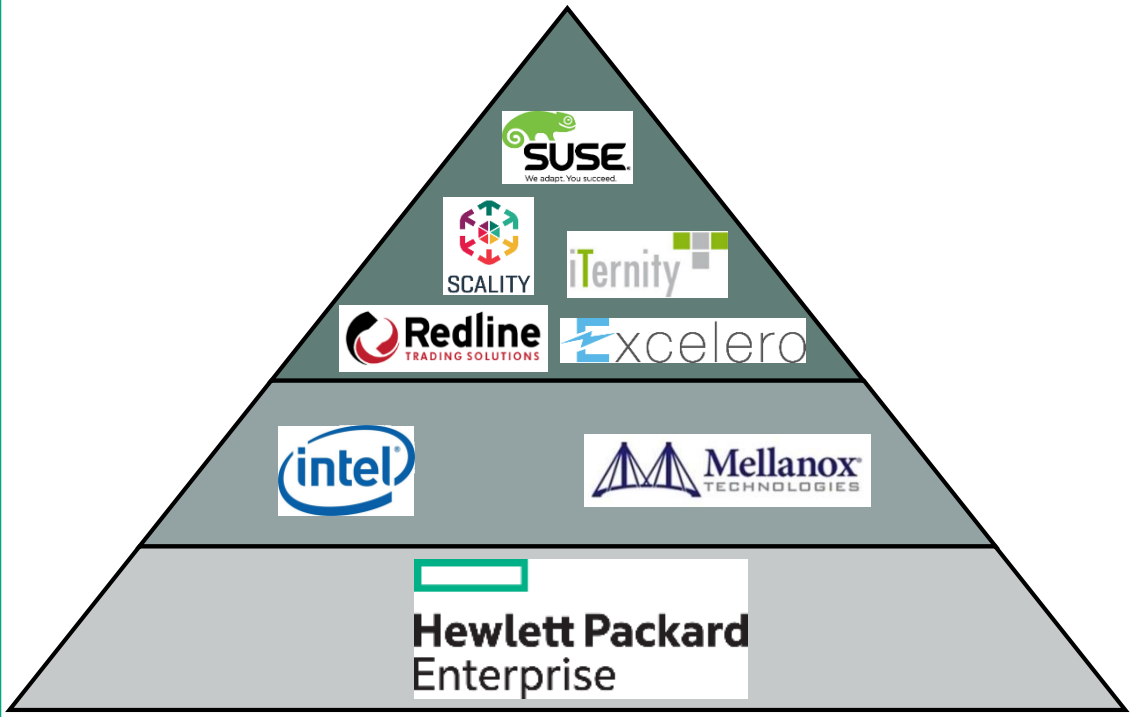


**End-2-End Solutions**

**Center of Excellence**

**Workload Optimized**

# HPE purpose-built portfolio for High Performance Computing

**HPC Industry Solutions**

Financial Services | Government & Academia | Life Sciences | Manufacturing | Oil & Gas, Energy | Weather & Climate Research

**HPC Services** — HPC Advisory Services, HPE Datacenter Care for HPC and Support Services, Flexible Capacity, Financial Services

## Supercomputing / Enterprise / Commercial HPC

**HPE SGI® 8600**

Liquid cooled, delivering industry leading performance, density, & efficiency

**HPE Apollo 8000**

Advancing the science of supercomputing

**HPE Apollo 6000 Gen10**

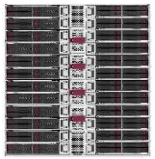

Extreme Compute Performance in High Density

**HPE Apollo 6000**

Rack-scale HPC

**HPE Apollo 2000**

The bridge to enterprise scale-out architecture

## HPE Performance Software Suite

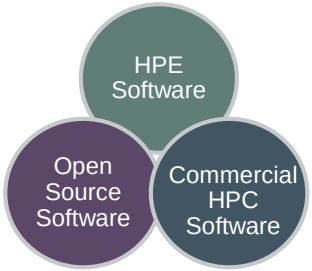


- **HPE Performance Software - Core Stack**
- **HPE Insight Cluster Management Utility**
- **HPE Performance Software - Message Passing Interface (MPI)**
- **HPE SGI Management Suite**

## Emerging HPC

**HPE Apollo 6500**

Rack-scale GPU Computing

**HPE Apollo SX40**

**HPE Apollo PC40**

HPC optimized industry standard servers supporting latest NVIDIA GPU technology.

**HPE Apollo KL20**

Intel Xeon Phi self hosted servers

## HPC Storage

**HPE Apollo 4520**

Purpose-built for HPC Storage

**HPE DMF**

HPC Data Management Platform

## HPC In-memory Compute / Networking

**HPE Superdome X**

**HPE MC990 X**

Scale up HPC

**Arista 7500E Series**

Low latency, high bandwidth

Hewlett Packard Enterprise | Intel | Mellanox | SUSE | Seagate

# Maximize High Frequency Trading speed and throughput for greater competitive advantage

Realize 20% faster trade analysis execution with the **HPE Trade and Match Server** solution***

## HPE Trade and Match Server Solution

### Optimized High Frequency Trading performance

Trade and Match Server for FSI with **HPE Apollo 2000**

### Speed

- Optimized for high frequency performance with overclocked processors
- Minimize cache coherent memory operations
- Minimize system latency

### Reliability

- Improved reliability and MTF with enterprise class infrastructure, qualification and services
- Reduce jitter for more efficient CPU utilization

### Costs

- Rightsized compute and storage
- No charge tuning and management tools
- No charge trusted advisory services, POC and deployment

## STAC N1 Public Report (Network Latency)

| STAC BENCHMARK COUNCIL | Latency (Microseconds) | | Best Results of all Submissions | |
|---|---|---|---|---|
| | Mean | Max | Mean | Max |
| **100K MPS** | 2.7 | 13 | Hewlett Packard Enterprise | Hewlett Packard Enterprise |
| **1M MPS** | 2.6 | 18 | Hewlett Packard Enterprise | Hewlett Packard Enterprise |

## HPE Trade and Match Server Benchmarking

- **#1 Latency** Performance; **65% improvement in max latency** for 1M messages per second (MPS)* (**33 microseconds faster** than Supermicro (#2))
- **12% improvement in data transfer speeds to decision engine****

Accelerated **access to data** for decision engines that deliver **faster trading insights** that increases **trading throughput** & **competitive position**

## Win the order execution race with the HPE Trade and Match Server Solution!

Hewlett Packard Enterprise

*As reported by STAC: based on audited STAC-N1™ results from the overclocked HPE Trade & Match servers using a Solarflare network stack. **STAC Report: HPE Overclocked Servers under STAC-N1** – 16 February 2017 https://stacresearch.com/news/2017/02/16/SFC170206.

** Internal HPE test results (SPECint 2006)

*** The HPE Trade and Match Server Solution utilizes specialized overclocking of the processor to deliver 20% higher performance than standard clock rate solutions

# *Open Source and HPC:*
## *Improved Quality, Speed, and TCO*

***Joseph George, Vice President of Solutions, SUSE***
***Twitter: jbgeorge***

Hewlett Packard Enterprise

SUSE

# Name That Project!

# Name That Project!

# Name That Project!

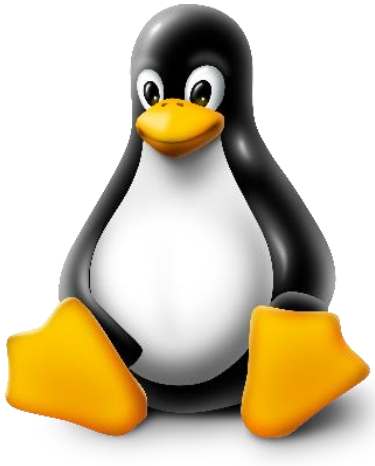

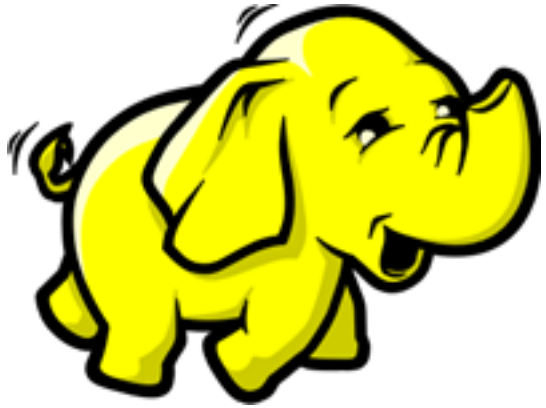

# Name That Project!

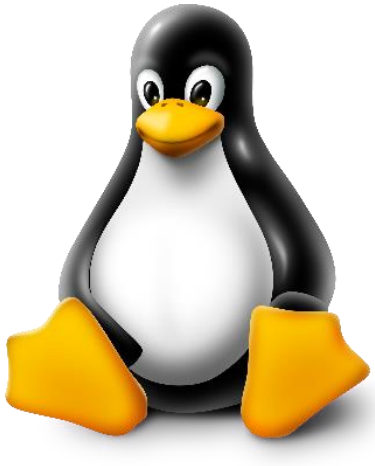

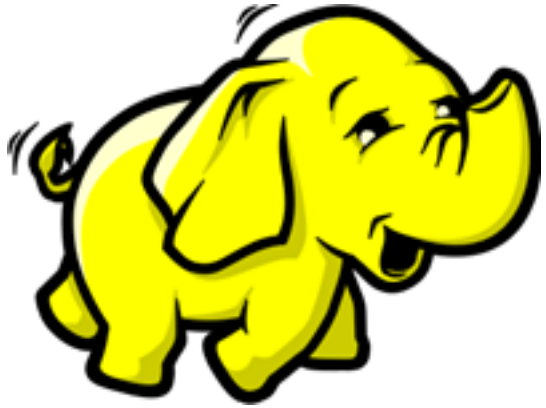

# Name That Project!

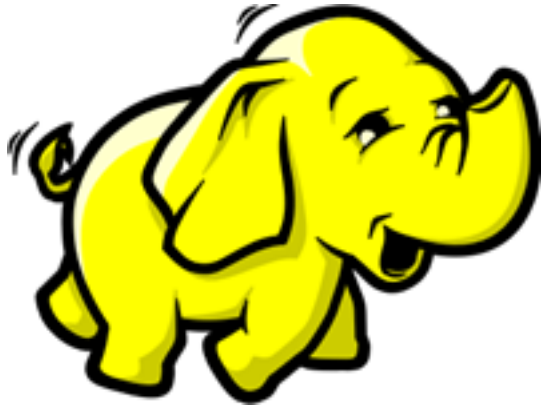

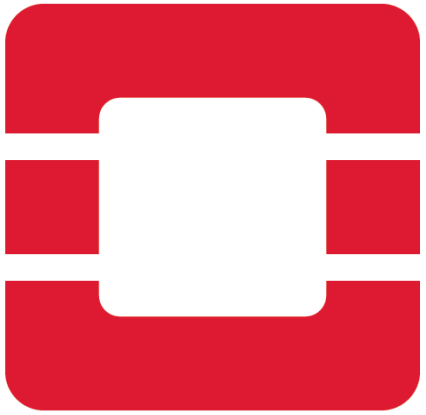

# Name That Project!

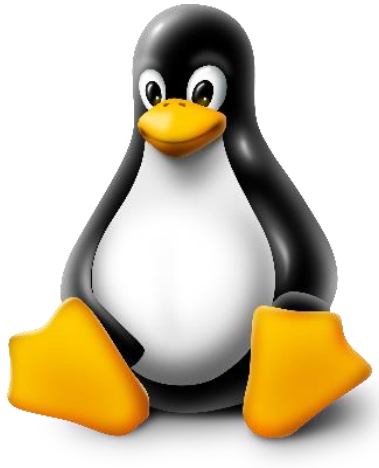

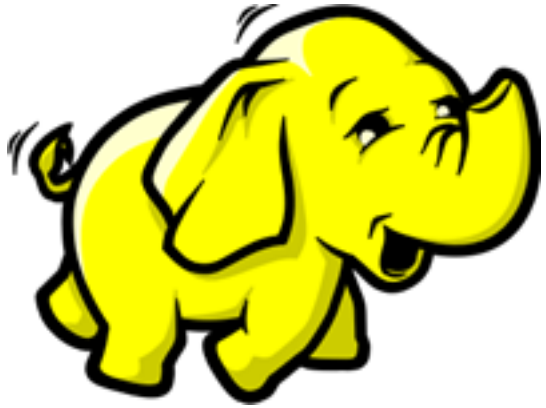

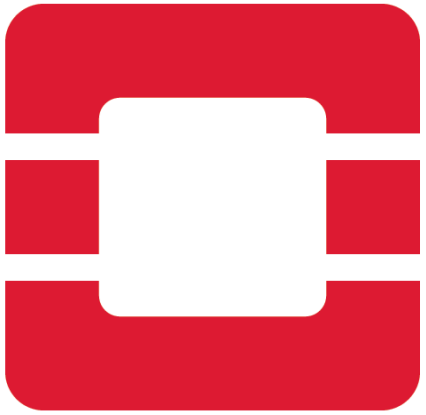

lustre™

# Name That Project!

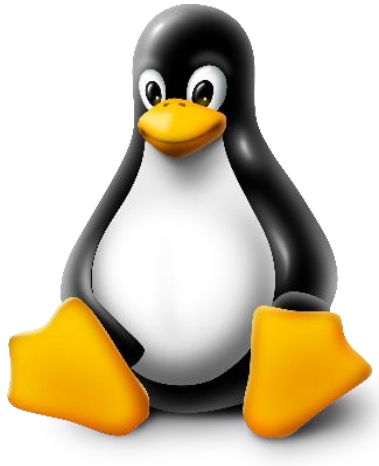

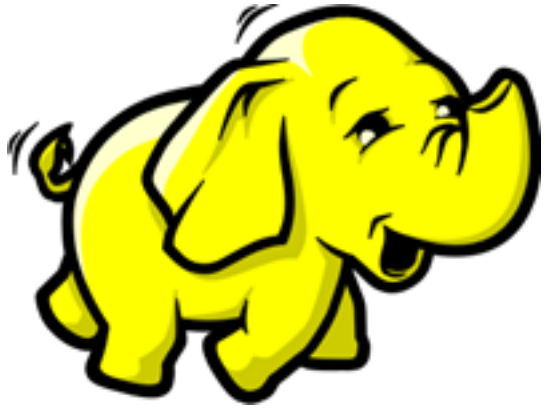

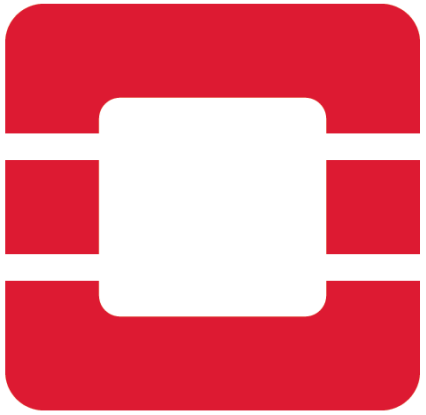

lustre™

# Name That Project!

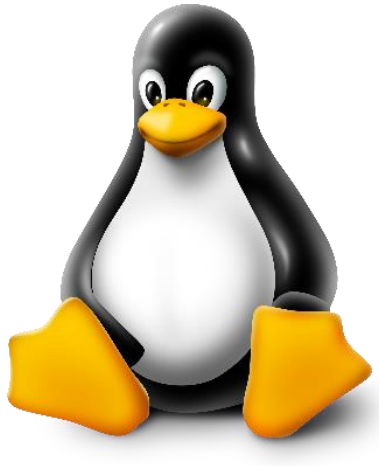

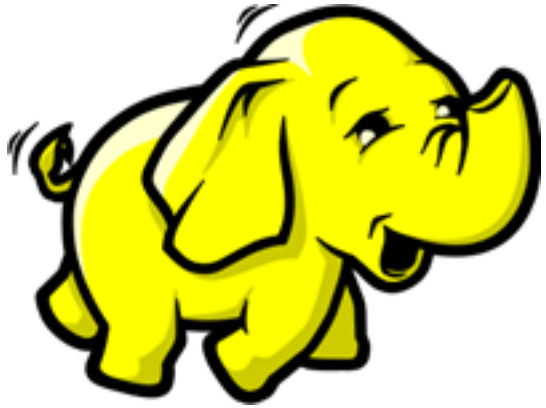

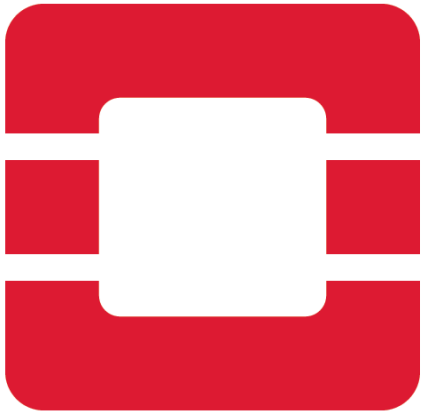


# Open Source Traction Today…

- **> 65% use OSS to speed application deployment**
- **> 55% use OSS in production infrastructure**
- **Reasons for Open Source Software (OSS) use**
  - Quality of Solutions
  - Competitive Features / Technical Capabilities
  - Ability to Customize / Fix
- **90% claim OSS improves efficiency, interoperability, and innovation**

BLACKDUCK


*Source: Blackduck 2016 Future of Open Source Survey Results*

# *Operating Systems*

## *(Linux in HPC)*

## Linux on 99.4% of the Top 500 Supercomputers

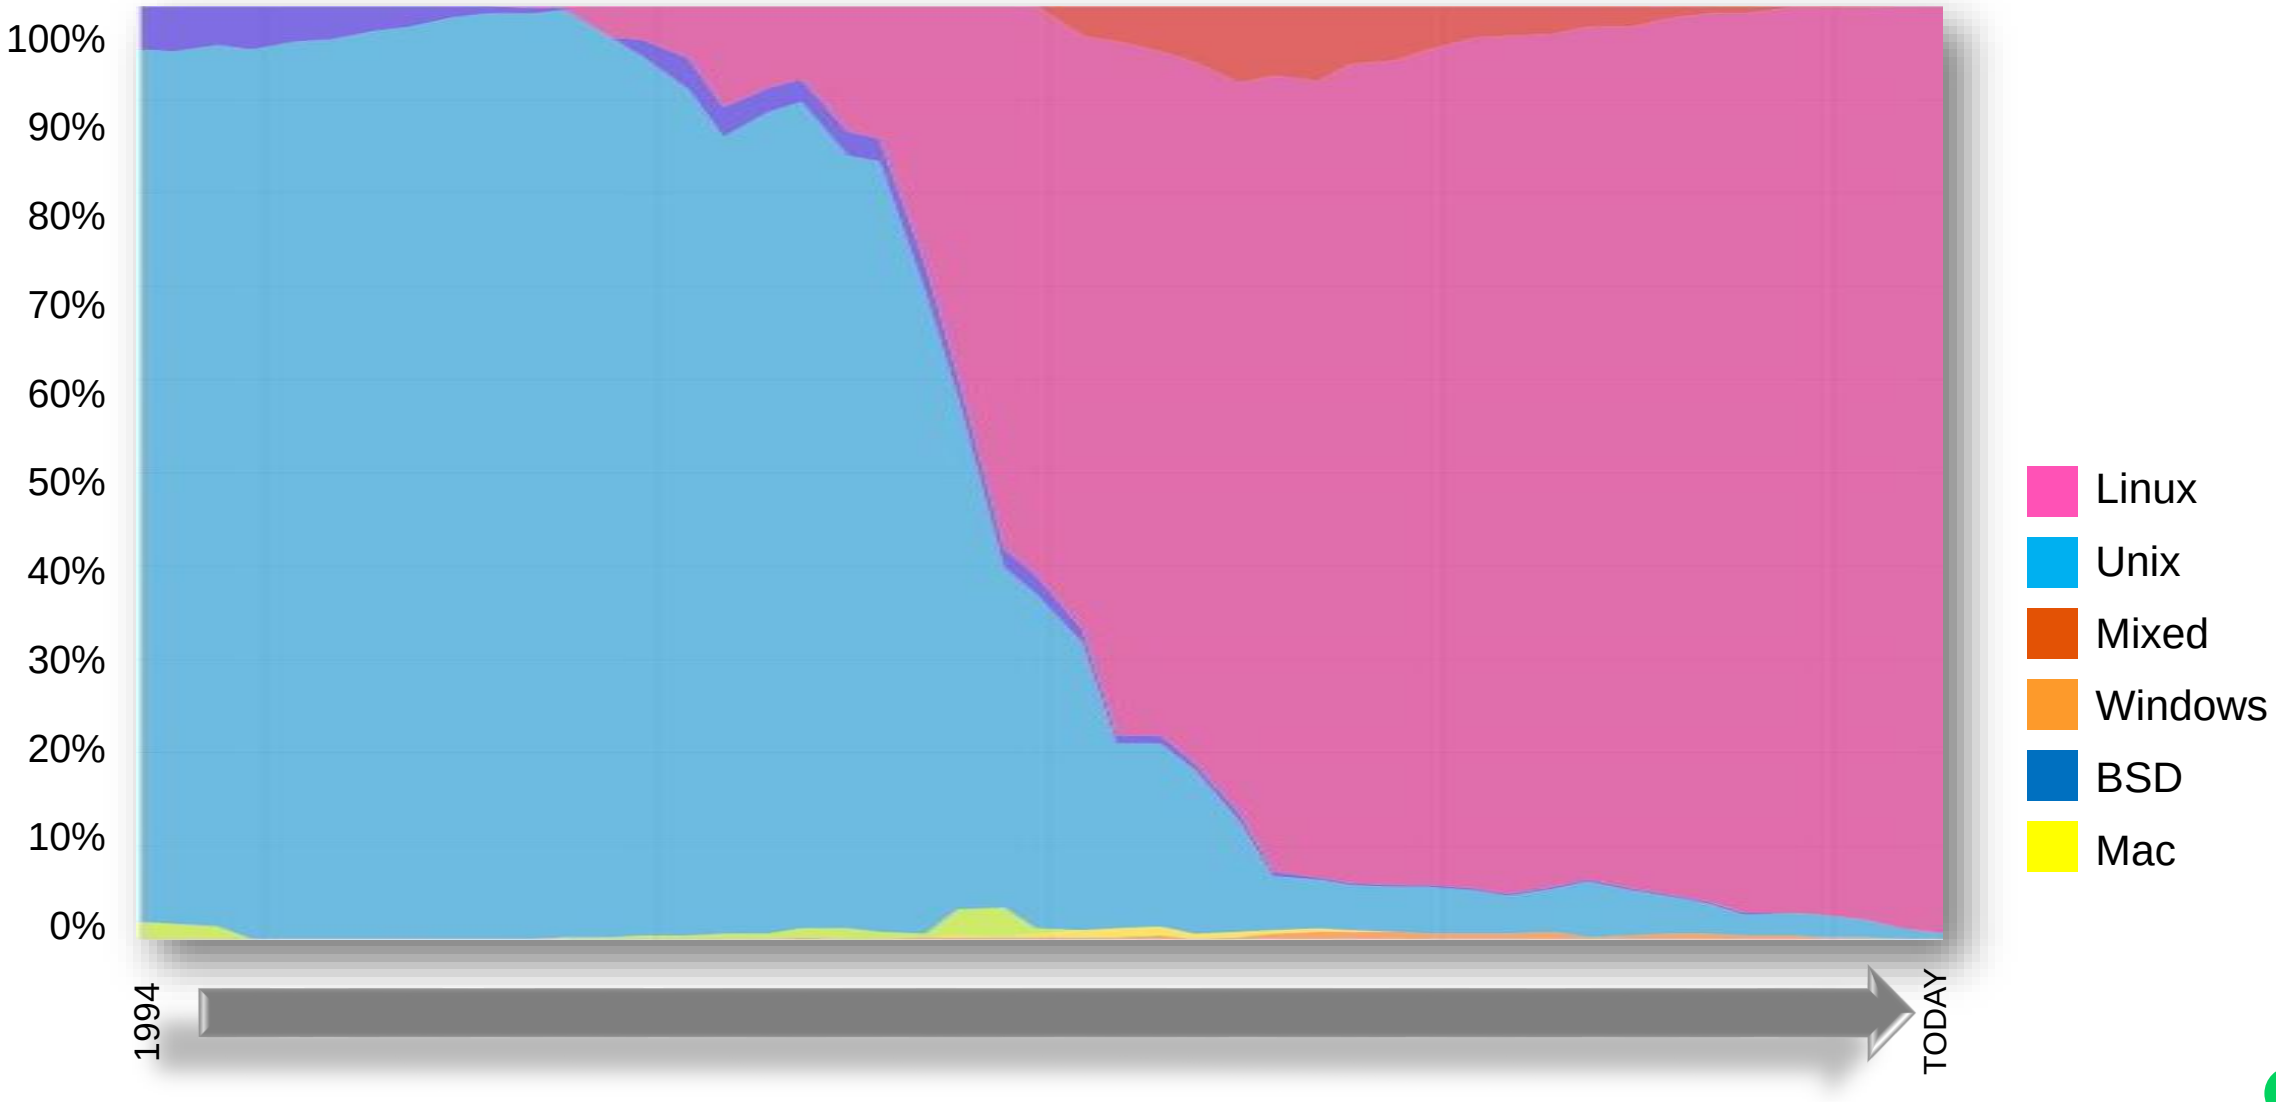

# OpenHPC Open Source Community

**openhpc.community**

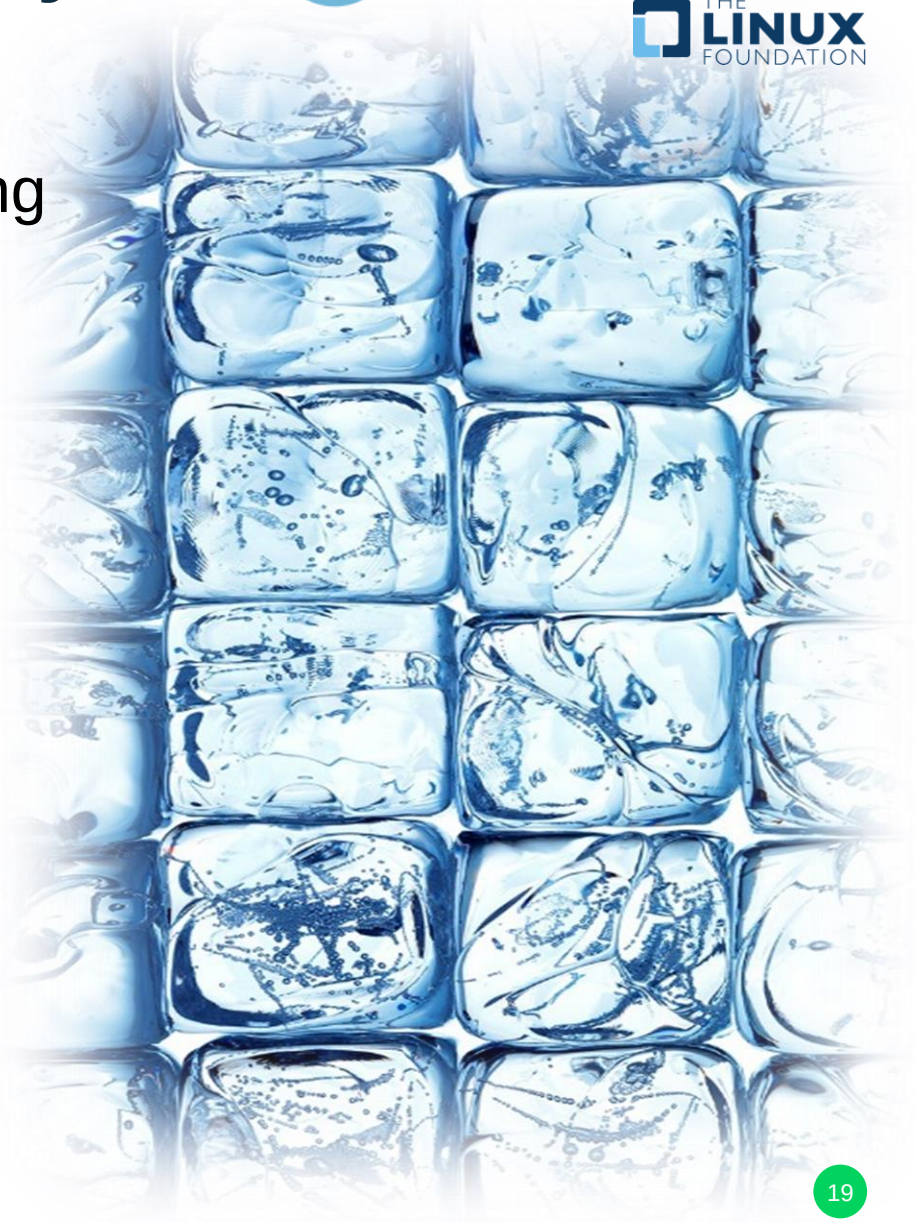

- Linux Foundation project – SUSE is a founding member (now 30+ full members)
- Provides common platform – standard HPC stack – for collaboration and innovation
- Provides full HPC stack (~300 packages) on top of SLES
- Build with SUSE build service: http://build.openhpc.community
- Simplifies installation, configuration, and maintenance of a custom software stack

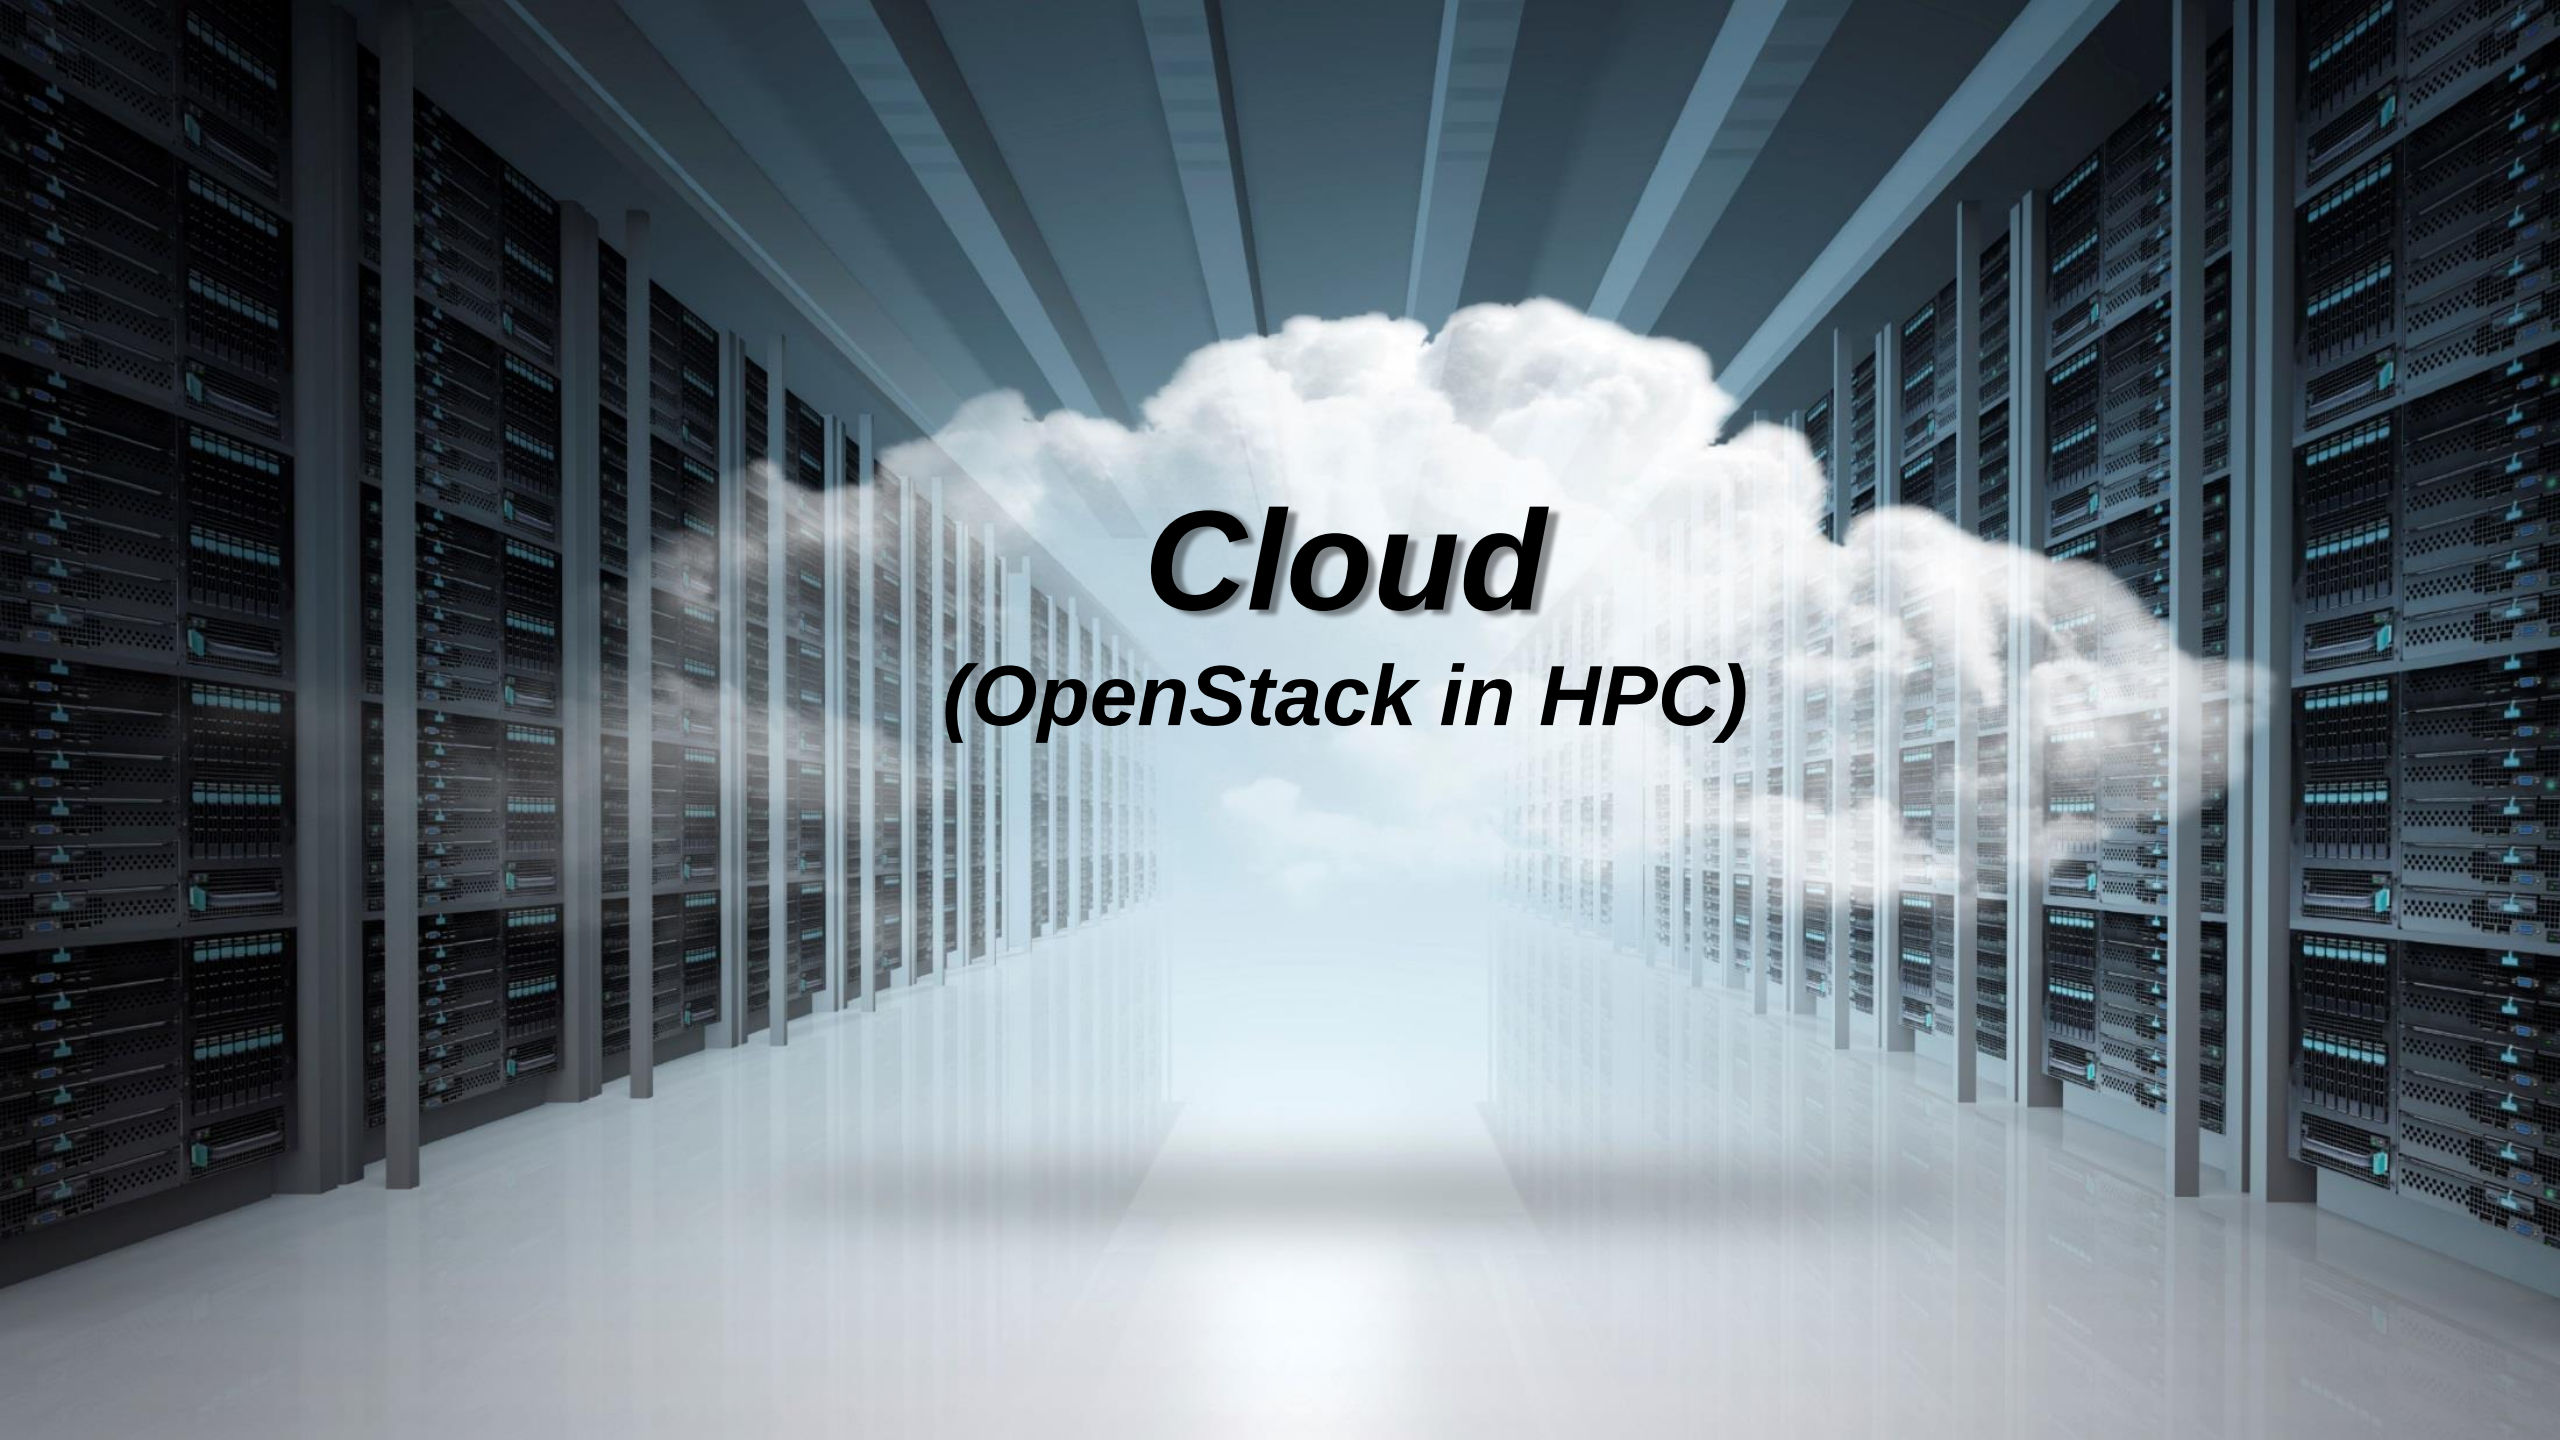
Cloud
(OpenStack in HPC)

# Cloud and HPC: Unique, Yet Similar

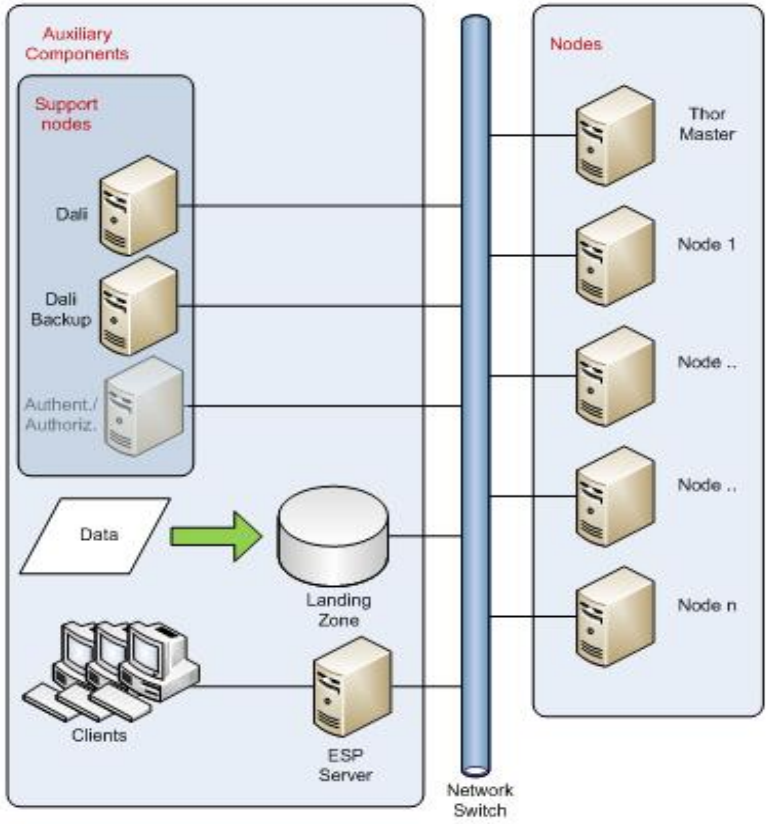


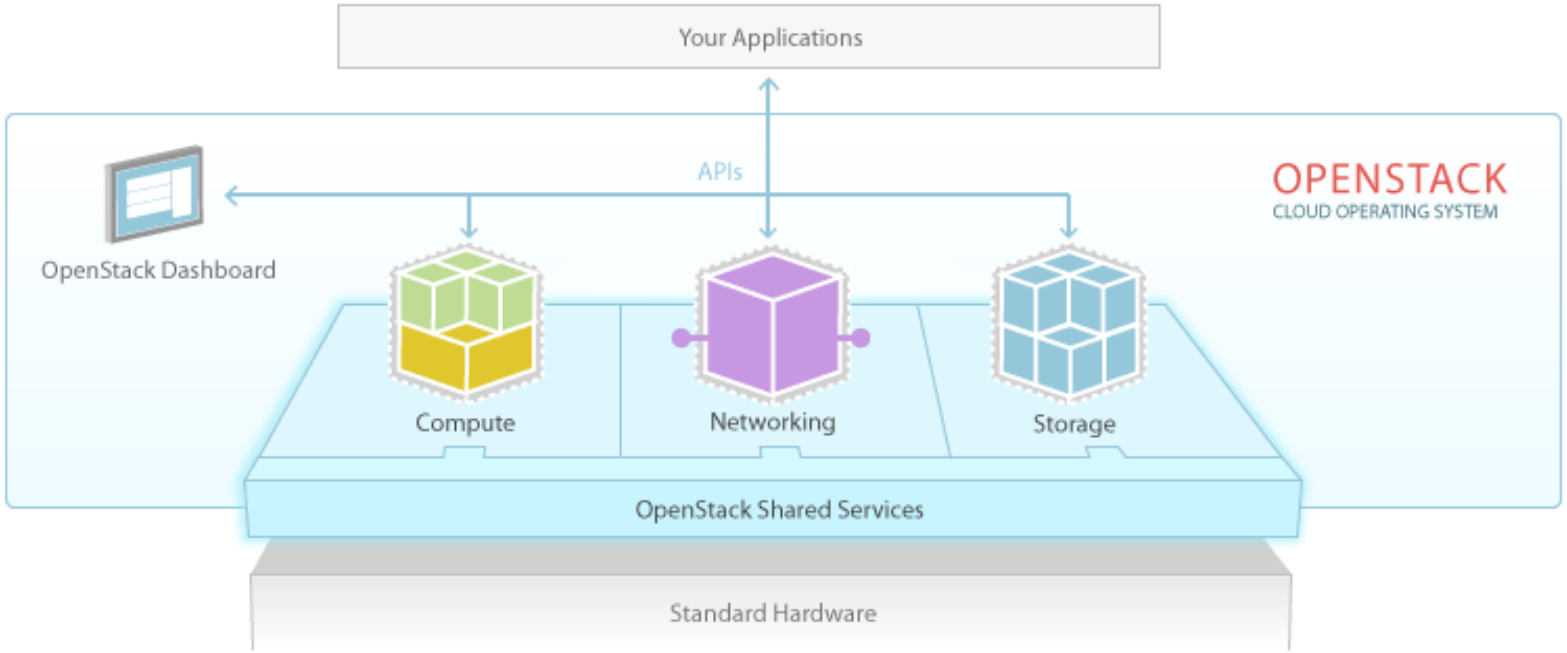


**Highly Distributed**

**Large Storage / Data Pools**

**Performance Management**

**Resource Management Key**

# HPCaaS in the Real World

GREGOR MENDEL INSTITUTE OF MOLECULAR PLANT BIOLOGY

- **The Challenge: Provide IT resources to scientists with strong high performance computing requirements rapidly with limited overhead**
- **The Solution: HPCaaS with SUSE OpenStack Cloud**
- **Addressing the Challenge: Scientists are now able to deploy applications by themselves, with IT adjusting resource allocation as needed**

- **Results: Users can now deploy services self service, with IT maintaining infrastructure**
- **Where are they now? Running in production for over one year with great success, now expanding into software-defined storage and containers**

YouTube *"The Gregor Mendel Institute of Molecular Plant Biology and SUSE OpenStack Cloud"*

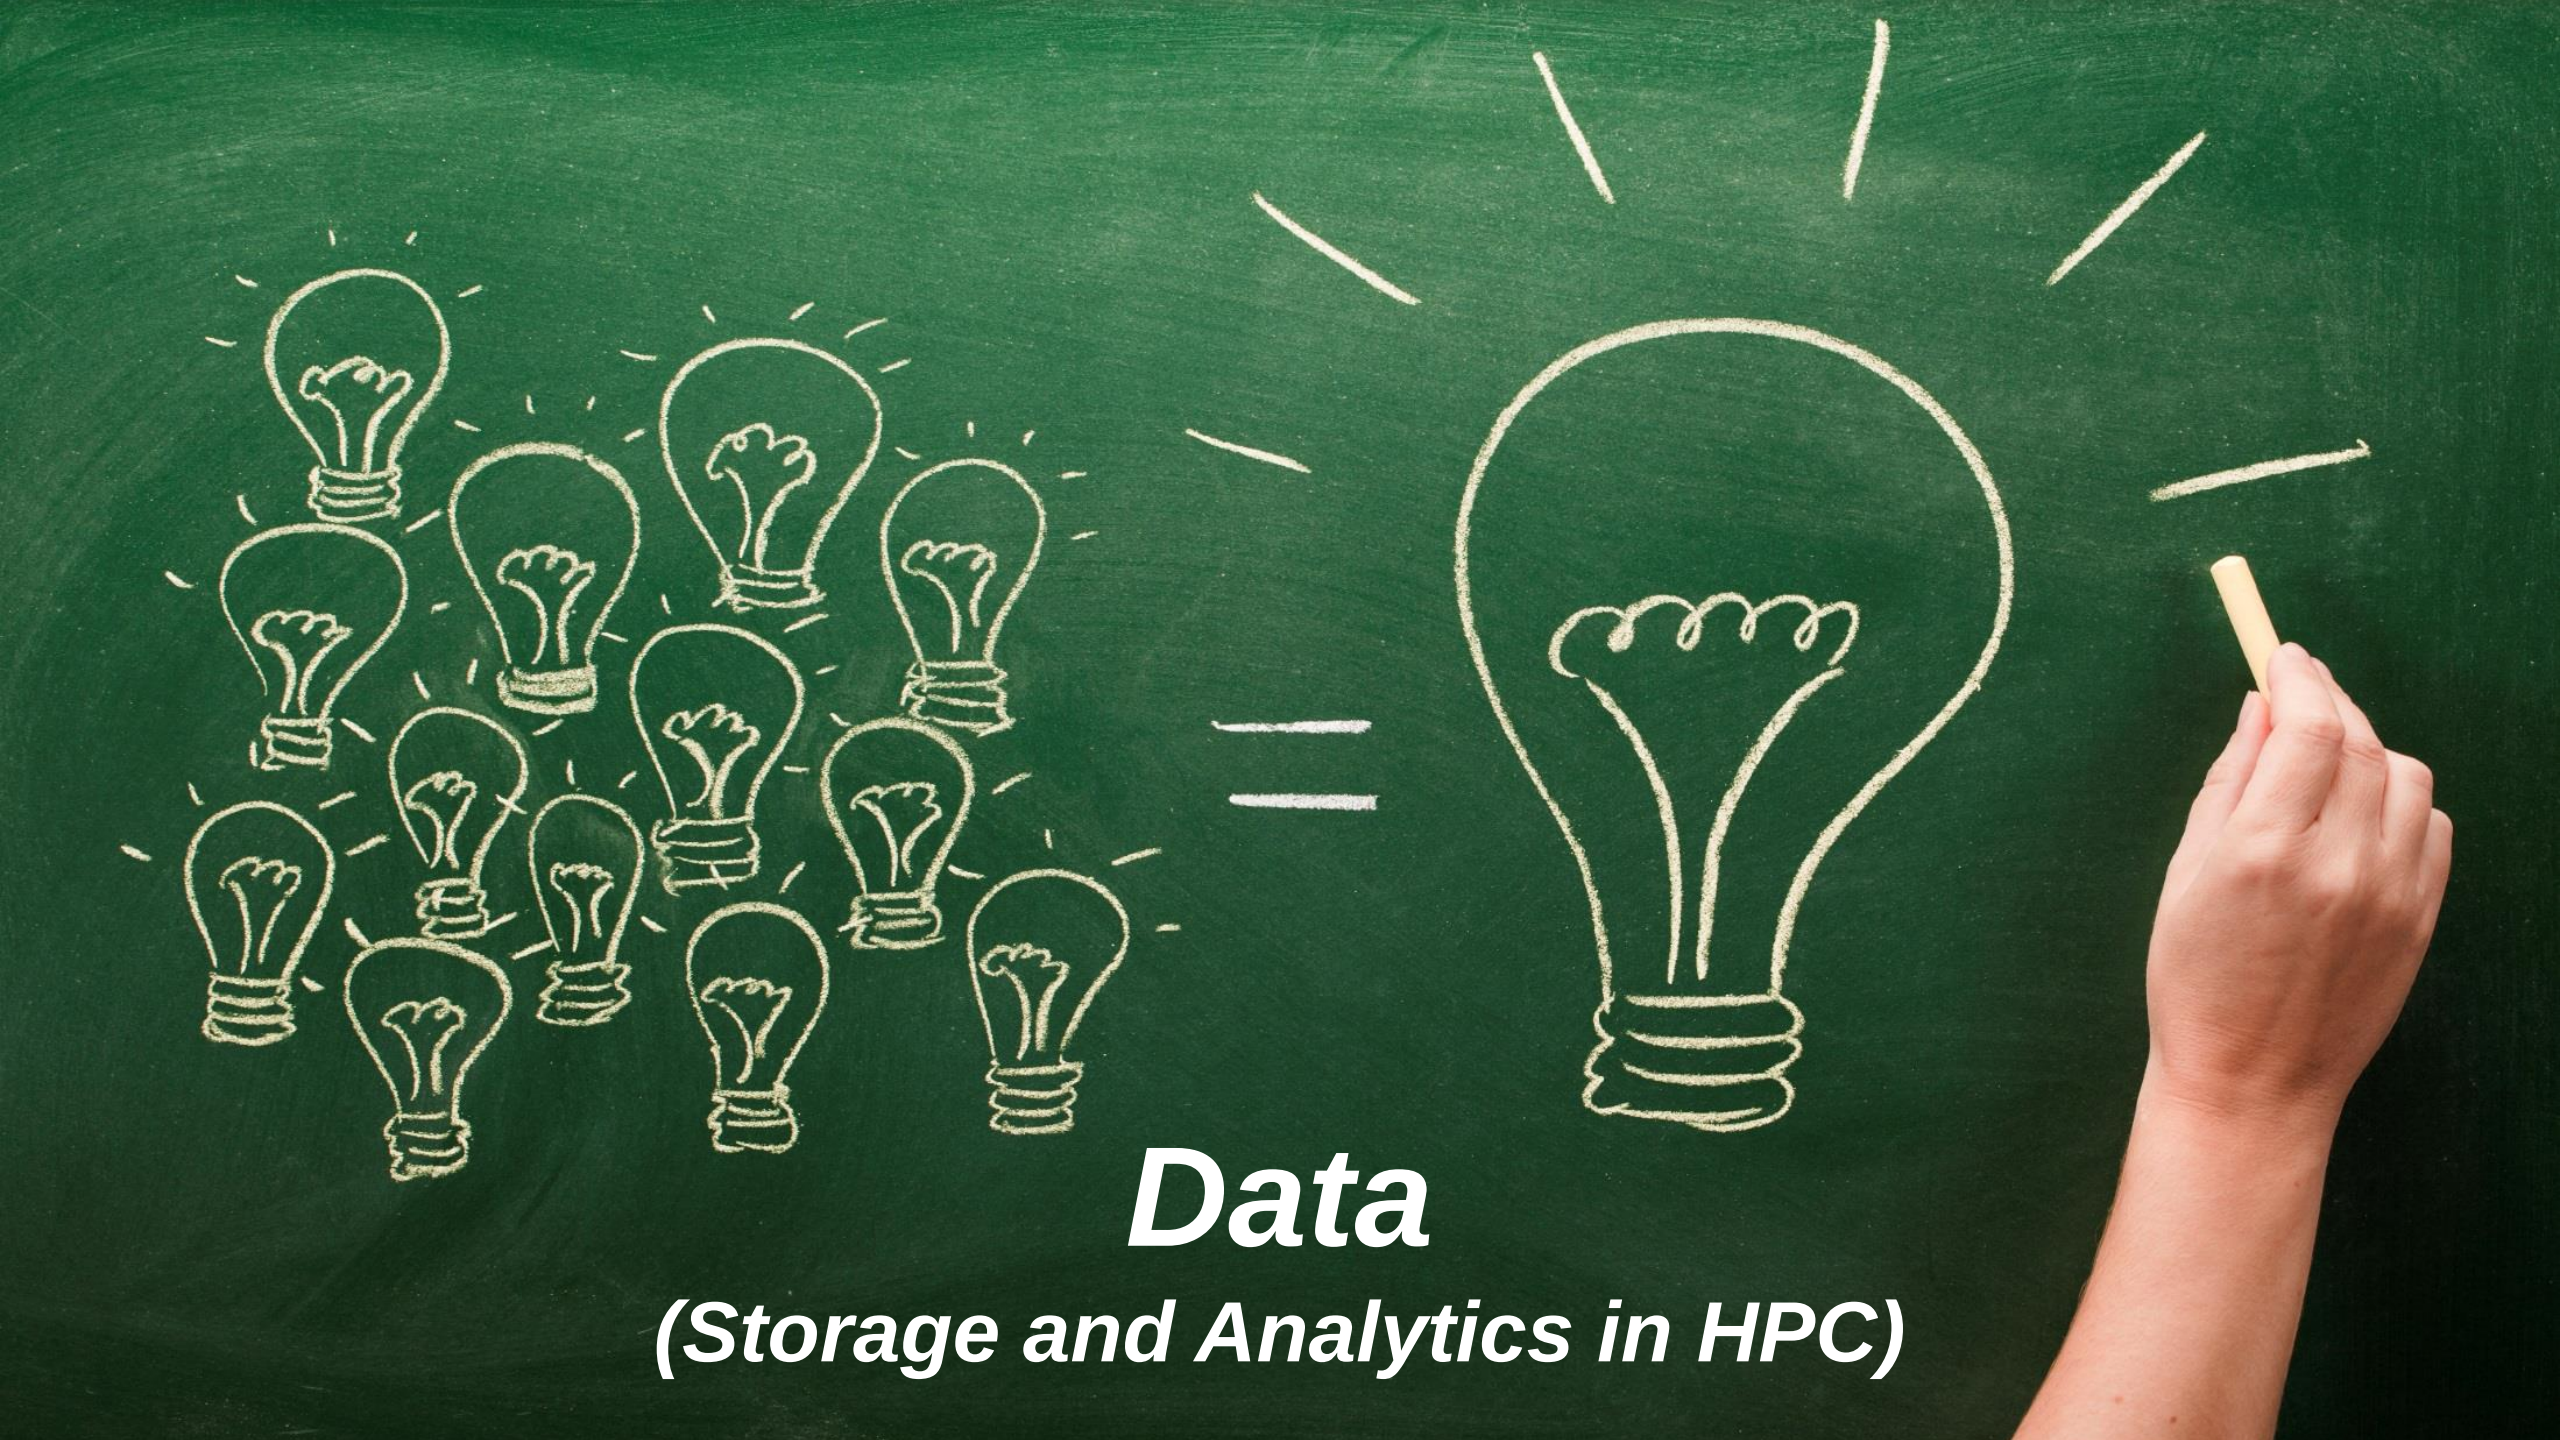
Data
(Storage and Analytics in HPC)

# SUSE Enterprise Storage

## Unlimited Scalability with Self Managing Technology

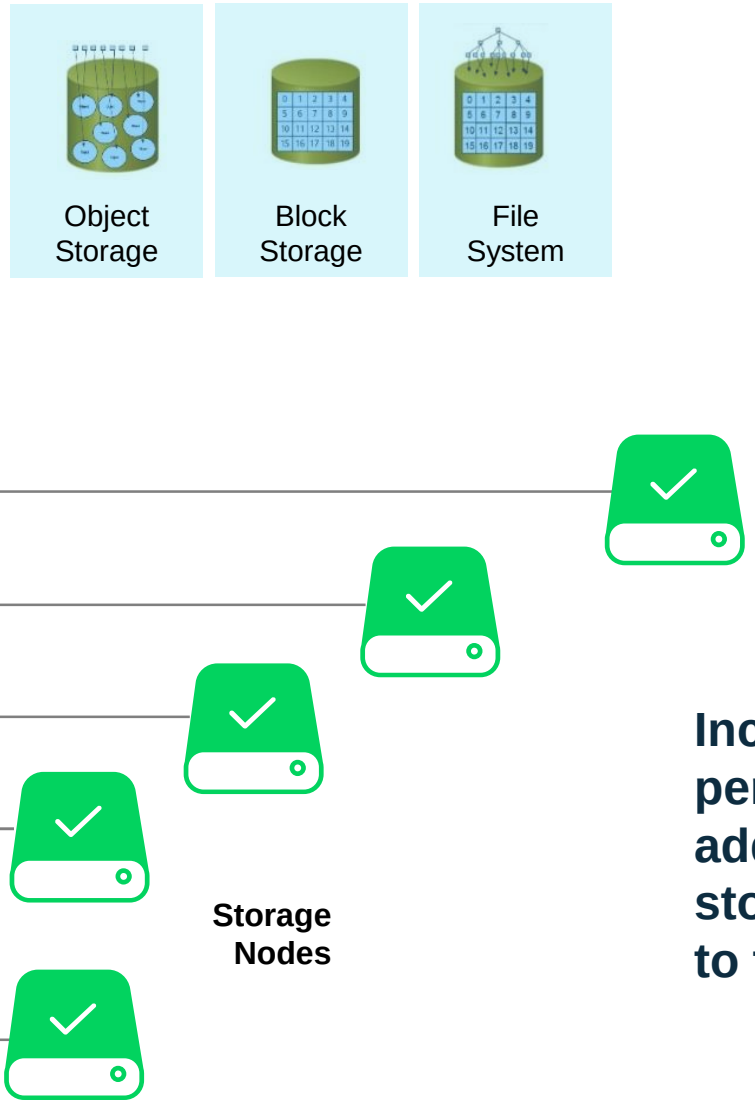


**Increase capacity and performance by simply adding new storage or storage nodes to the cluster**

# CASE STUDY: Orchard Park Police Department

## Overview

+ Read Story

*Home to the Buffalo Bills football team, Orchard Park is a town that is located in Erie County, New York. At the time of the 2010 census, the town's population was 29,054. The Orchard Park Police Department strives to keep town citizens safe and to serve in accordance with the values of integrity, respect, professionalism and community.*

| Challenge | Solution | Results |
| --- | --- | --- |
| • **Main challenge was supporting the data and storage obtained from the body cameras.**<br>• **Sought a storage solution to retain body camera footage that would fulfill legal mandates.** | • **Chose to implement SUSE Enterprise Storage because it is non-proprietary, scalable, flexible, cost-efficient and resilient.** | • **Maintains network performance of 400-500 IOPS**<br>• **Supports body cameras and critical surveillance tools**<br>• **Provides the ability to retain video data for legal purposes** |


**www.suse.com/products/server/hpc**
**www.suse.com/partners/alliance/hpe/**
**@jbgeorge**

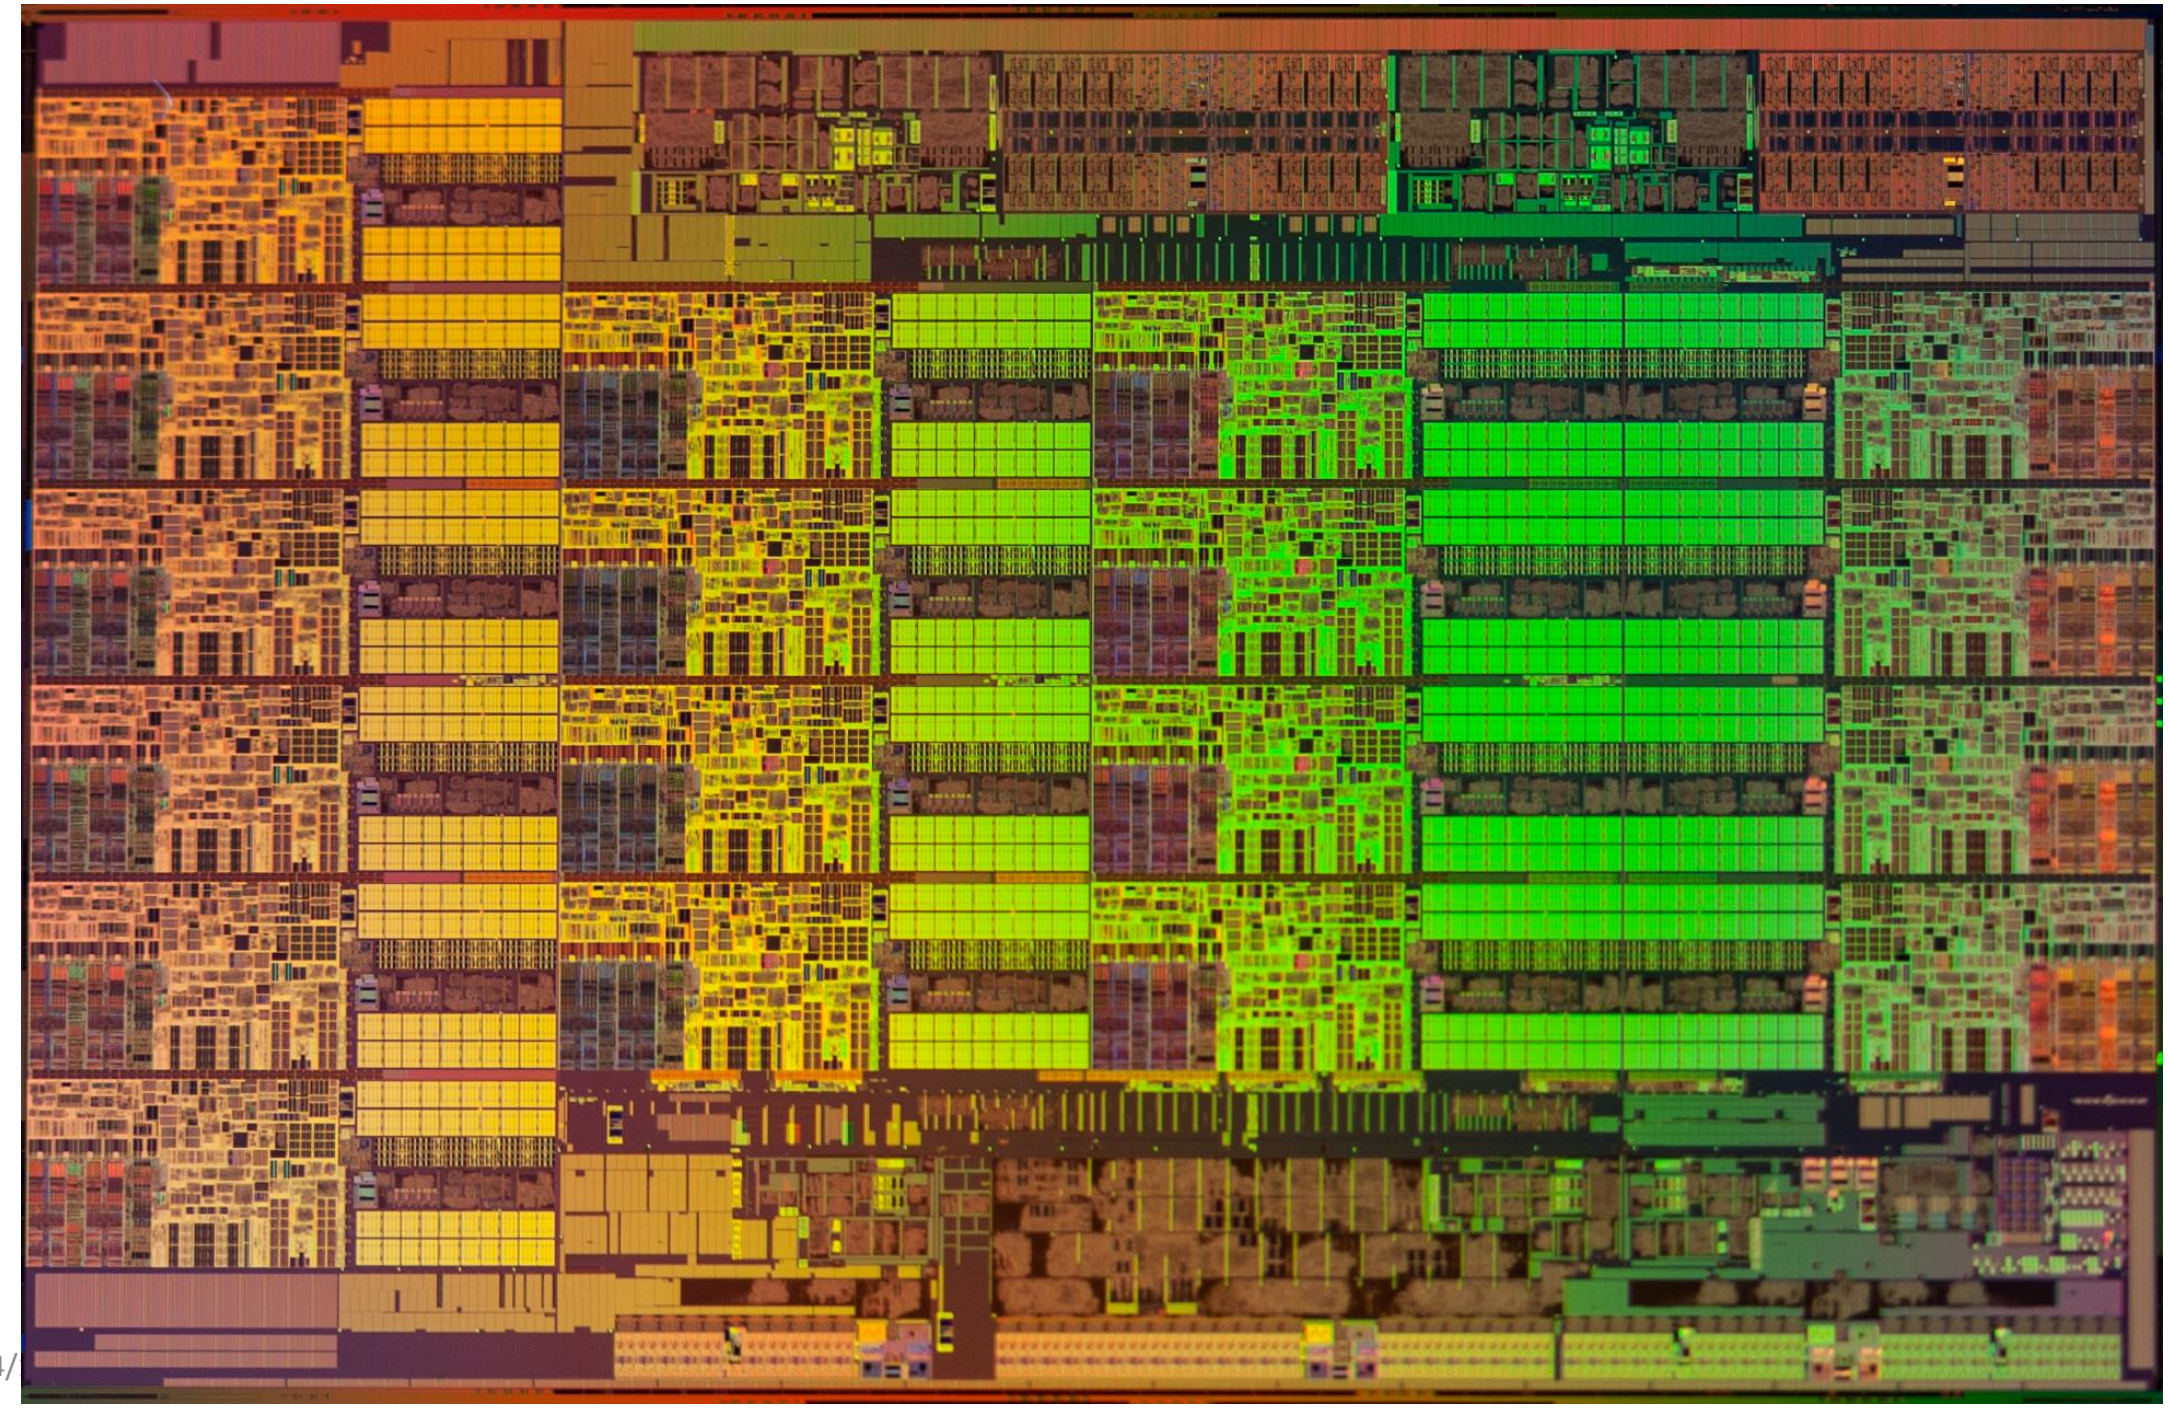

| Year | Cores | SIMD | L/C | LANES |
|---|---|---|---|---|
| 2007 | 8 | 128 | 4 | 32 |
| 2009 | 8 | 128 | 4 | 32 |
| 2010 | 12 | 128 | 4 | 48 |
| 2012 | 16 | 256 | 8 | 128 |
| 2013 | 24 | 256 | 8 | 192 |
| 2014 | 36 | 256 | 8 | 288 |
| 2016 | 44 | 256 | 8 | 352 |

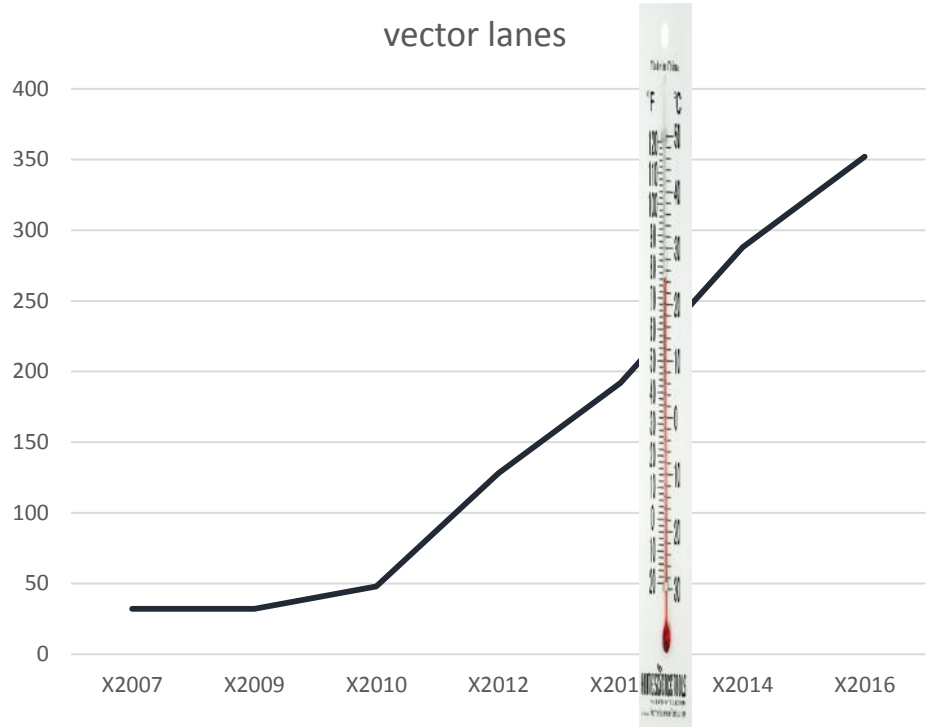


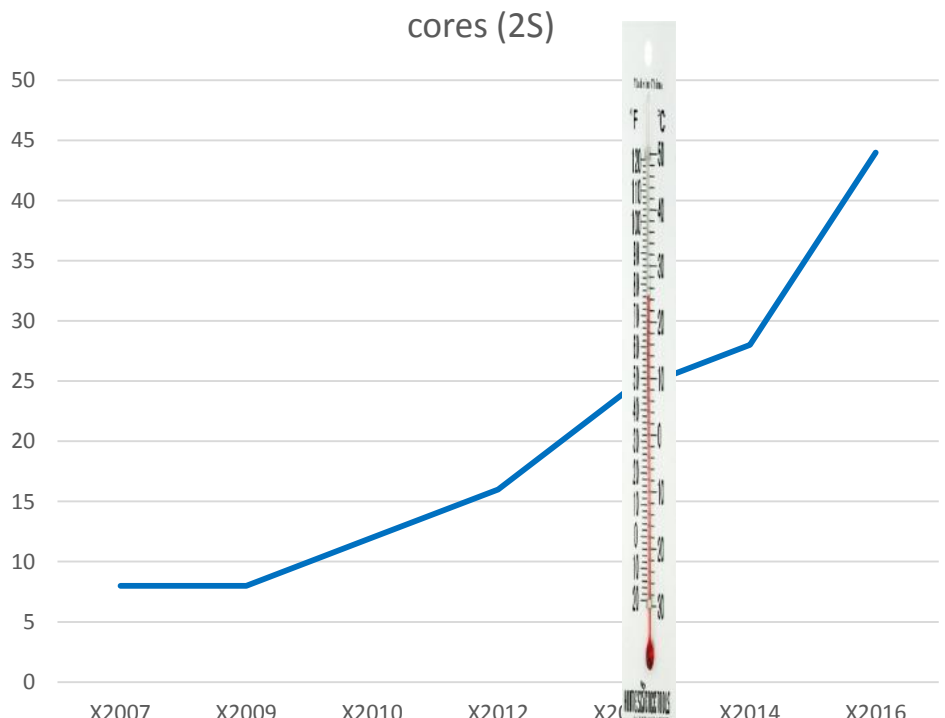

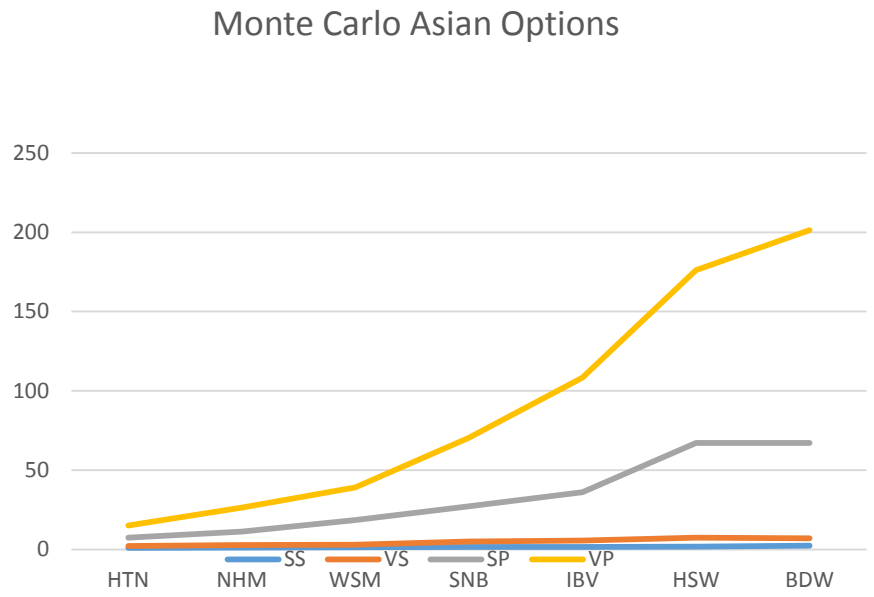
Monte Carlo Asian Options
250
200
150
100
50
0
HTN
NHM
WSM
SNB
IBV
HSW
BDW
SS
VS
SP
VP

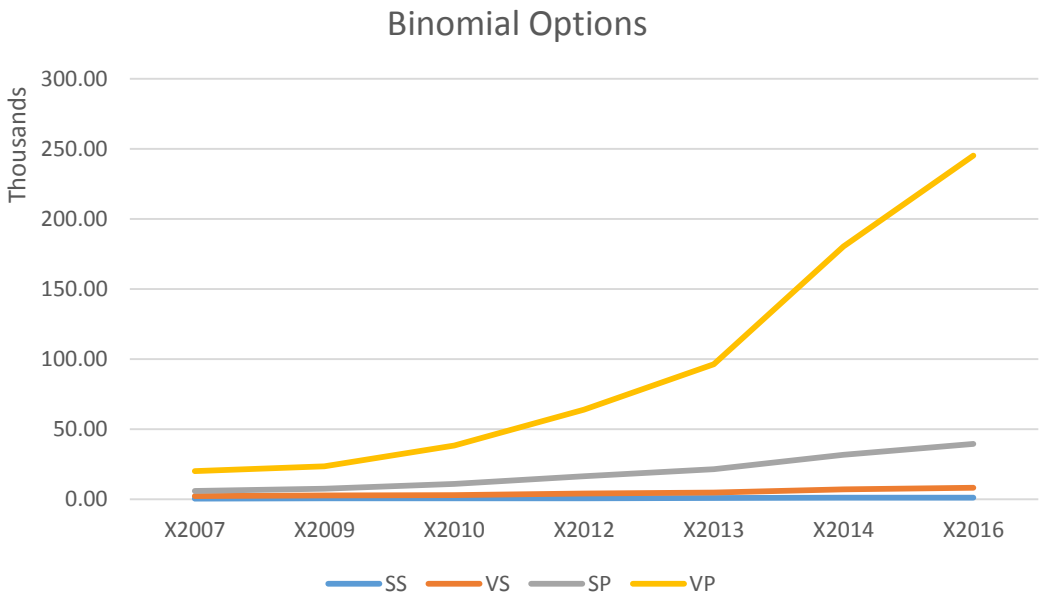
Binomial Options
Thousands
300.00
250.00
200.00
150.00
100.00
50.00
0.00
X2007
X2009
X2010
X2012
X2013
X2014
X2016
SS
VS
SP
VP

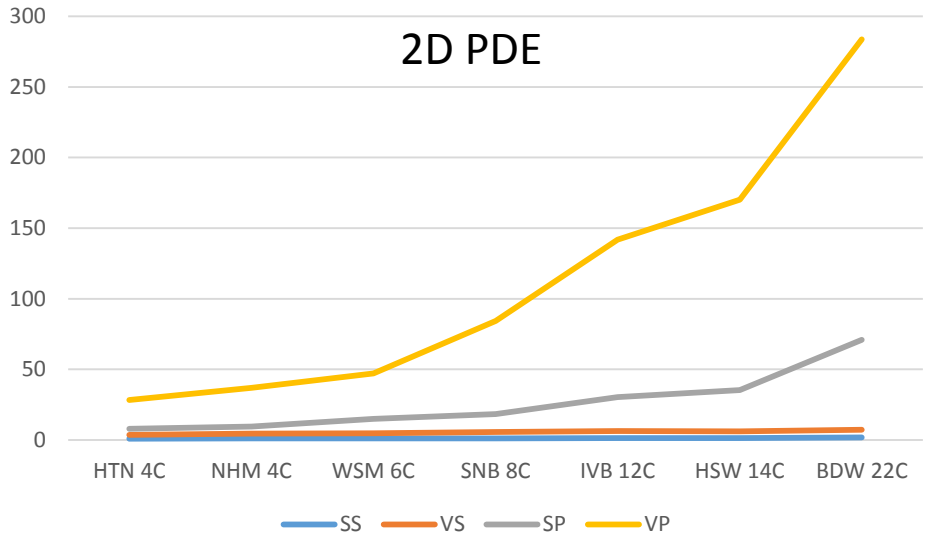
2D PDE
300
250
200
150
100
50
0
HTN 4C
NHM 4C
WSM 6C
SNB 8C
IVB 12C
HSW 14C
BDW 22C
SS
VS
SP
VP

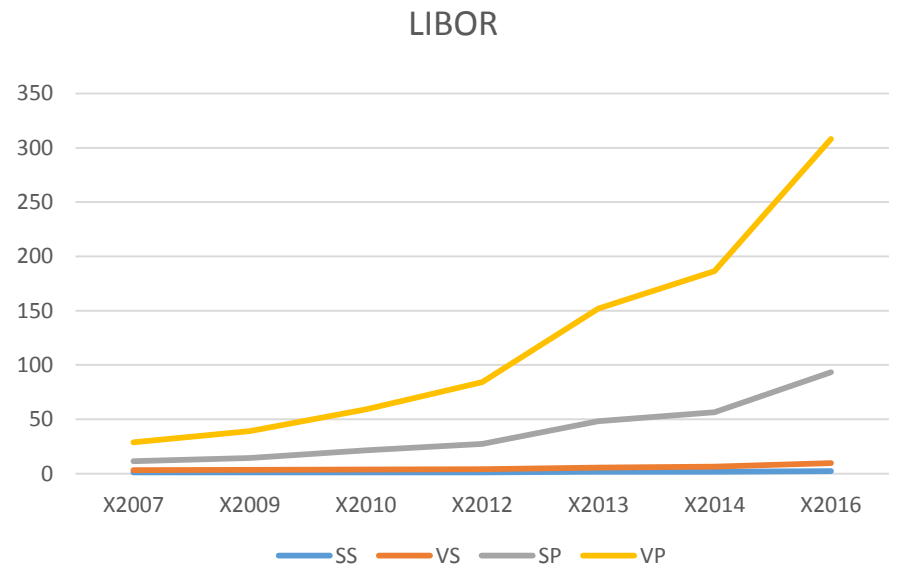
LIBOR
350
300
250
200
150
100
50
0
X2007
X2009
X2010
X2012
X2013
X2014
X2016
SS
VS
SP
VP

# Parallel Programming for CPUs

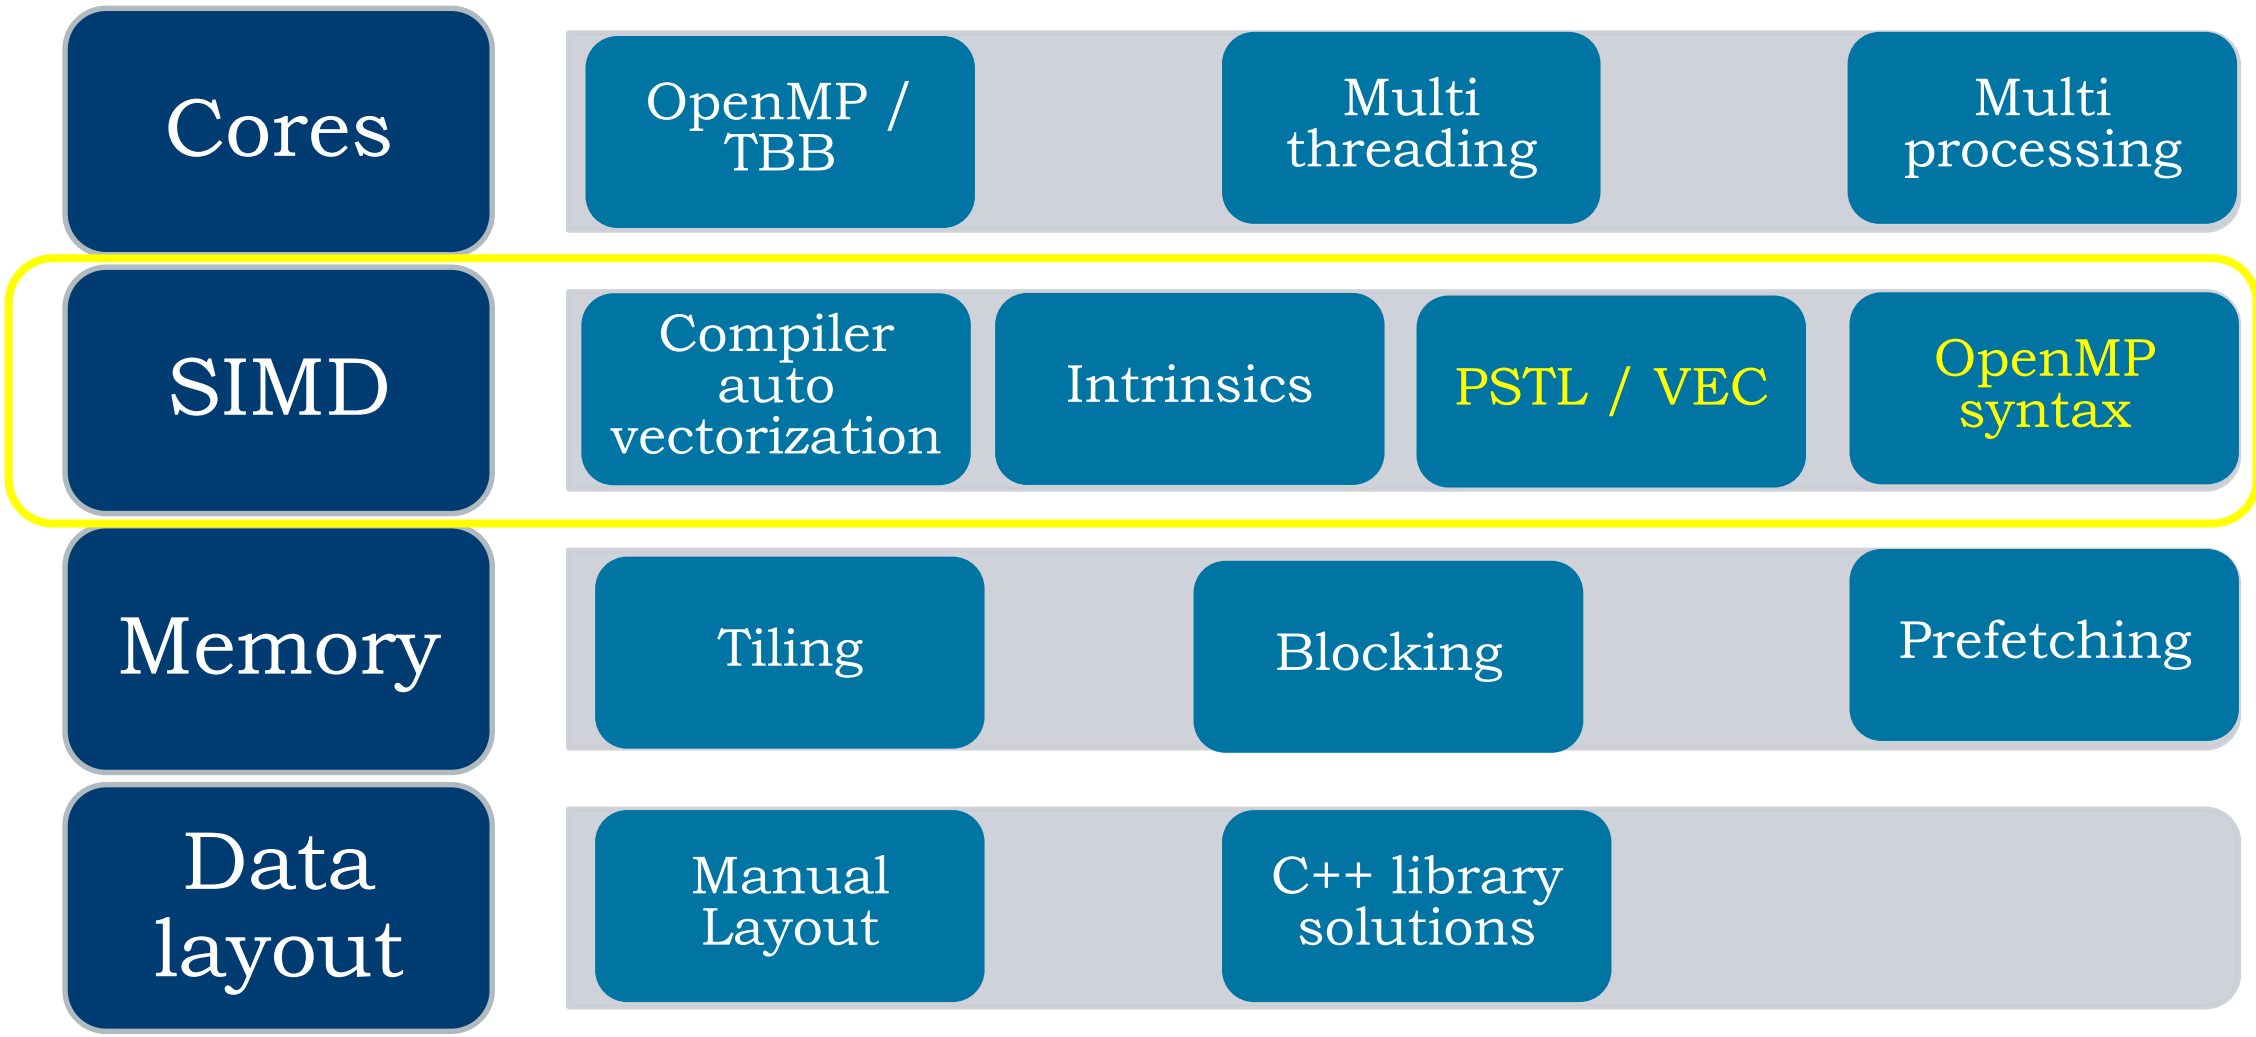


4 considerations to take care of when writing an efficient, unconstrained parallel program

# KNL Architecture Overview

## Tile

2 cores, Cache, memory HUB

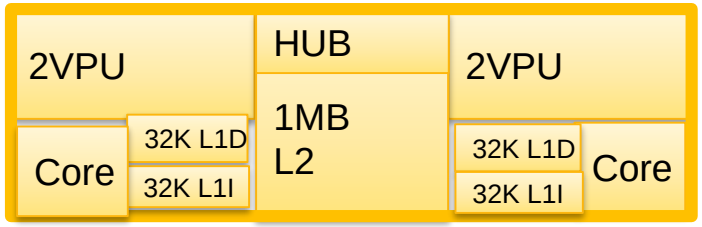


- **ISA**
  SSE4.2, AVX, AVX2
  AVX512: 32X512 registers
- 8 mask registers
- Floating point and int operations
- New semantics

- ✓ Silvermont based core
- ✓ Many per core enhancements
- ✓ Out-of-Order Cores
- ✓ 3X single-thread vs. KNC
- ✓ 1/3X single-thread vs Xeon

## Platform Memory

Up to 384GB (6ch DDR4-2400 MHz)

## On-package high bandwidth memory

Up to 16GB, ~460 GB/s STREAM at launch

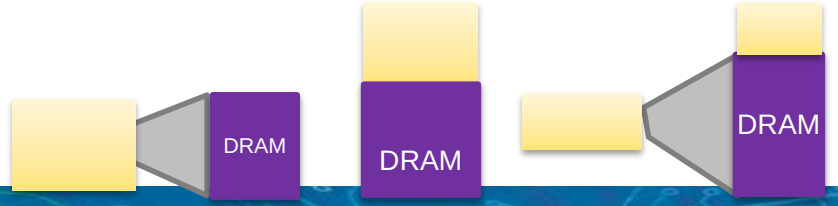


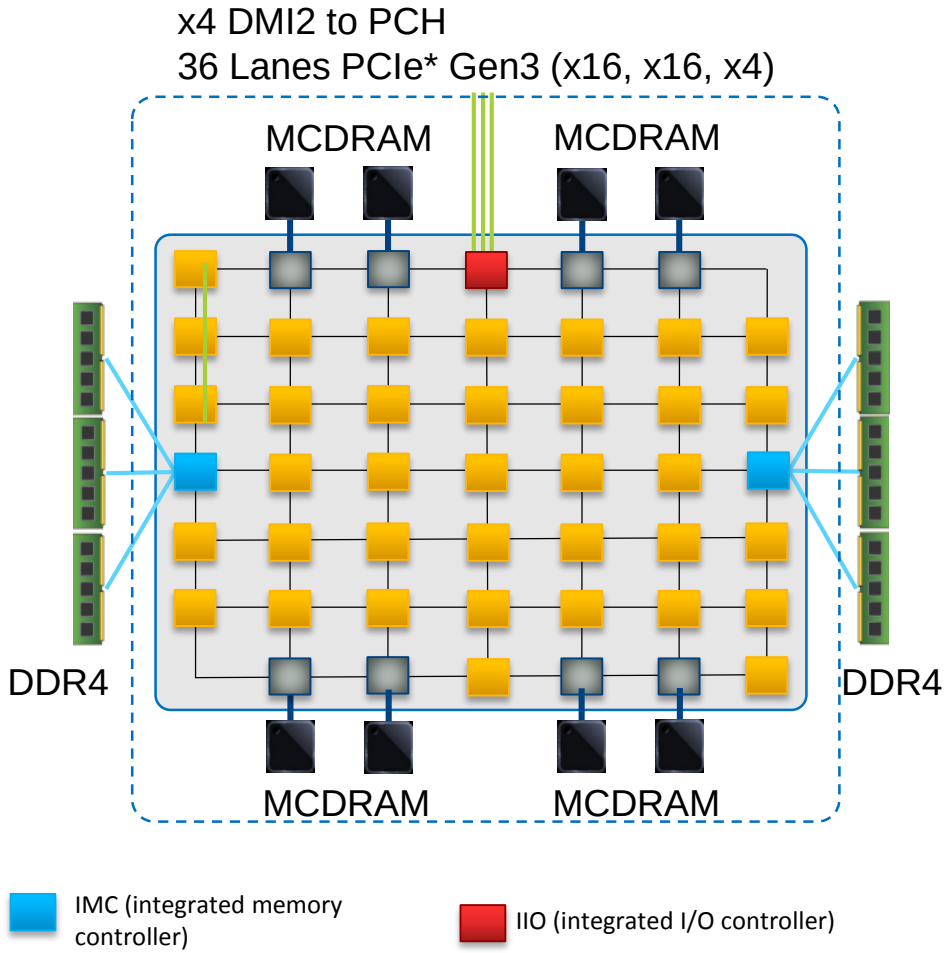

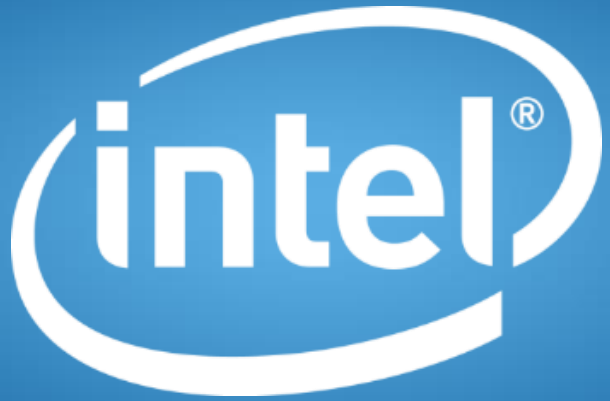
intel®

# KNL Architecture Overview

## ISA
Intel® Xeon® Processor Binary-Compatible (w/Broadwell)

## On-package memory
Up to 16GB, ~460 GB/s STREAM at launch

## Platform Memory
Up to 384GB (6ch DDR4-2400 MHz)

## Fixed Bottlenecks
- ✓ 2D Mesh Architecture
- ✓ Out-of-Order Cores
- ✓ 3X single-thread vs. KNC

TILE:
(up to 36)

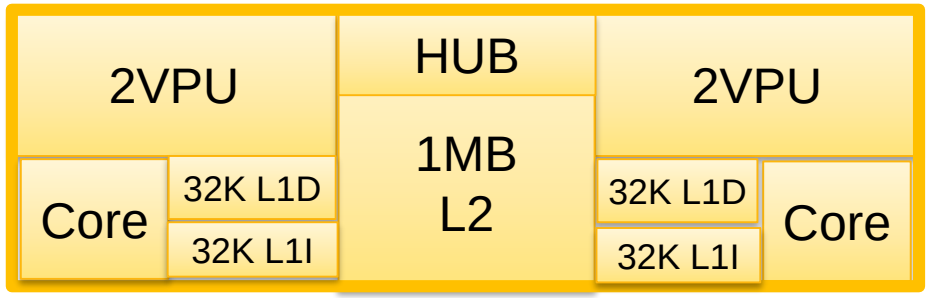


*Enhanced Intel® Atom™ cores based on Silvermont Microarchitecture*

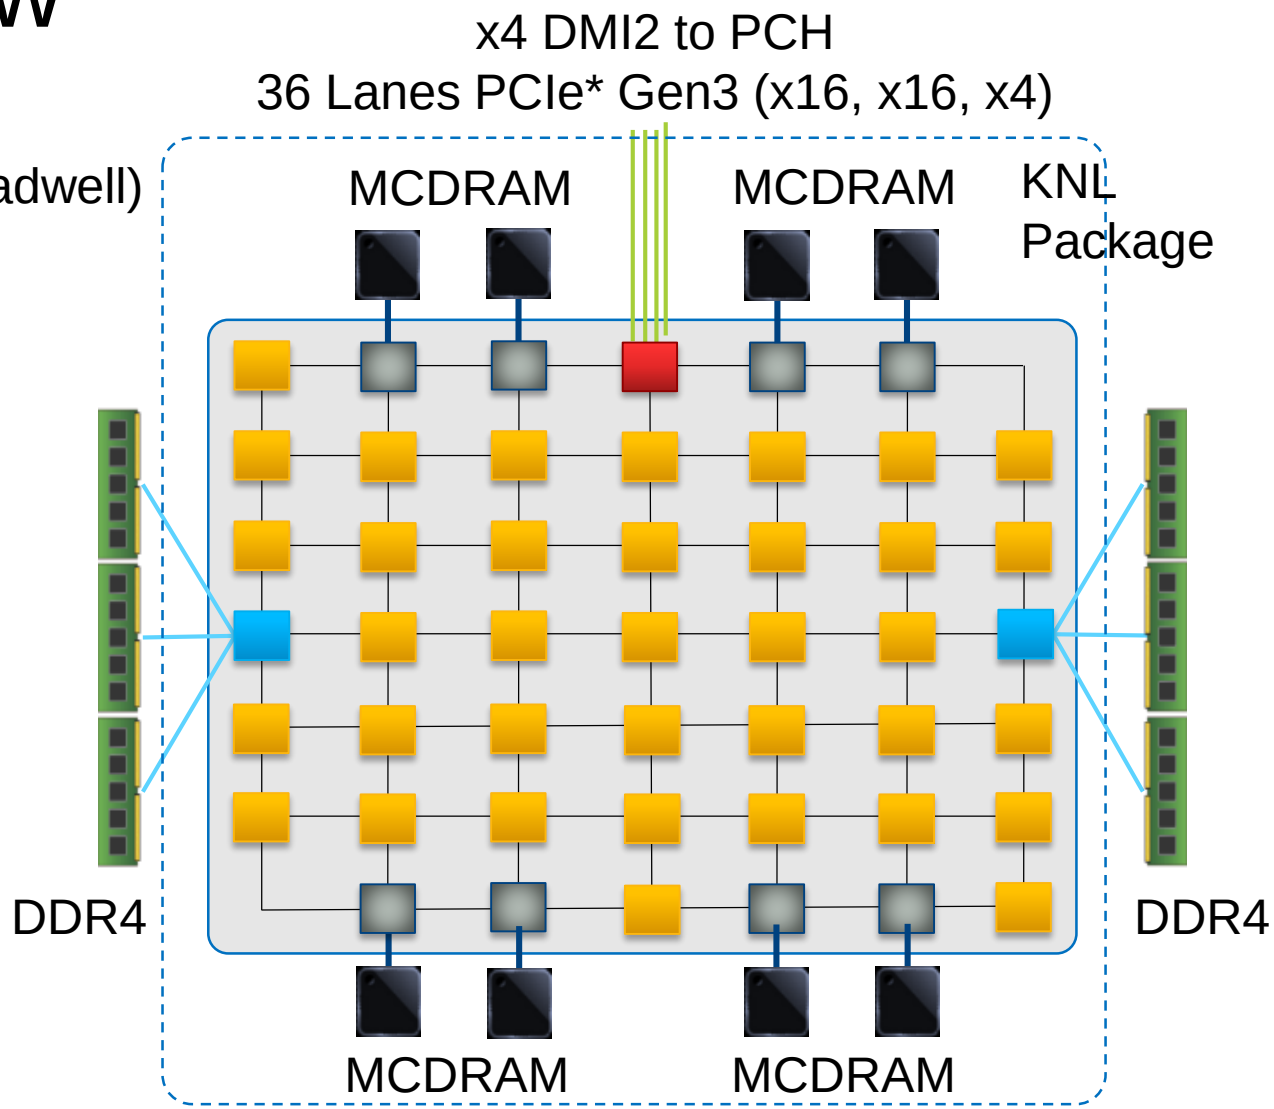


Tile | EDC (embedded DRAM controller) | IMC (integrated memory controller) | IIO (integrated I/O controller)

# #1 Best Practice in Parallelizing a Loop Hierarchy

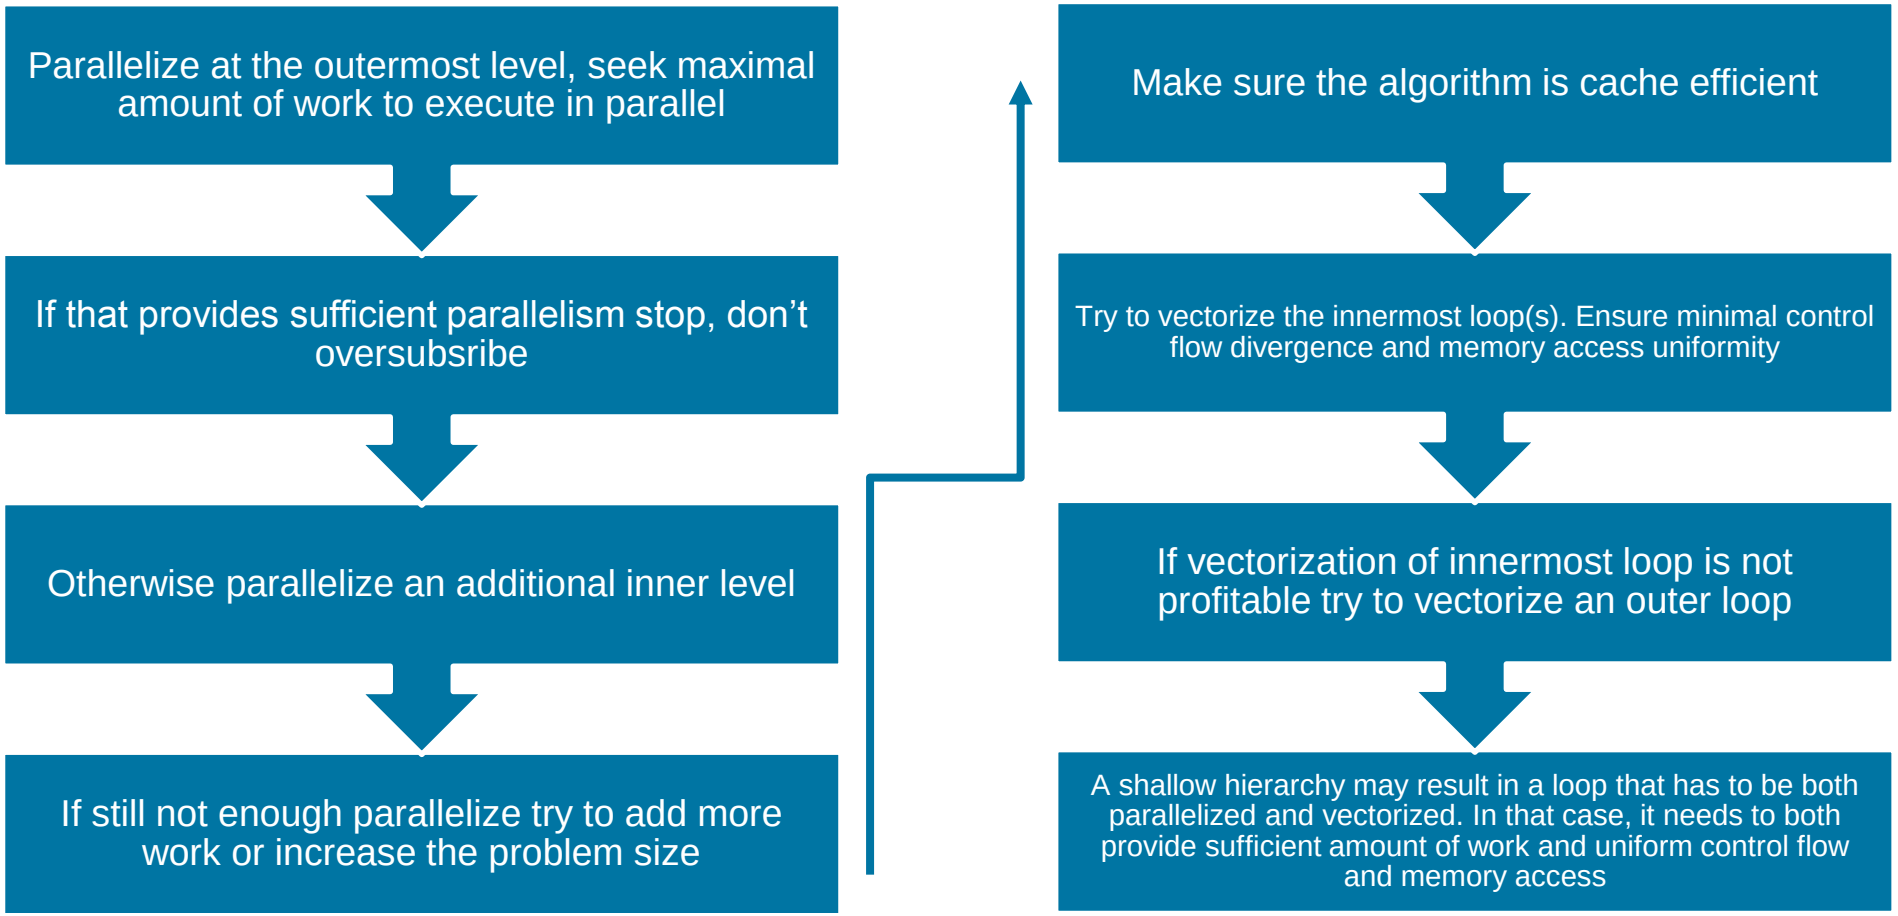


Vectorize Innermost, Parallelize Outermost (VIPO)

# Increments in HW architecture and programmability

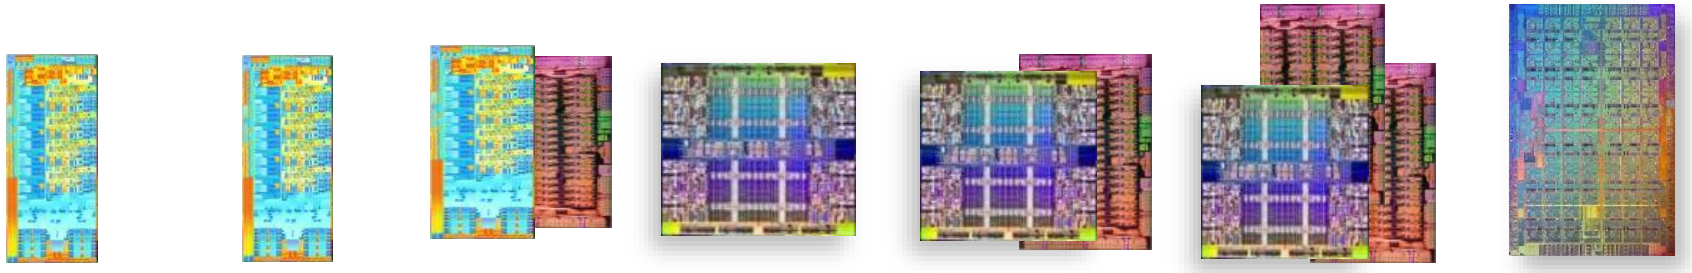

| | Intel Xeon processor E5 2697-V2 | Intel Xeon processor E5 2697-V2 | Intel Xeon E5 2697-V2 + Xeon Phi | Intel Xeon E5 2697-V3 | Intel Xeon E5 2697-V3+ Xeon Phi | Intel Xeon E5 2697-V3+ 2*Xeon Phi | Intel Xeon Phi 7290 |
|---|---|---|---|---|---|---|---|
| Report name | INTC130829 | INTC140507 | INTC140530 | INTC140814 | INTC140915 | INTC151018 | INTC161016 |
| Year | 2013 | 2014 | 2014 | 2014 | 2014 | 2015 | 2016 |
| cores | 24 | 24 | 24+61 | 36 | 36+61 | 36+122 | 72 |
| Threads | 48 | 48 | 48+244 | 72 | 72+244 | 72+488 | 288 |
| vectors | 256 | 256 | 256+512 | 256 | 256+512 | 256+2*512 | 512 |
| Parallelization | OpenMP | TBB | TBB | TBB | TBB | TBB | TBB |
| Vectorization | #SIMD | OpenMP | OpenMP | OpenMP | OpenMP | OpenMP | OpenMP |
| Heterogeneity | N/A | N/A | OpenMP | N/A | OpenMP | TBB | N/A |
| Greek.TIME(SEC) | 4.8 | 1.0 | 0.63 | 0.81 | 0.53 | 0.216 | 0.207 |

1st Heterogeneous Implementation

Dynamic Load Balancing between 3 devices

# Summary

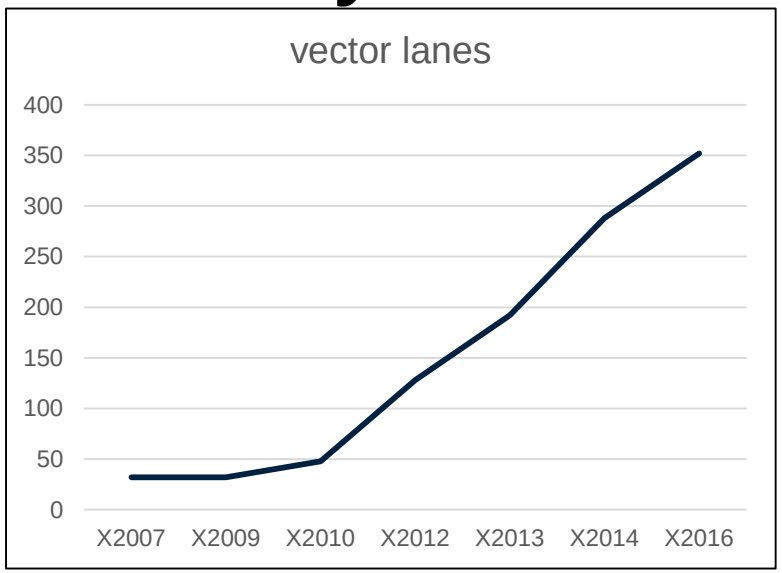
vector lanes
400
350
300
250
200
150
100
50
0
X2007
X2009
X2010
X2012
X2013
X2014
X2016

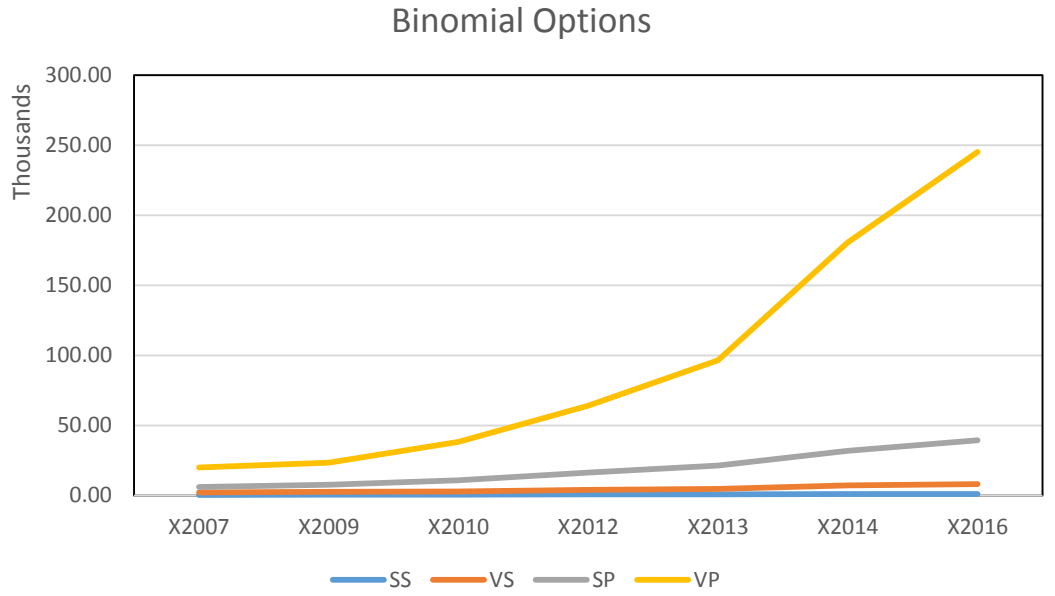
Binomial Options
Thousands
300.00
250.00
200.00
150.00
100.00
50.00
0.00
X2007
X2009
X2010
X2012
X2013
X2014
X2016
SS
VS
SP
VP

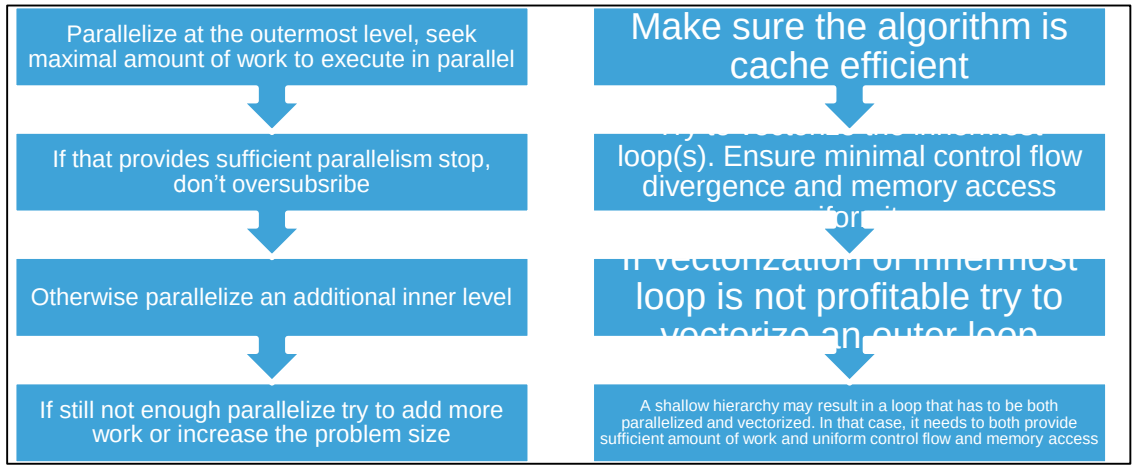
Parallelize at the outermost level, seek maximal amount of work to execute in parallel
If that provides sufficient parallelism stop, don't oversubsribe
Otherwise parallelize an additional inner level
If still not enough parallelize try to add more work or increase the problem size
Make sure the algorithm is cache efficient
loop(s). Ensure minimal control flow divergence and memory access
loop is not profitable try to
A shallow hierarchy may result in a loop that has to be both parallelized and vectorized. In that case, it needs to both provide sufficient amount of work and uniform control flow and memory access

Hewlett Packard Enterprise

Hewlett Packard Labs

# Blazingly Fast Monte Carlo to Accelerate Decision Making in Financial Services

Natalia Vassilieva, PhD
Senior Research Manager

# Monte Carlo simulations in finance

- Used to value and analyze instruments, portfolios and investments
- Enable better decision making
- Mostly used for
  - Derivatives pricing
  - Risk management

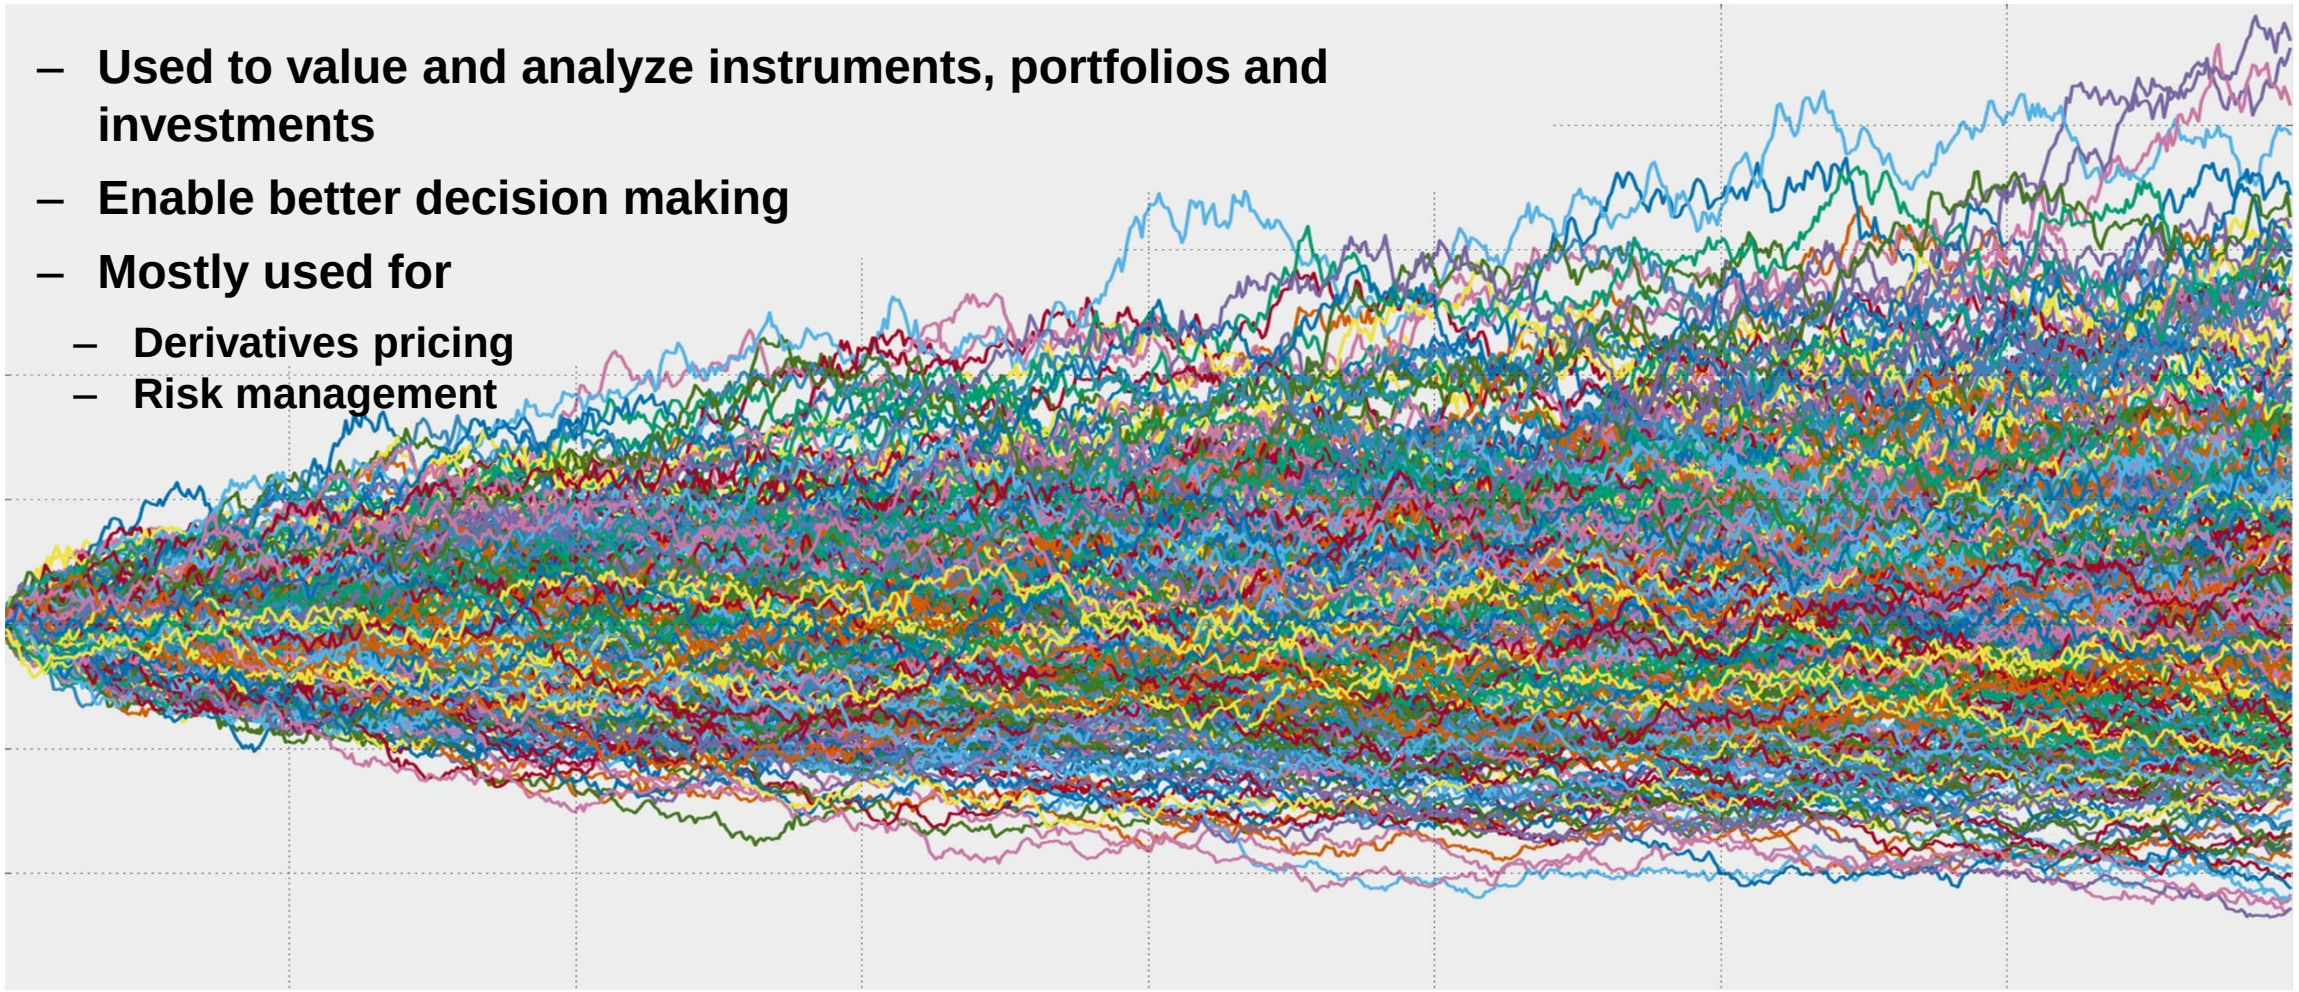

# The need for fast simulations is recognized



## Today's solution: use accelerators (GPUs)

Major trend is the need for "a massive acceleration in calculations". "**The use of GPUs has led to calculation speed increases of between 60 and 300 times**, " making it possible to manage [books of complex products] in quasi real time, instead of once or twice a day, and allowing enough Monte Carlo simulations [to take place] to get smooth gamma for better risk management."

Source: Murex Overall Technology ranking (https://www.murex.com/webdoc)

# Our solution: Memory-Driven Monte Carlo simulations

Leverage large memory to run Monte Carlo simulations **up to 10,000x faster**


## Traditional

Step 1: Create a parametric model $y = f(x_1,\ldots,x_k)$
**Step 2: Generate a set of random inputs**
**Step 3: Evaluate the model and store the results**
Step 4: Repeat steps 2 and 3 many times
Step 5: Analyze the results

## Memory-Driven

Capacity to store representative behaviors of pre-simulated model allows us to **replace** steps 2 and 3 with look-ups and simple transformations

Hewlett Packard Enterprise

# Example: empirical comparison with S&P 500 data

## Option pricing with Memory-Driven Monte Carlo

| Market data (true behavior): | Average Volatility | Option Price |
|---|---|---|
| | 20.31% | $39.56 |

| Model | Complexity | Average Volatility | Option Price | Mispricing | % Error | Time (ms) |
|---|---|---|---|---|---|---|
| Black Scholes | Low | 25.53% | $42.99 | $3.43 | 8.67% | 0.0017 |
| Heston | Low | 23.87% | $41.87 | $2.31 | 5.84% | 0.786 |
| VAR (with traditional Monte Carlo) | High | 23.23% | $41.41 | **$1.85** | **4.68%** | **24210** |
| VAR (with Memory-Driven Monte Carlo) | High | 23.23% | $41.44 | **$1.88** | **4.75%** | **3.18** |

The average volatility values are based on the S&P 500 index options averaged over a 9-year period.
The option value used in this example is $39.56 and the strike date is 10 days ahead.
The training data for the model fit are Wednesday call options, and the test data are Thursday call options.
The true option behavior is the actual volatility for the S&P 500 index for the mentioned period.
VAR model: http://fic.wharton.upenn.edu/fic/papers/09/0906.pdf

# Experimental comparison: Memory-Driven MC v.s. traditional MC

Option pricing and portfolio value-at-risk

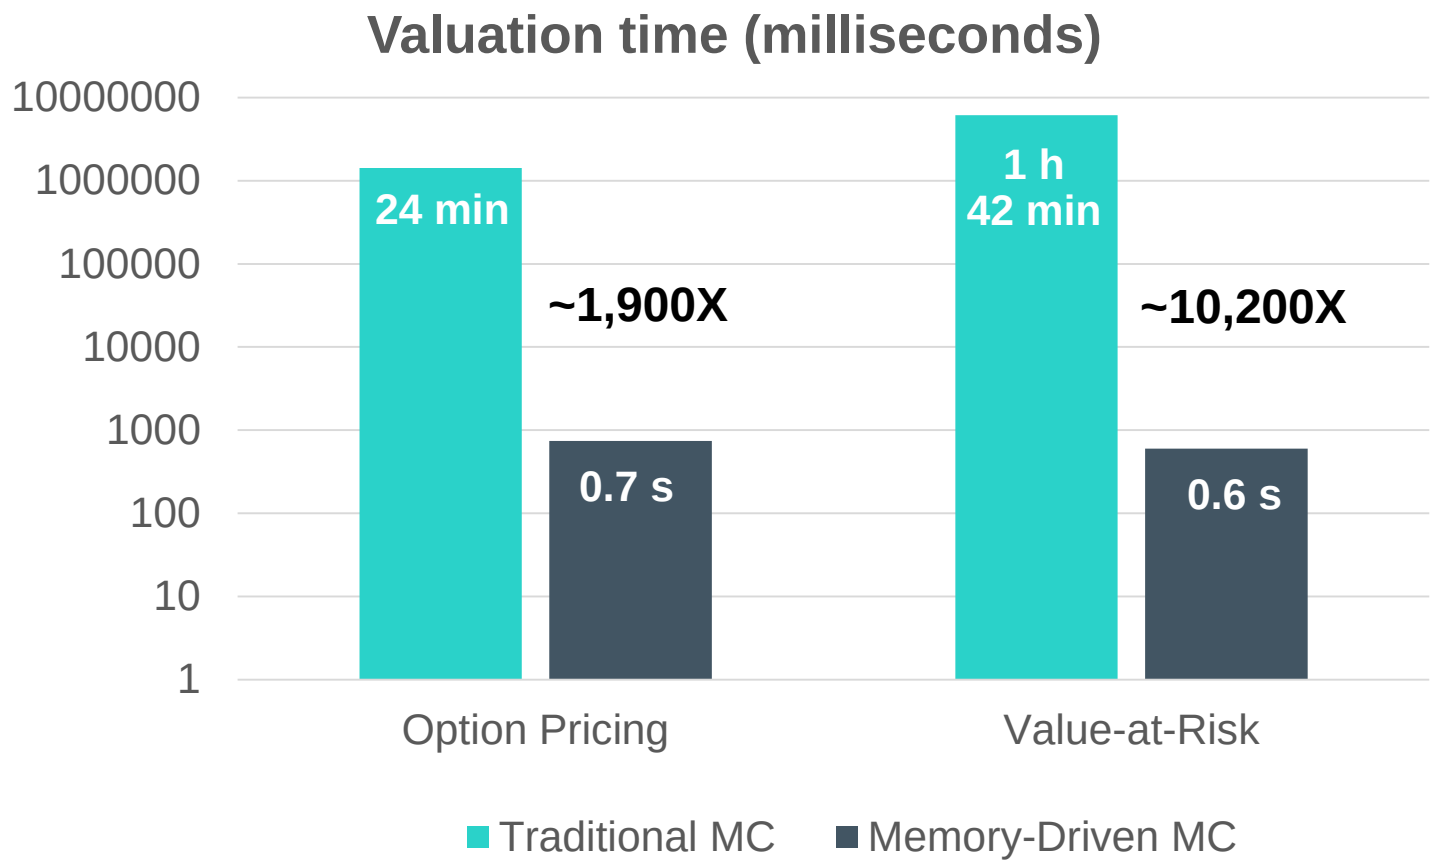


**Option pricing**

Double-no-Touch Option with 200 correlated underlying assets

Time Horizon: 10 days

**Value-at-Risk**

Portfolio of 10000 products with 500 correlated underlying assets

Time horizon: 14 days

# Experimental comparison: Memory-Driven MC v.s. traditional MC

## Local volatility vega profiles for Double No Touch

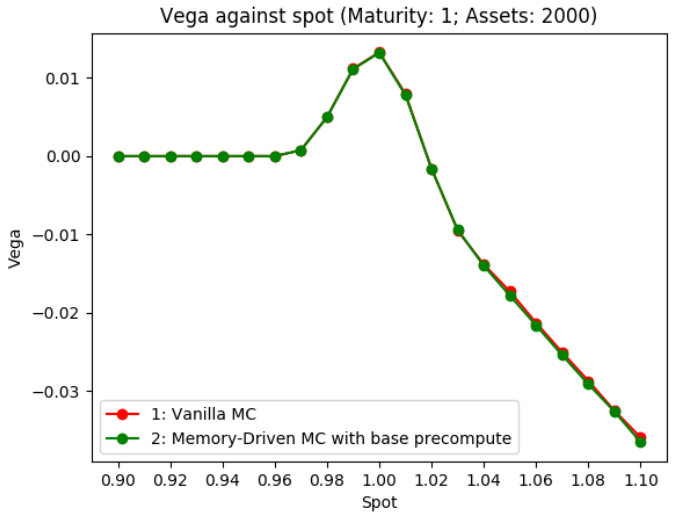


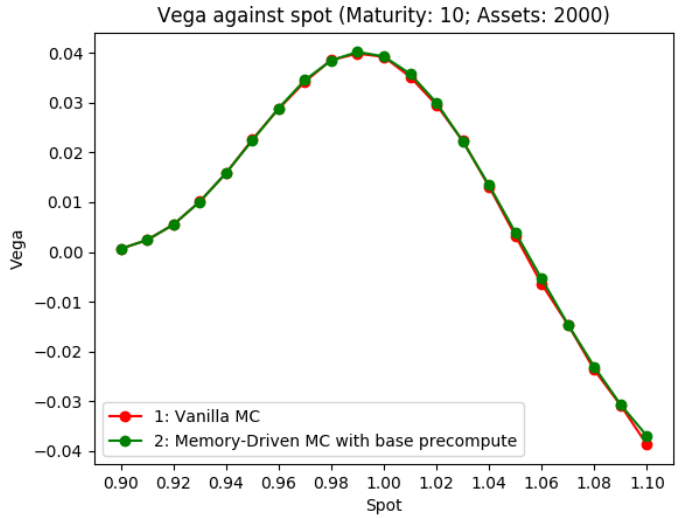


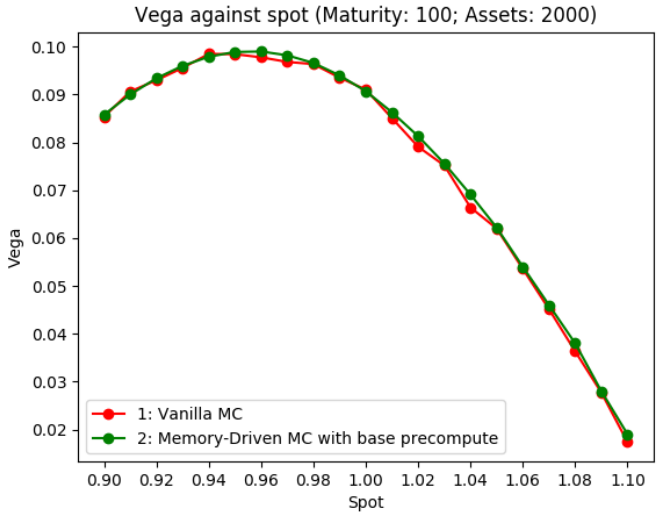


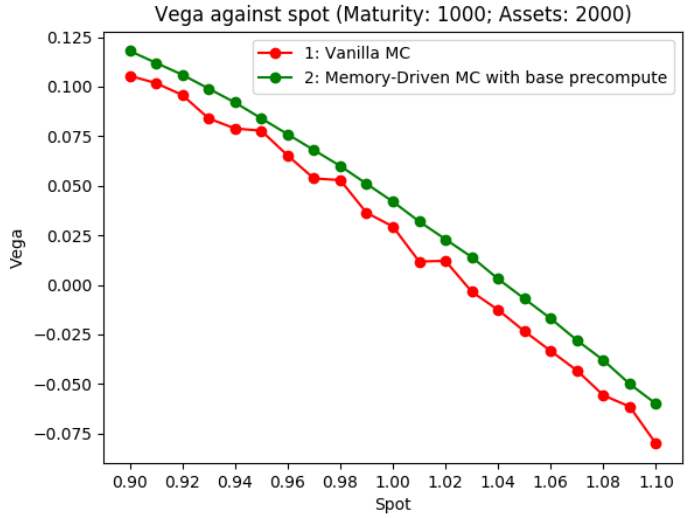


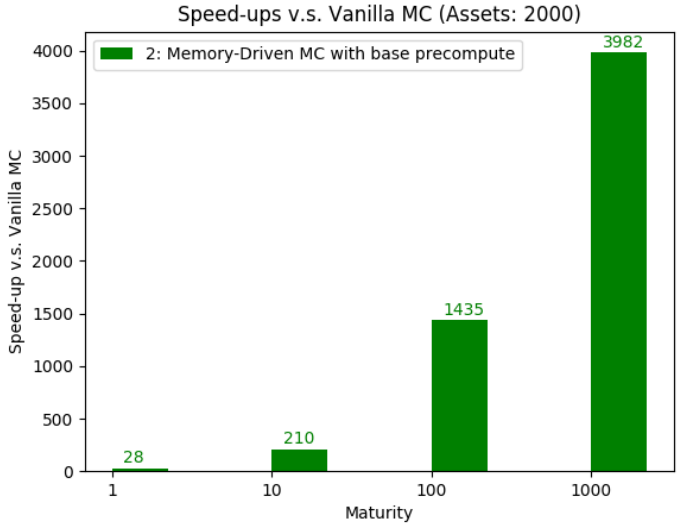


Hewlett Packard Enterprise

# Required infrastructure: ~100TB RAM

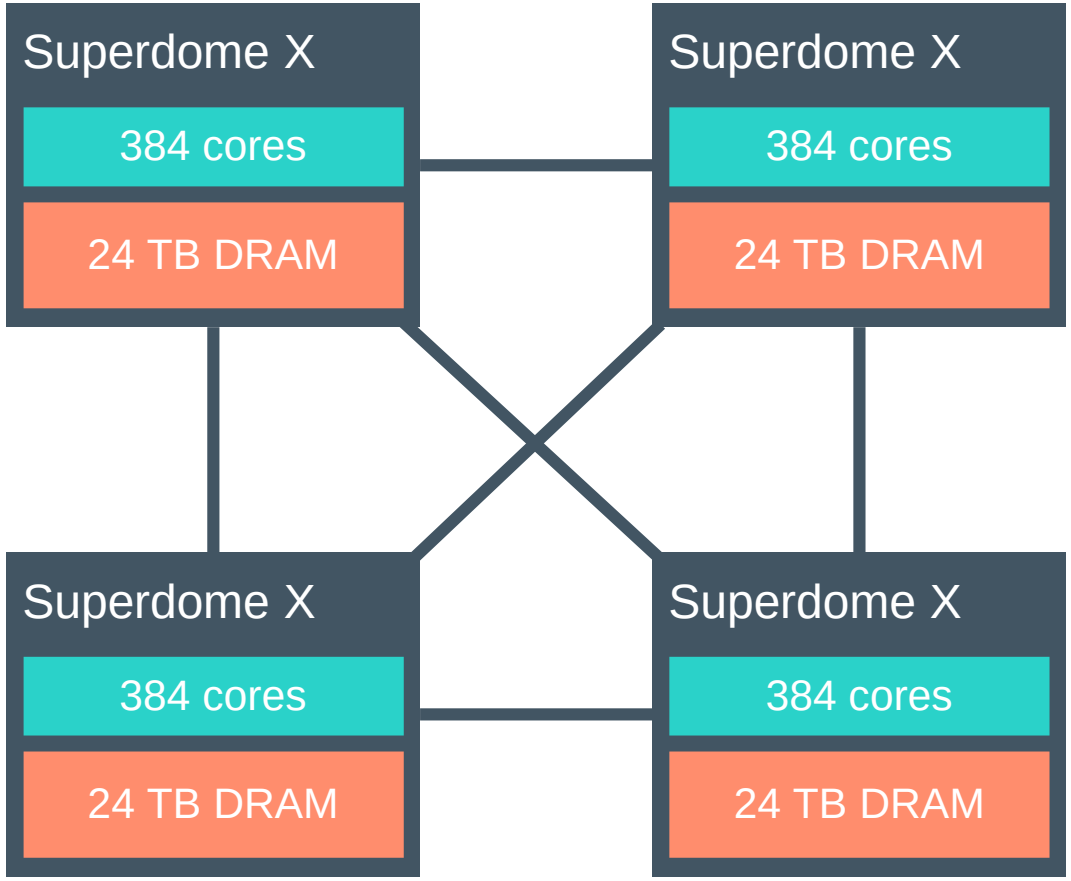


HPE Integrity Superdome X or SGI UV 300



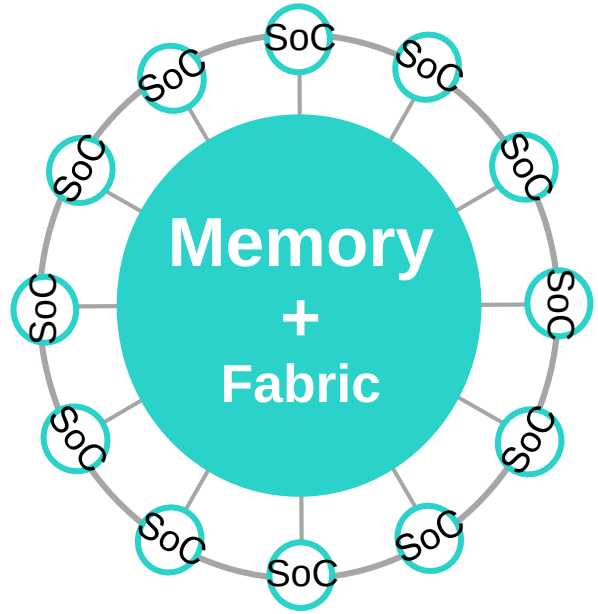


Memory Driven Computing

Hewlett Packard Enterprise

# From PoC in Labs to commercial solution

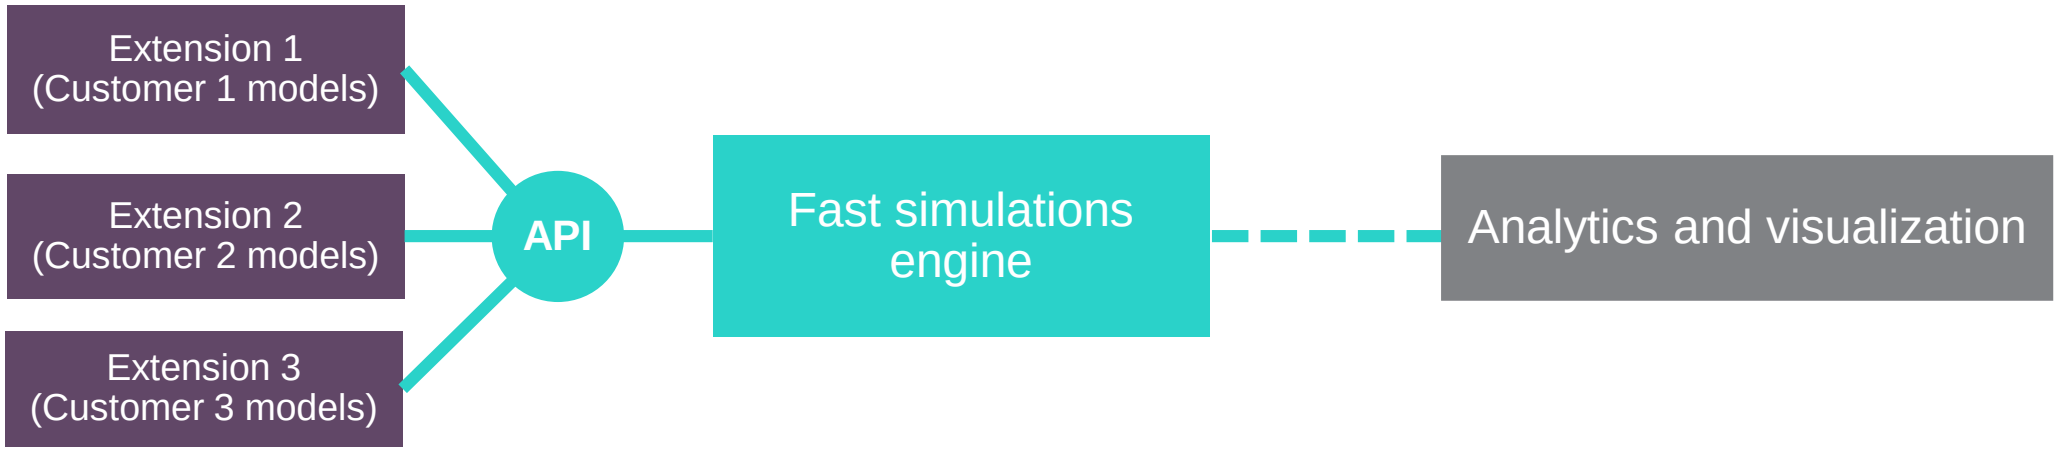


- Agnostic to your models: add extension which implements your model
- Mark-to-market and mark-to-future values in real time with Fast simulations engine
- Can provide integration with other common tools if necessary

# Memory Driven computing to revolutionize financial industry

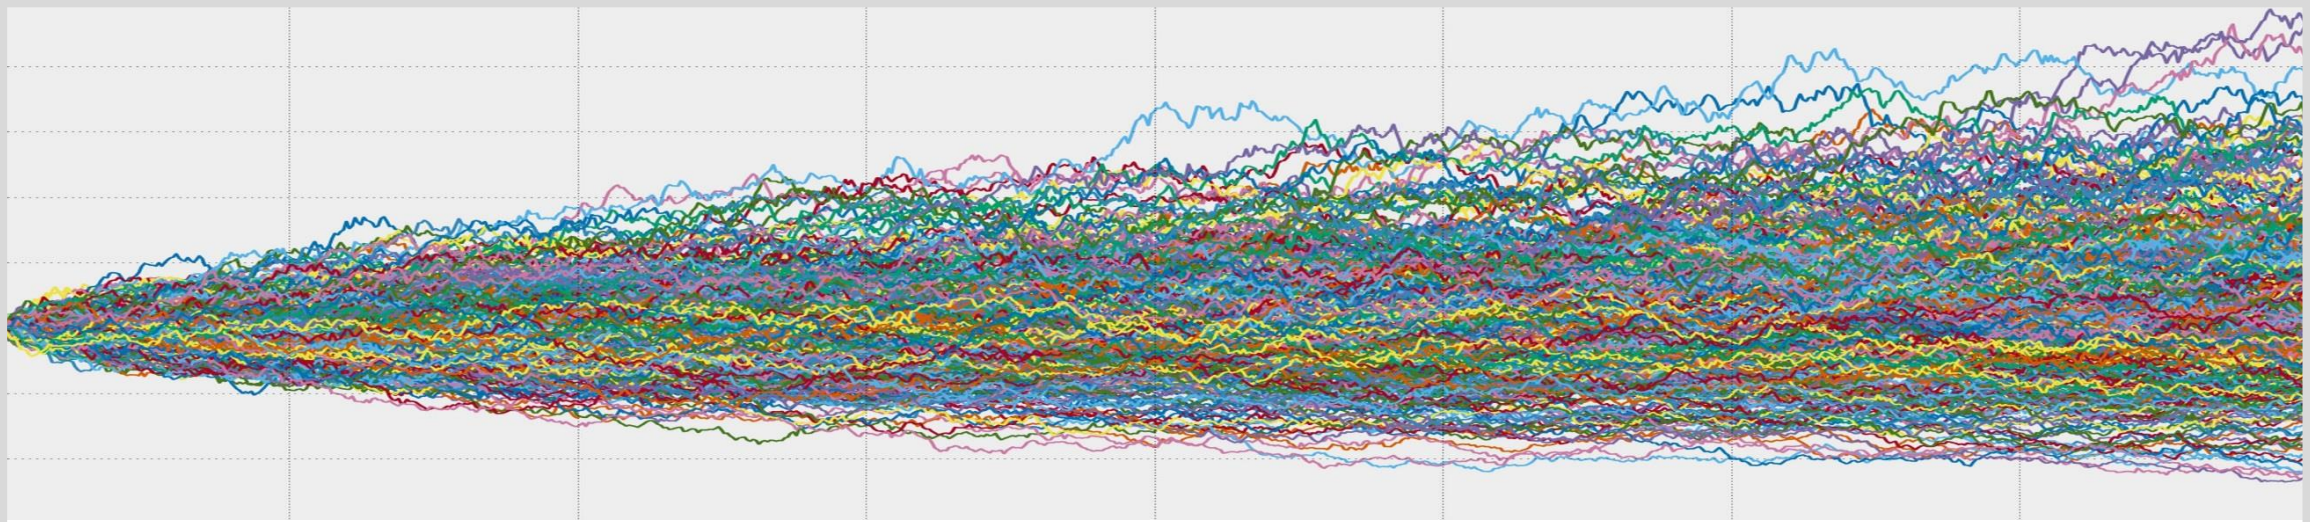

**Truly fast Monte Carlo simulations**

- Accurate pricing of complex deals in real time
- Portfolio risk estimation in real time
- Assessment of multiple scenarios

**Change the way you do investment decisions**

Hewlett Packard Enterprise

Hewlett Packard
Enterprise

# Thank you

Natalia Vassilieva
nvassilieva@hpe.com