

Components of a Research Computing Resource

Hussein Al-Azzawi
UNM Center for Advanced Research Computing
06/09/2021


<https://www.linkedin.com/in/hazzawi/>





CENTER FOR ADVANCED RESEARCH COMPUTING

- First public HPC Center in the State of New Mexico
- In 2000, UNM, in collaboration with IBM, built the first-ever Linux beowulf cluster in production
- Main research fields:
 - Nuclear
 - Materials science
 - Astrophysics
 - Genome sequencing





What is HPC?

What is a Cluster?





It Depends!

--Henry Neeman





What is a Cluster!

"... [W]hat a ship is ... It's not just a keel and hull and a deck and sails. That's what a ship needs. But what a ship is ... is freedom."




– Captain Jack Sparrow

"Pirates of the Caribbean"

Credit: Henry Neeman & Dana Brunson

http://www.oscer.ou.edu/acirefvirtres2016_talk_brunson_clusterdesign_20160811.pdf



“High-performance computing (**HPC**) is the use of ***parallel processing*** for running ***advanced application*** programs efficiently, reliably, and quickly.”

<http://searchdatacenter.techtarget.com/definition/high-performance-computing-HPC>

A cluster ***needs*** of a collection of small computers, called ***nodes***, hooked together by an ***interconnection network*** (or ***interconnect*** for short).

It also ***needs*** software that allows the nodes to communicate over the interconnect.

But what a cluster ***is*** ... is all of these components working together as if they're one big computer ... a ***super*** computer.

Credit: Henry Neeman & Dana Brunson

http://www.oscer.ou.edu/acirefvirtres2016_talk_brunson_clusterdesign_20160811.pdf








Computer Simulation

A computer simulation is a program that attempts to reproduce an abstract model of a particular system in the virtual space before building it in reality.

Parallelism

Performing computations in parallel by carrying out many calculations simultaneously. Utilizing thousands of general-purpose computers all working together to solve the same problem at the same time. This in fact is an excellent case for how modern supercomputers work.



Summit in its natural habitat - America's Fastest and Smartest HPC Cluster



image courtesy of Oak Ridge National Laboratory website



But, Why?

1

To resolve tomorrow's problems, which can't be studied using today's infrastructure, by using computer simulations

2

Cluster HPC helps bypassing the limitations of Moore's law

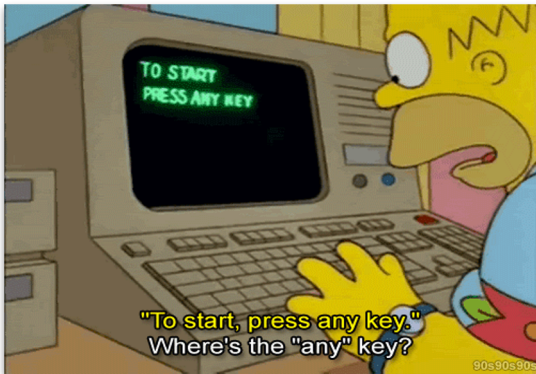
3

Increase the research outcomes by an order of magnitude by manipulating excessive data points and machine learning

4

Solve expensive and dangerous problems in simulations

HPC Challenges..



Staffing



\$\$\$



Power, Space, Cooling



Return on Investment (ROI)

“The Purpose of Research Computing is the Research, not the Computing”

Jonathan Dursi - <https://www.dursi.ca/post/purpose-of-research-computing>



Hats HPC staff “geeks” wear

- HPC Engineer
- Security/Network/Storage
- Research and Development
- CI Facilitator/Applications Scientist
- Negotiator and translator (Enterprise <-> Research)
- Heating, ventilation, and air conditioning (HVAC)



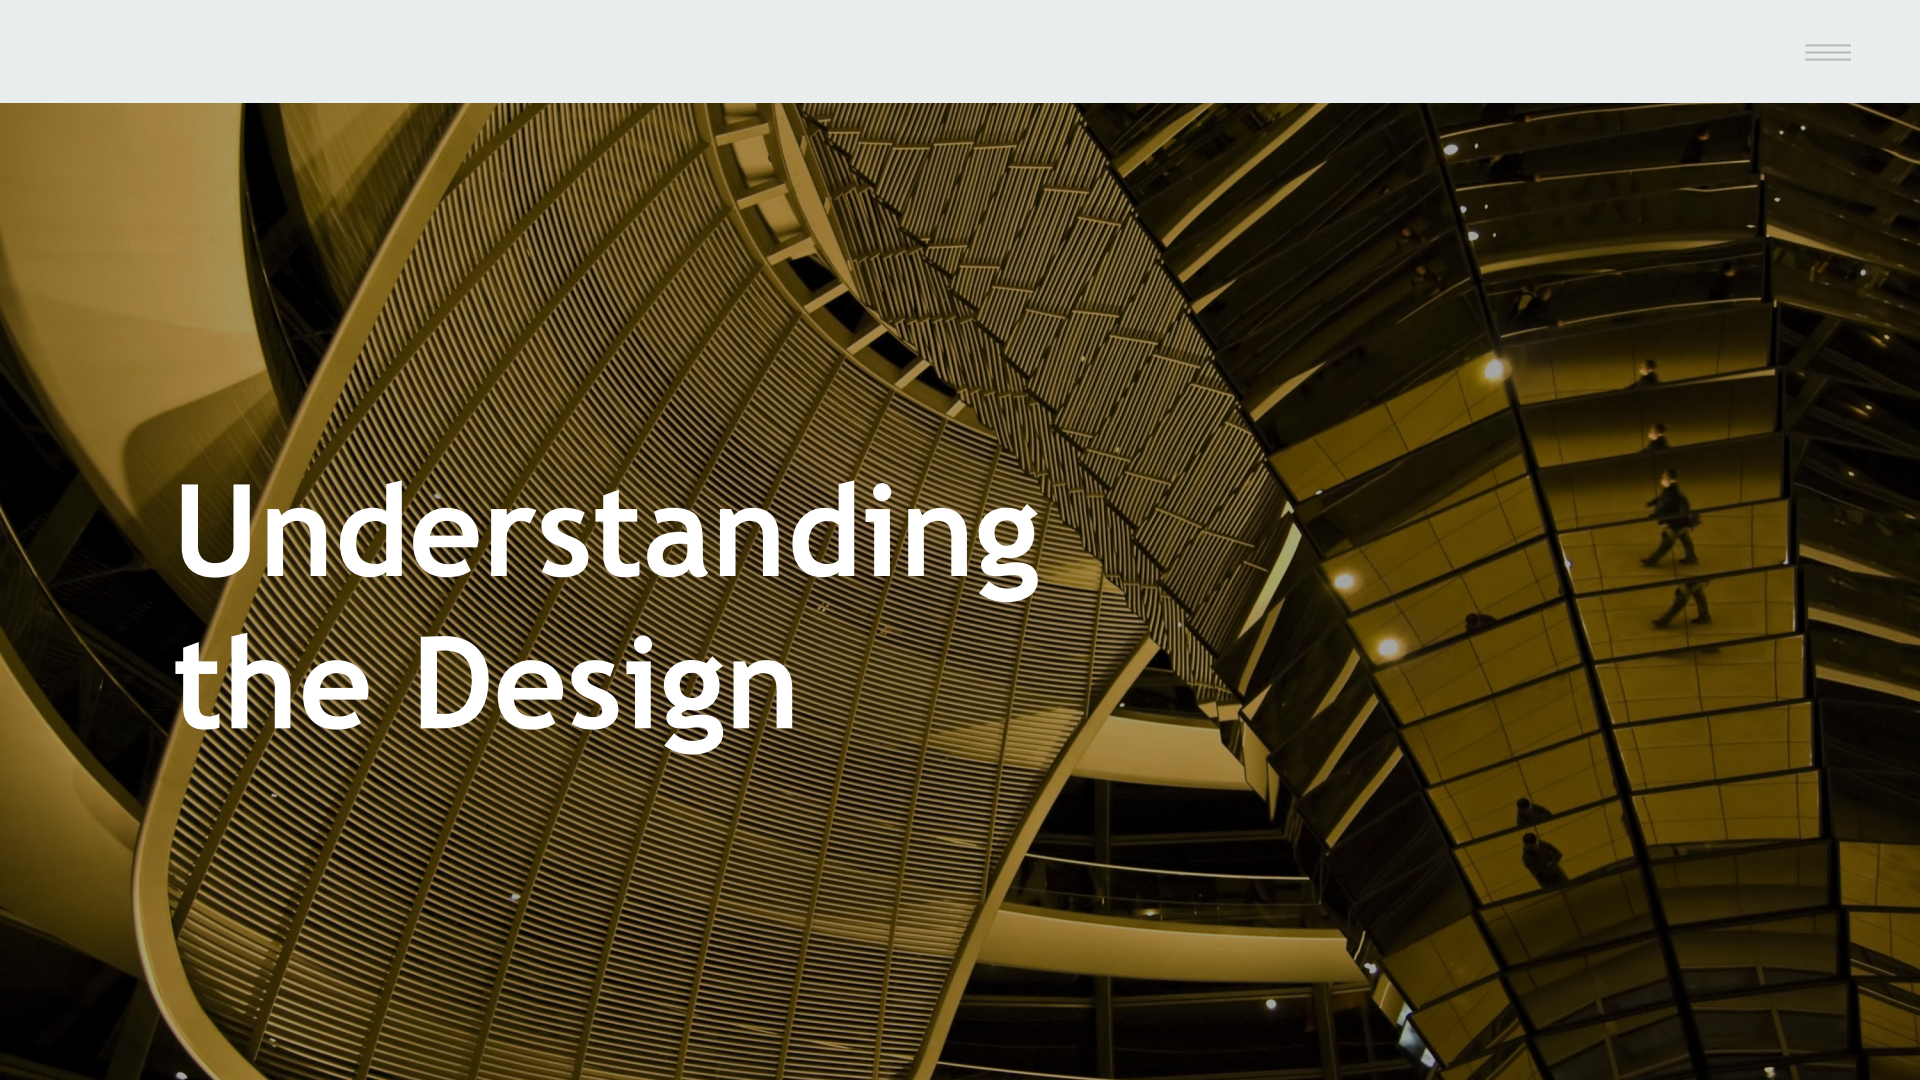
“Mile wide and half inch deep”



"It couldn't be a grad student; we had to hire a real person."

--Anonymous ;)





Understanding the Design


```
public class Program
{
    static void Main()
    {
        Console.WriteLine("Hello, World!");
    }
}
```

Left brain


I am the left brain.
I am a scientist. A mathematician.
I love the familiar. I categorize. I am accurate. Linear.
Analytical. Strategic. I am practical.
Always in control. A master of words and language.
Realistic. I calculate equations and play with numbers.
I am order. I am logic.
I know exactly who I am.

```
6*300111 // 6*300111
// declina 4d'eat'
// 1000000 is 1000000

16 1000111 // 16 1000111
// 1000000, 1000000, 1000000
state 10 1000000 // 10 1000000
pe 10 1000000 // 10 1000000

end

6*300110, 4 1000110, 6 1000110, 8 1000110, 10 1000110
6*300110, 4 1000110, 6 1000110, 8 1000110, 10 1000110
6*300110, 4 1000110, 6 1000110, 8 1000110, 10 1000110
```



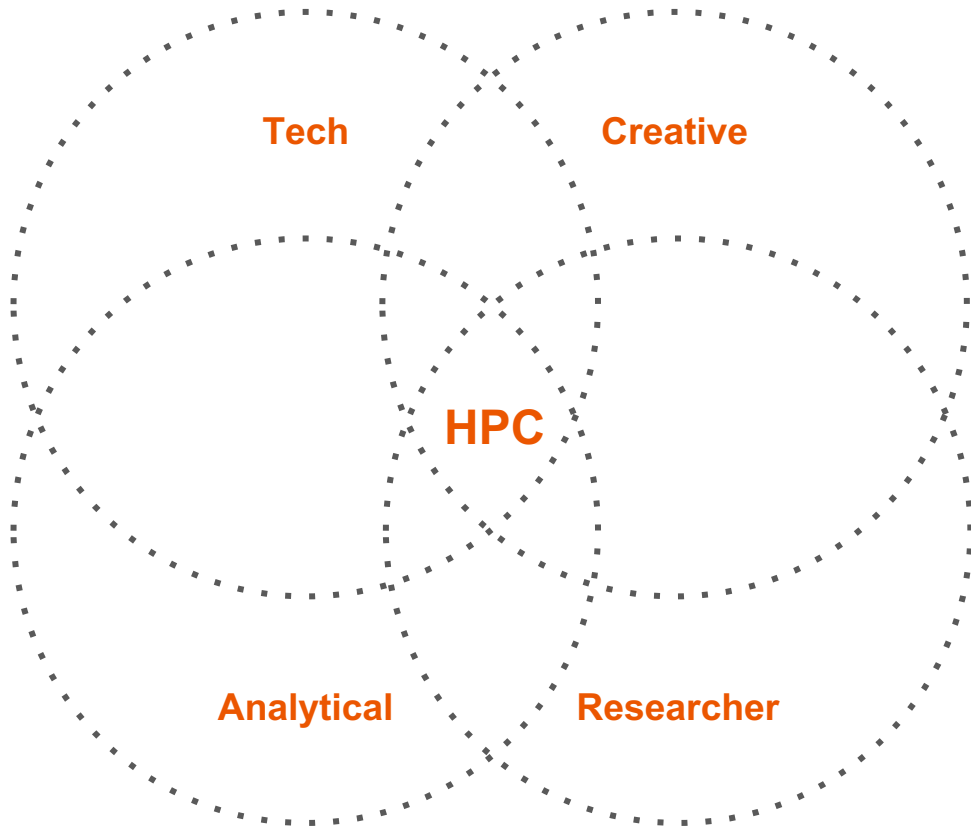
Right brain

I am the right brain.
I am creativity. A free spirit. I am passion.
Yearning. Sensuality. I am the sound of roaring laughter.
I am taste. The feeling of sand beneath bare feet.
I am movement. Vivid colors.
I am the urge to paint on an empty canvas.
I am boundless imagination. Art. Poetry. I sense. I feel.
I am everything I wanted to be.






Point of View




Cluster Components



Compute



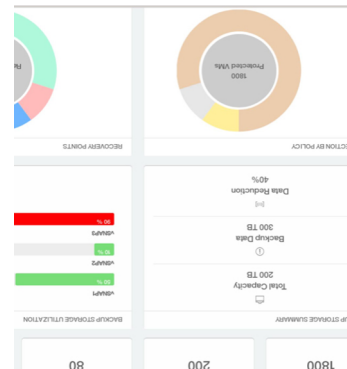
Network



Storage



Batch Scheduler



HPC SW

Compute

- CPU
- GPU
- Memory
- Interconnects
- Diskful/Diskless





Network Interconnects

Main factors:

- Latency
- Bandwidth

Types

- InfiniBand
- Ethernet



image courtesy of Mellanox Store website

Storage

Storage Usage

- Home directories
- Scratch space
- Project space

Storage Types

- Traditional file systems
- Parallel file systems



Resource Manager/ Scheduler

Resource Manager oversees the available resources of the cluster

- Warewulf, xCAT..

Scheduler determines what runs where..

- Slurm, Torque, Moab..

Me: Submits a new job
without checking input filenames


slurm.out: job terminated

Me:



HPC Software

- General usage
- Field dependant

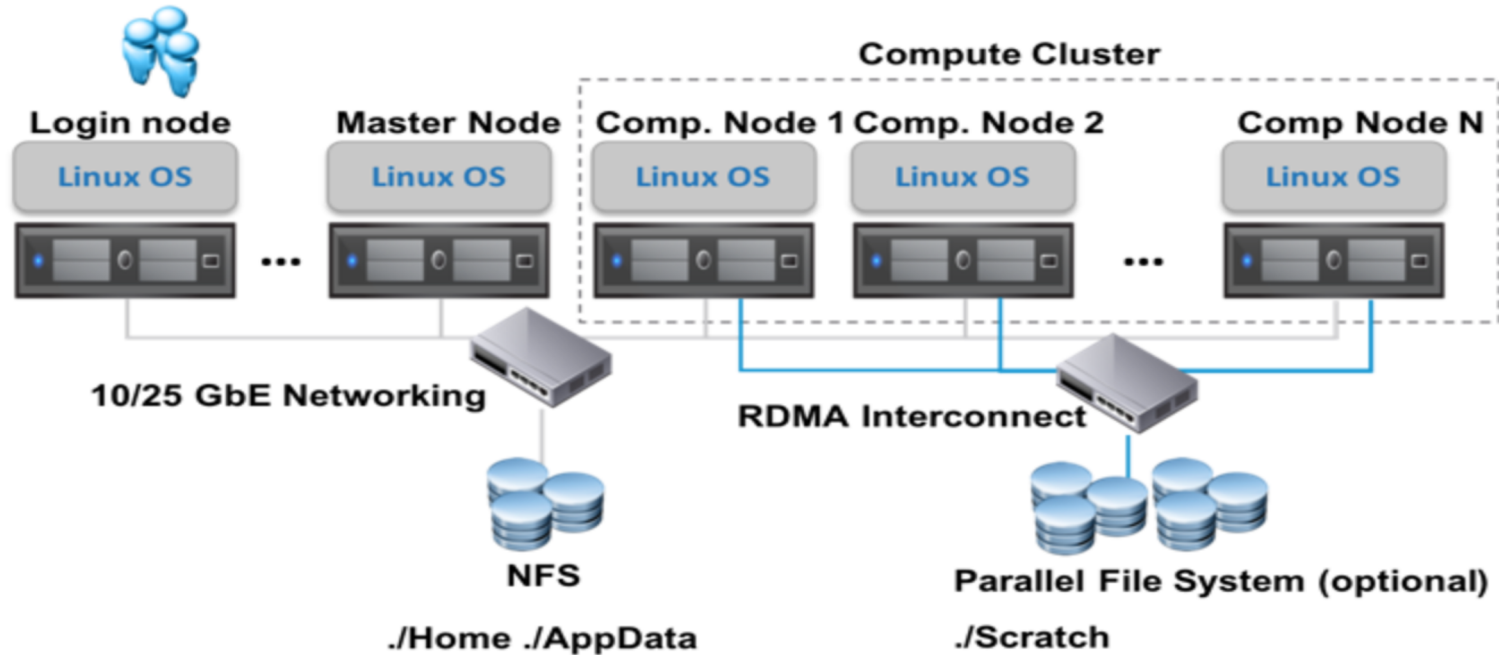


Other Components

- Management/Head nodes
- Login nodes
- Benchmarks
- Security and logs
- Monitoring utilities
- Backup systems
- Cooling systems
- Backup power
- Licensing
- Web interfaces/science gateways
- Data Transfer Nodes (DTN)
- Science DMZ



Traditional HPC Cluster





Acquisition start to finish

- Get Money
- Spec system based on what's needed
- Get bids (informally or formally)
- Buy
- Don't sign acceptance before you test everything with at least some of your workloads

Be warned: No matter what you do, the minute you send out the PO you'll think of something you should've done differently.

————— And it's okay, we all feel that way.

HPC in the Cloud

Pros


- Cost
- Easy to build
- Cloud bursting
- Huge business investments

Cons

- Cost
- Compliance
- Vendor specific
- Harder to track on the institutional level



Cloud/Virtualized HPC Cluster





Computational Science vs Computer Science



How long is a trillion seconds?

~32,000 years



Computers store numbers in floating-point format, so they are called floating-point numbers.

A single instruction like addition or multiplication is called an operation,

HPC speed is measured in terms of floating-point operations per second or Flop/s (Flops)

Rmax - Maximal LINPACK performance achieved

Rpeak - Theoretical peak performance

Ops per second	Scientific Notation	Metric Prefix	Unit
1 000	10^3	Kilo	Kflops
1 000 000	10^6	Mega	Mflops
1 000 000 000	10^9	Giga	Gflops
1 000 000 000 000	10^{12}	Tera	Tflops
1 000 000 000 000 000	10^{15}	Peta	Pflops
1 000 000 000 000 000 000	10^{18}	Exa	Eflops



Performance

CPU clock frequency in GHz
Clock speed in cycles per second (Hertz)
of floating-point operations per cycle

Example:

2.6 GHz processor

2.6 billion cycles per second

16 floating-point operations per cycle

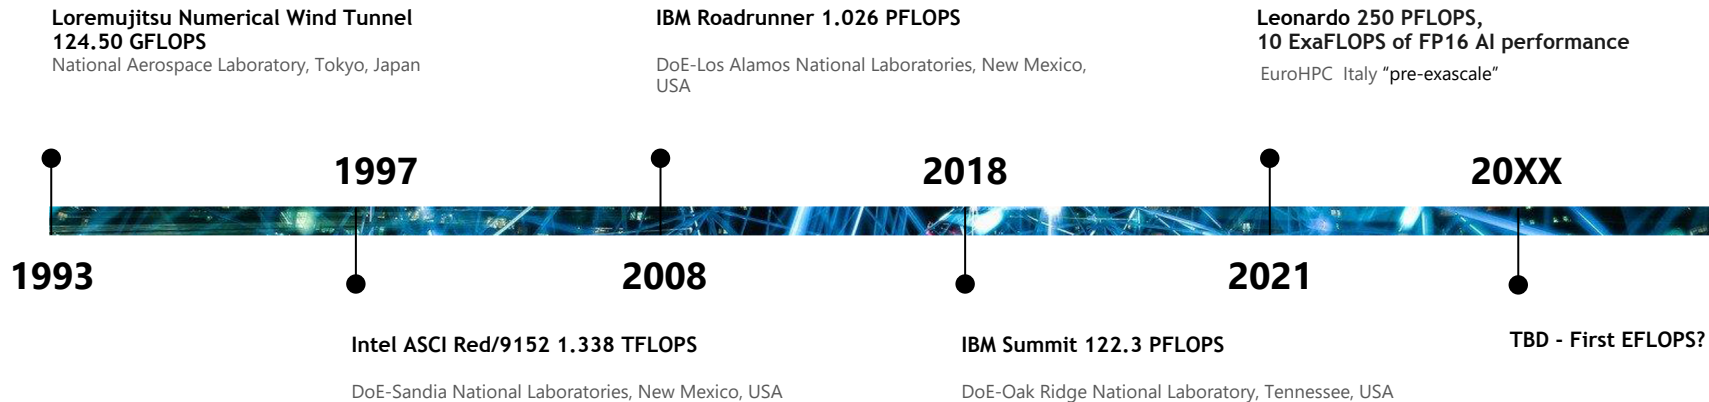
$16 \times 2.6 \text{ billion} = 41.6 \text{ Billion Flops}$

41.6 GFlops





Top 500 History - Peak Speed (Rmax)



<https://www.top500.org/>

<https://www.cineca.it/en/hot-topics/Leonardo-announce>



Quantum Computers


Transistors are already at the atom size,, what's next!

The use of quantum-mechanical phenomena such as superposition and entanglement to perform computations.

Qubits vs Binary bits

- 1 or 0 quantum state
- 1 and 0 superposition states

Using quantum algorithms to solve specific complex problems beyond the capabilities of a classical HPC



IBM Quantum Computer at SC18 floor

Great article by Scott Aaronson



“What Makes Quantum Computing So Hard to Explain?”

“

- What superposition really means is “complex linear combination.”
- “complex” in the sense of a real plus an imaginary number, while “linear combination” means we add together different multiples of states.
- So a qubit is a bit that has a complex number called an amplitude attached to the possibility that it’s 0, and a different amplitude attached to the possibility that it’s 1.
- These amplitudes are closely related to probabilities, in that the further some outcome’s amplitude is from zero, the larger the chance of seeing that outcome; more precisely, the probability equals the distance squared.

”



XSEDE HPC Systems

- 01 | Stampede2 | TACC
- 02 | Comet | SDSC
- 03 | Bridges-2 | PSC
- 04 | Jetstream | IU
- 05 | Open Science Grid | OSG
- 06 | Ranch | TACC



XSEDE

Extreme Science and Engineering
Discovery Environment

image courtesy of XSEDE website

<https://www.xsede.org/ecosystem/resources>

Thank You!

Hussein Al-Azzawi
UNM Center for Advanced Research Computing
06/09/2021

<https://www.linkedin.com/in/hazzawi/>

