
Workstations and Distributed Computing

Western PowerWorld User Conference

6/19/2019



Purpose

- What Do We Want?
 - More Processing Power!
- When Do We Want It?
 - Now!, or before this year's TPL effort.

Distributed Computing Add On

- Utilize your current system to full typology
 - Available for Powerflow analysis and Transient Analysis
- *** MUST BE ENABLED FOR EACH ANALYSIS TYPE *****

Testing Configuration	Total # of Cores	Powerflow Simulation Results				
		# of contingencies	Simulation Time (hrs)	Contingency / Min	Min / Contingency	% reduction
Desktop	1	170,820	44.9	63	0.01577	-
	3		17.1	167	0.0060	62%

Testing Configuration	Total # of Cores	Transient Simulation Results				
		# of contingencies	Simulation Time (hrs)	Contingency / Min	Min / Contingency	% reduction
Desktop	1	724	208.1	0.06	17.2	-
	3		71.1	0.17	5.9	66%

More Power Needed!

- Three Main Options
 - Cluster - Many small machines connected
 - VDI - Virtual Desktop Infrastructure
 - Powerful workstation

Evaluation of Options - Cluster

- Cluster of smaller computers
 - Complex implementation (many machines to take care of / maintain / purchase of software)
 - Physical storage location?
 - Usage of existing machines?
 - Who supports this?
 - Difficult to share

Evaluation of Options - VDI

- VDI – Virtual Desktop Infrastructure
 - Computing in the cloud (more like fake computer)
 - Multiple “computers” can be located on a single VDI.
 - PowerWorld software unable to dynamically add / remove random machines for distributed computing
 - PSE Infrastructure would required additional hardware for implementation (so what’s the point?)
 - This is the future, but depends upon distributed computing software limitations – Think using AWS

Evaluation of Options – Powerful Workstation

- Powerful Workstation
 - Easiest implementation
 - Need support from IT
 - Easy to share, via Remote Desktop
 - PSE did it, so it has to be the best, right?

Important Considerations

- Who is maintaining / supporting it?
 - Maintenance / Support drives hardware (Dell, Lenovo, etc)
- Does it need to be shared? By how many people? How often?
- What is the performance specification?
 - Trick question: Should be “How much money is available?”
 - Better question is : With funds available, how can I optimize the system?

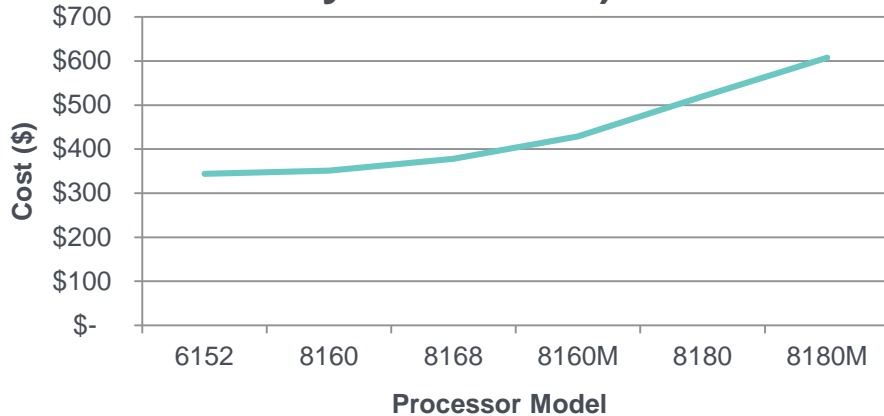
Basic CPU Terminology

- CPU – Central Processing Unit (physical device)
 - Speed - Processor Clock Frequency - GHz
 - Base Clock Speed vs Turbo Speed
 - Cores / Threads
 - Cores - number of processing units on CPU for each thread
 - Threads – single line of commands that get processed by core
 - Hyperthreading – Creates “Logical” cores to simulate multiple cores, allowing for more 2 threads per “Physical” core.

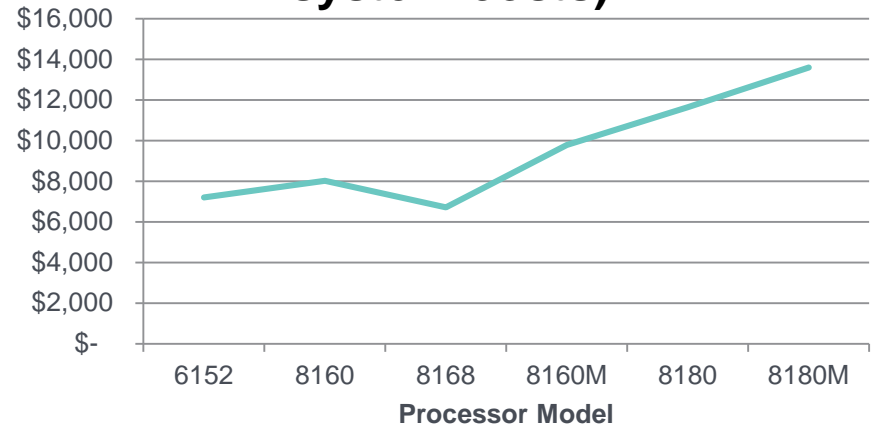
Make	Model	Speed		Per CPU		
		Base	Turbo	Cores	Threads	Cache
Xeon	6152	2.1	3.7	22	44	30.0
Xeon	8160	2.1	3.7	24	48	33.0
Xeon	8168	2.7	3.7	24	48	33.0
Xeon	8160M	2.1	3.7	24	48	33.0
Xeon	8180	2.5	3.8	28	56	38.0
Xeon	8180M	2.5	3.8	28	56	38.0

CPU Considerations – cost effectiveness

\$ / core (2018, based on total system costs)



\$ / Ghz, (2018, based on total system costs)



PSE Workstations

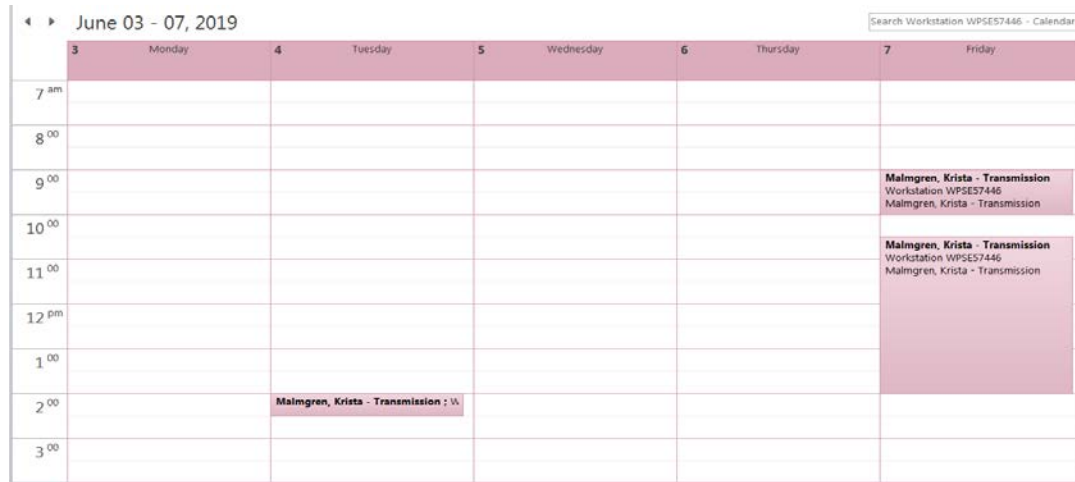
- 3 Workstations
 - 2 CPU's per workstation, Intel Xeon 8168 processors
 - 24 cores per CPU (48 per workstation, 144 total)
 - 256 GB of RAM (yes, that is correct)
 - 1 TB SSD internal hard drive (results not stored locally)
 - Headless configuration, accessed via remote desktop

PSE Workstation - Considerations

- Where are the workstations going?
 - MDF room in the Snoqualmie Data Center
 - Climate controlled, clean environment
- Who is supporting this?
 - IT supports it, response time can vary, issue for PSE
- How will results be stored?
 - Powerflow can be saved locally, Transient saved directly to shared drive

PSE Workstation - Considerations

- How will the workstation be shared?
 - Headless units require remote desktop for access
 - Easy way to share a computer, limits conflicts / interference
 - PSE created calendars for each computer, shared with 12 people, to help with coordination



Performance Results

Testing Configuration	Total # of Cores	Powerflow Simulation Results				
		# of contingencies	Simulation Time (hrs)	Contingency / Min	Min / Contingency	% reduction
Desktop	1	170,820	44.9	63	0.01577	
	3		17.1	167	0.0060	62%
One Workstation	47		1.9	1530	0.0007	96%
Three Workstations	143		0.7	3986	0.0003	98%

Testing Configuration	Total # of Cores	Transient Simulation Results				
		# of contingencies	Simulation Time (hrs)	Contingency / Min	Min / Contingency	% reduction
Desktop	1	724	208.1	0.06	17.2	
	3		71.1	0.17	5.9	66%
One Workstation	47		9.9	1.22	0.8	95%
Three Workstations	143		3.5	3.47	0.3	98%


- CPU**
100% 3.37 GHz
- Memory
23/256 GB (9%)
- Disk 0 (C:)
1%
- Ethernet
Not connected
- Ethernet
Not connected
- Ethernet
Not connected
- Ethernet
S: 8.0 R: 0 Kbps
- GPU 0
NVIDIA Quadro P
4%

CPU


Intel(R) Xeon(R) Platinum 8168 CPU @ 2.70GHz

% Utilization over 60 seconds

100%



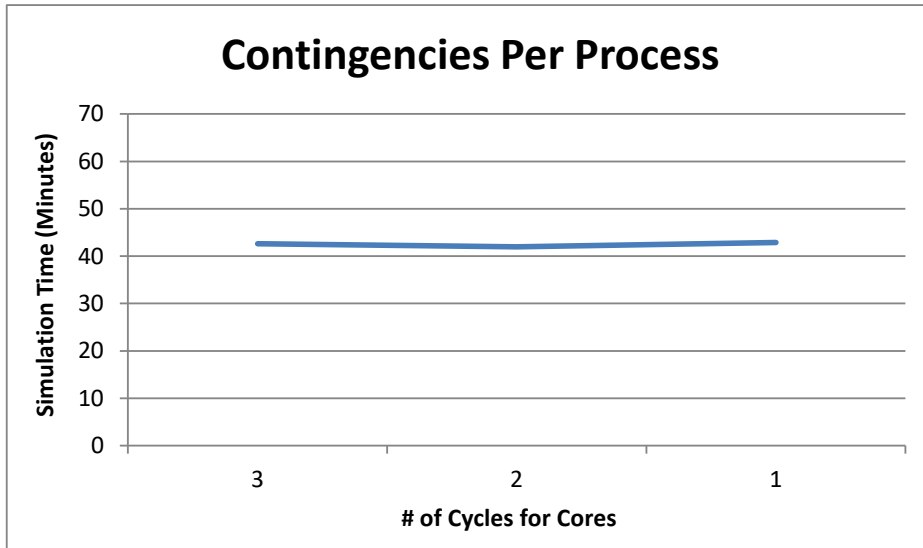
Utilization	Speed	Base speed:	2.69 GHz
100%	3.37 GHz	Sockets:	2
Processes	Threads	Cores:	48
291	4064	Logical processors:	48
Handles		Virtualization:	Disabled
99007		Hyper-V support:	Yes
Up time		L1 cache:	3.0 MB
0:00:10:29		L2 cache:	48.0 MB
		L3 cache:	132 MB



Actual Speed is 3.37 GH
 Turbo is 3.7 GHz
 Base is 2.7 GHz
 - don't spend money for turbo!

Powerflow Sensitivity Analysis

- Number of Contingencies Per Process (core)

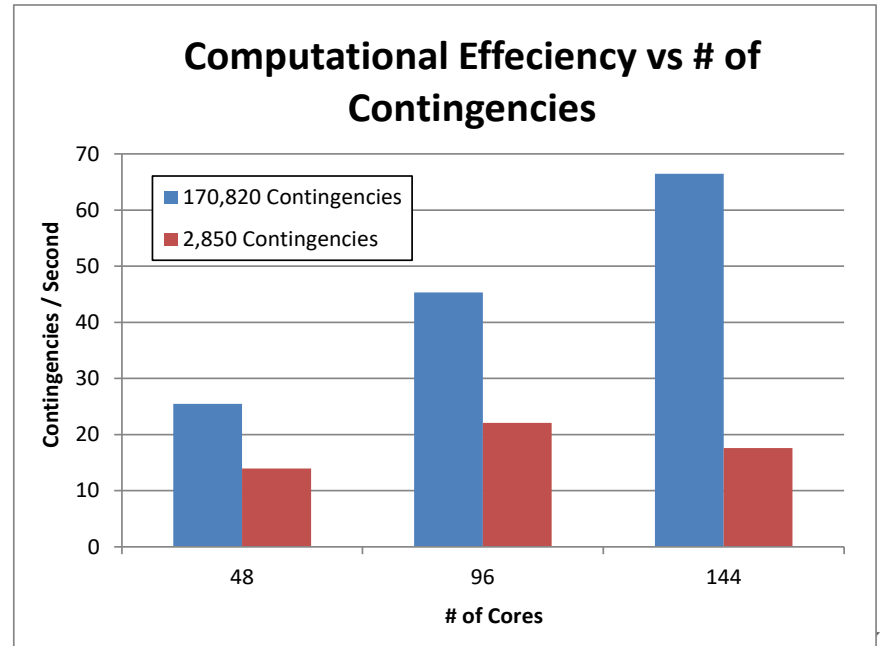


$$\left(\frac{\text{Number of Contingencies}}{\text{Number of Processes}} \right) / 2$$

Powerflow Sensitivity Analysis

- Number of Total Contingencies

Total # of Cores Used	# of contingencies per process	# of contingencies	Minutes	Cont / Sec
48	3634	170,820 Contingencies	112	25
96	1798		63	45
144	1195		43	66
48	61	2,850 Contingencies	3.4	14
96	30		2.2	22
144	20		2.7	18



Last Thoughts

- Get an understanding of budget / scope
- Be mindful of IT support requirements / needs
- Many groups can benefit from Workstation
 - Sharing is Caring.
- Please share your experiences!

Questions



David Meyer

david.meyer@pse.com