Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Benchmarking of cloud and local resources

Davide Grassano

Mid Term Review

21 August 2015

Outline

- Benchmarking
  - What is it?
  - Why is it needed?
- Benchmarks used
  - $t\bar{t}$bar_gensim
  - Hepspec06
- Cloud computing
  - What is it?
  - The cloud project
- Obtained results
- Prospect for the future

What is benchmarking?

- A series of tests to assess the capabilities of the hardware and software of the machine of interest
- Gives a metric to compare the performance of machines with different architecture
- Both standard benchmarks and specific ones (that mimics part of the workload of a full job) can be used

Why is it needed?

- Help to determine which machines are most suited to run your jobs on
- Allows the optimization of the allocation of resources
- Gives an estimation of how long the job will need to run

Benchmarks used

- $t\bar{t}$bar_gensim
  - Simulates the generation of a small (150) number of ttbar events.
  - Give metrics specific for a full CMS job based on the ttbar/s result
- hepspec06
  - Smaller package of the more notorious collection of benchmarks SPEC2006
  - Stress the CPU and compiler of the system
  - Gives a more standard metric for comparison of resources

What is Cloud computing?

- A cloud is a collection of services and infrastructure accessible from anywhere in the network
- Allows the sharing of resources for the users and the achievement of an economy of scale for the host
- The resources are not static, but dynamically reallocated on a ‘on demand’ basis, maximizing their efficiency.

What is Cloud computing?

The cloud project

- Allows a dynamic and more flexible allocation of resources on a on demand basis
- A local cloud can be used to share resources not in use with partner institutions
- A commercial cloud (like AWS) allows an extra allocation of resources for the execution of big workloads, with no upfront commitment of money

Obtained result AWS

Amazon	N_CC	CORE T	Speed	$ per hour	ttbar/s per	ttbar/s tota	ttbar per $/h	HS06 per cor	HS06 total	HS06 per $/h
m3.xlarge	4	on E5-26	2,50	0,266	0,0139	0,0557	0,209	14,3	57,1	215
m3.2xlage	8	on E5-26	2,50	0,532	0,0139	0,111	0,208	12,2	97,6	184
m4.xlarge	4	on E5-26	2,40	0,252	0,0201	0,0806	0,320	16,1	64,5	256
m4.2xlage	8	on E5-26	2,40	0,504	0,0191	0,153	0,304	15,1	121	240
m4.4xlarge	16	on E5-26	2,40	1,008	0,0198	0,317	0,315	13,5	217	215
c3.xlarge	4	on E5-26	2,80	0,105	0,0153	0,0611	0,582	14,9	59,4	566
c3.2xlage	8	on E5-26	2,80	0,210	0,0153	0,122	0,583	14,7	118	561
c3.4xlarge	16	on E5-26	2,80	0,420	0,0149	0,239	0,568	13,2	212	504
c4.xlarge	4	on E5-26	2,90	0,441	0,0228	0,091	0,207	17,5	69,9	158
c4.2xlage	8	on E5-26	2,90	0,882	0,0226	0,181	0,205	16,5	132	150
c4.4xlarge	16	on E5-26	2,90	1,763	0,0205	0,327	0,186	14,8	237	134

Obtained result local machine

	N_CORE	CORE TYPE	Speed(	ttbar/s per core	ttbar/s total	HS06 per core	HS06 total
cloudworker1148	8	ntel XEON X535	2,66	0,0179	0,143	8,32	66,5
fnpc2036	8	MD Opteron 238	2,90	0,0217	0,174	12,0	96,1
fnpc3000	16	MD Opteron 613	2,30	0,0173	0,277	9,92	159
fnpc4001	32	MD Opteron 612	2,00	0,0149	0,477	8,64	277
fnpc5009	32	MD Opteron 613	2,30	0,0162	0,520	9,45	302
fnpc6000	64	MD Opteron 637	2,30	0,0136	0,873	10,0	640
fnpc7024	64	MD Opteron 637	2,30	0,0136	0,868	9,49	607
Fermicloud134	1	E5-2660V2	2,20	0,0193	0,0193	17,9	17,9
Fermicloud148	1	E5-2660V2	2,20	0,0192	0,0192	17,8	17,8
Fermicloud149	8	E5-2660V2	2,20	0,0182	0,145	14,3	115
Fermicloud150	1	E5640	2,60	0,0229	0,0229	18,4	18,4
Fermicloud381	8	E5640	2,60	0,0217	0,174	15,2	122
prvm0189	1	ntel XEON X535	2,66	0,0173	0,0173	13,2	13,2
prvm0190	4	ntel XEON X535	2,66	0,0170	0,0680	11,4	45,5
FY2015 bid	48	Intel E2670V3	2,30	0,0195	0,9381		

For the future

- Run the $t\bar{t}$bar_reco benchmark to determine the efficiency loss for the multithreading
- Run full length CMS jobs

Thanks for your attention