

Our Team

Filled with passion and an incomparable latin flavor, we are a team absolutely interested and involved in the world of high performance computing. We are absolutely dedicated and we have an unstoppable desire to excel in the field and show the world that the raw talent that exists in Colombia can truly make a difference in this world.




Fig. 1: EAFIT team

This will be the second time that a team from Universidad EAFIT will compete in the student cluster competition, as last year we were part of a joint EAFIT-Purdue team. All of the experience that we acquired has carried over into our new team, making us not only extremely well prepared but also extremely motivated to outdo ourselves. All the members of our team come from very different backgrounds, but all of them share a common set of values, stemming from the values and virtues instilled in all of us by the University. These experiences, coupled to a constant exposure to different branches of Computer Science, will prepare our team for an extremely productive and fulfilling career.

Our Hardware

After measuring the performance and stability of the applications in many different scenarios, we've decided that, in our particular case, we would be best served utilizing only CPUs, with the added benefit of ensuring that we would not go over the power limit.

- 7 HP ProLiant DL380 Gen9, each with dual 14-core Intel Xeon E5-2695v3 processors at 2.3 GHz.
 - 1 master node with 192 GB of memory and 400GB SSD of shared storage.
 - 6 compute nodes with 64GB of memory.
- 3 HP ProLiant DL60 Gen9, each with dual 10-core Intel Xeon E5-2660v3 processors at 2.60 GHz and 64GB of memory.

Given that Trinity performs better if launched from a node with more memory, we increased the memory of one of the nodes to 192 GB.

Why our team will win

- Our undying desire to be the best, and our unending allegiance to both our team and our university makes us intensely motivated to give it our all in order to make both our university and the world proud of what we will achieve.
- We've studied, as deeply as we've been able to, each and every one of our applications, and we've done our absolute best to optimize them and to make them as stable as we can on our hardware.
- We've gone to the end of the earth to become one with our hardware, understanding it's ins and outs just as well as we know ourselves.
- We have given both our life and our soul to this competition, and our unbreakable resolve to become victorious, combined with the skills we've been building all this year, will allow us to confidently state that we are ready to bask under victory's sweet and shining glory.

How we prepared ourselves

- We were lucky enough to be able to count among our ranks three past SCC participants, and we have learned so much from them, their experiences and the challenges they've overcome.
- As a team, we've prepared with the utmost dedication, holding team meetings twice a week, working with every computational resource we've had available in order to learn as much as we can from both the apps and the hardware.



Fig. 2: Weekly meetings

- In order to balance the workload, but spread the knowledge, we've assigned each team member two of the applications, and, in doing so, we've assured a great sense of camaraderie and a shared sense of responsibility for each app.

	WRF	Milc	Repast	Trinity	Allinea	HPL	OS
Alejandro Gomez					•	•	
Daniel Salazar			•	•			
Camilo Correa		•		•		•	
Alejandro Salgado	•		•				
Johan Yepes	•	•					
Pablo Restrepo							•

Fig. 3: Workload distribution

- In addition, we've met with local experts that work in their day to day with some of the applications that we'll be using in the competition. Most notably, we went to the SIATA, an early warning system for meteorological phenomena in the city, that utilizes WRF. We also met with two professors from our university that work with HPC Repast and Trinity as part of their usual research activities.

Optimizing the apps

Being a team whose core is based on mutual communication and discussion, we devised a quite simple yet infinitely effective strategy for optimizing our applications: we split the workload of every application between two team members, and, as such, every decision, every analysis of a compilation result or allinea run has been discussed thoroughly. Therefore, every decision is born out of consensus, which without a doubt, is far more important than anything else we could've done.




Fig. 4: Applications

We've tested and used every compiler and MPI implementation that we have been able to find and adapt to our cluster, and we've read, twice over, every specification for all of these tools. With this in mind, we've worked to optimize as cleanly and precisely every as we've been able to application. Our main tool for this has been allinea, which has proven itself to be absolutely vital to understanding how our applications behave, and how we could make them behave.