

History of Research Computing Clusters at RIT

IBM Cluster – 1998 - 2010

The IBM Cluster was composed of 44 server nodes, each built from dual P3 1.4 GHz processors and 512 MB RAM. This linux cluster was designed to run parallel computing jobs that are tightly coupled and use the MPI over high speed interconnects (currently 1 Gigabit ethernet switched network). It also ran multiple serial jobs that did not need to communicate with each other but expect to run on dedicated processors.

A 'backfill' mechanism (Condor) was in place where completely opportunistic serial jobs could run in the background when other queued jobs were not utilizing all of the processors.

We inherited this cluster from Bioinformatics.

Solvay Cluster – 2008 - 2011

The solvay cluster was a 'condominium' cluster consisting of shares purchased by researchers. Each node had 32 cores with AMD Opteron 2.3 GHz processors and 64 GB of memory. There were originally three nodes for a core count of 96. A fourth node was purchased one year into the cluster's life which brought the total core count to 128. At this point the memory was also doubled to 128 GB per node. Each node was interconnected with 10 Gigabit links. The collection of nodes was managed by [Sun Grid Engine](#).

Solvay was designed to run parallel computing jobs that are tightly coupled using MPI. However, we found over the years that 99.6% of all computation was actually 'embarrassingly parallel' or 'massively serial' work.

Solvay took its name from the [Solvay Conference](#). Each of the compute nodes were named after a number of famous participants there (Einstein, Bohr, Curie, and Schrodinger).

The Solvay cluster was phased out in late 2011 and the surviving resources were rolled into the Tropos Cluster.

Werner Cluster – 2008 - 2011

At the same time we acquired the Solvay Cluster, an additional node was purchased for use by those Research Computing users who did not have a stake in the Solvay cluster. This node was separated off into its own 'cluster' with a different headnode and name the Werner Cluster (the compute node was called Heisenberg).

In 2010, the scheduler used on werner was changed from the [Sun Grid Engine](#) to the [Simple Linux Unified Resource Manager \(SLURM\)](#). This paved the way for the unification of the solvay and werner compute nodes in the Tropos Cluster under a more flexible SLURM configuration.

Tropos Cluster – 2011 - present

TODO