



IIT KANPUR SETS UP MISSION CRITICAL, HIGH PERFORMANCE COMPUTING POWERHOUSE ON RED HAT ENTERPRISE LINUX

32 NODE AND 96 NODE CLUSTERS POWER NATIONAL RESEARCH PROJECTS

FAST FACTS

Industry:	Academia
Challenges:	Modernize and scale proprietary, SMP-UNIX infrastructure. Set up a high performance computing environment that can handle cutting edge research projects. Simplify management and system administration.
Solution:	Software: Red Hat Enterprise Linux Hardware: AMD64 and Intel 32-bit servers
Benefits:	Improvement in performance and overall computation time through 64-bit x86 servers. Significant TCO reduction by eliminating the need for proprietary licenses and RISC based hardware. Freedom to modify and work with the source code

INTRODUCTION

Indian Institute of Technology (IIT), Kanpur is one of the premier institutions established by the Government of India. The institute has a goal to conduct original research of the highest standard and to provide leadership in technological innovation for the growth of the country. Established in 1959, IIT Kanpur boasts of a sprawling residential campus that spans across 420 hectares. With path-breaking innovations in both its curriculum and research, the institute is rapidly gaining a legendary reputation across the world.

CHALLENGES

In addition to offering formal undergraduate and post-graduate programs, IIT Kanpur is a leading contributor to national research & development.

The institute faces a formidable challenge in maintaining a state of the art research facility that enables both students and faculty to participate in national science and technology projects.

The projects typically involve development of complex applications in diverse areas like fluid dynamics, molecular modeling, speech recognition, etc. The conventional method employed by the institute to run these resource hungry applications was to use large RISC based, SMP servers running different flavors of UNIX.



Over the years, IIT Kanpur had amassed a significant number of standalone SMP servers that were quickly reaching their threshold limits. In addition to requiring the purchase of expensive compilers and proprietary software components, the flexibility offered in a closed environment was limited. The availability of local support staff was another concern, as the infrastructure had to be up and running on a 24x7x365 basis. In the event of a failure, sourcing support from outside the campus was a painstaking process.

A state of the art High Performance Computing Cluster (HPCC), capable of processing large amounts of parallel or sequential data became the need of the hour.

“For our high performance computing needs, we were looking for an operating environment that could deliver top performance and stability, while maintaining a high degree of sustenance for running complex jobs,” says Dheeraj Sanghi, Professor, CSE Department, IIT Kanpur.

“Some of the complex applications developed on campus require a continuous runtime in excess of a month. Imagine the frustration and loss if a crash were to happen on the Nth day. It would not only put the research project behind schedule, but would also mean a complete wastage of weeks of computation time. We required a popular platform that had campus wide support, which would reduce our dependence on outside vendors to a bare minimum,” adds Sanghi.

SOLUTION

IIT Kanpur’s tryst with Linux on the HPCC front began in late 2002. After setting up its first experimental Beowulf cluster with 16 Pentium III class servers on Red Hat Linux 7, the benefits of using an open source platform became apparent. Initially, the cluster was meant to be only a test bed for running parallel applications.

“Before this, we had used Red Hat Linux only on the non-critical services side to run applications like Web servers, mail servers, proxies, DNS, etc.,” claims Brajesh Pande, Senior Computer Engineer, Computer Centre, IIT Kanpur.

“However, the satisfactory results that we soon witnessed on the HPCC test bed justified the investment in state of the art Linux clusters for our research facility,” he adds.

In mid 2004, the institute setup its first Beowulf cluster for production use on the Red Hat Linux 9 platform, powered by 32 low cost, x86/32-bit servers. With the availability of low cost AMD 64-bit processors in the market and the significant 64-bit capabilities built into the Red Hat Enterprise Linux platform, the institute began to look at purchasing new 64-bit machines to add a second, more powerful cluster.

“Since Red Hat offered a stable, base OS that was compatible and certified across a wide range of 64-bit hardware, it was a natural choice for the second cluster project. Plus, it offered full support for different compiler platforms through GCC,” explains Sanghi.

In 2005, the institute setup its second, 96 node cluster with 98 AMD64 Opteron servers powered by Red Hat Enterprise Linux 3. Two dedicated engineers employed by the institute manage the two clusters 24x7x365.

“Instead of investing in expensive RISC-UNIX 64-bit servers, low cost AMD64 Opteron machines powered by Red Hat Enterprise Linux were a smarter choice for us. As Red Hat is the most popular Linux flavor in the market, everyone on campus was already familiar with its intricacies. In fact, we have been both using and actively following Red Hat’s progress right from the release of Red Hat Linux 5,” adds Sanghi.

BENEFITS

With proprietary UNIX-RISC servers, scaling the HPC infrastructure was an expensive proposition for IIT Kanpur. With Enterprise Linux running on non-RISC hardware, the institute has found a scalable, low cost solution that delivers the same performance, without any compromises.



“With Red Hat’s popularity on campus, both students and faculty now have the flexibility to develop applications using open standards on their Fedora or Red Hat machines. When these applications are ready, they can seamlessly move them to the HPCC environment for computation,” says Sanghi.

“Red Hat has also given us tremendous freedom to tinker with the OS, which wasn’t true in the earlier proprietary environment. In the event of a crash, skills are available throughout the campus to resolve issues immediately. With Red Hat, our dependence on support calls to proprietary vendors has been eliminated completely,” adds Sanghi.

“Another benefit of using a certified platform like Enterprise Linux is that third party HPCC software for resource allocation and node management run without any problems,” adds Pande.

A critical requirement at the time of setting up the two Linux clusters was scalability of usage. The institute currently has 100 users who access the HPC laboratory. In the next six months, the figure is expected to scale to 300.

“While designing the solution, we wanted to make sure that we could extract maximum throughput, to the last drop. Enterprise Linux coupled with our high speed servers have allowed us to support multiple applications and users simultaneously,” explains Sanghi.

With the new 64-bit HPCC environment powered by Red Hat Enterprise Linux, computation time has been reduced by 50%. “Applications that typically used to take 3-4 weeks to run earlier, now take less than two weeks,” he adds.

FUTURE PLANS

t With a 6 TB SAN already in place to support the 64-bit cluster, IIT Kanpur’s storage requirements are in the process of scaling rapidly.

“Red Hat’s Global File System would be a nice filesystem to have and we are actively considering it. Also, new technologies like Xen virtualization, Stateless Linux, SystemTap, etc. are exciting to look forward to in the next Red Hat releases,” adds Sanghi.

CONCLUSION

The freedom and flexibility offered by Red Hat Enterprise Linux has allowed IIT Kanpur to setup two of the most popular Linux HPC clusters in the country, which have contributed significantly to next generation research projects.



RED HAT SALES AND GENERAL INQUIRIES

Toll free numbers:

**Europe, Middle East and
Africa (EMEA)**
00800 7334 2835

Turkey
00800 448 820 640

Israel
1809 449 548

UAE
80004449549

E-Mail:

europa@redhat.com

www.europe.redhat.com

© 2007 Red Hat, Inc. All Rights Reserved. Red Hat, Red Hat Enterprise Linux, the Shadowman logo and JBoss are registered trademarks of Red Hat, Inc. in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds. All other trademarks are the property of their respective owners.

