



PURDUE UNIVERSITY

The Challenge

Purdue, already a computing powerhouse, needed even greater high-performance computing capacity. Specifically, the life sciences department needed to process a vast amount of complex data for work in drug discovery. The university didn't want to just build for this particular project – research leaders wanted the capacity to expand their computing capacity to other departments as their needs grew.

Like many universities, Purdue already had many individual computing clusters based in their own departments – the engineering school had its own, as did medicine, agriculture, and the rest. The issue with these individual clusters is that they're typically underutilized, sitting idle unless there is a specific project that requires high-performance computing.

The Solution

Two community-based, high-utilization clusters that would serve current and future needs. Partnering with PIER Group, Purdue installed Hewlett Packard Enterprise computer nodes with hardware from Mellanox and Cisco, emphasizing powerful high-performance computing that can assist campus researchers in their efforts to provide new breakthroughs.

The Results

At the time, the two supercomputers, Rice and Conte, both made the TOP500 list of supercomputers when they were built. They were built to grow along with Purdue's needs as the size of applications, the amount of data, and the demands on computing power continue to increase.

The two community-based supercomputers were built for approximately 25 percent of the cost of multiple department clusters. In addition, Rice and Conte both maintain a utilization rate of about 90 percent, boosting ROI every day, and the automatic allocation of computing capacity allows Purdue's central IT team to pursue ongoing, university-wide technology strategies.