



## High-Performing Servers on a University Budget? Totally Possible

Universities have unique computing needs and ever-growing data sets. Today's cutting-edge research calls for high-performing computer hardware to support it. However, since researchers often rely on grants and tight academic budget allocations to fund projects, they put particular importance on getting the maximum value for their investment—the highest performance in a server they can afford without any components or capabilities they don't need. And who can blame them?

When you search for a solution to your university's research application hardware needs, know that buying a preconfigured off-the-shelf (OTS) server will most likely offer you too much or too little. Sure, you may get lucky and find something that's just right, but that's rarely the case.

More often, choosing an OTS solution involves compromise. OTS server solutions rarely suit the highly specialized needs of university research. Universities almost always benefit from a custom server solution.

In the long run, a customized server configured exactly to your needs is likely to be more cost-effective and provide better outcomes. For example, if your research involves variable computing performance—or crunching massive amounts of data—a custom server configured to fit the project's exact high-performance computing requirements can offer the best bang for your buck.

### Server Trends

Three major trends are currently driving the server solutions market:

- Big Data: Rapidly increasing size of data sets
- Virtualization: Eliminating legacy systems and/or increasing use of existing servers
- Cloud Computing: Moving data to Internet-accessible private or public cloud storage

The real challenge? A university's skilled and persistent IT staff must keep pace with the very latest innovations in server products and services in addition to their other duties. Plus, they face countless variables pertaining to storage and redundancy and virtualization and software.

The solution? Work collaboratively with an integrator that has broad knowledge—one that has access to the manufacturers that are driving the technology and can really help you right-size your server solution.

## Big Data: The Data Doubling Effect

Since 2012, humankind has doubled the amount of data that previously took 2,000 years to amass. Data production across the globe now doubles about every 15 weeks, and the computer hardware that it takes to process that data is advancing rapidly. Ten years ago, for example, a server that held 250 gigabytes of data was considered a big server. Today it's 250 terabytes—a thousand-fold increase in one decade. That's remarkable growth.

Fortunately for you, the cost of storage continues to decrease. Moore's Law, as stated by Intel founder Gordon Moore, holds that data processors will be twice as fast and half as expensive every 18 months. That constant has held true since the 1960s.

The continual challenge for large organizations? Maintaining enough horsepower to manage, extract, and analyze such huge and growing data sets. Protecting and backing up the data is another ongoing challenge.

## Virtualization: Consolidating and Backing Up Data

Today, data sets are doubling, tripling, and even quadrupling in size. One increasingly difficult challenge for university research teams is to ensure that all data is backed up, preferably in multiple locations.

Due to the relatively low cost per terabyte of storage, one cost-effective option is to increase storage. But you have to implement safeguards. As you add storage, whether by using larger drives or more drives, make sure you have proper redundancy in your system. If and when a drive does fail, you want the reassurance that your system's data is backed up somewhere else—either on a different server or in a RAID environment within the existing server. This becomes a bigger challenge as you add multiple drives to that solution.

Drives are becoming so large that the time needed to rewrite data to a failed drive has gone from hours to days. If you are using early versions of RAID, your system is vulnerable; if another drive were to fail in that time, your entire system's data would be lost.

With the introduction of RAID 6, custom servers now have the ability to maintain double parity. This means your server can continue to serve data despite a two-drive failure. In other words, RAID 6 allows your server to sustain a second failure and continue to serve data, even while the first drive is being rebuilt—a safeguard your university's data deserves.

Another important consideration is the need for scalable redundancy in your server solution. Again, it is unlikely that an OTS solution will be scalable enough for university research purposes. As your data set grows—and you know it will—your hardware system should have the capacity to accommodate that growth and back it up. That's where a flexible custom configuration saves you money in the long run.

## Assessing Your Hardware Needs

How do you choose a server solution that will get the best performance out of an application? In addition to redundancy and scalability considerations, evaluate your organizational and project needs. Consider project requirements in each of the following areas:

- Operating system
- Type of storage
- Amount of storage
- Memory capacity
- Number of CPU cores
- Networking interface
- RAID level
- Processor model
- Form factor
- Number of hard drive bays
- Number of processors
- Need for application optimization
- Space for server racks
- Power, including redundant power sources
- Climate control
- 24/7 monitoring

Ask your researchers questions. What is the initial cost of such a system? What ongoing expenses are associated with maintaining it? What is the total cost of ownership (TCO)?

If the university team does not know how to address these questions, you should strongly consider consulting with a server expert before making any purchases. A good hardware provider will be deeply familiar with the technology, know what's on the market today and what's coming in the future, and be able to recommend the best hardware configuration for your university's needs.

## Integrating Cloud-Based Computing

Consider this scenario: Your researchers have a large data set that they want to crunch. They need 50 servers to do this crunching, but they don't have the budget for it. What can they do?

One partial or full solution might be to leverage either a private cloud within the greater university system or a third-party public cloud. This would give the researchers access to extensive resources without having to spend the money on the actual hardware or maintain the hardware themselves.

With cloud computing, you use the resources you need and then walk away.

## Choosing the Right Hardware to Maximize Your Budget

People often assume that custom is more expensive. More often than not, when university researchers get the exact configuration they need, they end up saving money.

The real challenge for buyers searching for the right server solution is often information overload. You have the option to customize a server configuration in many ways, from switching or adding an individual component to designing a fully customized server. Even within individual components, you have a wide variety of options.

Those options are actually a good thing, though, because they allow custom hardware providers the ability to configure a solution for your specific needs. Not all OTS systems have that same flexibility. In many cases, buyers either spend too much for components they don't need, or worse, they forego a potential component or feature they do need because they found a seemingly lower-priced solution without it.

## Collaborating with a Custom Hardware Provider

When all is said and done, what is the best way to find a server configured to provide the perfect balance of functionality, performance, and scalability? Work collaboratively with a custom server provider that has deep knowledge of the latest technologies available and experience customizing hardware configurations for a wide variety of organizations and application needs.

Look for a provider that is highly collaborative—one that will work closely with your university research team to understand the current and future needs of their program and its budget constraints. A reputable and trusted supplier will steer you, the buyer, away from costly purchasing mistakes and optimize the equipment to better match your organization's needs, reduce energy consumption, and maximize scalability with the latest technology and options.

Your hardware provider partner should ask you many questions, probe deeply into your current and projected computing needs, and then work with you on an ongoing basis to prioritize expenditures and stay within your budget.

You can get a high-performing server solution and control costs with the right partner by your side. The ideal situation is to start from the ground up—pick the chassis, the motherboard, the memory, the CPU, etc. With the guidance of a large-scale server solution expert, you optimize your ability to get as much value for your hardware investment as you possibly can.

Servers Direct is a premier provider of advanced build-to-order server and storage solutions for use in enterprise, data centers, education, government, and the private sector. For more than 10 years, IT professionals have relied on Servers Direct expertise to help develop custom matched solutions with uncompromised quality and support.

20480 East Business Parkway  
City of Industry, CA 91789  
1.800.576.7931  
[www.serversdirect.com](http://www.serversdirect.com)

 **SERVERS DIRECT**<sup>®</sup>  
© 2014 Servers Direct. All rights reserved.