

Aligning Business Organizations

Sun's approach for building energy-efficient datacenters



Summary

- Over time, business cycles of expansion, contraction, mergers, and acquisition cause IT and real estate inefficiency.
- Inefficiency is difficult to address because the CIO doesn't pay the power bill, and the facilities organization doesn't know future power requirements.
- To bridge the gap between facilities and its IT customers, Sun established a new, centralized organization with the authority and funding to purchase IT equipment, build datacenters, and reduce power usage.
- Using its centralized group, Sun reduced IT-related expenses in its San Francisco Bay-Area facilities, substantially decreasing real estate and utility costs.
- The solution was a standard "pod" design with modular power, cooling, cabling, and power metering that helps enable scalable, intelligent growth.



Real companies expand and contract. They invent, develop, test, and deploy new products and services. They merge, acquire, and orchestrate an ever-changing workforce. Over time, these activities can lead to energy and real estate inefficiency that affects the bottom line — no matter how successful the company.

When Sun focused on its own San Francisco Bay-Area portfolio, it managed to compress its datacenter space by more than 60 percent and retire more than 5,000 compute devices — saving on space, power, cooling and taxes. It did all of this in just 12 months, without impacting the business. And to top it off, Sun's efforts garnered nearly \$1 million in utility rebates and efficiency awards.

Opportunities at Sun

Sun is not unlike most companies in that its own business cycles had left it with a world-wide patchwork of real estate and IT assets.

- When Sun assessed its own situation, it found that it was supporting a wide range of datacenter spaces, many of which were poorly managed and maintained.

- The majority of rooms were running old, inefficient hardware that was still in use due to limited budgets and a culture that encouraged groups to retain capital equipment.
- Some of Sun's campuses were only half full, but could not be vacated due to the cost of replicating complex technical infrastructure that supported the company's R&D and business activities.

Despite the fact that Sun delivers billions of dollars worth of products every year, it faced a situation that was similar to those faced by many of its customers: inefficient, expensive, and inflexible datacenter space that limited how quickly the company could respond to new, emerging markets.

Origins of inefficiency

Space and power inefficiency grows over time due to a number of factors that are common to most companies.

Business cycles

Normal cycles of growth and contraction, acquisitions and spin-offs, cause churn in space and IT infrastructure.

Corporate culture

In most companies, groups retain old equipment because it is considered “free” when it reaches zero book value. Since most groups do not pay the real estate or power bills out of their budget, they have a limited understanding of what space and utilities cost the company. With frequent reorganizations and product or service changes, old equipment can become orphaned and continue to consume space and power while satisfying no business need.

The language gap

Core stakeholders tend to have different priorities and goals and speak different languages. The CIO, IT organization, and engineering organization, those who create and run enterprise applications, care about application performance and service levels, while the facilities organization cares about real estate and power expenses.

Lack of visibility

The lack of visibility from one organization into another can result in unnecessary expense. The facilities organization doesn’t always understand or even believe IT requirements. This lack of visibility causes them to react to the power and cooling requirements of each new generation of high-density equipment as

it arrives rather than having a long-term, cost-effective plan for scaling power and cooling infrastructure to meet future requirements. Conversely, the IT organization understands new generations of equipment, but doesn’t usually understand the time and expense it takes to implement large, high-density data-center power and cooling infrastructure. Both of these blind spots can cause reactive behavior that is expensive to the company.

Results of inefficiency

Business expansion and contraction contributes to inefficiencies in real estate, which is Sun’s second-highest cost. Real estate consumes both capital and operating expenses.

The culture that says “capital is important” costs in terms of energy and taxes. Equipment that engineers see as “working” and “paid for” costs more in real estate space, energy, and taxes than new equipment. An older, fully configured Sun Fire™ V880 server can consume \$2,600 in power per year while idle, while a fully configured Sun Fire T2000 server powered by Sun’s UltraSPARC® T1 processor with CoolThreads™ technology can consume as little as \$200.

Organizations saddled with aging, fixed IT infrastructure are the ones most limited by their space, power, and cooling constraints. These constraints place real limitations on a company’s ability to adapt to changing business conditions, including rolling out new applications, expanding into new markets, handling mergers and acquisitions, and consolidating real estate. The solution is to move toward a more dense and more flexible infrastructure that can be changed quickly to embrace new technology and new business opportunities.

Finally, business changes drive new datacenter requirements. When these changes in the technical infrastructure portfolio are not managed, duplication and inefficiencies can become very expensive, very fast. The solution is to have a group responsible for standardizing and managing this portfolio worldwide.

Aligning stakeholders

Sun found that the key to changing patterns causing power and space inefficiency is to better align the organizations and stakeholders in a company. This must be done at a corporate level to ensure a companywide perspective is taken.

The CFO, with ultimate visibility into current costs and potential benefits, must provide the incentive and funding needed to restructure the company’s technical infrastructure and compute resources. The CIO and CTO must understand how their technology choices impact the company’s power and space consumption. The facilities organization needs a road map from which it can plan future requirements, helping it to take a proactive rather than a reactive stance. Every organization with lab or datacenter space is a stakeholder and needs to be part of the process.

Sun created an independent organization with both the authority and the financing needed to create and execute a technical infrastructure strategy. Sun began the process by characterizing its technical infrastructure portfolio in order to have a well-quantified starting point from which it could measure improvements.

Sun learned through the development of this program that it must have realistic goals that are fully supported by the stakeholders. In Sun's culture, dictates are not effective. The key to the program's success was collaboration and buy-in on identifying the problem and creating solutions.

Managing cultural issues

One of the first cultural issues that the new organization grappled with was that its initial goal of complete centralization was in conflict with the real needs of engineers, many of whom need to have frequent, hands-on access to their equipment. The organization addressed these needs by characterizing each existing lab and datacenter space by the kind of access its users required. In some cases, remote access was sufficient. In other cases, immediate proximity to engineering space was required.

With the high cultural value of capital equipment, it would have been difficult to get organizations to give up their outdated systems without incentives. The cost savings of using new servers is so great, however, that Sun included server replacement as part of the central group's budget. Internal Sun groups were offered new equipment in exchange for their outdated servers and storage — just the enticement organizations needed.

To solve finance and tracking concerns, the equipment and depreciation budget was transferred to each organization. This allowed them to track the assets through their normal processes and have the depreciation expense included in their budget for the new fiscal year.

Sun's approach

Sun formed a centralized group called Global Lab & Datacenter Design Services, which had both the authority and the budget to design and execute a technical infrastructure strategy.

Sun's shared datacenter operating model

After completing an inventory of Sun's global technical infrastructure portfolio, a sharing model was designed to reduce Sun's power and real estate footprint.

Sharing technical infrastructure means consolidating and increasing density, offering an opportunity to drive down costs by using more efficient power and cooling solutions. Sun's operating model is based on a stack that includes physical, network, naming, compute, and storage layers. All layers are dependent on each other, so it was crucial to start with the foundation. The first goal was to implement the physical layer and succeed in getting organizations to share space, power, cooling, and cabling, thus reducing overall costs.

Datacenter and lab construction standards

Sun's central organization used the sharing model to create generic design and construction standards. A consistent set of technical infrastructure standards could be repeated worldwide. This would simplify and speed the process of building next-generation environments. The standards had to be generic, yet support scalability and mobility so that Sun could easily reconfigure its spaces as needed.

Hardware replacement program

The hardware replacement program was an integral part of Sun's strategy. The program offered incentives to organizations willing to replace their aging equipment with Sun's current generation of energy-efficient servers. The program, described in detail in the brief "*Creating an Energy-Efficient Datacenter*," was so successful that it succeeded in achieving more than an 80-percent reduction in square footage and a 60-percent reduction in utility power, without down time or application porting.

Deploying pods

Sun's operating model and construction standards pointed to the use of "pods," self-contained groups of racks and/or benches that optimize power, cooling, and cabling efficiencies to facilitate rapid and simple replication throughout the datacenter.

Base power and cooling parameters can be adjusted up and down in 4 kW-per-rack increments.

Choosing a smaller, rather than larger, pod size maximized power, cooling, and cabling designs that could be used in both the small and large datacenters that Sun needs to support. When an organization requires a larger amount of space, the pod design can be replicated as many times as necessary to provide the total aggregate space and infrastructure required. Smaller pods also act as a standard increment by which lab and datacenter space can be scaled up and down as business cycles — including reorganizations, acquisitions, expansions, consolidations, and new product development efforts — dictate.

The pod design defines the datacenter's shell as well as how it is organized internally. Using innovative yet simple techniques for power, cooling, cabling, and energy metering and monitoring, Sun was able to create spaces that can be reconfigured and adapted to different uses without having to design each one from scratch. The innovations internal to the pod design are described in a set of solution briefs that include the following:

- *Powering Sun's Energy-Efficient Datacenters* — Sun standardized on a hot-pluggable overhead busway to distribute power to each rack. There is no need to engage an electrician to change breakers and receptacles or pull new cable when rack configurations change. Now, the modular busway allows power drops to be installed and moved on demand. This makes the space much more flexible and eliminates large amounts of copper and waste.
- *Cooling Sun's Energy-Efficient Datacenters* — The pod design specifies that in-row, overhead, or underfloor cooling solutions can be used independently or in conjunction to neutralize the spot heat loads generated by high-density equipment. There is no single product that solves every problem, rather the pod design allows for multiple solutions that can adapt to the business or environmental challenge.
- *Connecting Sun's Energy-Efficient Datacenters* — Increased density puts higher demands on cabling infrastructure, which can substantially increase cost and have an impact on other aspects of datacenter design when traditional approaches are taken. The pod's simplified cabling architecture cuts cable costs by more than 50 percent, saving copper, improving airflow, and simplifying reconfiguration for new, high-density equipment.
- *Accurately Measuring Datacenter Power Efficiency* — The ability to control and monitor remote environments is essential to running an efficient datacenter. This includes monitoring temperature and humidity as well as being able to power equipment on and off at the plug level. Sun's power distribution architecture uses in-rack power metering that can be monitored locally and remotely. This technology enables datacenters to trend their datacenter power usage at a highly granular level.

The bottom line

Sun's shared model reduced both real-estate and utility costs in 12 months. Rather than constraining the company, the new, modular, efficient datacenter design improved Sun's agility and contributed to its operating margin last year.

Learn More

For more information on how to align your business organizations for better space and power efficiency, or to obtain the additional solution briefs mentioned in this document, please contact your Sun representative or visit sun.com/eco.

New datacenter designs can accommodate today's requirements and scale into the future while significantly reducing costs and deployment times, with only an incremental investment. Sun customers can increase their own datacenter density and simultaneously cut costs by learning from Sun's experience and leveraging Sun's strategy and technology in their own datacenters.