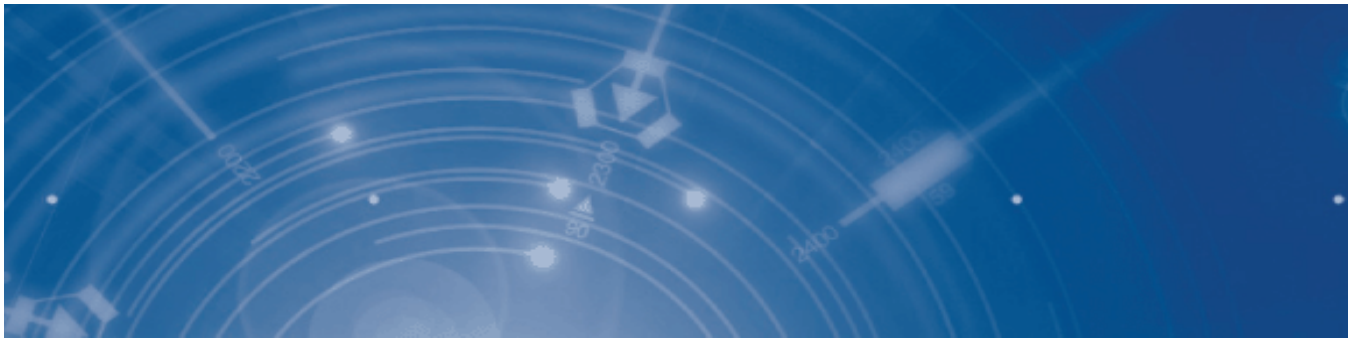

U.S. Data Centers: The Calm Before the S

25 September 2007

Rakesh Kumar

Gartner RAS Core Research Note G00151687

U.S. enterprise data centers are facing considerable disruption during the next three or more years. Er demands and new technology implementations will cause changes and increase costs.

Overview

This research outlines why energy and floor space constraints will cause considerable disruption to U.S. enterprise data centers during the next three years. Users must plan early and change rapidly to avoid costly delays.

Strategic Planning Assun

By 2011, more than 70% of U data centers will face tangible related to floor space, energy and/or costs.

During the next five years, th center hosting providers for c services will grow rapidly in th

During the next five years, m enterprise data centers will sp energy (power and cooling) a hardware infrastructure.

Key Findings

- The U.S. data center collocation market will grow rapidly during the next few years and should be considered as an alternative to internal data centers.
- Upward-spiraling infrastructure demands and increasing energy costs mean that the energy proportion of IT costs could double by 2012.
- Advanced monitoring and modeling tools will change the nature of data centers from static entities to dynamic "living organisms" during the next five or more years.

Recommendations

- Users must start monitoring the energy consumption of their data centers, and modeling the floor capacity available for future growth. This information should be used quickly to make strategic decisions about the nature and importance of the data centers.
- Projected major changes to the capital or operational part of an IT budget must be quantified quickly and appropriate provisions established, such as securing funding to build a new data center.
- Technologies such as virtualization and dynamic workload

management must be implemented quickly to improve the use of assets and defer the procurement of new hardware as a mechanism to optimize floor space and energy costs.
- Where data center refurbishment or new builds are required, users must ensure that they and their engineering partners use the evolving modeling and monitoring tools.
- If hosted data center space is required, then users should move quickly to avoid rapidly rising costs, and should secure mid- to long-term contracts (three to six years).

Analysis

CIOs of large U.S. organizations must prepare for a period of rapid changes in their data centers. The disruption will be accompanied by a significant increase in capital and operational expenditures. Failure to respond quickly and appropriately to the changing market conditions and technologies will result in needlessly high energy bills, expensive service contracts and delays in implementing new technologies.

During the next two or more years, three main issues will come together and cause the disruption:

- Legacy data centers won't have sufficient power and cooling requirements for the next generation of high-density server and storage equipment.
- The volume growth of IT infrastructure will exceed the available data center floor space for most organizations.
- The need to manage upward-spiraling energy costs through optimization tools and modeling techniques.

Gartner estimates that more than 70% of the world's Global 1000 organizations will have to modify their data center facilities significantly during the next five years. The U.S. has the biggest concentration of large (greater than 50,000 square foot) data centers, the majority of which were built more than seven years ago. Their design envelope is insufficient to handle servers' current and future energy needs.

These legacy data centers typically were built to a design specification of about 35 to 70 watts per square foot. Current design needs can vary from between 150 to 200 watts per square foot, and by 2011, this could rise to more than 300 watts per square foot. These figures for energy per square foot represent just the energy needed to power the IT equipment; they don't include the energy needed by air-conditioning systems to remove the heat generated by this equipment. Depending on the tier level and future equipment density plans in the data center, these cooling needs can increase the overall power requirements by an additional 80% to 120%. For more details about these types of calculations, see "The Impact of High-Density Server Growth on Data Center Power and Cooling Metrics."

The implication is that most current data centers will be unable to host the next generation of high-density equipment, so CIOs will have to refurbish their established sites, build new ones or look for alternatives (such as using a hosting provider).

Strategic Planning Assumption: By 2011, more than 70% of U.S.

enterprise data centers will face tangible disruptions related to floor space, energy consumption and/or costs.

In August 2007, the U.S. Environmental Protection Agency (EPA) reported that U.S. data centers consumed 61 billion kilowatt hours (kWh) in 2006 (1.5% of all U.S. energy consumption) and cost $4.4 billion to operate. The report also estimated that this could increase to 100 billion kWh and $7.4 billion by 2011. The EPA's assumptions are rather conservative, however; we believe that a projected 15% CAGR (compound annual growth rate) of the high-volume, high-density servers could add another 10% to 15% to these figures.

The EPA's report also stated that the large enterprise data centers are consuming most of the electricity (38% in 2006). Again, we believe that as site consolidation accelerates in the U.S., the proportion of these large data centers will grow, placing additional pressure on energy consumption and supply.

During the next three or more years, one of the most-important changes to the U.S. data center landscape will be midsize and large users' increasing propensity to use data center collocation services. Traditionally, the U.S. market has been reluctant to embrace the notion of leasing space and running IT services from that location. However, during the past nine months, Gartner has detected a shift in attitude that will accelerate during the next few years. We believe that the fiscal equation of an expensive capital cost for a new, owned data center — as opposed to the much more inexpensive, ongoing operational costs of leased space — will encourage companies to explore the use of collocation space.

Moreover, the perceived issues of lack of control and weaker security with collocation players generally haven't manifested themselves. The net result is that leasing space from a well-designed, modern data center provider can yield financial and operational benefits. Currently, the market is supply-constrained, resulting in a rapid increase in costs. Gartner expects this market to become very active during the next few years, so users should move quickly to secure good prices.

Strategic Planning Assumption: During the next five years, the use of data center hosting providers for core data center services will grow rapidly in the U.S.

Return to Top

Managing the Disruption

Users must develop a set of prioritized actions as a mechanism to manage the immediate issues and the fallout that will result from market changes during the next few years. The following four suggestions encapsulate the key actions:

Quantifying and Communicating the Issues: Users must establish models for gaining a granular view of their energy costs, floor space and infrastructure topologies. As a result, users will understand where the

power is going (servers, cooling, and uninterruptible power supplies or power distribution) and how much it costs. This must be performed as a primary task in data centers, and then it should be rolled out to the appropriate office locations.

It's also important to develop projection models based on the growth of IT equipment and changes to the data center or office layout. This will enable the consideration of different scenarios of equipment deployment and internal financial management (such as chargeback). For example, a model using computational fluid dynamic (CFD) analysis will help determine how many more servers can go into a data center, where they should be located and whether any changes must be made to an air conditioning system.

Infrastructure Optimization and Rationalization: Users must systematically consolidate their current machines and workloads to use up all available spare capacity. In this way, users will maximize the returns on their investments in server technology. Moreover, they'll delay the need to purchase more (of the newer, high-density) hardware and, therefore, push back the manifestation of energy problems by months (if not years).

It's also important for new servers to be procured on the basis of running them about 60% to 70% of the time. This will ensure that users can better manage the power and cooling issues at the start of new hardware deployments.

Strategic Planning Assumption: During the next five years, most U.S. enterprise data centers will spend as much on energy (power and cooling) as they will on hardware infrastructure.

To achieve these levels of use, choose appropriate software tools. These include virtualization software (from companies such as Microsoft and VMware) and better workload tools. However, these tools won't be of much use unless organizations change their operational processes to benefit from the software. For example, having multiple, virtual machines in a single server should ensure that production partitions can be run next to test partitions in the same box. This requires a change in the application production/acceptance and testing processes.

Large mainframe groups also may benefit from using this platform as a very energy-efficient platform for traditional workloads, and also as an alternative for new Linux applications. In this sense, users should evaluate the energy costs of using multiple, x86 high-density servers for new Linux applications, or using the specialist Integrated Facility for Linux (IFL) mainframe engine for the same applications.

Future Data Center Options: IT organizations must address the long-term, strategic nature of their data centers' relevance. This will become relevant and poignant when the issue of new data center builds or major refurbishments presents itself. The capital costs of a new build are very high and the design specification is extremely variable.

For example, the energy per square foot can range from 100 watts to 300 watts, with extreme environments using 450 watts or more. This creates a huge variation regarding the type of facilities equipment to be used, as well

as an extremely uncertain operational cost model. As a result, Gartner predicts a rapid uptake of collocation space as a short-term measure to evaluate other options, or as a longer-term IT service delivery strategy. However, this rapid increase also will drive up the cost of collocated space quickly, thus reinvigorating the build vs. collocate decision for many companies. In such cases, it's important for IT organizations to balance the immediate (energy or floor space) problems with the more-strategic issues of risk management, the cultural posture of internal/external data centers, and costs.

Using Evolving Tools and Techniques: When building a new data center or refurbishing an older one, users must embrace the latest modeling and monitoring tools. Users also must select an engineering company that has the skill set to use these emerging technologies and to adapt company processes to take advantage of these tools as they become available. In addition, during the early design phases, consider a tiered floor-space model, creating "cool zones" for specific equipment types and designing power and cooling solutions for those locations, rather than designing for the extreme case across the entire data center.

For example, it will become very important to integrate elements of the building management software with the core system management tools. It also will become mainstream to use modeling tools (such as CFDs) in the initial design and subsequent maintenance of data centers.

Return to Top