Level 3 COMMUNICATIONS

Connecting and Protecting the Networked World℠

WHITE PAPER

CONTINUITY, CAPITAL AND SECURITY: THE IMPACT OF HIGH DENSITY DATA CENTERS

Executive Summary

Big data is changing the world, particularly for enterprises, service providers and government entities that need to deploy bigger, faster servers to capture and process an expanding avalanche of data from multiple sources. This rapid growth can put a huge strain on existing facilities as they struggle to increase processing power and storage capacity within data centers that were built before the introduction of today's high-density, high-capacity systems. To accommodate these new technologies, many organizations are turning to off-site data centers, which provide infrastructure specifically designed to handle increased server densities and to provide higher levels of reliability and redundancy.

There are many reasons why organizations use off-site data centers, but the most compelling reason is often to support a business continuity and disaster recovery (BC/DR) capability. Other reasons include data center centralization, geographic expansion, peak load handling and secure data storage. Flexibility is also key because long-term planning can be challenging when software applications are constantly adapted to meet new requirements and hardware architectures continue to evolve.

One reason for using offsite facilities that is often overlooked is the financial benefit of leasing space instead of building and owning a private facility. The high costs, long planning cycles, consumption of scarce financial capital and continuing operation expenses of a private center can easily exceed the cost of leasing space on a recurring, flexible basis.

This paper begins with a discussion of offsite data center applications and benefits, followed by a discussion of some of the financial parameters surrounding the lease vs. build decision. The implementation process is then addressed, along with some considerations for choosing a service provider. Two sidebars are included: a data center planning checklist and a discussion of the benefits of using a remote facility as the primary data center for an enterprise.

Introduction

During the past two decades, data storage and analysis requirements for virtually every organization have increased dramatically, as more services are migrated online and more data is gathered and stored. This trend towards big data shows no sign of stopping, as public Internet services, private intranets and extranets, video, mobile and business-critical enterprise applications continue to increase the demand for processor power, data storage, and rapid response times to user requests. Accompanying this trend is a move toward increased system density, as servers get smaller, disk drive capacities grow, and communication links become quicker and shorter.

To improve system performance, many organizations are installing servers with multiple, closely-packed server blades, copious amounts of memory, and high-bandwidth network connections. By increasing system density, the response times within systems can be faster, and placing systems in close proximity allows the communication latencies between them to be reduced, giving a further boost to processing speed. However, as system density increases, so does the demand for power and cooling in ever smaller spaces.

Modern organizations also increasingly require highly reliable, around-the-clock access to critical enterprise data. As more employees take advantage of mobile devices and flexible working hours, they have grown accustomed to near-instantaneous access to all of their files and programs from anywhere around the globe. Plus, with the growth of intranet-based services, users may no longer have local disk data storage, so when Cloud servers or data networks are down, more employees find it difficult to do their jobs. These forces combine to have a significant impact on new data center designs, with an increased emphasis on reliability and redundancy for all key system elements and infrastructure.

High-density data centers (such as Level 3's new centers in San Francisco and Omaha) have been specifically designed to accommodate high-density,

high-reliability systems with redundant power and cooling systems. A typical data center built a decade or more ago used an average power density of ten watts per square foot, whereas new, high-density centers are built to support twenty to thirty times this power density (200-300 W/ft2), with even higher densities available for specific applications. In order to achieve high system reliability at this greater level of overall power density, all of the key support systems, including power conditioning and generation, cooling, data communications and wide-area connectivity have to be engineered appropriately. This requires a thorough reworking of all the design assumptions for the facility, starting literally from the ground up.

Data Center Applications

Leased facilities can host server and storage systems to support a wide range of organizational functions. In many cases, two or more of the following strategies will be of benefit to entities seeking cost reduction, flexibility, or protection, whereas other organizations may focus on a single major application for their first off-site deployment. Regardless of size, the flexibility offered by off-site, leased data centers allows for future changes in the roles that these facilities can play.

Business continuity and disaster recovery

Modern data centers require a complex set of inputs to operate, including power, environmental controls and communications links. Having two or more sites for critical data infrastructure (i.e., primary and backup) can help ensure that business operations will not be significantly impacted by natural disasters or major outages. Ideally, the backup facility should be located far enough away from the primary facility that the chance of a single outage or disaster affecting both locations is small or nonexistent.

Data center consolidation

Combining and centralizing data centers can often result in substantial cost savings. A single, large data center is easier to manage, maintain, and operate than several smaller ones, particularly when the costs of synchronizing data sets across multiple storage clusters are taken into account. Many organizations find that is it more efficient and cost-effective to locate centralized servers within a high-reliability data center, instead of designing and building their own facilities.

These are powerful arguments for centralizing organizational data centers. But, there are other reasons that some organizations chose to have multiple data centers. One obvious one is for greater resilience towards natural or man-made disasters. Another is to provide better service to users (including web-based customers) who are spread over large geographic area.

Global organizations may opt to deploy regionally consolidated data centers, if circumstances warrant. For example, in some jurisdictions, the regulations surrounding the storage and transfer of personally identifiable information are quite strict. These rules can have an impact on where the servers containing such personal information are located, forcing equipment deployments in certain regional or country-specific locations.

Data center extension vs. build outs

Transitioning from an existing hardware configuration to a new generation of technology can be challenging, particularly when both old and new systems have to coexist during new application testing and data migration. Since technology tends to be refreshed more often than facilities (every three to five years for computing resources versus every 10 to 20 years for facilities), organizations with fixed facilities can find it difficult to introduce new technologies.

Geographic data center expansion

When organizations need to serve users that are spread across a wide geographic area, locating web servers in data centers that are close to users often makes sense. Proximity reduces the network transmission delays between client and server, allowing faster page load times and higher-performance file and video downloads. Even simple web pages with multiple elements will load more quickly when delivered by a

server that is within a few hundred kilometers of the user, instead of being housed on a different continent.

Peak load handling

Designing a data center to handle peak traffic loads can be punishingly expensive, particularly if traffic peaks are large and infrequent. The problem is, if a system is designed to handle a peak load that is twice its average load, then approximately half of the system's resources will be sitting idle most of the time. A centralized data center can accommodate peak loads that correspond to activity levels that ebb and flow during the day for organizations with facilities in multiple time zones around the globe.

Secure data storage

For many organizations, data security is a must. But beyond being a good business practice, minimum levels of data security are mandated by government regulations in a variety of industries, including healthcare and finance. Creating a facility that can not only meet all the requirements but also comply with the rigorous audit requirements can be a challenge. Organizations can simplify and streamline their compliance procedures by leasing space within a facility that has successfully passed audits related to, say, financial transactions.

Using Remote Site for Production, On-Premises for BC/DR

In a traditional disaster recovery scenario, an off-site system is used to store critical data and host processing power that can be used in an emergency in case the primary production system becomes unavailable due to natural disaster, major component failure or human error. However, some organizations are turning this arrangement on its head, and locating their primary data centers within a remote-hosted facility while keeping their disaster recovery facilities on-premises. Why is this? Well, there are several reasons:

- Off-site servers and storage can be housed in high-reliability data centers, with levels of security and redundancy that can be cost-prohibitive in a private data center. Hosting primary production at one of these centers could increase overall system availability and reduce the probability of a natural or man-made disaster taking the production servers off-line.
- In the (unlikely) event that a primary facility experiences a catastrophic failure, actions will need to be taken to bring the backup capabilities into service. This may be easier to accomplish when the backup facility is situated for easy access by IT staffers.
- When on-site servers are not being used for disaster recovery, they can be used for prototyping and testing new services and software. This arrangement allows application development staff to have easier, quicker access to the systems being used for development and troubleshooting. Plus, this system can be used as a test environment for evaluating new hardware and software vendors.

Benefits of Outsourcing and Leasing Data Center Space

Using leased space in a data center facility can provide a variety of benefits to organizations and their clients or customers. These centers can leverage the latest advances in power conditioning/backup, cooling and communications technology to deliver increased efficiency and lower costs. In addition, costs for essential staff such as security and building maintenance can be spread over multiple clients, supporting high-quality service without the overhead costs that each organization would have to face with individual data centers.

To achieve corporate environmental objectives, modern data centers can be designed to minimize the use of environmentally expensive inputs, such as electricity for cooling. With the new technologies that have recently become more widely available, such as LED lighting and ambient air cooling, these data centers offer not only significant savings in power consumption, but they also feature a lower environmental footprint and reduced greenhouse gas emissions per kilowatt of power consumption. These innovations can help reduce electric power bills, a major component for any processing installation.

Even though modern data centers are designed to run autonomously, without direct human intervention, they still require ongoing monitoring and maintenance to ensure optimal performance. The costs of hiring staff

to perform building and equipment maintenance and to manage third-party suppliers can significantly add to operational expenses. These functions also consume funds that could be used to hire staff to perform tasks that are more aligned with strategic organizational goals. In a shared data center, maintenance costs can be spread over a multiple clients, helping to keep individual expenses lower.

High-reliability data centers also usually include advanced security features. It is typical to have 24 x 7 on-site security, particularly to supervise access by client and third-party personnel performing hardware installations, repairs, and upgrades. Modern data centers also offer access control that is localized to specific equipment racks as well as comprehensive video surveillance.

Another benefit of leasing space within a shared facility addresses the challenges associated with replacing an outdated set of technology with a new generation, as is done every three to five years within many organizations. For organizations that own a fixed data center, finding space to house the new technology while it is being assembled and tested can be difficult. For customers of leased data centers, it is relatively easy to lease new space where the new technology system can be installed, while leaving the legacy system intact until cutover. Then, the old system can be removed and the space that it occupied can be relinquished to the facility provider to use for other clients.

Financial Considerations

Performing a thorough financial analysis is also important when deciding whether to build a new data center or lease space within an existing one. The cost of providing space to house data processing equipment involves a tradeoff between operational expenses (OPEX) and capital expenditures (CAPEX). In many companies, CAPEX expenditures (especially at the magnitudes associated with large data centers) require significant planning, review and high-level approvals. Large, fixed CAPEX expenditures, which must be amortized over a decade or more, can limit the ability of an organization to adapt to changing business circumstances or changes in technology.

The biggest capital investment for a high-performance data processing center is actually building the data center that will house servers, storage, and datacom equipment. These specialized buildings — with large amounts of power, air conditioning and redundancy — can consume a large amount of capital for any organization.

One example of the costs of a high-reliability data center is available for public inspection, provided by the Uptime Institute in their Total Cost of Ownership calculator.[1] Presented in a white paper, this model is based on a facility with 20,000 square feet of active device area (where servers, disk drives, tape drives and datacom equipment are housed) along with another 20,000 square feet allocated to support equipment such as cooling, power conditioning, uninterruptable power supply (UPS) and other related facilities. The capital cost of this facility was calculated to be $72 million with an estimated life of 15 years, giving an annualized capital cost of $7.9 million assuming a real discount rate of 7 percent. Annual operating expenses for this model facility amount to $2.9 million excluding electric power and communications circuits. These expenses, totaling nearly $11 million, are a fixed annual cost throughout the 15-year life of the facility, meaning that they won't change whether the data center is full of equipment or only partially used.

Leased space within a shared high-reliability data center offers organizations a flexible alternative to this fixed expense by providing rack space, power, cooling, security and communications that can scale up or down as the needs of the organization change. One major benefit of this flexibility is that it allows organizations to introduce new technologies, such as lower-power servers or higher-density storage, without worrying about violating the assumptions that were made in designing a fixed facility that may be only partway through its expected lifecycle. Another advantage is the availability, on relatively short notice, of additional

space that could be provided if the organization experiences significant growth or if additional room is needed to deploy a new generation of technology.

Data Center Transition Planning Checklist

Creating any new server/storage installation is a complex endeavor, whether within leased or owned facilities. By developing a set of project plans in advance to cover each of the major tasks, organizations can help minimize the time and expense of the transition. The following list provides a high-level view of some of the key topics that need to be considered in any data center plan:

- New hardware and software acquisition
- Installation logistics, including equipment staging and staffing needs
- Transition plan for existing hardware and software
- Temporary and permanent connectivity requirements
- Security architecture and practices for physical facilities and personnel
- Internet and intranet security and firewalls, including update procedures
- Local and remote technical support and professional services plan
- Data backup and restoration procedures, data mirroring
- Service level agreements with carriers and other providers
- Communications redundancy and diversity
- System testing and turn-up procedures
- Ongoing hardware and software maintenance
- System and network monitoring
- Remote access to device console ports and diagnostics
- Disaster recovery and business continuity plan

Typical Implementation Process

Moving into an off-site, high-reliability data center requires several activities on the part of both the center's provider and the organization that will be using the center. With the correct plan, this transition can be made in a cost-effective and efficient manner.

A good first step is developing a budget and a schedule for the transition. Fortunately, the timeframe for moving into leased space at an existing data center can be much more condensed than the time that would be required to plan and build a new data center, particularly considering the building permits and other approvals that need to be obtained. A typical move into a data center can be accomplished in six to nine months, although quicker deployments are certainly feasible. This is in contrast with the multi-year process typically required to design, obtain permits, construct, outfit and test a new standalone data center.

Since the actual computing and storage equipment is furnished by the organization leasing space within the data center, the hardware design effort is virtually identical to the process that would be used to install an in-house system. Suitable equipment from reputable vendors must be selected and purchased from suppliers. The equipment must be shipped to the data center, along with adequate quantities of data and power cables and connectors. One extra component often installed in off-site data centers is a mechanism that allows remote access to the console port on servers and routers that also has the ability to reboot devices remotely.

During the installation phase, staff must be provided to "rack and stack" the equipment, connect the appropriate cabling between the devices, and install the required software. This includes making connections to local area network switches and routers, plus configuring those devices. In addition, wide area network connections (often via private circuits) must be installed to allow users within the organization to access the systems, and Internet connectivity is also frequently required. Some high-reliability data centers (such as those from Level 3) have built-in connectivity to major telecommunications providers, including on-net access to the Level 3 backbone. This can provide a

significant cost savings, since local loops do not need to be purchased to make connections from the data center to a long-haul carrier's PoP (point of presence). Once the equipment and software are installed and the necessary internal and external networks connected, the new system needs to be tested. This testing does not necessarily need to be done on-site at the remote facility; provided that the network connections are operating properly, the tests can be conducted under remote control from the organization.

Choosing a Data Center Provider

After the decision has been made to investigate using an off-site high-reliability data center, a provider needs to be selected. There are a wide range of choices on the market, but most organizations will benefit from having a provider that offers at least the following capabilities:

- Full redundancy for all critical systems, including power, cooling, and communications
- High power densities — at least 200 watts of usable power per square foot
- Adequate standby power capacity, with a plan for handling long-duration power outages
- Robust security, including on-site personnel, building and cage-level access control
- Location that has a low incidence of natural disasters
- In-building connections to a variety of communication providers
- Multiple communication access paths that are homed to diverse carrier networks
- Favorable references from existing customers
- Certification from Uptime Institute and other governing bodies

Organizations that deal with financial transactions frequently require two industry certifications: the PCI-DSS (Payment Card Industry Data Security Standard) and the SSAE 16 (Statement on Standards for Attestation Engagements Number 16). Both of these certifications rely upon periodic audits to ensure that data processing and storage facilities are in compliance with security and resource availability requirements, as well as certify that mechanisms are in place to protect the confidentiality and privacy of customer data. Providers such as Level 3 have designed facilities with an eye towards meeting these certifications and also have direct experience with these audit processes, which can greatly simplify the process for client organizations.

Conclusion

Today's off-site data centers can provide a better level of redundancy and can support higher server densities than facilities that were designed as recently as a few years ago. These advances, coupled with the lower power consumption and greater flexibility inherent in facility leasing, combine to create a powerful incentive for organizations of any size to seriously consider hosting their data centers in facilities owned and managed by a leading communications carrier.

The newest facilities from Level 3 in Omaha and San Francisco offer extremely high levels of redundancy and security. The Scott Data Center in Omaha, Nebraska, has received Tier III certification from the Uptime Institute as being Concurrently Maintainable. These new facilities join a network of hundreds of other Level 3 data center hosting facilities around the world that can provide high-quality facilities to host data centers for all types of organizations.

For more information on Level 3 data center services, please visit our website, www.level3.com. See for yourself how Level 3 can help enhance system performance, security and flexibility with an attractive package of leased facilities and services.

Connecting and Protecting
the Networked World℠

ABOUT LEVEL 3

We build, operate and take end-to-end responsibility for the network solutions that connect you to the world. We put customers first and take ownership of reliability and security across our broad portfolio.

1.877.2LEVEL3
INFO@LEVEL3.COM
LEVEL3.COM

DATA NETWORKS | SECURITY | CONTENT DISTRIBUTION | DATA CENTERS | APPLICATION PERFORMANCE | VOICE | UCC

13DCS03