

WHITE PAPER

Data Centre: “Build vs. Buy”

By Laura Stevenson,
Director, Product Marketing, BLACKIRON Data



BLACKIRONDATA

I EXECUTIVE SUMMARY

When considering different options for enhancing data centre performance in response to increased technology or business requirements, organizations must ensure that TCO comparisons take all relevant cost factors into consideration, including the obvious factors as well as the not-so-obvious.

Traditional data centres that have been in operation and that have remained unchanged for the past several decades have pretty much become functionally obsolete, either because they are unable to handle the higher power densities of newer IT equipment or simply because they are not energy efficient and could not be upgraded cost-effectively once they went online.

Faced with the prospect of having to build a new data centre, or refurbish or expand existing facilities, senior corporate management and IT decision-makers are now able to consider numerous options, including continuing to build and operate their own data centre capabilities (i.e., “Build”) or outsourcing their data centre requirements to a third-party colocation services provider (i.e., “Buy”).

Making an informed “Build vs. Buy” decision requires a comprehensive and in-depth analysis of the costs, risks and opportunities associated with each option. This paper focuses on the cost aspect of the decision and offers a data centre Total Cost of Ownership (TCO) model as a framework for examining the full range of Build costs, including capital equipment costs – upfront as well as periodic over time – and ongoing operating expenses.

The paper also presents some insights into aspects of data centre cost that are often overlooked or underestimated that might result in an incomplete picture of the Build costs and thus an inaccurate comparison against the cost of outsourcing equivalent data centre services.

For discussion purposes, consider using a 900 square foot, 20-cabinet data centre built to Uptime Institute’s Tier II standards. Traditional analysis suggests it would take roughly five years for the return on an investment in the Build option to break even with the cost of Buying equivalent data centre services from a top-tier data centre services provider. However, we believe that traditional analysis does not take into account a number of critical TCO items that are often overlooked, and 11 years is more realistic. It could even be longer, depending on various data centre design decisions and trade-offs.

The very real question this raises is “Who can tell what the future holds for data centres, and computing in general, in five years, 11 years or even more!”

We believe the Buy option to be the answer, as it provides the best protection against increasing data centre complexity, cost and risk, and eliminates the need to worry about uptime, technology obsolescence and future requirements. It also preserves valuable capital that can be invested in core business initiatives.

COMPREHENSIVE TCO ANALYSIS FUNDAMENTAL TO DATA CENTRE “BUILD VS. BUY” DECISION

Legacy Data Centres No Longer Adequate

A January, 2011 Gartner survey of more than 2,000 CIOs worldwide yielded the accompanying list of technology priorities.¹

What’s interesting to note about this is where most of the priority items live – they live in the data centre. Perhaps even more interesting to note, however, is that the majority of today’s data centres were built 15 years ago or more, when power requirements and processing densities were much lower. Now, consolidation, virtualization and high-density hardware installations, for example, are stretching current data centres to the limit.

RANK	TECHNOLOGY PRIORITIES
1	Cloud Computing
2	Virtualization
3	Mobile Technologies
4	IT Management
5	Business Intelligence
6	Networking, Voice, and Data Communications
7	Enterprise Applications
8	Collaboration Technologies
9	Infrastructure
10	Web 2.0

SOURCE GARTNER

In their current state, data centres simply cannot accommodate developments such as increased processing power, virtualization, new server technology and requirements for low-latency data connectivity. For example, according to AFCOM, the leading association of data centre management professionals, power failures and limits on power availability will interrupt data centre operations at more than 90% of all companies over the next five years.

New Dimensions Drive-Up Data Centre Complexity, Cost & Risk

As a result of rapidly changing technology, organizations can no longer expect to build a data centre with a 20-year life expectancy, at least not without planning on a major upgrade about every five years. In addition, a number of new dimensions are driving complexity and cost through the roof and increasing the risk associated with building and operating today’s data centres. These dimensions include:

- **Data Centre Consolidation:** straining power, space, and HVAC, as well as increasing risks associated with system uptime and application availability;
- **Technology Convergence:** servers, SAN and networking now need to be architected as a single, converged infrastructure platform;
- **Virtualization:** addressing many data centre problems, but adding to infrastructure complexity and manageability challenges — a single server could be running 30 virtual machines/application processes and driving 60+ network cards on multiple VLANs across secure prioritized segments;






- **Cloud Computing:** offering cost and agility benefits, while increasing complexity of security, identity and access management by extending applications and data beyond the corporate firewall;
- **Workforce Mobility:** driving new, and more complex requirements around remote application access, performance, authentication and identity management;
- **Branch Office Growth:** demanding an effective branch-computing model compatible with re-centralization;
- **24x7 Business:** requiring zero infrastructure downtime and ‘extreme availability’ for core business applications, thus necessitating improved business continuity and disaster recovery;
- **Web-Based Services:** providing service-delivery alternatives (i.e., SaaS), while creating data privacy, security and governance challenges;
- **Data Growth/Management:** affecting data centre design and infrastructure choices around storage architecture and information lifecycle management (ILM);
- **Process Automation:** reducing manual systems management by introducing dashboards for infrastructure performance metrics and automation such as ITIL-based process automation templates and the Configuration Management Database (CMDB); and
- **GreenIT:** quest for environmental sustainability driving data centre strategy, including technology upgrade, better power management, improved asset utilization and ‘e-waste’ management.

Organizations facing these challenges often find themselves in a position of needing to decide whether to build and operate their own data centre or buy infrastructure services from a third-party data centre services provider. **This dilemma — the so-called “Build vs. Buy” decision — is not restricted to just new ‘green field’ builds, however, but also applies to typical scenarios related to existing data centres, such as:**

- An existing data centre that is reaching the end of its useful life, perhaps because the IT equipment (i.e., servers, storage, networking) or the supporting facilities equipment (i.e., power, HVAC, smoke/fire, security) are too old and cost too much to maintain/upgrade;
- An existing data centre that no longer has the capacity headroom — space, power, cooling, processing, bandwidth, staff — to support the additional demands that accompany continued, and sometimes unexpected, business growth;
- An existing data centre that does not have the capability to support the technology, interconnectivity and power density characteristics of today’s newer server, storage and networking technologies;
- An organization whose IT systems are becoming even more critical and thus sees value in increasing the investment in its current data centre to ensure greater system reliability, availability and uptime; and
- An organization that wishes to make a fundamental shift in its data centre strategy or computing and service delivery paradigm, such as migrating to cloud computing.



Separating Myth from Reality

For those already familiar with the many potential colocation benefits previously listed, it is not unreasonable to ask, “If data centre outsourcing is so great, why hasn’t everyone done it?”

The answer given by 81% of the respondents to the Info-Tech survey who chose not to colocate was their concerns regarding privacy and security. Interestingly enough, however, 84% of surveyed organizations that successfully colocated their data centres stated that the facility security standards of their colocation partner actually exceed the current security capabilities of their own organizations.²

This apparent paradox highlights the fact that there are many factors at play in the “Buy vs. Build” decision, not the least of which are emotional issues such as fear of the unknown, fear of job loss and concerns over loss of control. These concerns are often fueled by a lack of education as to what colocation service providers actually provide.

The “lack of security” response may simply be a smokescreen for these emotional issues, especially when you consider the levels of physical security in play in a modern data centre hosting environment. This often

includes hardened facilities made of bullet-resistant materials, multi-factor access authentication with biometrics, access turnstiles and mantraps, full surveillance coverage with 24x7 monitoring, and more. While there may in fact be some data centres that aren’t as committed to security, the top-tier data centre service providers are continually going through various certifications (i.e., PCI Compliance, SSAE 16 Compliance), subjecting their facilities, their security technology and procedures to close scrutiny by third-party auditors.

Even in the case of industries that have historically been colocation averse, such as Healthcare, Financial Services and Government that require additional layers of security and privacy to protect data, the Info-Tech survey found that data centre colocation can, in fact, be successful regardless of the industry.²

This paper will help organizations separate the myth from the reality of data centre colocation, at least in terms of the most critical dimension of the ‘Build vs. Buy’ decision — Total Cost of Ownership.



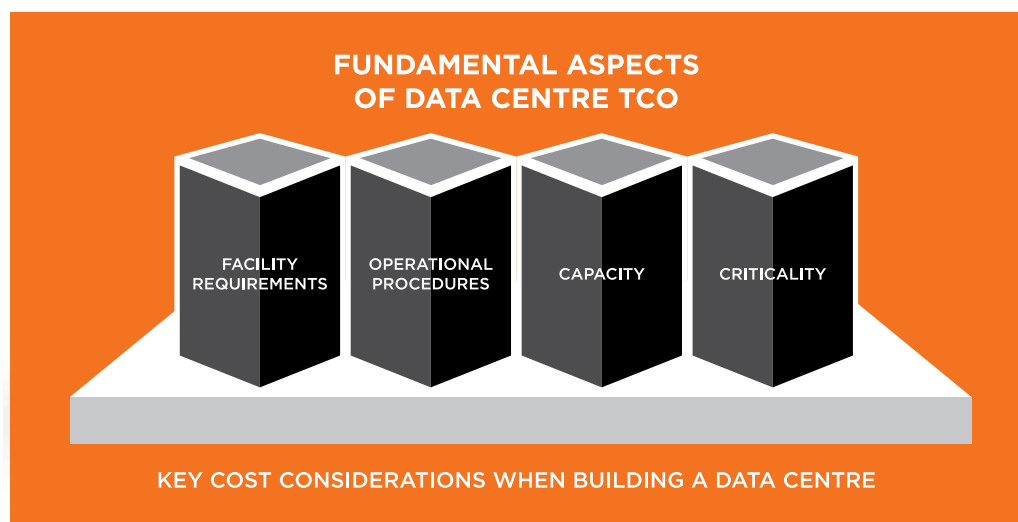
Build vs. Buy — Understanding the True Costs

Making an informed decision about whether to build or buy data centre services requires a thorough understanding of all the costs involved on both sides of the issue so that a reasonable comparison can be made.

The “Buy” scenario is relatively easy — the data centre services provider will typically help an organization identify its current and future needs, assist with the architecture and design of an effective colocation solution to those needs and then provide a detailed cost quotation for all the services that will be delivered and the Service Level Agreements (SLAs) governing those services.

The “Build” part, on the other hand, requires analysis that is more complex and far-reaching than most would expect or have experienced. As a result, true data centre TCO is almost always grossly underestimated, leaving decision-makers without complete and accurate information with which to make an informed decision.

The four-pillar data centre TCO model presented here has been developed to give organizations facing data centre decisions a better understanding of the nature, complexity and extent of the cost-related issues associated with building, operating and maintaining an in-house data centre capability.



While this paper is not intended to be an in-depth treatise on every aspect of this model, a short summary of each of the four cost pillars, along with several illustrative examples, will help to set the tone for better data centre decision-making. In all likelihood, it will also provide the understanding that executives need to better justify and feel more confident about a “Buy” decision, as well as give them valuable cost data against which to measure and compare the colocation services provider’s cost quotation.

Facility Requirements — conspicuous & easy to understand

This first pillar consists primarily of two sub-categories of upfront costs that together represent 60% of the total data centre investment:

1. Physical Facility Costs

i.e. walls, doors, raised floor, entrances, turnstiles, mantraps, security/reception station, etc.

Designing, building and operating a data centre is considerably different than it is for other buildings and requires a highly-specialized level of expertise and experience. Even very experienced data centre operators may not have the design and build resources necessary to undertake such a major building project. The nuances and subtleties of a sophisticated data centre design are somewhat similar to other types of specialized facilities, such as an oil refinery or chemical plant, which require specialized design skills and experience, often acquired through engineering firm engagements that may not be part of an organization's core skill sets.

2. Computing Environment Costs

i.e. cabinets, power, HVAC, smoke/fire detection and suppression, etc.

The initial capital costs, as well as the ongoing costs, are directly related to the levels of redundancy of the power and cooling equipment — Tier II, III or IV, based on Uptime Institute standards. The higher the level of redundancy, the more equipment is needed in order to prevent infrastructure equipment failure from impacting the operation of the resident computing equipment.

A common pitfall for organizations looking for ways to cut costs is choosing less-than-best products with which to build and outfit their data centre facility. Compromising on quality to save cost is a bad move to make when you are in the data centre business, and the potential risks include degradation of equipment performance and a reduction in the Mean Time Between Failure (MTBF) of critical facility components, which can have serious impact on the performance and uptime of IT equipment residing in the data centre.

Designing and building a data centre requires careful scrutiny of costs associated with physical facility construction and the purchase of major capital equipment, both of which have a significant impact on final cost. One advantage enjoyed by larger, top-tier data centre services providers is their purchasing power and ongoing relationships with construction firms and equipment vendors. Since a data centre build project is likely to be a once-in-a-career event for most CFOs and CIOs, they typically don't have access to the same special pricing deals and volume discounting as a data centre services provider who buys and uses these products repeatedly, either while regularly updating existing facilities or building additional data centres across the country to accommodate additional customers and those wanting multiple data centres for business continuity reasons. Colocation vendors get better quality services and products at a lower cost.

Operational Procedures — things get a bit fuzzy here

Some aspects of this TCO category are obvious and fairly well understood by most, while others are less so and often greatly underestimated. This pillar is characterized by the wide variety of cost dimensions it encompasses and by the fact that it includes an initial upfront investment, which represents about 15% of total data centre investment, as well as ongoing additional costs of 20% to 30% yearly — both capital costs and operational expenses.

The biggest components of this category are likely to include:

- **Ongoing/preventive maintenance of facilities equipment**

(i.e., power backup/generation, HVAC);

The cost of preventive maintenance, and ensuring consistency in procedures for maintenance and support of mission critical infrastructure, contributes significantly to overall data centre TCO. To cite one example, consider the ongoing maintenance of an Uninterruptible Power Supply (UPS). An Emerson Network Power study on the impact of preventative maintenance on UPS reliability revealed that the Mean Time Between Failures (MTBF) for units that received two preventive maintenance visits per year is 23 times better than a UPS that has had no preventive maintenance. Being aware of, and acting upon, this type of knowledge has a direct bearing on TCO.

- **IT equipment — initial purchase/installation, ongoing maintenance, periodic refresh/upgrade**

(i.e., servers, storage, network components, telephony, peripherals, software licences);

- **Power/electricity costs**

The cost associated with power consumption is steadily increasing as the cost of electricity increases and as the power density of technology also increases. In fact, numerous studies have shown that energy cost is now the largest expense related to operating a data centre, and based on current and foreseeable trends, is expected to continue to rise over time. Designing and building for energy efficiency is now an operational necessity.

This is a direct cost that needs to be fully understood by specialized data centre technical staff with the ability to measure power consumption at key points throughout the data centre, determine the expected power consumption by various infrastructure combinations and forecast changes in consumption — and thus the impact on cost — in lock-step with changing technology and changing business needs over time. The increased power densities of today's new server processors are quite unlike those of previous years, for example, and power supplies, which historically were only efficient when operated at peak power load, can now deliver much-improved performance when operating at only 40% to 60% of peak load. In addition, new power supplies can support 40% to 50% greater load at lower cost — delivering more for less.

It takes an expert in power management to understand and manage today's complex power environment and to ensure the continuous, reliable delivery of clean power throughout the data centre, based on ample redundant UPS and diesel-generated backup power capabilities. These are the challenges that colocation services providers continually deal with when working with multiple customers and data centre sites every day.

■ Staffing for numerous data centre roles

(Including security Maintenance & support for both facilities and IT equipment, network & performance monitoring, help desk and others).

The 24x7 availability of technically-qualified support and maintenance staff for the mission-critical systems, such as UPS, power distribution, backup generators and cooling systems, is also key to ensuring continuous availability. Certified electrical and mechanical personnel are good examples. The importance of this factor should not be underestimated, and there are significant ongoing costs of hiring and retaining these highly-specialized and qualified personnel.

The cost impact of staff churn, which can be a significant TCO factor, must also be taken into account — there is an active job market for experienced IT professionals these days, so it must be determined how to retain such staff in spite of other lucrative offers, and what impact their leaving would have on operations and TCO.

Data centre services providers, on the other hand, have a much easier time attracting and retaining specialized data centre expertise because of the multi-customer, multi-site nature of their business and the higher volume of activity and interesting and unique challenges this creates. Because of multi-tenancy, staffing costs can be amortized across multiple customers, reducing the effective cost to any one customer.

■ Additional costs, often hidden

(Including rent, capital depreciation, asset management, development of operational policies and procedures, development and ongoing testing of security, backup and disaster recovery plans, and numerous others).

The effort and cost for disaster recovery planning and periodic testing, for example, often turns out to be more than what was originally budgeted for. One reason is that any changes that have taken place in the data centre must be taken into account in terms of updating the disaster recovery plan, updating or adding to the redundant infrastructure and updating recovery procedures. Companies that build their own data centres often ignore disaster recovery planning and investments in redundant infrastructure — at least until they have experienced their first disaster. With their minds on cost avoidance, they may simply have underestimated the value of investing time and money in this area — perhaps another misguided case of sacrificing quality to cut cost.

Let's Check “Build vs. Buy” ROI So Far...


It is likely safe to say, that most organizations are at least aware of the basic cost considerations addressed by the first two pillars of the data centre TCO model — Facilities Requirements and Operational Procedures.

Considering only the costs of those first two pillars, and assuming a 900 square foot, 20-cabinet ‘data centre model’ built to Tier II Uptime Institute standards, it takes approximately 11 years before an investment in an in-house data centre would equate to the cost of a comparable colocation solution. In other words, **it takes 11 years before the return on the investment in building an in-house data centre reaches the break-even point with the cost of buying equivalent data centre services from a third-party colocation services provider.**

While this calculation includes initial build and ongoing operational costs, as well as some increases in electricity costs, it does not account for any additional outlay required for equipment refresh or replacement, which in all likelihood would be required a couple of times over an 11-year period. This makes the “Buy” model more attractive for an even longer period of time.



“BUILD VS. BUY” ROI COMPARISON




Data Centre Capacity — the “stranded or capped” dilemma

The capacity of a data centre is described as being either “capped” or “stranded”. As the name suggests, a data centre with capped capacity has been built to meet today’s needs and is not able to support any future growth without further investment in expansion or upgrade. On the other hand, a data centre with stranded capacity has been over-built in anticipation of future growth; it has unused capacity today, and the capital that was invested to build that capacity is also considered to be stranded and thus not available to be used for more strategic business initiatives.

This dichotomy creates a lose-lose investment dilemma for organizations that choose to design and build their own data centres.

Those that choose to build a bigger facility upfront that does not limit their capacity for growth are left with stranded capital, and the ROI break-even point for “Build vs. Buy” can easily extend beyond the original 11 years to 15 years, or more. Those choosing to avoid stranded capital by building a smaller facility upfront that more closely matches current needs often find themselves faced with the need to invest additional capital to expand their facility before they have even reached full capacity. As indicated by the diagram on page 11, this can push the “Build vs. Buy” ROI break-even point out even further, perhaps to as far as 20 years.

“BUILD VS. BUY” ROI COMPARISON




The utilization of a data centre changes over time and can range from 10% to 90%. Capacity, utilization, and the variation in utilization over time, are important factors to consider when performing a data centre TCO analysis to support a “Build vs. Buy” decision.

However, who really has any idea what their needs are going to look like or how technology will change in 11 years, much less 15 or 20, and whether a “Build” investment will ever actually break-even against the compelling economics of a “Buy” decision. Buying colocation services requires no ‘crystal balling’ — buy what you need today, and then expand (or contract) when you need to. From the customer’s perspective, their outsourced data centre capacity is never capped or stranded.

Data Centre Criticality – uptime is king, but there’s a cost

With the continued increase in reliance on information technology systems to support critical applications and business processes, a single downtime event now has the potential to significantly impact the profitability – and, in extreme cases, even the viability – of an enterprise. In fact, **a worldwide study found that 73% of respondents believe that an outage of up to 24 hours could prove fatal to their organization, and that 32% believe that an outage of less than four hours could have this same effect.**⁴

As a result, CIOs no longer have any choice but to invest in uptime and in robust, reliable systems. Just how much should be invested, however, depends on the value the organization places on uptime, or conversely, what the cost of downtime is to the organization. Determining this is an important aspect of a data centre TCO analysis that is often overlooked. In fact, in a survey conducted by the Neverfail Group, 54% of respondent companies did not know the cost of their downtime.



The cost of downtime can vary greatly, based on the organization's dependency on its computing systems. In a study of 41 large data centres (i.e., 2,500 square feet minimum) by Emerson Network Power in conjunction with the Ponemon Institute, the average cost of downtime was found to be \$5,600 per minute, with the average cost per downtime event being \$505,500. These costs are based on a variety of factors, including but not limited to data loss or corruption, productivity losses, equipment damage, root-cause detection and recovery actions, legal and regulatory repercussions, revenue loss and long-term repercussions on reputation and trust among key stakeholders. Business disruption and lost revenue were cited as the most significant cost consequences of downtime.³


In another study by the Fibre Channel Industry Association, the cost of downtime for financial services institutions and information technology companies averaged more than \$1.3 million per hour, while the cost for healthcare and professional services organizations averaged \$550K/hour, and \$75K/hour for other enterprises such as Airline reservations.

This final TCO pillar or category of data centre cost corresponds to the amount that needs to be invested in infrastructure redundancy and other forms of downtime mitigation to ensure an appropriate level of uptime to meet business requirements. Interestingly, the leading – and most costly – root causes of downtime reported by respondents were directly related to vulnerabilities in data centre power and cooling infrastructures.

Apart from equipment failures that cause systems to go offline, emergency repairs and even planned maintenance on systems may require them to be taken offline intentionally. An investment in infrastructure redundancy plays an important role here in a concept known as “concurrent maintainability”, which enables systems to be bypassed, and thus taken offline without impacting data centre operations. This is a key criterion in Uptime Institute's Tier III and Tier IV data centre certification levels.

For our sample data centre model of 20 cabinets occupying roughly 900 square feet, it will cost about 20% of the initial upfront capital investment to ensure Tier II standards for uptime of 99.9% (a.k.a. “3 nines”). However, this still equates to 43.2 minutes of downtime monthly, which could have a significant impact on business. One way to further reduce downtime is by investing in additional redundancy, perhaps based on data centre design standards that approach Uptime Institute's Tier III classification. This would equate to additional data centre costs of approximately 30% of the original upfront investment and a 25% increase in ongoing operational costs. When this is added to our original “Build vs. Buy” ROI comparison, these increased costs push the break-even point out to 16 years. Investing in “criticality” means investing in greater uptime, which in turn carries with it additional upfront and ongoing costs.

“CRITICALITY”... HOW MUCH UPTIME IS ENOUGH?



99.9% = **43** minutes
UPTIME OF DOWNTIME



“Build” Option Requires Crystal Ball Gazing

Making an informed “Build vs. Buy” decision requires a comprehensive and in-depth analysis of the costs associated with each option. Depending on the organization’s type and size, performing an honest self-evaluation of core strengths and the level of data centre design-build expertise and experience of internal resources can also be an important factor in making the best decision.

A prospective colocation services partner will happily provide the cost to “Buy” whatever data centre services are needed. Since their core business – the business they focus exclusively on every day – is designing, building and operating data centre facilities for their customers, they can be very cost effective, as well as being able to rapidly deploy data centre capabilities based on a well-proven design.

For the “Build” side of the analysis, the data centre TCO model presented herein hopefully provides a useful and effective framework for examining the full range of costs associated with a comparable “Build” option. Depending on the scope and complexity of their data centre requirements, each organization must decide for itself how heavily it should weight each of the four pillars of the model when conducting its TCO analysis.

Without a complete understanding of all the costs involved in a “Build” scenario, what may initially calculate out to be an 11-year ROI, can quickly turn into 15, 20 or even 25 years. **How comfortable do you feel that you can look that far out and predict the future for data centres, cloud computing and the state of infrastructure services in general?**

“While there are a multitude of factors to be considered, in the end you will have to ask if building, owning and operating your own data centre is a strategic advantage to your business, or just a burden on internal resources and capital, perhaps best outsourced to a qualified external organization.”

Julius Neudorfer, CTO and founder of North American Access Technologies, Inc.

1 Gartner Press Release, “Gartner Executive Programs Worldwide Survey of More Than 2,000 CIOs Identifies Cloud Computing as Top Technology Priority for CIOs in 2011”, January, 2011

2 Info-Tech Research Group, “Develop a Colocation Strategy: Successful data center colocation strategies are a function of cost and vendor selection” at <http://www.infotech.com/research/ss/develop-a-colocation-strategy>

3 Emerson Network Power, “Understanding the Cost of Data Center Downtime: An Analysis of the Financial Impact on Infrastructure Vulnerability”, 2011

4 Osterman Research, “Key Issues in Selecting an E-mail Disaster Recovery and Business Continuity Solution”, May, 2009