

Data Center Physical Realities in a Software-Defined And Virtualized World

Optimizing IT Infrastructures for Today's Business Demands

Overview

As our businesses become increasingly digital, we tend to think about technology in non-physical terms. Our IT infrastructure becomes “the cloud.” Our servers and storage become “virtual.” Our networks become “software-defined.”

The reality, however, is that information technology (IT) *always* depends on physical infrastructure. Without hardware there is no software. So to lose sight of technology's physical realities is to jeopardize the wisdom with which we manage its physical underpinnings. And unwise management of those physical underpinnings leads to both higher operational costs and higher operational risks.

In fact, the great irony of virtualization is that it makes the physical realities of IT infrastructure even more important by dramatically increasing density and scale of workloads in the data center—and by dramatically increasing the consequences of physical infrastructure failures.

This white paper addresses five key aspects of IT that are inextricably tied to computing's physical realities, even as that computing becomes more virtualized, software-defined, and cloud-based:

- **Technology economics**
- **Business risk**
- **Operational scale**
- **Agility and innovation**
- **Security and compliance**

By diligently addressing these physical realities, data center managers can better enable the organizations they serve to succeed in an increasingly tech-centric world—while they continue to pursue the evolving opportunities offered by virtualization, the cloud, and software-defined IT infrastructure.

More Data = More Physical

Whether you're managing massive amounts of capital or making it easy for people to flag a ride downtown, your business depends on software and data. That's why companies are making massive investments in Agile development, Big Data, and other disciplines that help them get the upper hand when it comes to digital customer engagement and analytic insight

However, every organization's ability to successfully use software and data in any form—including databases, transaction processing systems, mobile apps, and IoT networks—depends on healthy, efficient underlying physical infrastructure. In fact, the more your business depends on software and data, the more it depends on the uninterrupted physical flow of electrons across silicon chips, spinning disks, solid-state media, and network cables made of copper or optical fiber.

The physical realities of computing thus absolutely constrain business performance in several specific ways, including:

- **Accurate Delivery of the Right Power to the Right Device**

You can't compute without power. So it's essential to deliver the right power to every device all the time. This has become a non-trivial engineering challenge as virtualization intensifies power demands and as next-generation infrastructure—including high-density OCP servers, converged systems, and DC-powered drives—introduce greater variability into power requirements across the data center.

More Data = More Physical (Cont.)

- **Appropriate Temperatures and other Environmental Conditions**

Computing equipment generates a lot of heat, especially as data centers get denser. It is thus essential to always apply sufficient cooling everywhere across your data center infrastructure to avoid the overheating that can cause equipment failure—while also avoiding wasteful overcooling. It is also essential to protect your equipment from other threats such as excessive humidity, vibration, and airborne particulates.

- **Access for Maintenance and Troubleshooting**

Disks crash, fans short out, and power supplies fail. So physical equipment requires physical service. These tasks must be performed quickly and reliably by authorized technical staff so as not to excessively consume their limited time. Again, this is a non-trivial challenge given the increasing density of equipment in the data center—as well as the way the relationship between that equipment and the critical IT services

it supports keep becoming more dynamic and complex.

- **Access for Maintenance and Troubleshooting**

Disks crash, fans short out, and power supplies fail. So physical equipment requires physical service. These tasks must be performed quickly and reliably by authorized technical staff so as not to excessively consume their limited time. Again, this is a non-trivial challenge given the increasing density of equipment in the data center—as well as the way the relationship between that equipment and the critical IT services it supports keep becoming more dynamic and complex.

These physical aspects of technology management should not be minimized or overlooked. Done right, they enhance the performance of IT and, by extension, the business as a whole. Done wrong, they undermine that performance—with potentially serious adverse consequences to both the top and bottom line.



SOUND PHYSICAL INFRASTRUCTURE FOR DIGITAL BUSINESS SUCCESS

The success of every business is ultimately built on the physical foundation of data center infrastructure—whether that infrastructure is on-premise, off-premise, or hybrid.

Five Critical Impacts

Business leaders have historically thought of physical data center management as a tactical challenge for IT. But as business itself becomes more pervasively software- and data-centric, this is no longer the case. The underlying physical infrastructure of every business now directly impacts the business in five critical ways.

Critical Impact 1: Technology Economics

Technology costs money to acquire and operate. At one time, these costs were a relatively small percentage of overall business spending. As the work of building, marketing, selling, and supporting products and services is increasingly done with software and data, however, technology costs have become more significant.

And those costs can't be measured by IT budgets alone. Many groups outside IT now engage in their own initiatives around applications and analytics. These groups bear the "above the line" costs for their initiatives. But "below the line," such initiatives continue to drive physical infrastructure costs through workloads that drive consumption of server, storage, and network capacity.

Physical infrastructure drives costs in three main ways:

- **Electricity Costs**

Information technology accounts for about 10% of the world's total electricity generation, according to a study by Digital Power Group. In the data center, these costs are driven by equipment, cooling, and UPS devices that ensure business continuity. As digital engagement with customers continues to expand, these costs all increase.

- **Labor Costs**

Companies need skilled IT staff to monitor, maintain, and troubleshoot physical data center infrastructure. Increasingly, these skills have to be available on a 24x7 basis to ensure round-the-clock digital engagement with customers. Many sites

also require security staff to maintain 24x7 security surveillance of the data center.

- **Floorspace Costs**

Floorspace can be a significant cost even for data centers outside urban areas. And some percentage of data center floorspace is always non-productive, since companies must maintain room for expansion. The acquisition of additional floorspace is a particularly large cost event, since it must typically be acquired in large increments. And the cloud doesn't always offer a cheap alternative—because floorspace remains a cost component of any cloud service.

Data center infrastructure management best practices significantly reduce these costs by ensuring that power, people, and floorspace are used with maximum efficiency. At large enterprises, these savings can be substantial. Improved data center economics are important to small-to-midsize enterprises as well, since their costs are proportionally greater due to lesser economies of scale. And for service providers, these savings translate directly into increased profitability (see Service Provider Challenge section, p.10).

Critical Impact 2: Business Risk


As businesses become increasingly dependent on technology, the risks associated with physically based failure of their technology infrastructure become more problematic. Data center downtime no longer *just* costs companies a bit of employee productivity. It halts revenue, jeopardizes long-term customer relationships, and can badly undermine brand value.

Five Critical Impacts (Cont.)

Every business's financial well-being is contingent on the physical well-being of its data center infrastructure. But here are just a few key threats to consider:

- **Power Failures/Brownouts**

Loss of and/or insufficient electrical power to the data center represents a major risk to every business. Extreme weather, in particular, is dramatically increasing the frequency and magnitude of regional power outages (see chart)—with the five-year annual average of outages *doubling* every five years. Power can also be interrupted by everything from human error by grid operators to supply-demand disparities in summer months.



Businesses are more likely than ever to suffer loss of power to the data center due to an extreme weather event.

- **Cloud/Internet Connectivity Disruption**

A common misconception is that businesses insulate themselves from physical disruption of their IT services when they adopt cloud, SaaS, and other Internet-based computing resources. But often, just the opposite is true, because physical disruption of a single edge networking device—due to a power issue, an overheated wiring closet, or some other physical problem—can cut a business

off from all of its external IT resources. So the adverse consequences of that single physical point-of-failure are actually much greater, rather than much less.

- **Device Blindness**

When IT staff lacks clear, accurate, and up-to-date visibility into all physical infrastructure in the data center, two bad things can happen: 1) it can take them longer to fix many types of technical problems that require the right device to be modified in the right way, and 2) they can inadvertently make mistakes that disrupt critical IT services. Both of these risks increase as data center infrastructure becomes denser and more complex.

Business activity is threatened by many other risks as well: HVAC failure, flooding due to extreme weather or broken pipes, tremors and earthquakes, etc. So every business must take appropriate steps to mitigate these risks based on both their probability and their financial impact on the business.

Critical Impact 3: Operation Scale

Growth is an almost universal business goal. But, as any experienced business leader can tell you, growth presents its own set of challenges—especially if it happens quickly. And as business growth increasingly requires companies to grow their IT operations, data center managers must increasingly be prepared to scale infrastructure capacity.

This operational scalability is closely tied to physical factors, including:

- **Floor/Rack Density**

Because the amount of space in any data center is finite, IT must make optimum use of existing rack footage. This is ideally done from the outset—so that IT staff doesn't have to scramble at some

Five Critical Impacts (Cont.)

future point of crisis to eke out more room. Plus, as noted above, increased density makes it essential that IT staff maintain accurate documentation of data center assets and cabling—even as moves, adds, and changes occur over time.

- **Power Distribution and Monitoring**

Proper power capacity planning is a must for any growing business. Without such planning and real-time monitoring, critical IT services can be threatened by an inability to deliver sufficient power to a growing number of more power-thirsty devices. Also, as power needs grow, it becomes more economically worthwhile to monitor and optimize data center efficiency metrics (PUE, DCiE, PI, etc.).

- **Intensified Cooling Requirements**

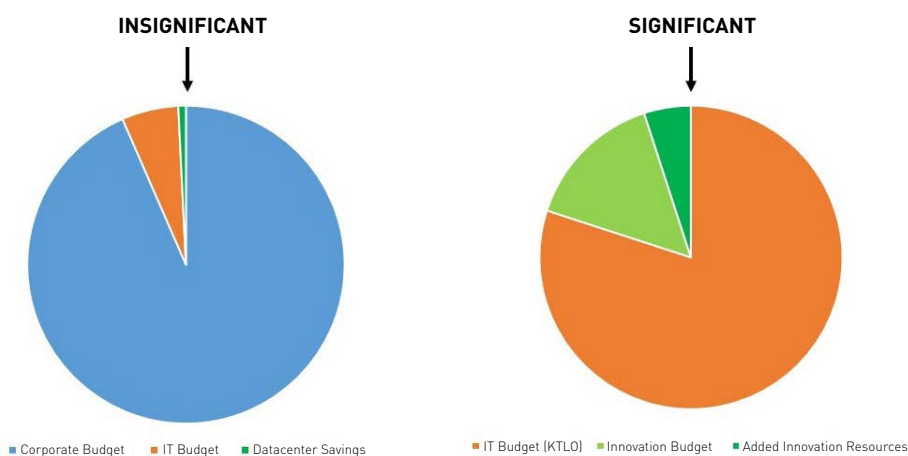
As noted above, increased equipment density can make it difficult to properly cool the data center to required temperatures. Scaling IT operations to meet the growing needs of the business therefore requires a variety of measures, including investment in cooling capacity, optimized equipment

layout for temperature regulation, and additional environment sensors.

Many business and IT decision-makers are attempting to insulate themselves from these scalability issues by using the cloud—which promises elastic, adaptable IT capacity. The cloud alone, however, is an insufficient scalability solution for most organizations. That's in part because most organizations for a variety of reasons cannot run all their systems in cloud and in part because the economics of the cloud are often not very favorable when it comes to extended growth.

Critical Impact 4: Agility and Innovation

Technology change nowadays entails more than just mere expansion of scale. Businesses also have to nimbly adopt new technologies as they become available and as they become relevant to the business. These technologies can include solid-state storage systems that dramatically improve performance over traditional spinning disks, specialized security



HOW DO DATA CENTER SAVINGS IMPACT THE BUSINESS?

Data center savings appear insignificant when viewed only as a percentage of total corporate spending. [e.g. 3% reduction of an IT budget that represents 6% of total corporate spending = .0018%]

The significance of data center savings is better reflected in how it frees IT budget for innovation. [e.g. shifting 3% of IT spending from the 85% currently allocated for “keeping the lights on” (KTLO) to the 15% currently allocated for innovation represents an increase of 20%]

Five Critical Impacts (Cont.)

appliances, and converged systems that provide tight integration between purpose-configured compute, storage, and network components.

IT staffs must be able to accommodate the accelerating flow of these infrastructure technologies in and out of the data center. This flow puts pressure on all aspects of data center infrastructure management already mentioned above—including power, cooling, and monitoring. Converged infrastructure in particular can add to data center density, since it packs more components, power supplies, etc. into a smaller form factor.

It is also important to note that business agility can include scaling down as well as scaling up. So when periods of lower infrastructure utilization occur, IT staffs have to be able to detect and act on opportunities to consolidate VMs, power down devices, etc. in order to save on operating costs.

Critical Impact 5: Security and Compliance

Cybersecurity is a central concern for every organization today. Security breaches can not only halt operations, interrupt sales, and cause irreplaceable data loss. They can also do permanent harm to customer relationships and brand reputation.

Unfortunately, while IT leaders focus on cybersecurity measures such as authentication and encryption, they often fail to properly address basic physical security in the data center. This leaves businesses vulnerable to data theft via server USB drives, equipment theft, sabotage and other threats.

Historical monitoring of physical data center access is also often necessary for regulatory compliance, since auditors may require such documentation as proof that sensitive data was not exposed to unauthorized individuals—and/or that the business has exercised due diligence in preventing same.

The Service Provider Challenge

IT service providers come in all shapes and sizes: traditional colo/hosting facilities, cloud-based infrastructure and platform solutions (IaaS and PaaS), pure-play SaaS, and a growing ecosystem of vendors providing databases, analytics, security, and more on an on-demand basis.

From a purely physical perspective, these service providers face the same challenges as enterprises. They have to optimize power efficiency, keep equipment cooled, mitigate risks, and adapt to constantly changing requirements.

There are, however, three reasons why service providers need to be especially diligent when it comes to optimizing physical management of their data center facilities:

- **Operating Margins**

Data center operations represent a fraction of a fraction of enterprise IT budgets. Not so with service providers—for whom data center floorspace, power, cooling, and operations represent the third largest cost behind equipment and HR. Increased physical efficiencies thus translate directly into higher profits.

- **Cost Visibility**

Enterprises can choose whether or not to actually chargeback or showback data center costs. Service providers have no such luxury. They need to know exactly how much it costs them to deliver every service for every customer. Without that insight, they can't accurately price their services or measure the real profitability of their active accounts.


- **Scalability/Adaptability**

In enterprises, changes in the data center are regularly necessary to support the business. For service providers, changes in the data center essentially are the business. Any constraint in the speed and ease with which those changes can be made pose a threat to growth, operating margins, customer value, and competitive differentiation.

For these reasons and others, service providers need to be even more diligent about optimizing physical management of their facilities.

CLOUD MAKES DATA CENTER INTEGRITY EVEN MORE CRITICAL

As enterprises embrace cloud, data center integrity remains critical—and may become even more so—because uninterrupted access to external resources depends on healthy network infrastructure.



Achieving Physical Fitness

Given the impact of physical infrastructure on business success, what steps should business leaders take? How can data center owners best act to optimize the physical fitness of their enterprise infrastructure?

While a complete inventory of data center best practices is beyond the scope of this paper, here are three actions every company can take immediately to address the physical realities of business in an increasingly tech-centric world:

1. Elevate Responsibility for Physical Infrastructure Strategy

Many companies leave it to front-line hands-on techs to let them know if “there is anything wrong” in the data center. But this if-it-ain’t-broke-don’t-fix-it mentality doesn’t best serve the business, because front-line techs are typically only empowered to address immediate uptime and performance issues—and not to make strategic decisions about operational efficiency, corporate carbon footprint, and the like.

Business leaders must therefore ensure that the right people at the right level of the organization own physical infrastructure strategy and are accountable for its KPIs.

2. Re-Assess and Re-Inventory the Data Center

A lot has changed in the past few years. Data center infrastructure density has increased. Customer-facing application components have become more distributed across hybrid cloud. Climate change is a growing C-level concern. Yet many IT organizations have not taken a long, hard look at their physical data centers in light of these dramatically different conditions.

Business leaders should address the major changes in enterprise IT with a corresponding re-evaluation of current physical state of the data center, along with associated operations.

Achieving Physical Fitness (Cont.)

This re-evaluation should include:

- How equipment is distributed in the data center and how decisions are made about where to place new equipment as it arrives
- Ensuring proper use of PDUs to optimize power distribution and provide accurate data about data center energy consumption—in tandem with best-practice processes for analyzing and acting on that data
- Investigation of how network equipment connecting the enterprise to the cloud is protected in the data center and remote offices
- Control and monitoring of physical access to the data center
- Collaboration with CCO, CFO and other stakeholders regarding IT's role in carbon footprint initiatives

3. Tap Outside Expertise

Businesses are constantly leveraging external technical resources. They outsource application development and mobile interface design. They hire ethical hackers to perform penetration testing. They work with DR consultants to rigorously safeguard

business continuity. Yet, for some reasons, many IT organizations assume that they fully possess all the experience and knowledge necessary to optimize physical management of the data center.

This is probably not the case. In fact, most IT organizations possess minimal in-house expertise in electrical and thermal engineering. Nor do they usually stay abreast of the latest development in DCIM best practices.

It therefore makes sense for business leaders to tap outside expertise in these critical areas—especially considering how much business success depends on reliable, efficient underlying physical infrastructure.

Simply put, businesses that take their physical infrastructure for granted spend more, achieve less, and expose themselves to greater risk than those that exercise more diligence in the management of power, cooling, operational efficiency, and security. That's why every business leader should come to terms with the physical realities of their increasingly technology-centric world—and take the time to ensure that those realities are working in their organization's favor.

About Raritan

Raritan, a brand of Legrand, is a trusted provider of rack power distribution units, branch circuit monitors, transfer switches, environmental sensors, KVM-over-IP switches, serial console servers, and A/V solutions for data centers and IT professionals. Established in 1985 and based in Somerset, N.J., Raritan has offices worldwide serving customers in 76 countries. In more than

50,000 locations, Raritan's award-winning hardware solutions help small, midsize, enterprise, and colocation data centers to increase efficiency, improve reliability, and raise productivity. And provide IT departments with secure, reliable remote access tools needed to manage mission-critical environments. For more information, visit us at Raritan.eu