



Increase Power Density To Decrease Data Center TCO

White Paper

Table of Contents

Modern Trends in IT Infrastructure Deployment and Management	1
Disadvantages and Limitations of Low-Density Data Centers	3
Cost and Operational Benefits of High-Density Data Centers	6
Case Study: How High-Density Data Centers Deliver Reduced TCO	7
Conclusion and About Colovore	11

EXECUTIVE SUMMARY

Modern IT user demands are driving new developments in software applications. Efficient commodity and converged hardware is being widely deployed to support this systems development thrust. As a result, the total power consumption in data centers is increasing to levels that legacy, low-power density facilities cannot easily support.

Consolidating applications using virtualization isn't the only way to save IT OpEx. This whitepaper will show that consolidating servers currently in sparsely occupied low-density cabinets into fully-occupied high-density cabinets can provide significant savings only when the data center is designed to be high-density from day one.

MODERN TRENDS IN IT INFRASTRUCTURE DEPLOYMENT AND MANAGEMENT

IT environments have grown complex, sprawling, and power-hungry in recent years due to the proliferation of systems, devices, and storage necessary to support the incredible growth in IP-based connectivity. This complexity is being tamed by advances in modern server technologies and software infrastructure management. Both clearly provide efficient and cost-effective ways for companies to organize and improve their IT infrastructure.

The cause of this trend is easy to see, and explain. End users demand more software features. This drives the need for more hardware deployed on an increasingly larger scale. To save time, software developers want access to compute, storage and memory resources without having to worry about the underlying disparate hardware systems.

This has led to the growth of virtualization and new demands to provide ways to improve server resource utilization while simultaneously consolidating new and legacy applications in their own secure, independent environments. Hardware is not just becoming more heavily utilized, it is also becoming more of a generic commodity often having multiple functions consolidated into a single converged chassis. For example, higher-end blade server chassis often come with network, and storage capabilities built in.

Many companies are also experimenting with CapEx efficiency by standardizing on a limited range of desired platforms and developing their own provisioning systems. Others are exploring ways to distribute compute loads more evenly throughout the day, cabinet or data center to further improve efficiency.

This desire is the foundation of cloud computing, both private and public. The trend is guaranteed to continue because hardware commoditization, virtualization and converged infrastructure solutions from major OEMs including Cisco, Dell, HP, and SuperMicro enable corporations to decrease IT overhead and operational costs.

Virtualization Benefits and Considerations

This trend in IT infrastructure is both clear and powerful:

- Research firm Gartner forecasts that 82.4% of total operating system deployments will be virtualized by 2016
- Research firm Frost & Sullivan estimates that more than 70% of midsize and large businesses in the US are already using server virtualization
- Research firm IDC reports that blade servers now account for nearly 17% of all server revenues

The benefits of virtualization are many:

- Consolidation of server infrastructure into fewer, more powerful systems
- Aggregation of electrical load into larger, more easily manageable units
- Decreased capital costs (ie, standardized server hardware, cabling, rack switches, software licensing, and installation)
- Simplified maintenance requirements and operational costs

According to IDC, blade servers on average consume 40% less power than traditional rack servers and take up 37% less space. Given that power can consume up to 70% of a company's data center operating costs, efficiencies in electrical load management can yield considerable savings. With its newly-launched Moonshot cartridge-based server, HP claims a 77% cost reduction per cartridge versus a traditional server and estimates that the savings in TCO for a company deploying a 16-cartridge server deployment versus a traditional 16-server rack system are \$212,000, or nearly 50%, over a 3-year timeframe.

Even so, the HP Moonshot and similar converged infrastructures can be challenging to provision. The Moonshot, for example, hosts 45 server cartridges in the space occupied by just over 4 traditional 1 rack unit (RU) servers, making its deployment options in traditional data center facilities and cabinets limited due to its overall power consumption. Moonshot's specifications show that the chassis is designed to support the consumption of 500W to 1kW of power per rack unit based on half-to-full service load.

This is a very practical and important consideration companies must be aware of as they seek to deploy virtualized environments and blade servers-- how is your current data center equipped to handle this infrastructure? Simply put, older data centers can present meaningful challenges and in some cases impede the deployment of modern server infrastructure. This is unacceptable, especially when IT is increasingly satisfying customer demands. It is imperative that companies understand the capabilities of their existing data center environments in order to take advantage of server trends without the datacenter capacity tax.

DISADVANTAGES AND LIMITATIONS OF LOW-DENSITY DATA CENTERS

When the limited capabilities of legacy data centers are involved there can be a number of challenges when modern user needs clash with the software and hardware requirements to meet them. We'll explore this in more detail next.

Virtualization and Its Impact on the Data Center Cabinet

As mentioned, one of the primary advantages of deploying virtualization and blade server infrastructure is a reduction in the amount of power required to run those servers through the efficient use of shared components. But IT managers need to exercise due diligence when evaluating blade systems. Hardware manufacturers are simultaneously taking advantage of component miniaturization to maximize the number of virtual machines (VMs) per chassis.

Technological developments in traditional servers are seeing increased unit efficiencies as well, but increased on-board, CPU, RAM and storage capabilities in the same physical footprint are driving up their power consumption.

So while individual server power draw can be reduced upon the implementation of newer architectures, power requirements for those servers *at the data center cabinet level increase considerably*.

Using standard off the shelf hardware it is now possible for servers to draw up to 1 kW peak load per rack unit. With 40 - 50 rack unit cabinet sizes, this means that the power required per cabinet increases by orders of magnitude when compared to legacy datacenter capabilities, and can easily grow to 20 kW+ per cabinet. Also, directly related, the amount of cooling required to ensure proper server function also increases at the cabinet level. This can place a significant strain on a data center's infrastructure and operating capacities, and in many cases mature data centers are not be able to support the higher-density loads.

The Evolution of Data Center Cabinet Power Requirements

Traditional data centers were built in many cases in the late 1990s and early 2000s. At the time, facilities were designed to supply 50-75 watts of electricity per square foot. As a cabinet usually occupies 25 square feet of space when walkways are included, these facilities could only support total power draws of 1-2 kW per cabinet. As servers advanced to require data centers with 100 and then 200 watts of electricity per square foot, data center operators upgraded or built new data centers to accommodate the increased electrical requirements. 4-5 kW per cabinet densities emerged as the standard in data center architecture, and today the vast majority of data center floor space still features this 4-5 kW per cabinet standard. It is quite expensive to retrofit mature data centers to significantly upgrade the power and cooling infrastructure necessary to deliver higher densities, not to mention the risk of business continuity and service delivery interruption during times of intense renovations.

The current challenge is that server power requirements have continued to grow and grow. As previously mentioned, today an off-the-shelf, 1U Dell R620 server can draw up to 1000 watts to run—an increase of 20x since the year 2000. This has major ramifications for traditional low-density data center operators and ultimately leads to inefficiency, both from a cost and operations perspective for the end user, and potential obsolescence.

The Challenges of Modern Servers in Low-Density Cabinets

So how does a low-density data center provider handle a customer who wants to deploy modern server infrastructure? There are two strategies. The first forces the customer to distribute their IT infrastructure across multiple unfilled cabinets and non-contiguous floor sections. This is often referred to as “spreading the load,” and it is necessary because the data center does not have enough power and cooling to adequately support cabinets with more than 4-5 kW of draw. The end result is

significantly higher operational costs, with cabinets only partially-full with servers and networking equipment. Specifically:

- half-empty, power-starved, under-utilized racks still cost the same amount per cabinet per month to rent, leading to wasteful monthly cabinet fees
- more expensive cross-connects and top-rack switches are required to connect non-contiguous space or cabinets
- cabling costs increase as a result of longer runs, and cable management becomes complex and often unwieldy
- maintenance costs and operational complexity increase as a result of widely-distributed, non-contiguous infrastructure sprawl
- server operating performance can suffer if blanking panels are not adequately deployed
- there is a greater operational risk of tripping smaller-capacity power circuits with high-power usage server chassis
- a single server chassis can consume more than 50% of a 15A, 120V power circuit which leads to less efficient HVAC utilization

In addition to spreading the load, the second strategy low-density data center providers employ is a re-distribution of power and cooling infrastructure to provision a “high-density” section of the floor. However, the customer still has to pay for the real estate that has been robbed of power to provide the capacity elsewhere. Customers are therefore billed for these “virtual” or “phantom” cabinets. This unused space can be physically disguised as caged-off walkways or ad-hoc staging, storage and office space areas around cabinets. While this may serve as a stop-gap measure for existing clients wishing to deploy modern infrastructure, there are important issues to assess both in the near-term and longer-term:

- are the existing data center racks deep or wide enough to handle the increased volume of power cables, or even physically support the new server form factors?
- as IT equipment is refreshed in the same set of racks, cables, power circuits, switches and breakers will need to be added, modified, and pulled—if the wrong unit is disconnected, service delivery can be materially disrupted
- what remaining capacity exists in the data center to grow in the future? In many cases, little-to-no new power and cooling have been added to the plant, but providers have simply reallocated existing resources. Long-term IT scalability can thus be severely limited, presenting risk to your IT roadmap and company's growth

- where will your new and future infrastructure be physically located? Typically legacy data centers feature limited or no new contiguous space, so you will be forced to spread your servers across the floor in disparate racks, which decreases operational efficiency

As discussed, there are many disadvantages and limitations with housing modern equipment in legacy datacenters. This paper will now explore the impact this can have on a company's bottom line.

COST AND OPERATIONAL ADVANTAGES OF HIGH-DENSITY DATA CENTERS

This section will provide IT managers with a simple model to contrast the cost of low density colocation with that of purpose built high density facilities.

The High-Density Data Center

Just as there have been major advances in processing capabilities and performance at the server level, there have been significant advances in core data center infrastructure. New data centers are able to take advantage of improvements in modern electrical and mechanical distribution systems to significantly increase energy efficiency, power density and cooling capacities per square foot—all of which allow for material reductions in operating costs for both data center operators and their customers. Power densities of 20 kW per rack are now achievable and critical to a company's IT success.

Legacy data centers typically feature dispersed, air-based cooling systems which result in the inefficient use of power and delivery of cooling. For a data center operator, a key statistic is called PUE, or Power Usage Effectiveness. The statistic measures the overall energy efficiency of the data center and is a ratio determined by dividing the amount of power entering a data center by the power used to run the computer infrastructure within it. A ratio of 2.0 would indicate that for every 2 watts coming in to the utility meter, only 1 watt is delivered out to the IT load on the data center floor, and the other watt is used for cooling and other infrastructure overhead. According to the Uptime Institute, in a 2011 survey of 500 data center operators, the average PUE was 1.8. However, legacy data centers feature PUEs greater than 2.0, meaning that much of the power is lost to servicing the building and not being delivered to the customer racks. The end result is that the data center operator is forced to charge more per kW of power delivered to the customer, in order to compensate for its high expense of power distribution and cooling.

Modern data centers, in contrast, can now be built to feature PUEs in the 1.2 – 1.4 range (some companies such as Google, Facebook, and Microsoft have reported PUEs in their data centers in the 1.07 – 1.2 range, but this can be misleading because those data centers are custom-built for their own in-house applications and infrastructures). With this improvement in operating efficiency, modern data center providers can operate with lower power costs and therefore reduce expenses for their customers. They can also deliver much higher power density per rack across their entire data center floor, which we will discuss next.

Case Study: How High-Density Data Centers Deliver Reduced TCO

As PUEs are reduced, the ability to deliver more cooling to the data center rack increases. This drives power density per cabinet considerably, and enables many operational and cost advantages. In short, the higher the power density, the lower the TCO for a customer.

To illustrate this, we have provided a case study for a typical retail colocation customer featuring 60 kW of aggregate load per month.

First, let's assume some typical cost and infrastructure deployment figures for the end user:

- monthly power costs \$300 per kW
- monthly cabinet rental costs \$600
- monthly cross-connect costs \$250
- Layer 3 network switch costs \$5,000

Second, let's assume that a typical legacy data center can provide 4 kW per rack of power density, while a high-density operator can provide 20 kW per rack.

Now let's look at a representative deployment for both of these data centers:

Low-Density Data Center Deployment

- 60 kW total load @ 4 kW per cabinet = 15 cabinets required
- Two Layer 3 network switches to manage the network of cabinets

High-Density Data Center Deployment

- 60 kW total load @ 12 kW per cabinet = 5 cabinets required (this also allows for 8kW of available expansion space per cabinet)
- Two Layer 3 network switches to manage the 4 cabinets

At 60kW the TCO difference between the two scenarios is significant as will be shown.

Example 1 - 60 kW Implementation

In this example, let's compare the two deployments over a 24-month contract:

Low-Density Data Center TCO

One-time:		\$10,000 in top-rack switches
Monthly:		\$18,000 in power costs
		\$9,000 in cabinet rentals
		\$500 in cross-connect costs
Total Monthly Cost	=	\$27,500
Aggregate 24-month TCO	=	\$670,000

High-Density Data Center TCO

One-time:		\$10,000 in top-rack switches
Monthly:		\$18,000 in power costs
		\$3,000 in cabinet rentals
		\$500 in cross-connect costs
Total Monthly Cost	=	\$21,500
Aggregate 24-month TCO	=	\$526,000

Using the same infrastructure, a 60 kW customer in a high-density colocation facility can save \$144,000, or 22% in total TCO, as compared to a legacy low-density facility. The savings are material!

Furthermore, for growing companies there are additional cost and operational advantages in using high-density racks. Let's assume that:

- our customer doubles in size, growing to 120 kW per month of electrical load
- in the high-density scenario the customer now uses 17kW of power per rack

Now we'll see how this new scenario continues to provide greater savings.

Example 2 - 120 kW Implementation

As before we'll compare the two deployments over a 24-month contract:

Low-Density Data Center TCO

One-time:		\$20,000 in top-rack switches
Monthly:		\$36,000 in power costs
		\$18,000 in cabinet rentals
		\$1000 in cross-connect costs
Total Monthly Cost	=	\$55,000
Aggregate 24-month TCO	=	\$1,340,000

High-Density Data Center TCO

One-time:		\$15,000 in top-rack switches
Monthly:		\$36,000 in power costs
		\$4,200 in cabinet rentals
		\$750 in cross-connect costs
Total Monthly Cost	=	\$40,950
Aggregate 24-month TCO	=	\$997,800

Using the same monthly cost assumptions as in the first phase, the total cost savings increase to \$342,200 over 24 months - a 25% reduction in TCO.

This is because in the high-density data center, the customer only had to add 2 new cabinets despite doubling its power draw. The 8kW of initial reserve capacity per cabinet allowed the customer to grow vertically within its rack space for some time without additional cabinet costs. With only 2 new 17 kW cabinets required, the 60 kW of power increase was easily absorbed.

In the legacy data center, the customer now has to manage 30 cabinets as the additional 60kW translated into 15 additional cabinets. In other words, lower density causes the customer to pay for 22 unnecessary cabinets and related network infrastructure, as well as the management time and hassle that the distributed cabinet layout requires.

Notes on method and assumptions: The analysis assumes that a low power density deployment will require 30% more switch infrastructure due to the physical distribution of equipment. Additionally, it assumes one cross connect per switch.

Other Advantages of High-Density Data Centers

In addition to the meaningful reduction in TCO, a high-density data center deployment yields several other important advantages in terms of day-to-day management and operating efficiency:

- More streamlined, simplified cable management: with many fewer cabinets, cross-connects, and top-of-rack switches deployed, the ability to physically manage cabling and connections is much improved, leading to operational stability
- Smaller, contiguous footprints to manage: with fewer cabinets, companies can manage their infrastructure in a more compact layout, saving time and the anguish of trying to support multiple nodes in real-time all over the data center floor in disparate locations
- Facility consolidation: with fewer cabinets the need to have multiple data center facilities can be reduced. In colocation scenarios this means more streamlined vendor management, from account management to legal administrative support
- Efficient growth: IT managers can plan to add more server chassis to cabinets with less need to run additional power circuits because of the high-power capacity circuit initially provisioned

This case study provides factors to consider when evaluating legacy and higher-density data centers. The step-by-step evaluation process should be helpful in any due diligence exercise.

CONCLUSION

Modern business needs drive new innovative applications that require increasing quantities of more powerful hardware. Though hardware efficiencies are improving, the total usage per rack unit is increasing at a faster rate, which causes aggregate power consumption to rise. Legacy data centers cannot accommodate this trend. High-density facilities offer a better total cost of ownership and help future-proof businesses.

IT executives, you must incorporate data center power density into your thinking. It is quickly becoming the defining variable in IT infrastructure management and success.

Peter Harrison, CTO and Co-Founder

December 2013



About Colovore

Colovore is the Bay Area's leading provider of high-density colocation solutions. Located in Santa Clara, CA, Colovore's modern data center features wall-to-wall power densities of 20 kW per rack, leading to significant reductions in TCO for our customers while also ensuring long-term IT scalability. Our facility features 9 MW of power, is LEED-Platinum certified, and operates with a PUE below 1.3. In addition to colocation space, we also feature a range of managed services to make the daily life of an IT professional better, including Ethernet connectivity, managed firewalls and VPNs, rack-and-stack, and cable management services.

Consolidate, Converge, Colovore.

www.colovore.com