Shared Infrastructure: Scale-Out Advantages and Effects on TCO

A Dell™ Technical White Paper

Dell | CIO Insight Series

The power to do more

Table of Contents

Introduction

The massive data needs of our connected age are now moving toward even greater density and efficiency in medium to large data centers. Applications like High Performance Computing (HPC), Web 2.0, online gaming, and cloud technologies are beginning to overlap and require similar features in the data center: high density, serviceability, and power efficiency. The first of these applications, HPC, has been moving from large, monolithic supercomputers to clustered environments that use commodity hardware. This trend is causing many hardware providers to offer system-optimized hardware for scale-out installations. Like the HPC community, administrators of Web 2.0, online gaming, and cloud data centers have found the economics of scale-out data centers using commodity hardware to be very attractive and overall more compatible with their usage models. Regardless of the specific needs of the application — whether computing efficiency, I/O performance, or large memory support — new and open source technologies to address networking fabric limitations, distributed file systems, and distributed compute management have made the implementation of these deployments both easier and more cost effective.
The growth of scale-out deployments has, however, come with an entirely new set of interrelated challenges. Each of these challenges has a direct effect on the total cost of ownership (TCO) of the data center. The largest of these are power, cooling, and physical space. Depending on the ownership model employed in the data center, each of these challenges will affect the cost in different ways. To address each challenge and offer systems that are affordable and efficient, the concept of shared infrastructure has been introduced.

What is Shared Infrastructure?

The concept of shared infrastructure has been used widely in the industry with products like blade servers. Recently, however, this technology has been extended to other uses. There are now servers that share some of the same resources that blade servers use (such as fans and power supplies), but forgo the integration of others (such as network fabric switching). Each of these shared infrastructure implementations has very specific pros and cons based on the specific usage model being adopted in the data center. Those considerations, which are all based on a customer's specific needs, are outside the scope of this paper.

For the purposes of this paper, a shared infrastructure system is defined as any system that pools resources in a single chassis and spreads them among independent server nodes within that same chassis to achieve optimized performance. Dell's latest addition to the shared-infrastructure landscape is the PowerEdgeTM C6100, a 4-node-in-2U server that shares chassis, power, and cooling with the other nodes in the server and nothing else. Systems like this have the potential of greatly reducing TCO for large clustered deployments.

TCO Advantages of Shared Infrastructure Computing

The advantages to total cost of ownership on a shared infrastructure system are focused in three main areas: power efficiency, system scaling efficiency, and

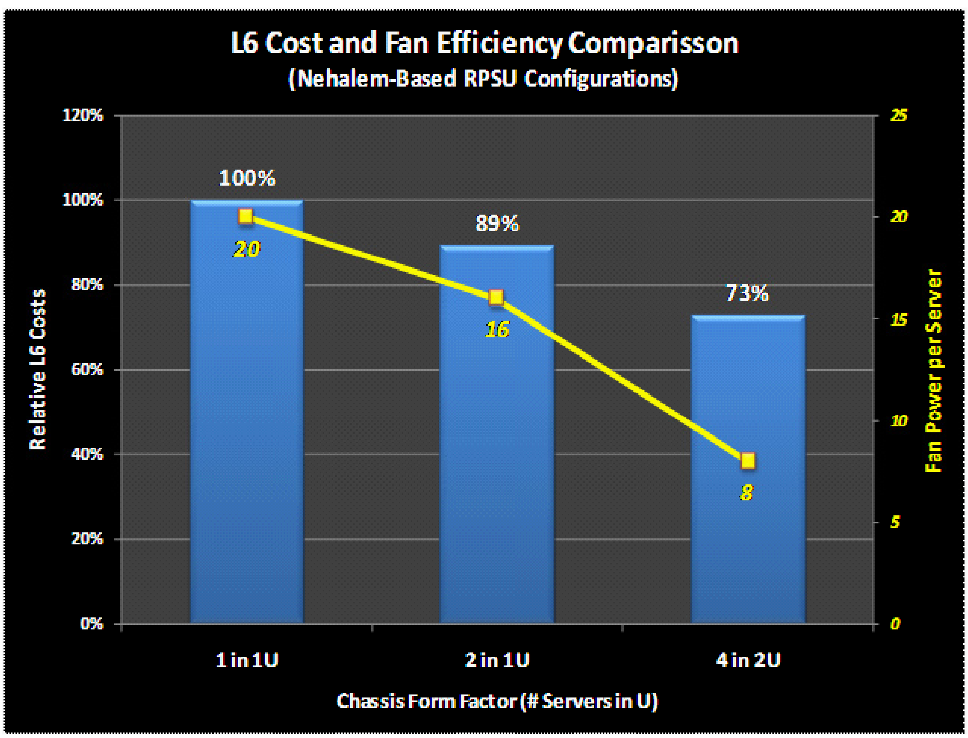

Figure 1: L6 cost and fan efficiency comparison between chassis form factors

compute density. Each of these features is not only focused on one aspect of owning and maintaining the data center, but rather focused on driving down TCO.

Most modern cluster-focused systems have been optimized to draw as little power as possible at the node level, but shared infrastructure systems have the advantage of sharing powered devices across multiple nodes. Firstly, these systems can be effectively cooled by fewer larger fans. This is due to a mix of greater thermal properties in a larger chassis as well as the higher efficiency of larger fans. In studies completed in Dell's labs with various shared infrastructure systems, a 60 percent decrease in fan power per server has been achieved when moving from a 1-node-in-1U configuration to a 4-node-in-2U configuration. (Note: L6 refers to hardware that includes only chassis, motherboard, and power supplies.)

In addition to the efficiency of the fans, other shared components, such as power supplies, will lead to greater efficiency in the shared infrastructure system. Not only are the larger high-efficiency power supplies more cost effective than those of smaller form-factor systems, but the fact that they are shared among more nodes allows them to maintain higher average utilization than a power supply in a single system.

Most power supplies are inefficient at very low utilization, so maintaining a higher average utilization will help the power supply to stay in a more efficient range of the operating curve. These factors all combine to provide lower per-node average electricity draw, as shown in Figure 2, when comparing a reference 4-in-2U system to a 1-in-1U system in Dell's labs. The results are an average power savings of 11 percent per node over a 1-in-1U system at the same performance level.

In traditional data center environments, each component of the server scales directly with the number of nodes. In shared infrastructure, however, as you add more nodes per chassis, you avoid the need to repurchase power supplies, power distribution logic, chassis infrastructure/components, and fans for cooling. As we refer back to Figure 1, we see a 27 percent improvement in per-node cost, based on acquisition cost in Dell reference systems.

In addition, there are a number of advantages that lie outside of traditional TCO calculations. When the entire lifetime of the solution is considered, the shared infrastructure system drives further savings. These manifest themselves as savings in network switching infrastructure, integration services, cabling, and shipping. If the entire solution is planned carefully, the shared infrastructure system can allow greater utilization of the ports for rack switches. In some situations the user can avoid running cables between multiple racks or purchasing more switching gear than required if more systems are consolidated in fewer racks, as they can be with shared infrastructure.

During integration, the ability to consolidate nodes into a single box cuts down the number of racks to integrate as well as the number physical units to load into the rack. In addition, this tighter integration cuts down on both number of cables and cable lengths. Even in shared infrastructure systems that only share power, chassis, and fans, reduced power cabling can result in savings.

Finally, whether the systems are shipped to the customer fully integrated into racks or shipped as single units, fewer racks and higher density means fewer boxes and less shipping weight. As a total solution, all of these savings can not only accumulate into great financial savings, but also great savings to

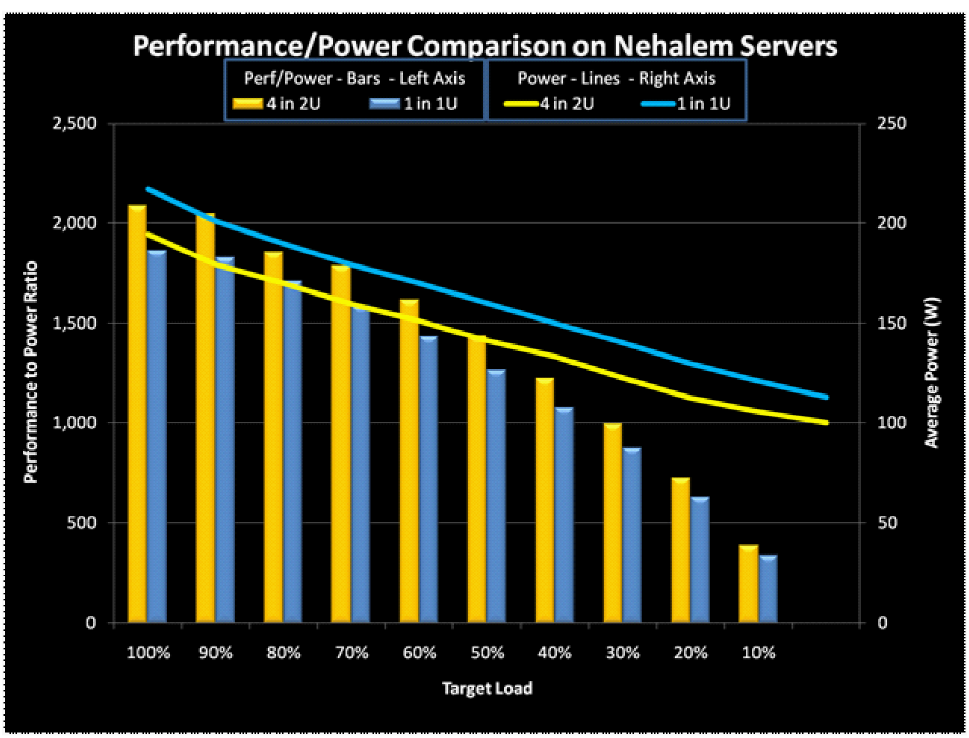

Figure 2: Shared Infrastructure Performance/Power Comparison

the environment through reduced metal, shipping material, fossil fuels, and carbon emissions.

Data Center Economic Models

The advantages discussed above manifest themselves in different ways depending on the economic properties of each administrator's data center. Though there are many models for owning and maintaining a data center, the two main models discussed in this paper are data center ownership and co-location. Both share the same challenges, but the priority and effect of each of the data center's challenges are different between these two models.

The first, and more traditional, model is to own the entire data center. This model is still predominant in the HPC arena, where universities, government agencies and large corporations have their own dedicated data center space and budget. In addition, large Fortune 500 corporations looking to create their own private clouds or public cloud services sometimes opt for a wholly owned data center. In this model, space is limited to the floor space the company has available for data center use, all power is metered for both the servers and for cooling, and server management/maintenance is handled by in-house personnel.

The second model of co-location has been growing and is a more cost-efficient model for Web 2.0 and public cloud startups, as well as smaller corporations in need of private clouds. In a co-located data center model, the owner of the computer equipment leases data center space from a third party. The lessee is charged for a circuit capped at a certain wattage/current draw, a fixed number of racks, cages to physically secure hardware, and data uplinks to outside networks. Many co-located data centers will charge differently for different combinations of the items listed here, but the overall concept is that each additional circuit or rack causes greater overall cost. Therefore, in this model many of the same challenges exist as in the owned data center model but affect the overall cost in subtly different ways. Both models are still in wide use, and this paper's aim is to show how a shared infrastructure system can solve the problems of both models.

TCO Advantages of Shared Infrastructure for Owned Data Centers

For users who own a data center, the issues addressed with shared infrastructure are overall power draw, cooling, physical space, and server maintenance. The efficiency, scaling, and density features of a shared infrastructure system combine to minimize operational cost in this environment. Based on basic TCO models, the savings in an owned data center can be significant when moving from a traditional server architecture to a shared infrastructure model. When comparing the total cost of a 500-node data center that uses standard 1U compute nodes versus 4-node-in-2U shared infrastructure systems like the PowerEdge C6100, the TCO savings are significant. The results, as displayed in Figure 3, show a total savings of 19.8 percent to the cost of running the data center.

The results of the TCO comparison show first that the higher power efficiency of the shared infrastructure system drives down the cost to power the systems. The roughly 19 percent power savings drives a 19 percent reduction in power cost. This power savings is also reflected in the power-to-cool cost. With an assumed power usage effectiveness (PUE) of 1.8, the higher efficiency of the 4-in-2U system drives a 19 percent savings in Power to Cool Cost as well.

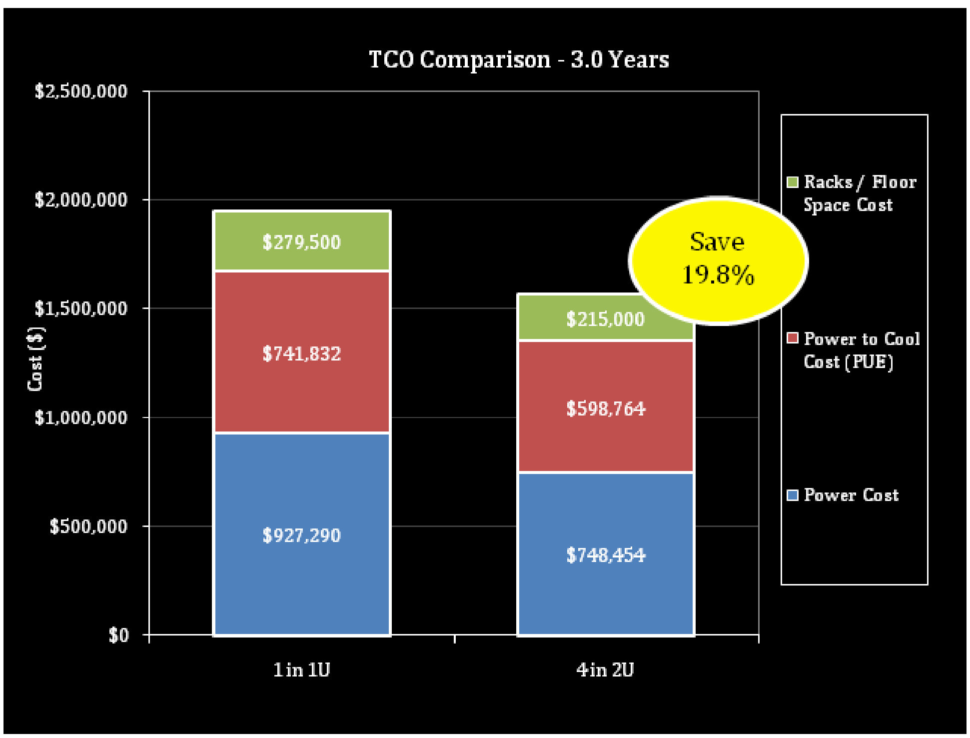

Figure 3: Owned Data Center TCO Comparison

The largest savings to TCO, however, is due to the reduced number of racks needed with the shared infrastructure system. When assuming a 15kW power maximum per rack and 40 usable U of rack space, the 1-in-1U system can fill the 40U of rack space in each rack, while the 4-in-2U system fills 26U to stay within the power budget. However, the 40U of 1-in-1U servers per rack includes only 40 nodes, while the 26U of 4-in-2U servers includes 104 nodes. That is a 160 percent improvement in nodes per rack compared to the 1-in-1U system. Therefore, to complete the 500-node data center, we only need 10 racks for the shared infrastructure systems compared to 13 for the 1-in-1U systems. This reduces both the overall cost racks, as well as the cost of floor space within the data center, resulting in a 92 percent reduction in the cost of racks and floor space.

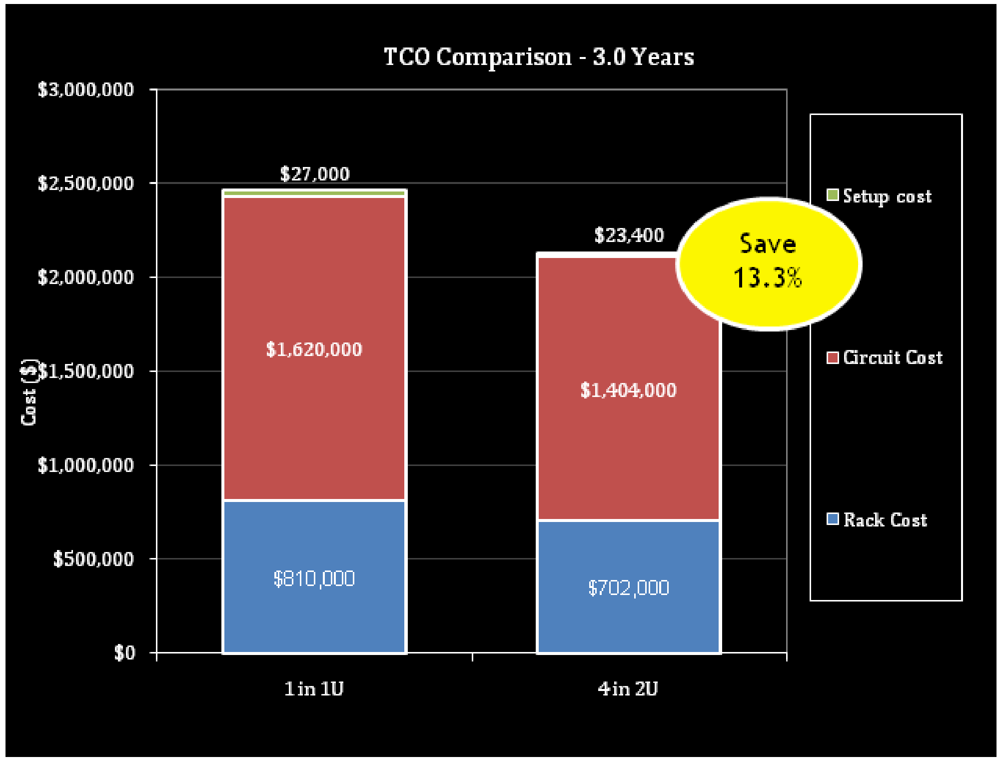

Figure 4: Co-located Data Center TCO Model — with Power Options

A

TCO Advantages of Shared Infrastructure for Co-located Data Centers

Though the co-located data center spreads out costs differently, the advantages of shared infrastructure still can have a major effect on overall cost in a co-located data center environment. This can be best analyzed by looking at the different items paid for in a co-located environment and how each is addressed by moving to a shared-infrastructure system rather than a traditional server system.

The main concerns in most co-located data center environments are the number of racks and the size of the power circuit per rack. Some providers give tenants a broad array of circuit options to meet their rack power needs for added cost, but some providers will limit all racks to a fairly low overall power draw. Therefore, if the servers are not optimized for power, more racks may be purchased than are actually necessary, because filling the rack would exceed the circuit's power limit. This problem can be easily addressed with the power efficiency available in a shared-infrastructure system.

shared-infrastructure system allows more nodes to fit within a rack's specific power footprint as well as potentially reducing the size of the power circuit required on each of the racks. In a co-located arrangement where there are various power circuit options, the

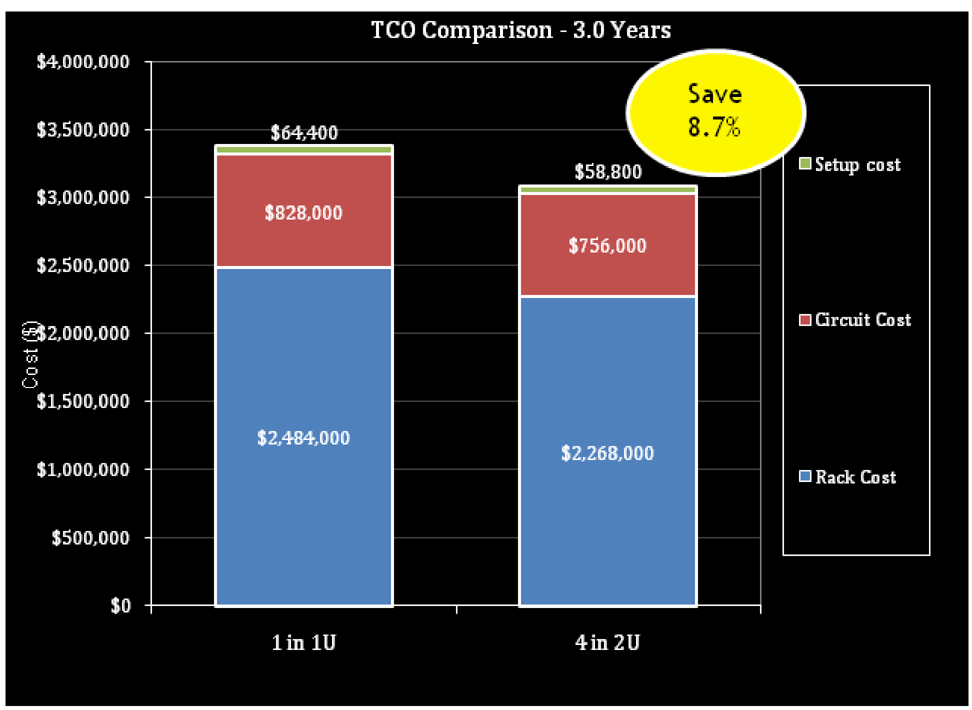

Figure 5: Co-located Data Center TCO Model — Only 3kW Circuit

density of a shared infrastructure system compared with a traditional 1-in-1U system will reduce the number of total racks required. In estimates using Dell development systems, this can result in savings of more than 15 percent, shown in Figure 4. The TCO model in Figure 4 assumes various power circuit options from 3kW up to 12.5kW, which allows the 1-in-1U system to reach 35 nodes per rack (assuming 40 usable U per rack), while the 4-in-2U system reaches 40 servers per rack in 20U. This increase in density reduces the number of racks and circuits required to setup and monitor monthly.

Some data centers, like those in Asia, do not provide options for circuit sizes and the available circuits are quite small. The output of the TCO model in Figure 5 shows that even in this scenario, the increased efficiency of the shared infrastructure system still reduces the total cost of ownership by more than 8 percent. In this scenario the high energy efficiency of the shared infrastructure system allows 12 systems per rack compared with 11 for the 1-in-1U configuration. This seemingly small difference adds up to great savings in a 500-node application.

Best Practices and Learnings with the New Breed of Shared Infrastructure Systems

Though systems that integrate multiple independent nodes into a single chassis without integrating networking and management, such as the PowerEdge C6100, are just gaining momentum in the mass data center market, Dell has extensive experience in both building and integrating these systems via our Data Center Solutions division. The currently shipping PowerEdge C6100 follows years of custom design and improvements based on customer needs across tens of thousands of installed systems. The result is a third-generation system available now.

The learnings from years of installing shared infrastructure systems have led to improved serviceability, improved efficiency, and higher density. The technology currently available in the PowerEdge C6100 allows serviceability at a single unit level in a 4-node-in-2U system. Earlier generations of systems may require bringing down additional units and multiple tools to repair a downed node, but our third-generation design allows the user to service a single node with a single spring latch. This incremental improvement across generations has yielded a focused feature set and fine-tuned components that drive greater power efficiency and maximum density.

Though the majority of the content in this paper focuses on the advantages of shared infrastructure, the decision to implement a shared infrastructure system in the data center should be a holistic one. While there are some savings that almost any implementation can reap, the true savings come when the entire solution is based on full consideration of every factor. Factors like the power needs of each node, the expected failure domain of servers, number of disks per node, type of disk per node, and number of nodes per rack all must be carefully considered, or the user can run the risk of negating the advantages of the shared infrastructure system.

Ultimately, any cost/benefit analysis must be done at the complete integrated and delivered solution level to realize the full benefits. As a rule of thumb, the questions you should ask are:

- What is my power budget per node?
- How homogeneous are my hardware needs per node?
- How much storage do I need per node?
- How fine is the level of control I need over the power of each node?
- What type of switching infrastructure do I need?

The answers to these questions may very well lead you to the great savings of shared infrastructure.

Conclusions

In the sections above, we reviewed key points about shared-infrastructure, its advantages, and its effects on different data center cost models. Regardless of an organization's data center environment, the reduced cost, power efficiency, and server density drive overall cost down. Though these savings are generally favorable for all applications, the areas where they really shine are scale-out environments. In scale-out environments, the many advantages of shared-infrastructure can be spread over hundreds to thousands of servers, yielding significant TCO savings.

To learn more, please visit Dell.com/PowerEdgeC

Appendix A. Data Center TCO Model Assumptions

The tables below contain the assumptions used for the TCO models presented in this whitepaper.

Table 1: Global Assumptions for Owned DC (Figure 3)

Global Assumptions for Owned DC	
Number of Servers	500 nodes
Useful Life	36 months
kW Available per Rack	15 kW
Energy Cost per kWH	0.20 per kWH
PUE Ratio	1.8
Cost of Space	$500 per Sqft
Space Requirement per Rack	14 Sqft per Rack
1 in 1U Configuration	
Form Factor	1U
Nodes per System	1
Processor	2 x Xeon® X5650
Chipset	Intel® 5500
Memory	6 x4 GB 1333MHz DDR3
HDD	1x160GB, 7.2K RPM, 3.5" SATA
Add-In Card	QDR IB Daughtercard
Wattage Required	352.5W
Workload	High Performance Linpack
1 in 1U Configuration	
Form Factor	2U
Nodes per System	4
Processor	2 x Xeon® X5650
Chipset	Intel® 5500
Memory	6 x4 GB 1333MHz DDR3
HDD	1x1TB, 7.2K RPM, 3.5" SATA
Add-In Card	QDR IB Daughtercard
Wattage Required	284.8 W per each node
Workload	High Performance Linpack

Table 2: Global Assumptions for Co-located Data Center with Power Options (Figure 4)

Global Assumptions for Owned DC	
Number of Servers	500 nodes
Useful Life	36 months
Cage Cost (monthly/rack)	$1,500
Cage Cost (setup/rack)	$1,000
PUE Ratio	1.8
Cost of Space	$500 per Sqft
Space Requirement per Rack	14 Sqft per Rack
Available U per Rack	40
Circuit wattage per rack (kW)	12.5
Circuit Cost (monthly)	$3,000
Circuit Cost (setup)	$800
1 in 1U Configuration	
Form Factor	1U
Nodes per System	1
Processor	2 x Xeon® X5650
Chipset	Intel® 5500
Memory	6 x4 GB 1333MHz DDR3
HDD	1x160GB, 7.2K RPM, 3.5" SATA
Add-In Card	QDR IB Daughtercard
Wattage Required	352.5W
Workload	High Performance Linpack
4 in 2U Configuration	
Form Factor	2U
Nodes per System	4
Processor	2 x Xeon® X5650
Chipset	Intel® 5500
Memory	6 x4 GB 1333MHz DDR3
HDD	1x1TB, 7.2K RPM, 3.5" SATA
Add-In Card	QDR IB Daughtercard
Wattage Required	284.8 W per each node
Workload	High Performance Linpack

Table 3: Global Assumptions for Co-Located Data Center with Only 3kW Circuit (Figure 5)

Global Assumptions for Owned DC	
Number of Servers	500 nodes
Useful Life	36 months
Cage Cost (monthly/rack)	$1,500
Cage Cost (setup/rack)	$1,000
PUE Ratio	1.8
Cost of Space	$500 per Sqft
Space Requirement per Rack	14 Sqft per Rack
Available U per Rack	42
Circuit Wattage per rack (kW)	3.1
Circuit Cost (monthly)	$500
Circuit Cost (setup)	$400
1 in 1U Configuration	
Form Factor	1U
Nodes per System	1
Processor	2 x Xeon® X5650
Chipset	Intel® 5520
Memory	6 x4 GB 1333MHz DDR3
HDD	2 x 3.5" SATA
Wattage Required	1064W for 4 servers; 266W for each server
Workload	SPECpower @ 70% load
4 in 2U Configuration	
Form Factor	2U
Nodes per System	4
Processor	2 x Xeon® X5650
Chipset	Intel® 5500
Memory	6 x4 GB 1333MHz DDR3
HDD	2 x 3.5" SATA
Wattage Required	943W for 4 nodes; 235.75 W for 1 node
Workload	SPECpower @ 70% load