

Strategies for Deploying Blade Servers in Existing Data Centers

By Neil Rasmussen

White Paper #125



Executive Summary

When blade servers are densely packed, they can exceed the power and cooling capacities of almost all traditional data centers. This paper explains how to evaluate the options and select the best power and cooling approach for a successful and predictable blade deployment.

Introduction

Blade servers offer significant advantage over traditional servers – improving the processing ability while consuming less power per server. However, with their smaller footprint, blades can be much more densely packed, resulting in racks that can require up to 20 times the electrical power and generate up to 20 times the heat, compared with what the average traditional data center was designed for. This can stress the capabilities of existing power and cooling systems. To effectively deploy blade servers, the power and cooling infrastructure in a data center must be upgraded, or the blade server loads must be spread out over multiple racks.

There are a number of strategies that can be used to deploy blade servers. This paper provides guidelines for determining the appropriate power and cooling strategy, based on the needs and constraints of a specific installation.

The Core Challenge

The core challenge related to blade server installation for most existing data centers is related to power and cooling **distribution**. Most data centers have raw power and cooling capacity but do not have the infrastructure to deliver this capacity to a high density area. Unfortunately, many users do not even understand they have this problem until they attempt deployment. This happens because virtually no data centers are documented or instrumented to provide the operators with information regarding the density capability of the data center in a particular area of the facility. The technical reasons for these problems are described in detail in the white papers and application notes referenced at the end of the paper, but are summarized here:

Insufficient airflow: Blade servers require approximately 120 cfm of cool air per kW of power rating.

Most traditional data centers only provide 200-300 cfm of air per rack location, which **is 10 times less air than fully populated rack of blade servers** and limits the average rack power to below 2kW. A blade server which does not receive enough cool air will end up ingesting its own hot output air and overheat. **This is by far the biggest challenge in blade deployments.**

Insufficient power distribution: Blade servers draw much more power than typical data center power distribution systems were designed for. This problem shows up in three forms: 1) Insufficient number and / or wrong type of power wiring installed under floor or overhead, 2) Insufficient nearby Power Distribution Unit (PDU) capacity, and 3) Insufficient quantity of breaker positions. Any of these problems can prevent the ability to deliver high density electrical power.

Note that of the two key problems described, the cooling distribution problem is the main constraint. For this reason, the primary focus of this paper is on selecting a cooling architecture. The power architecture will follow from the selected cooling architecture and is dependent on the specific brand of blade server. For specific details consult references 2, 5, 6, and 7 listed at the end of this paper.

The 5 Different Methods for Deploying Blade Servers

There are 5 basic approaches to cooling blade servers. Once an approach is chosen, there are a variety of different products and techniques that can be used to implement it. These approaches are described in detail in APC White Paper 46, "Cooling Strategies for Ultra-High Density Racks and Blade Servers" and summarized in **Table 1**.

Table 1 – Application of the five approaches to cooling high density enclosures

Approach	Advantages	Disadvantages	Application
1 Spread the load Split equipment among enclosures to keep peak load down	Works anywhere, no planning needed Essentially free in many cases	High density equipment must be spread out even more than approach 2 Uses more floor space Can cause data cabling issues	Existing data centers, when high density equipment is a small fraction of the total load
2 Borrowed cooling Provide average cooling capability with rules to allow borrowing of underutilized capacity	No new equipment needed Essentially free in many cases	Limited to about 2X the design power density Uses more floor space Requires enforcement of complex rules	Existing data centers, when high density equipment is a small fraction of the total load
3 Supplemental cooling Provide average cooling capability with provision for supplemental cooling equipment	High density where needed and when needed Deferred capital costs High efficiency Good floor space utilization	Limited to about 10 kW per enclosure Racks and room must be designed in advance to support this approach	New construction or renovations Mixed environment Location of high density equipment is not known in advance
4 High density area Create a special high density row or zone within the data center	Maximum density Optimal floor space utilization High density equipment does not need to be spread out High efficiency	Need to plan a high density area in advance, or reserve space for it Must segregate high density equipment	Density 10-25 kW per rack When there is a requirement to co-locate high density devices New construction or renovations
5 Whole Room Provide high density cooling capability to every rack	Handles all future scenarios	Extreme capital and operating costs of up to 4X alternative methods May result in extreme underutilization of expensive infrastructure	Rare and extreme cases of large farms of high density equipment with very limited physical space

To deploy blade servers, one of the methods must be selected. The selection is based on constraints of the current installation as well as the needs and preferences of the user.

The Blade Deployment Process

The process of preparing the physical environment to support the deployment of blades includes the following key elements:

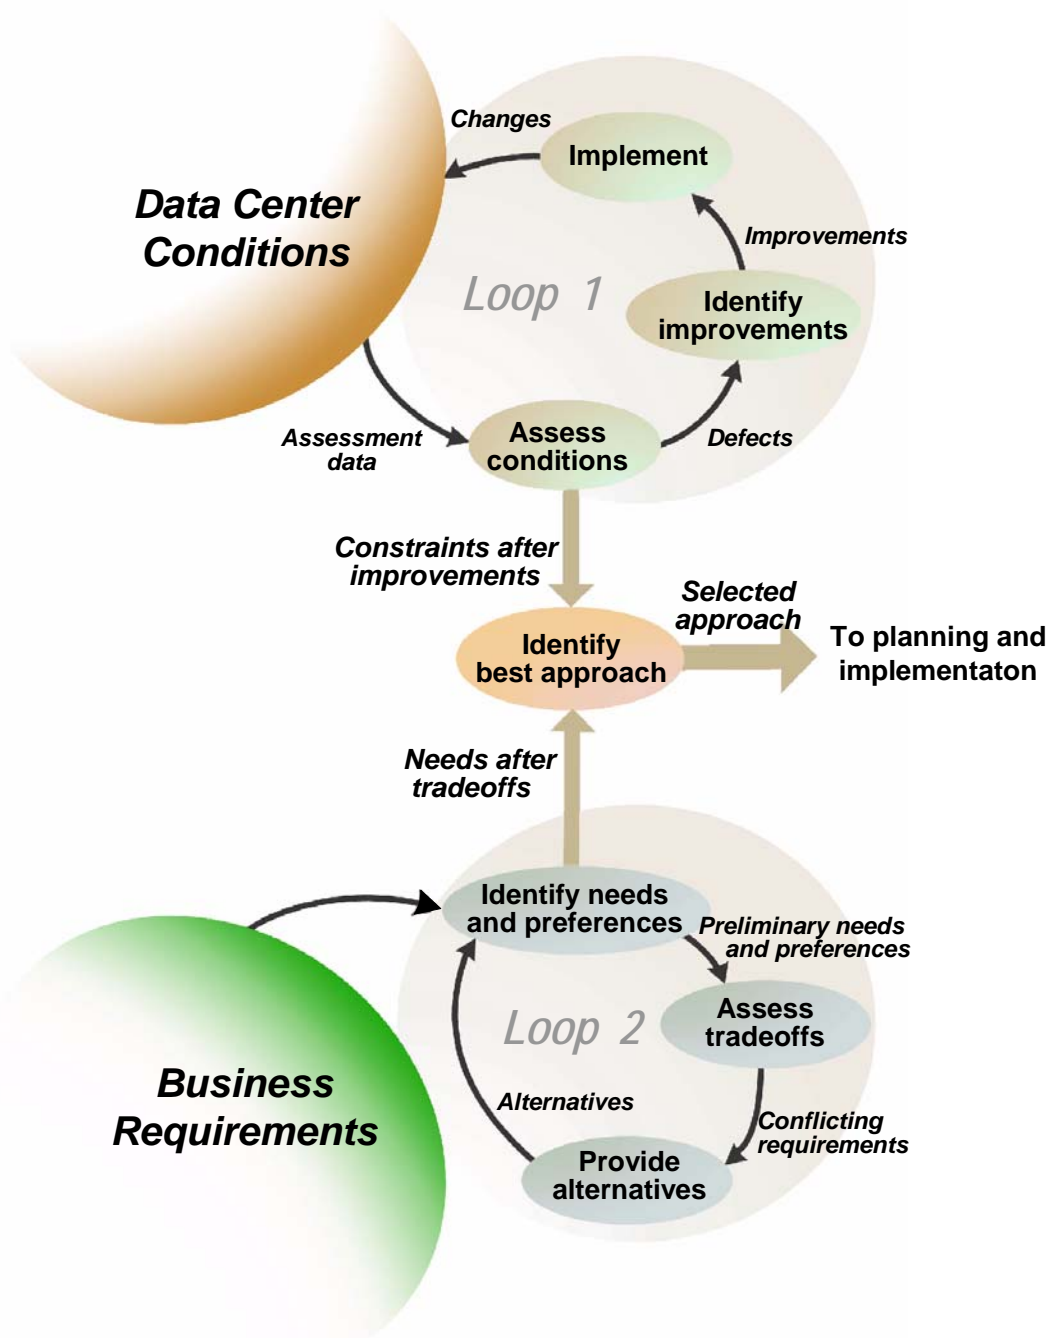
- Identifying the constraints of the existing facility
- Identifying user needs and preferences
- Determining the appropriate design approach for power and cooling
- Designing and subsequently implementing the design

A map for this process is provided in **Figure 1**. The figure is a process flow map showing the various process steps and the resulting data at each step. The process includes two key loops at the front end, where the constraints and the user needs and preferences are determined via iteration. This is essential in order to allow proper adjustments and tradeoffs to be made. Typically the initial constraints and preferences change after a review of the situation and the associated tradeoffs. In the most common example, the preference or requirement to densely pack blades is often relaxed when the consequences of this approach are fully understood. This analysis occurs in Loop 2 in the process map.

Another common situation is where an assessment of the current installation identifies problems which can be easily corrected and increases the ability of the data center to handle blade power and cooling requirements. These adjustments occur in Loop 1 of the process map.

In the next sections, the various processes that contribute to the selection of a design approach are discussed in more detail.

Figure 1 – Process map for determining the appropriate deployment method for blade servers in an existing data center



Identifying the Constraints of the Existing Facility

Existing data centers have various hard constraints that cannot be changed. These constraints are as follows:

Precision power capacity. The data center may not have sufficient excess UPS capacity to power a proposed blade server installation.

Precision cooling capacity. The data center may not have sufficient excess precision cooling capacity to cool a proposed blade server installation. This limitation refers to the raw capacity of the computer room air conditioners, and not the air distribution system.

Floor space limits. The total floor space in the data center may be constrained, or the floor space available for the blade deployment may be constrained. If severe enough, these constraints may force certain design approaches.

No ceiling plenum. The room may not use or have a ceiling return air plenum. The room may be height constrained so that no ceiling plenum is possible. This constraint may eliminate some design options.

Raised floor restrictions. The existing raised floor, if any, may be less than 2 feet in height and / or be partially filled with wires or piping. This may constrain the air distribution capability of the raised floor which may preclude some design options.

Weight restrictions. The data center floor may have floor loading limitations, particularly when raised floors exist. This may preclude some design options.

Frequently, the constraints in an existing data center are not documented and are not obvious, and an assessment of the conditions must be done.


Assessment of existing conditions

An assessment of the existing conditions in the data center is essential to blade deployments. This assessment may be superficial if the number of blade servers is on the order of one rack of blades or less. However, for deployments above this number, the depth and detail of the assessment must increase substantially.

During an assessment, various data on the capacity of the power and cooling systems is collected, including the nameplate capacity and, more importantly, the actual capacity as implemented. In addition, the existing load conditions must be assessed to determine the magnitude and physical distribution of the loads. Most importantly, the power and cooling distribution systems must be studied to quantify the ability of the system to deliver power and cooling to high density loads.

In some cases where the complexity of the deployment is high, it is desirable to simulate the data center using computer models, both to determine the “as-is” conditions and more importantly to provide verification of the proposed design. An example of data from such a model is shown in **Figure 2**.

Figure 2 – 3D Computational Fluid Dynamics (CFD) model of a data center showing airflows and temperatures (supplied by APC Professional Services)



It is advisable for all data center operators to have a rudimentary knowledge of assessing a data center. For complicated, high cost, or high risk installations, it is recommended that specialists be used to perform these assessments. APC and other vendors provide professional data center assessment services.

Identifying improvements – Basic data center hygiene

The existing conditions of a data center often include a number of weaknesses that should be identified and corrected before any further steps are taken, because they can affect the data which serve as a foundation for the blade server deployment. These problems can include:

- Lack of use of blanking panels
- Leaks in the raised floor or air supply system
- Improper configuration of air returns

- Improper configuration of vented floor tiles
- Unused under-floor wiring that can be removed
- Improper set-points on air conditioners

A more detailed explanation of these issues is provided in APC White paper #42, "Ten Cooling Solutions to Support High Density Server Deployment" and #49, "Avoidable Mistakes that Compromise Cooling Performance in Data Centers and Network Rooms".

Identifying User Needs and Preferences

In addition to hard physical constraints on the facility, customers often have soft constraints or preferences. These constraints may be absolute, or they may be relaxed if the cost to comply with them is too high. These needs or preferences may preclude some blade deployment options and suggest others. These needs include:

Uninterrupted operation. The most important need may be that the installation be minimally intrusive on existing data center operations, with minimal risk to the operating IT equipment. For example, no available scheduled down time may be possible.

High availability of resulting system. The most important need may be that the resulting system be of the highest possible availability. This would require that power and cooling systems be redundant, and that the system be tested to ensure redundancy.

Co-location of servers (densely packed). There may be a very strong desire or requirement to pack the blades at the maximum possible density. Some of these reasons for this include:

- System is a showcase / demonstration system
- Desire to conserve floor space
- Regulatory or legal requirement to keep all servers in one small location
- Simplify data cabling
- Desire to have a logical grouping of IT equipment (i.e. all web servers co-located)
- Different owners for different areas of the data center
- Simplifies administration of equipment (i.e. upgrades)
- A perception (usually wrong) that this will save money

Note that packing at full density can be extremely costly and require intrusive construction and modification of an existing data center. It is strongly suggested that alternatives that involve spreading be considered before the decision is made to densely pack blade servers.

Prepare for follow-on deployments. This may be the first of a series of blade server deployments, in which case the current deployment should lay the foundation for future deployments, and should not preclude or interfere with future deployments.

Time. There may be a requirement to deploy blade servers rapidly. If this is the case, planning, contracting, and construction may be undesirable.

Cost. The primary preference may be to deploy blade servers at minimal cost. This provides a clear direction.

Selection of Deployment Method

Once the constraints of the existing facility are well understood and the appropriate tradeoffs have been completed against the various user needs and preferences, the selection of the deployment method from among the 5 basic methods can be completed. The deployment method is selected on the basis of cooling issues since these issues are the primary constraint on practical systems. After the deployment method is determined, the power issues are resolved.

The key variable that affects the deployment method is the density of deployment. Many customers assume or prefer that blade servers be deployed at their maximum density. This is often not an appropriate assumption when fitting blades into an existing environment. In fact most blade servers use a modular chassis structure and can be deployed at a lower density than the maximum rack density. For example the IBM BladeCenter™ consists of independent chassis that can be deployed at increments of between 1 and 6 per rack. While it may appear to reduce the benefits of blades to spread them, in fact the cost, system availability, and speed of deployment may in fact be improved by spreading, particularly when attempting to install blades in an existing environment.

Many existing environments were designed for a power density of 2 kW per rack or less. When deploying blades at 10-30 kW per rack in such an environment, the blade cabinets disproportionately consume the power and cooling infrastructure, such that the data center eventually ends up with extra space which cannot be used when the power and cooling supplies are all utilized. For this reason there typically is no real benefit to conserving space during a blade deployment in most existing data centers. It is this fact that makes it practical and cost effective to spread blades out in existing data centers. **Deploying blades at full density is typically only cost effective in new facilities specifically designed to support high density, when the size of the deployment is large, or when there is a very severe constraint on space.**

Therefore, the central decision in blade deployment is the degree to which the blade chassis are spread among racks – that is, how many blade chassis will be installed per rack. The actual brand and model of blade server chosen may limit the practical ability to spread blades; for example some blade servers use independent chassis which are easy to spread, while other blades use a backplane system which makes it impractical to spread blades except in certain deployment increments. For a more complete discussion of these issues, refer to APC Application Notes related to specific brands of blade servers. When different

blade chassis deployment densities are mapped to the five key blade deployment methods described earlier, the result is **Table 2**.

Table 2 – Map showing blade server deployment criteria for different combinations of blade chassis density and blade deployment method, indicating preferred combinations

# Chassis per rack	Spread the Load	Borrowed Cooling	Supplemental Cooling	High Density Area	Whole room
1	Most data centers can accommodate	All data centers can accommodate	All data centers can accommodate. Adjacent blade racks allowed	Not cost effective compared with alternatives	Not cost effective compared to alternatives
2	Only if data center has unusually high cooling distribution capacity	Most data centers can accommodate, use of adjacent racks may be restricted	All data centers can accommodate. Adjacent blade racks allowed	Not cost effective compared with alternatives. A higher density target should be set for new zones or rows.	Not cost effective compared with alternatives. A higher density target should be set for a whole room.
3	Impractical: power density exceeds typical data center capacity	Most data centers can accommodate, but adjacent racks are not practical in most cases	Requires hot air return plenum or ductwork. Adjacent blade racks allowed	The maximum limit for well designed raised floor cooling systems	Not cost effective compared with alternatives. A higher density target should be set for a whole room.
4	Impractical: power density exceeds typical data center capacity	Data center must have unusually high cooling distribution capacity, rules are strict	Depends on the specific combination of blade server and supplemental cooling solution	Hot air scavenging systems are needed	Hot air scavenging systems are needed. Total room rebuild required.
5	Impractical: power density exceeds typical data center capacity	Impractical: power density exceeds typical data center capacity	Impractical: power density exceeds capability of known supplemental cooling devices	Hot air scavenging systems are needed	Hot air scavenging systems are needed. Total room rebuild required.
6	Impractical: power density exceeds typical data center capacity	Impractical: power density exceeds typical data center capacity	Impractical: power density exceeds capability of known supplemental cooling devices	Only if there is a severe area limitation. The cost may be extreme to achieve this density over a sustained area. May require rules.	The cost may be extreme to achieve this density. Total room rebuild required. Hot air scavenging systems are needed.

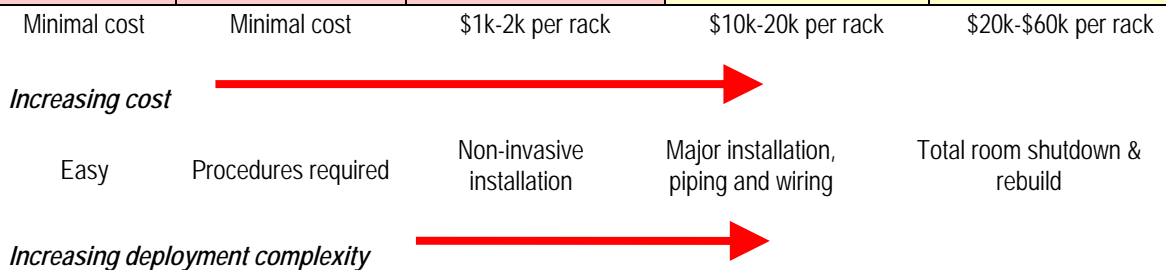


Table 2 shows that for the 30 possible combinations of six spreading density levels and 5 deployment methods, there are approximately 11 preferred combinations and another 7 marginal combinations for a total of 18 practical deployment combinations. To select the best alternative, thousands of combinations of user preferences, practical constraints, and existing condition data must map to these 18 deployment combinations. This mapping requires extensive analysis and rules and can be implemented as a software algorithm, but its full description is beyond the scope of this paper.

While developing tools to perform this analysis, APC has identified some key observations:

- If the fraction of racks of blades to be deployed is more than 25% of the total rack locations in a room, it is possible that an existing room will require a total rebuild of its power and cooling systems. This suggests that for any deployment of this magnitude, a new room be built, unless it is possible to shut down the data center for a period of time.
- For existing data centers where the deployment of 1-5 racks of blades is planned, it is attractive to spread the blades out at 25% to 50% of their full density (i.e. less than 3 chassis per rack), in order to minimize the impact on the data center operations and reduce the cost of deployment. For most data centers, the cost of achieving very high density is much greater than the space cost associated with a few extra rack locations.
- For the common case of an existing data center where the bulk cooling and power capacity exists, supplemental cooling increases the deployment density for a low cost while providing a predictable result.

Approaches not recommended

The following list is a set of approaches and actions that are routinely taken by data center operators but are flawed. These approaches do very little to help and often make matters worse.

Reducing air temperature. One of the easiest, and worst, actions a user can take is to reduce the air temperature set-point on the computer room air conditioners to attempt to solve data center hot spots. Taking this action will reduce the capacity of the air conditioners, dramatically increase humidifier water consumption, and dramatically decrease the operating efficiency of the data center (and consequently significantly increase the electrical bill). All of this will happen and it will NOT even solve the problem, since the problem is an airflow problem and NOT an air temperature problem.

Floor grates. Another apparently logical action is to replace the vented tile in a raised floor with a tile that has less air resistance. Such tiles often look like grates instead of the familiar perforated tile. This approach can help in the case of an isolated rack but has severe side effects, particularly when used in larger numbers. The use of these tiles in the typical data center will cause the airflow in other areas to decrease, but more importantly these tiles cause the significant and unpredictable variations to occur in airflow between tiles. This problem is described in more detail in APC White Paper 46, "Cooling Strategies for Ultra-High Density Racks and Blade Servers".

Top-of-rack fans. The use of fan trays installed in the top of racks is very popular, despite the fact that these fans provide no benefit in a properly designed IT rack. The problem of servers overheating is NOT due to hot air inside the rack. It is due to hot air at the air intakes of the servers which are located in the front. These fans just make more heat, and can even reduce cooling capacity in a well designed data center. Many customers specify fan trays based on old legacy specifications without an understanding of their purpose. There are some effective fan assist devices that couple to the rack; these are described in greater detail in APC White paper #42, "Ten Steps to Solving Cooling Problems Caused by High Density Server Deployment".

Isolating racks. Isolating racks away from rows in an area open on all sides is sometimes used in an attempt to reduce the density in an area and to allow more vented floor tiles to be associated with a rack. However, this approach allows hot exhaust air to return around the sides of the rack to the server intake. The overall effect is no benefit. It is much better to keep racks in a hot aisle cold aisle arrangement and use unloaded racks with blanking panels between blade racks, wider cold aisles, supplemental cooling devices, and/or hot aisle containment systems to boost performance.

Conclusion

The deployment of blade servers results in significant improvement in processing ability, however it can stress an existing data center's power and cooling systems. There are a variety of approaches to powering and cooling blade servers. The best approach for a specific installation will depend on the constraints of the existing design and the needs and preferences of the data center operator.

This paper outlines the issues and choices involved in blade server deployment. A process is described for selecting a deployment method based on constraints and needs.

Most users do not understand the cooling constraints of densely packed blade servers. When the options and their advantages are considered, deployments involving the spreading of blades will be attractive for many existing facilities because of the savings in cost and time, and the reduction in interference with data center operations.

References

- 1) **APC White Paper #46, “Cooling Strategies for Ultra-High Density Racks and Blade Servers”**
- 2) **APC White Paper #29, “Rack Powering Options for High Density”**
- 3) **APC White Paper #49, ” Avoidable Mistakes that Compromise Cooling Performance in Data Centers and Network Rooms”**
- 4) **APC White paper #42, “Ten Cooling Solutions to Support High Density Server Deployment”**
- 5) **APC Application Note #76 “Configuring Data Centers to Support IBM BladeCenter Servers”**
- 6) **APC Application Note #75 “Configuring Data Centers to Support HP BladeSystem p-Class Servers”**
- 7) **APC Application Note #74 “Configuration of InfraStruXure for Data Centers to Support Dell PowerEdge 1855 Blade Servers”**

About the Author:

Neil Rasmussen is a founder and the Chief Technical Officer of American Power Conversion. At APC, Neil directs the world's largest R&D budget devoted to power, cooling, and rack infrastructure for critical networks, with principal product development centers in Massachusetts, Missouri, Denmark, Rhode Island, Taiwan, and Ireland. Neil is currently leading the effort at APC to develop modular scalable data center solutions.

Prior to founding APC in 1981, Neil received his Bachelors and Masters degrees from MIT in electrical engineering where he did his thesis on the analysis of a 200MW power supply for a Tokamak Fusion reactor. From 1979 to 1981 he worked at MIT Lincoln Laboratories on flywheel energy storage systems and solar electric power systems.