APC

APPLICATION NOTE #76

Configuring Data Centers to Support IBM BladeCenterTM Servers

Abstract

The IBM BladeCenter operates at a power density 10X greater than the capability of a typical data center, which demands special attention to the power and cooling system. Success in deploying these servers – there are five basic strategies – requires an understanding of several inputs, including the business needs driving the deployment, user preferences, blade requirements, and constraints. This Application Note considers these inputs, and guides the reader to:

- Select the best deployment method from the five basic methods
- Choose the appropriate rack, cooling, power, and management solutions
- Apply these solutions to fit individual needs

Needs Assessment

BladeCenter servers have environmental requirements that exceed those of traditional rack-mounted servers, and can overload their environment if not properly integrated. Therefore, it is important that the IT manager take the time to conduct a needs assessment to understand business needs and preferences and identify facility constraints. This is particularly important if the blade servers are to be deployed in an existing data center. Key needs, preferences, and constraints are listed below, and explained in further detail in APC White Paper #125, "Strategies for Deploying Blade Servers in Existing Data Centers."

Business requirements and user preferences include:

- Existing data center uninterrupted operation – Is scheduled maintenance downtime a possibility?
- High availability of resulting system – Are redundant power and cooling required?
- Dense pack; co-location of servers – Do the blade servers need to remain in a specific area or can they be spread out?
- Preparation for follow-on deployments – Is this blade server deployment going to be phased?
- Time – Are there any time constraints to the deployment?
- Cost – Are there budget constraints?

Constraints include:

- Precision power capacity – Is there remaining power capacity available to support the blade server deployment?

APC Legendary Reliability®

- Precision cooling capacity – Is there remaining cooling capacity available to support the blade server deployment?
- Weight limits – What are the data center floor weight restrictions?
- Floor space limits – What are the space constraints within the data center?
- Ceiling plenum restrictions – Do room height constraints limit the ability to have an effective return air ceiling plenum?
- Raised floor restrictions – If a raised floor exists, are there any wire / piping obstructions or is the floor less than 2 feet in height?

APC and others offer professional services to help customers determine the proper NCPI (Network-Critical Physical Infrastructure) components to support blade server deployment. These services can include a site walk-through and needs assessment, turnkey installation service, server installation / migrations, and network integration. APC Professional Services Group can provide an end-to-end turnkey solution to ensure proper deployment of blade servers in an existing or new data center.

Planning

There are five basic alternative methods (listed in the table below) that can be used to install blade servers. The needs and constraints identified previously, with specific attention paid to cooling issues, will determine which approach is best for each application. Based on the deployment method selected, APC solution recommendations are made for each approach later in this Application Note. These five methods are explained in greater detail in APC white paper #46, “Cooling Strategies for Ultra-High Density Racks and Blade Servers.”

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 method 2 Typically limited to one BladeCenter chassis per rack 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	Density limited to about 2X the average power density Typically limited to 1 - 2 BladeCenter chassis per rack Uses more floor space Blade racks cannot be adjacent 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 and when needed Deferred capital costs High efficiency Good floor space utilization	Typically limited to 1 - 3 IBM BladeCenter chassis per rack 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 area 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

The table below, from APC White Paper #125, summarizes the blade server deployment criteria for different combinations of blade chassis density and blade deployment method, indicating preferred combinations. There are 11 preferred combinations (noted in green).

# 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.
	Minimal cost	Minimal cost	$1k-2k per rack	$10k-20k per rack	$20k-$60k per rack
	Increasing cost →				
	Easy	Procedures required	Non-invasive installation	Major installation, piping and wiring	Total room shutdown & rebuild
	Increasing deployment complexity →				

Rack Specification

The rack enclosure must physically support the BladeCenter chassis and have provision for power and data cabling to support the servers while allowing proper airflow into and out of the rack. The following attributes summarize the rack enclosure specifications for the IBM BladeCenter.

IBM BladeCenter Attribute	Value	Notes
Number of power supplies per chassis	(4) hot-swappable power supplies	Can be configured to supply redundant power to each of the blade chassis components
Standard power supply	2000 watts	Models prior to August 2004 may include 1200, 1400, or 1800 watt power supplies
Dual feed compatible	Yes	
Voltage requirement per power supply	200-240V 50/60 Hz, single phase	
Inlet connector of each power supply	IEC C20	Models prior to August 2004 may have IEC C14 inlets
Number of power supplies in fully configured rack	24	Each power supply has its own power cord
Nameplate power consumption of fully configured BladeCenter chassis	5400 watts	
APC-measured power consumption	4050 watts	This equates to a power consumption of 24,300 watts for a heavily configured 42U rack. Measurements taken on a heavily configured BladeCenter chassis operating at 100% CPU utilization
Mid-plane of BladeCenter chassis	Divided into 4 separate power buses with each one fed by a single power supply	Power supply modules 1 & 2 provide redundant power to the management modules, the switch modules, the blower modules, and processor blades 1 - 6. Power supply modules 3 & 4 provide redundant power to processor blades 7-14.
Number of power supplies in fully configured rack	24	Each power supply has its own power cord

APC answers the need for rack-based power distribution with its single and three-phase rack-mounted PDU products (RMPDU). APC-metered rack-mounted PDUs have an input twist-lock line cord and multiple output receptacles plus a built-in web/SNMP interface and a local user interface for current monitoring. Up to four rack-mount PDUs can be tool-lessly installed in the APC NetShelter SX rack. These products provide 14.4 or 25 kW worth of redundant power to the rack, depending on the version used.

APC provides power distribution, and UPS options to accommodate the power required by IBM BladeCenters. The table below presents the rack level power distribution solution(s) based on these attributes, for varying numbers of chassis per rack, and assumes the use of dual feeds to each rack. In many cases, the power infrastructure (UPS and main power distribution) exists. In cases where it doesn't, InfraStruXure UPS and main power distribution solutions (which are site specific) are readily available from APC.

In addition to these solutions, it is suggested to deploy the blades using the following best practices:

- Use metered PDUs to allow monitoring of branch circuit currents and ensure AC power redundancy (i.e. branch circuits should always be loaded to less than 50% of capacity).
- Evenly distribute the load across all 3 phase pairs. This balances the load and prevents breaker trips.

All recommended configurations assume a fully configured blade chassis with four redundant 2000 watt power supplies. In order to capitalize on these redundant power supplies, a 2N architecture should be implemented by providing a separate UPS / PDU power train for each input source. A minimum of two rack-mounted PDUs are supplied to provide breaker redundancy for the power supplies (feed A & feed B). BladeCenter power supplies 1 and 3 should be supplied by feed A, and power supplies 2 and 4 should be supplied by feed B. See **Figures 3 and 4** on page 11 and 12.

Number of Chassis per Rack	Solution	Part Numbers
1	(2) Metered Rack PDU, 20Amp, ZeroU, 3 phase, 208V	(2) AP7864
2 - 3	(4) Metered Rack PDU, 20Amp, ZeroU, 3 phase, 208V	(4) AP7864
4 – 6*	(4) Metered Rack PDU, 50Amp, ZeroU, 3 phase, 208V	(4) AP7868

*Also recommended for 1-3 chassis per rack when growth to 4 or more chassis is likely or if preference is to have no more than two whips per rack

Management Specification

Having visibility to many data points is required for the reliable operation of Network-Critical Physical Infrastructure (NCPI) and specifically the availability needs of IBM BladeCenters. Critical data points that should be monitored include current monitoring per branch circuit, inlet temperatures to racks, humidity levels, AC voltages, and UPS and battery status. APC's InfraStruXure power, rack and cooling components have built-in remote management capabilities ensuring access to individual devices is both easy and economical. Each browser-accessible device is quickly accessed from anywhere on the network without the need to install client software. Identifying problematic trends before they escalate is made easy with an exportable data log. An event log enables to the operator to pinpoint the timing and sequence of events leading up to an incident. SNMP trap forwarding allows individual devices to be integrated with an enterprise management system.

The suggested management solution(s) for all blade deployment approaches is the same – APC InfraStruXure Manager (part # AP9420 for 25 nodes, AP9421 for 100 nodes, AP9422 for 500 nodes, and AP9423 for 1000 nodes). APC's InfraStruXure Manager enables centralized management for up to a maximum of one thousand APC devices located throughout a network. A private IP network provides the option of monitoring of up to 253 APC devices from a single IP address. InfraStruXure Manager enables the quick assessment of your present situation and notifies the appropriate personnel should situations that threaten availability occur. Analysis features help to plan for changes in availability, power, runtime or cooling requirements. SNMP trap forwarding is used to send data to enterprise management systems, and the Modbus RTU protocol enables alarms and data points to be sent to your building management systems.

Close management of a high density NCPI increases the overall system availability by allowing the user to identify any power and thermal anomalies, and correct them before they cause a failure. APC recommends the use of an InfraStruXure Manger, metered RM PDUs, and an environmental monitoring unit with two temperature probes (top and bottom) on the front (air inlet) of each rack.

Conclusion

This Application Note describes how to successfully and predictably deploy IBM BladeCenters using the APC InfraStruXure system. For new data centers, designing high density zones or rows is typically the preferred approach. This is because designing an entire data center for high density is generally costly, inefficient, and actually wastes floor space.

For existing data centers, the use of supplemental cooling or borrowing rules is preferred for deployments in which BladeCenters are a small fraction of the total installation. This is because these are the least intrusive installation methods with the least risk of interfering with an operating data center. When the number of BladeCenter racks being deployed in an existing data center begins to exceed 5 racks or 10% of the total racks, high density zones or rows are the preferred approach. This is because spreading begins to become unmanageable and less effective with this quantity or fraction of blade racks. When the number of BladeCenter racks being deployed begins to exceed 10 or 25% of total racks, building a new room is recommended.

For powering the IBM BladeCenter, APC offers two metered rack-mounted PDUs, which mount in special tracks in the NetShelter SX rack and do not interfere with cabling to the BladeCenter or consume "U" height. The 20 amp rack PDUs suggested can be used for configurations of racks up to 3 chassis, and the 50 amp rack PDUs can be used for racks with 4 or more chassis.

For cooling the IBM BladeCenter, the key is to manage airflow effectively. For lower density deployments, APC offers the Rack Air Distribution Unit or Rack Air Removal Unit. For deployments at maximum density, APC suggests the APC NetworkAIR In-Row precision air conditioner with Hot Aisle Containment.

APC uses computer aided tools to design zones, rows, or complete data centers for blade servers. Consult an APC sales representative for more information on these tools. APC also offers professional assessment and planning services that can assist with blade server deployments.

Diagrams

Figure 1 – APC Rack Air Distribution Unit (ADU) and Rack Air Removal Unit (ARU)

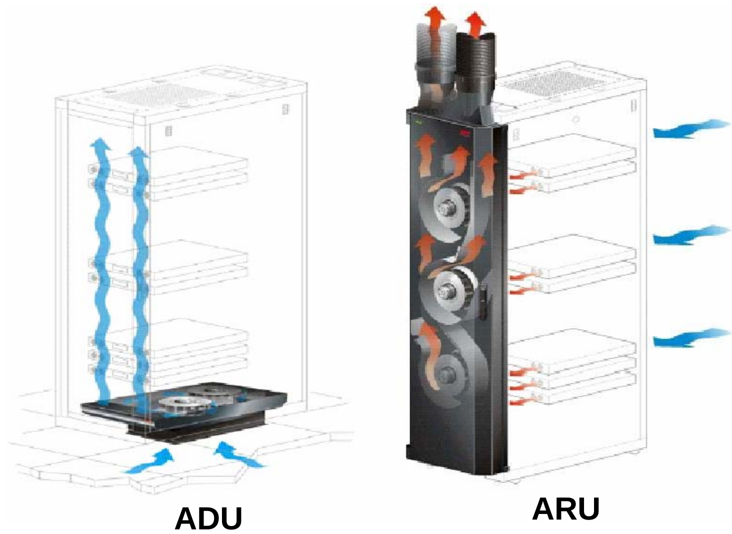

Figure 2 –NetworkAIR In-Row Cooling with hot aisle containment system

Figure 3 – Recommended PDU configuration for 1-3 BladeCenter chassis per rack

1 Chassis - 2x AP7864 20A, 3PH Metered PDU
2-3 Chassis - 4x AP7864 20A, 3PH Metered PDU

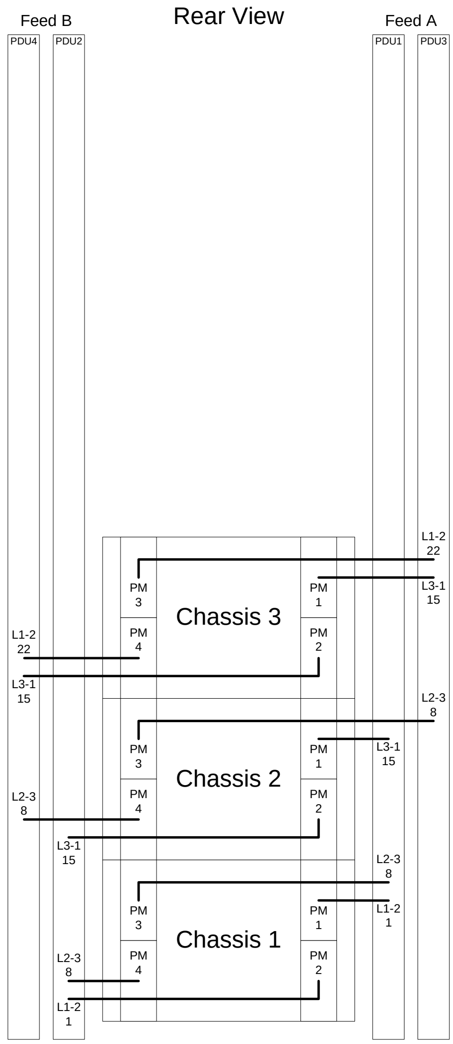

Note: Outlet numbers shown are for power supplies with C20 inlets. If using power supplies with C14 inlets, add 1 to the outlet numbers above.

Figure 4 – Recommended PDU configuration for 4-6 BladeCenter chassis per rack

4-6 Chassis - 4x AP7868 50A, 3PH Metered PDU

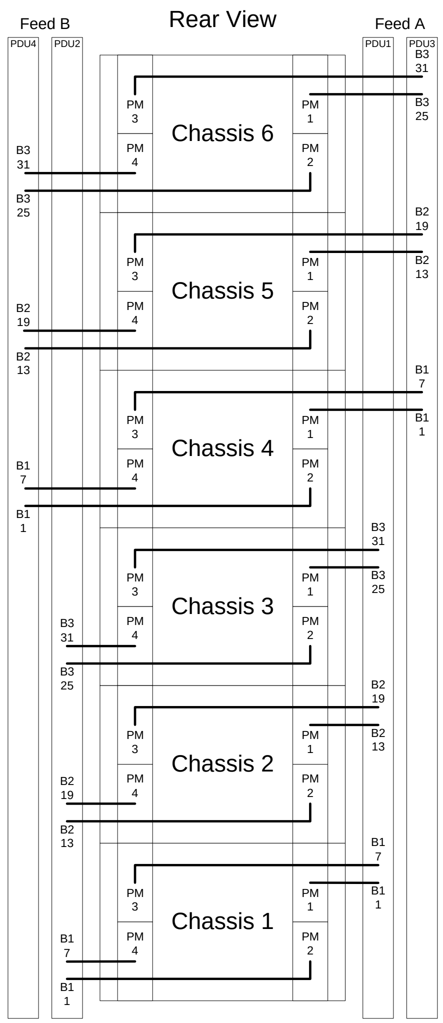

Note: Outlet numbers shown are for power supplies with C20 inlets. If using power supplies with C14 inlets, add 1 to the outlet numbers above.