APC APPLICATION NOTE #74

Configuration of InfraStruXure for Data Centers to Support Dell PowerEdge 1855

Abstract

Many companies are making plans to incorporate Dell PowerEdge Blade Servers into their data center applications. Using blades can pose many challenges to IT and Facilities manager's Network Critical Physical Infrastructure (NCPI). Simply put, NCPI is the environment consisting of power, racks, cooling, management and services that supports the technology layer of a business hierarchy. Traditional ways of NCPI deployment may be ineffective when used with blade servers. Blade deployment will require careful up front planning for new NCPI applications, and adaptation of existing NCPI.

Introduction

Many companies are making plans to incorporate Dell PowerEdge Blade Servers into their data center applications. Using blades can pose many challenges to IT and Facilities manager's Network Critical Physical Infrastructure (NCPI). Simply put, NCPI is the environment consisting of power, racks, cooling, management and services that supports the technology layer of a business hierarchy. Traditional ways of NCPI deployment may be ineffective when used with blade servers. Blade deployment will require careful up front planning for new NCPI applications, and adaptation of existing NCPI.

Planning and Needs Assessment

It is important to evaluate your intended application before implementation begins. This will ensure that each rack has the proper power and cooling provided for it. Blade Servers have environmental specifications that exceed traditional rack mounted servers, and can overload their environment if not properly integrated.

Things to consider:

1. Are the blade servers to be deployed in an existing or new datacenter environment?
2. How many blade servers will be deployed?
3. Are you planning on fully utilizing your rack footprint (put as many chassis in a rack as you can)?
4. For existing installations, how much power and cooling capacity is available within the datacenter?

APC and others offer professional services to help customers determine the proper NCPI components to support their blade server deployment. These services can include a site walk through and needs assessment, turn-key installation service,

server installation / migrations, and network integration. APC Professional Services group can provide an end-to-end turn key solution to ensure proper deployment of blade servers in an existing or new NCPI.

Rack Recommendations

The Rack that is to be used must physically support the blade server chassis(s), have provision for power and data cabling to support the servers, and allow adequate air flow to cool the blade server chassis. Either a vendor neutral or Dell specified rack should be used. The rack must also accommodate convenient places to mount enough rack mounted power distribution units to power the blade server chassis, and be able to support the physical load. It is important that each rack mounted power distribution unit is easily accessible and there is no visual obstruction to its receptacles/plugs. A fully configured Blade Server Chassis takes up 7U and weighs 130lbs (59kg); a rack fully configured with Blade Centers will have to support 780 lbs (354kg). The rack being used must also have adequate open space on the front and rear doors to accommodate proper cooling. APC recommends it Netshelter SX rack (APC Part Number AR2100BLK). The Netshelter SX has vendor neutral mounting, can support 2000 lbs (907kg) of load, has accommodations for up to eight vertical mounted (zero U) rack mounted PDU's (RMPDU) capable of delivering > 50kW of power per rack, and has over 800 sq inches (0.516 sq m) of open surface area (66%) on the doors to the accommodate proper cooling. The Netshelter SX also has options for overhead data and power cable distribution, rack cable management, and air distribution, which are all described later in this application note and can be configured via the APC InfraStruXure Build-Out Tool.

Power Architecture

Blade Center Severs can consume up to 757W / U. Each blade server chassis has slots for (4) hot-swappable power supplies, which can be configured to supply redundant power to each of the blade chassis components. 1200 and 2100 Watt power supply options are available and are determined by application. Each hot swappable power supply requires 200-240V single phase AC power. Each power supply has an IEC input connector (C13 for 1200W / C19 for 2100W), and its own power cord. A fully configured blade center chassis will have (4) IEC C14 / C20 power cords. Chassis configured with the 1200W power supply option require at least three power supplies to power a fully configured blade chassis, and four power supplies for redundancy. Chassis configured with the 2100W power supply option require at least two power supplies to power a fully configured blade chassis, and four power supplies for redundancy.

A fully configured 42U rack can accommodate as many as (6) Blade Server Chassis, require (24) IEC outlets (C13 or C19 depending on power supply used) and 31.8kW worth of 208-220V, single phase AC Power (2100W power supply including internal electrical losses). This poses many challenges to the people who design and support NCPI. APC answers the need for rack based power distribution with its single and three-phase rack mounted PDU products (RMPDU). Up to eight RM

PDU's can be tool-lessly installed in the APC Netshelter SX. This will provide 11, 25, or 58kW worth of redundant power to the rack depending on the version used.

APC rack mounted PDU's are available in basic, metered, and switched versions.

Basic – Has an input twist lock line cord and multiple output receptacles.

Metered – Same electrical features as the basic with a built in web/SNMP card and local user interface.

Switched – Same features as the metered version with the added capability to turn output receptacles on or off via a local serial interface or remotely across the Ethernet. Both methods are password protected.

The chart below in Figure 1 shows the part numbers and configuration of each APC recommended rack mounted PDU option:

Blade Center Recmmended Rack Mounted PDUs - North America / Latin America					
APC Part Number	**Input Line Cord**	**Output Receptacles**	**Rating**	**Type**	**Max Qty / Rack**
AP7564	NEMA L21-20P	(36)C13 & (6)C19	5.6kW	Basic	4
AP7864	NEMA L21-20P	(36)C13 & (6)C19	5.6kW	Metered	4
AP7961	NEMA L21-20P	(21)C13 & (3)C19	5.6kW	Switched	4
AP7568	Hubbellock 50A	(36)C13 & (6)C19	12.5kW	Basic	4
AP7868	Hubbellock 50A	(36)C13 & (6)C19	12.5kW	Metered	4
AP7968	Hubbellock 50A	(21)C13 & (3)C19	12.5kW	Switched	4
AP7567	Hubbellock 50A	(6)C19 & (3)C13	14.4kW	Basic	8

Blade Center Recmmended Rack Mounted PDUs - Europe / Asia / Latin America					
APC Part Number	**Input Line Cord**	**Output Receptacles**	**Rating**	**Type**	**Max Qty / Rack**
AP7851	IEC309 2P,3W,16A	(20)C13 & (4)C19	3.6kW	Metered	4
AP7553	IEC309 2P,3W,32A	(20)C13 & (4)C19	7.4kW	Basic	4
AP7853	IEC309 2P,3W,32A	(20)C13 & (4)C19	7.4kW	Metered	4
AP7953	IEC309 2P,3W,32A	(21)C13 & (3)C19	7.4kW	Switched	4
AP7555	IEC309 4P,5W,32A	(6)C19 & (3)C13	13.8kW	Basic	8

Figure 1, APC Recommended Rack Mounted PDU's for Blade Center Applications

The number of blade servers that are to be installed in each rack will determine which of the above rack mounted PDU's should be used. All recommended configurations assume a fully configured Blade Chassis with redundant power supplies. A minimum of four rack mounted PDU's should be configured to provide breaker redundancy (each power supply of the chassis should be fed from a different rack mounted PDU) if the 1200W power supplies are configured as shown in figure 2.

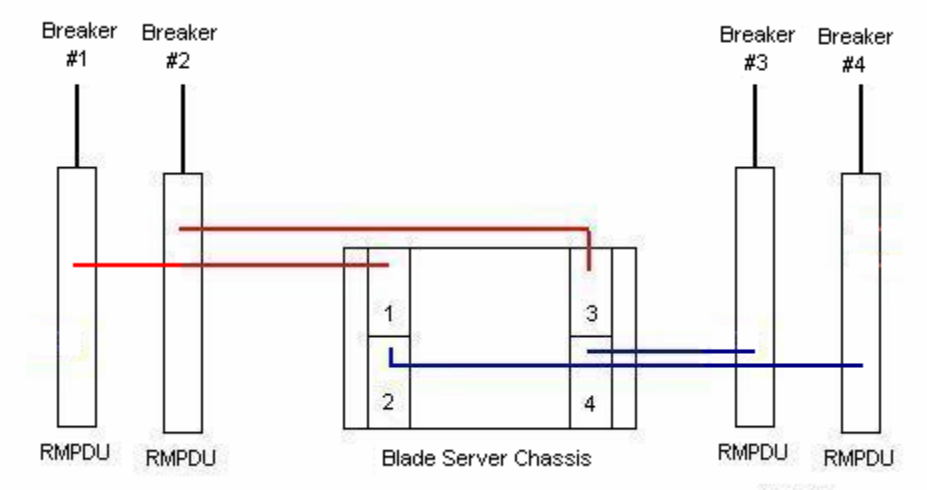

Figure 2, Power Cord Routing with 1200W Power Supplies

A minimum of two rack mounted PDU's should be configured to provide breaker redundancy (two power supplies should be fed from each rack mounted PDU) if the 2100W power supplies are configured, as shown in figure 3. Figure 4 shows a matrix of the recommended APC RMPDU solutions based on numbers of server chassis per rack.

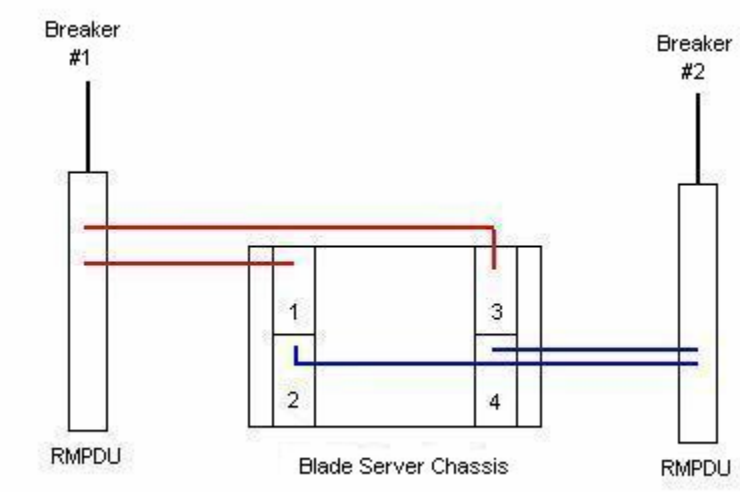

Figure 3, Power Cord Routing with 2100W Power Supplies

1200W Power Supply				2100W Power Supply			
Recommended Configurations - North America / Latin America				**Recommended Configurations - North America / Latin America**			
Blade Chassis Per Rack	**RMPDU Option #1**	**RMPDU Option #2**	**RMPDU Option #3**	**Blade Chassis Per Rack**	**RMPDU Option #1**	**RMPDU Option #2**	**RMPDU Option #3**
1 to 2	(4) AP7564	(4) AP7864	(4) AP7961	1	(2) AP7564	(2) AP7864	(2) AP7961
3 to 6	(4) AP7568	(4) AP7868	(4) AP7567	2	(4) AP7564	(4) AP7864	(4) AP7961
				3 to 5	(4) AP7568	(4) AP7868	(4) AP7567
				6	(6) AP7567	N/A	N/A
Recommended Configurations - Europe / Asia / Latin America				**Recommended Configurations - Europe / Asia / Latin America**			
Blade Chassis Per Rack	**RMPDU Option #1**	**RMPDU Option #2**	**RMPDU Option #3**	**Blade Chassis Per Rack**	**RMPDU Option #1**	**RMPDU Option #2**	**RMPDU Option #3**
1	(4) AP7851	N/A	NA	1	(2) AP7553	(2) AP7853	(2) AP7953
2 to 4	(4) AP7553	(4) AP7853	(4) AP7953	2	(4) AP7553	(4) AP7853	(4) AP7953
4 to 6	(4) AP7555	N/A	N/A	3 to 5	(4) AP7555	N/A	N/A
				6	(6) AP7555	N/A	N/A

Figure 4, Recommended RMPU configurations based on number of server chassis per rack

APC provides power distribution, and UPS options to accommodate the power required by blade servers. The chart in Figure 5 shows several power options based on total number of Blade Center servers that can be powered by each configuration.

1200W Power Supplies				2100W Power Supplies			
Qty of Blade Chassis	**Power Required**	**APC UPS**	**APC PDU**	**Qty of Blade Chassis**	**Power Required**	**APC UPS**	**APC PDU**
1 to 4	19kW (Note #1)	ISX20kW (Note #2)	N/A	1 to 3	16kW (Note #1)	ISX20kW (Note #2)	N/A
5 to 8	38kW (Note #1)	Symmetra PX 40	ISX 40kW PDU (Note #2)	4 to 7	37kW (Note #1)	Symmetra PX 40	ISX 40kW PDU (Note #2)
9 to 16	77kW (Note #1)	Symmetra PX 80	ISX 80kW PDU (Note #2)	8 to 15	80kW (Note #1)	Symmetra PX 80	ISX 80kW PDU (Note #2)
17 to 31	149kW (Note #1)	2 x Silcon 160kW	ISX 150kW PDU (Note #2)	16 to 28	148kW (Note #1)	2 x Silcon 160kW	ISX 150kW PDU (Note #2)
Notes:							
1.) Power levels do not include fan power of any air distribtution products that may be needed. Please use the InfrStruXure Build-Out Tool for proper configuration							
2.) Distribution Conductors are configured using the ISX Build-Out Tool based on the number of enclosures and RMPDU's used							

Figure 5, APC Recommended UPS and PDU options

For improved availability, 2N architecture may be utilized to provide a separate UPS / PDU power train for each input source. Power distribution troughs may be installed on the top of each Netshelter SX rack. This allows the use of pre-configured, flexible overhead cabling to deliver power from the UPS / PDU to each RMPDU that is configured within the system. The conductor lengths and floor layout of the racks is configured using the InfraStruXure Build-Out Tool.

Cooling Architecture

Providing a proper cooling architecture is the real challenge to effectively deploying Dell PowerEdge Blade Servers. The first imperative is to install the racks in a hot aisle / cold arrangement with the air return plenums for the CRAC located in the hot

aisles. This will minimize the amount of hot air that is allowed to mix the cold supply air to the servers. The hot aisle should be at least 3ft wide and the cold aisle at least 4ft. All unused U positions within the racks should be filled solid blanking panels to keep air from re-circulating between the aisles.

Raised Floor Environments

In a raised floor environment, the cold air is distributed under the floor and provided to the front of the racks through vented floor tiles. 3kW per rack is the highest average density that can be supported in this environment. Rack densities in excess of this can be supported as long as these racks are spread out across the datacenter and allowed to "drink" from their lower density neighbors. Air distribution techniques can be incorporated to ensure that the proper amount of CFM is delivered to the >3kW load racks and to minimize the hot air / cold air mixing. Even with these measures in place the average rack density of the room is still limited to no more than 3kW (lower density racks make up the difference). APC has two air distribution products that can be used as point products to increase rack capacity, which are shown in figure 6.

The APC Air Distribution Unit (ARU) (APC Part Number ACF002) is a vendor neutral rack device that can be utilized on racks that dissipate up to 4kW of heat load. It is a fan tray that occupies the bottom 2U of the racks and draws air through a flexible plenum from under the floor (a 1sq ft hole must be cut in the tile directly below the rack) and forces this cold air up the front the rack. This increases the available CFM to accommodate one blade server per rack.

The APC Air Removal Unit (ARU) (APC Part Number ACF103BLK) is a device that can be utilized on racks that dissipate up to 8kW of heat load. It is used in place of the back doors of a Netshelter SX rack. It uses fans to draw the hot air out of the back of the rack and direct it out the top. The ARE also removes static pressure in rack and increases airflow to accommodate up to two blade servers per rack. The top plenum kit (APC Part Number ACF120) then directs the hot air either down in the hot aisle or directly into a ceiling hot air return plenum for the room CRAC or CRAH(s).

Figure 6, ADU and ARU

The best way to avoid the negative effects of high density racks is to set maximum allowable rack density levels (based on the cooling environment maximums) and avoid deploying the high density racks all together. Money and resources can be saved by simply spreading the blade chassis across the entire datacenter space rather than compressing them all into a small area.

In-Row Cooling

APC has introduced an In-Row cooling concept that can support average rack densities up to 5kW with peak densities as high as 12kW (share from its low density neighbors). Peak racks should always be placed directly behind the Computer Room Air Handler (CRAH) to ensure limited air mixing in a hot aisle / cold aisle arranged room, a specially designed CRAH is placed in the row of racks. It draws hot air in from the rear of the rack (hot isle), cools it and expels the cool air into the cold aisle. Increased rack density is achieved because the architecture is not limited to the amount of CFM that can be handled by the floor tiles.

An In-Row system can cool > 5kW per rack if is outfitted with a hot aisle containment system, which includes a roof and row end door kit that will fully contain the hot aisles of the datacenter. This eliminates all air mixing and can accommodate up to 20kW of average rack density. Figure 7 shows an InfraStruXure System incorporating In-Row Cooling with a Hot Aisle Containment system.

Figure 7, In-Row Cooling Concept with hot aisle containment system

Recommended Cooling Solutions

Figure 8 shows several different APC recommended cooling solutions based on rack load level and type of cooling environment used.

Raised Floor Environment			
Blade Chassis Per Rack	**APC Recommended Solution**	**Blade Chassis Per Rack**	**APC Recommended Solution**
1	ACF002	1	ACF002
2	ACF102BLK & ACF120 (note #1)	2	ACF102BLK & ACF120 (note #1)
In-Row Cooling Environment			
Blade Chassis Per Rack	**APC Recommended Solution**	**Blade Chassis Per Rack**	**APC Recommended Solution**
1	In-Row Cooling Solution (note #2)	1	In-Row Cooling Solution (note #2)
Up to 2	In-Row Cooling Solution (note #2,3)	Up to 2	In-Row Cooling Solution (note #2,3)
Up to 5	In-Row Cooling with hot aisle containment (Note #2)	Up to 5	In-Row Cooling with hot aisle containment (Note #2)

Notes:
1.) Assumes a 40F (22C) temperature rise across server
2.) Configured using the InfraStruXure Build-Out Tool
3.) Only with APC design assistance

Figure 8, APC Recommended Cooling Solutions based on the number of Blade Chassis configured per rack

Management

Proper management of the NCPI layer of your network is essential to supporting the availability needs of your blade servers. With problems such as overloaded distribution circuits, aging batteries and insufficient cooling a regular occurrence, fault notification of critical issues via email/pager or SNMP traps (events), ensures crucial situations are dealt with in a timely manner. Many APC 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 software. Identifying problematic trends before they escalate is made easy with an exportable data log. An event log enables to you 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.

You can also choose to reduce the complexity of your network-critical physical infrastructure by managing your APC devices from a single console specifically designed for network-critical physical infrastructure. APC's InfraStruXure Manager enables centralized management for up to a maximum of one thousand APC devices throughout your network. A private IP network provides the option of monitoring of up to 253 APC devices from a single IP address. It enables the quick assessment of your present situation and notifies the appropriate personnel should situations that threaten availability occur. Analysis features help you 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 will increase 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 PDU's, and an environmental monitoring unit with two temperature probes (top and bottom) on the front (air inlet) of each rack.

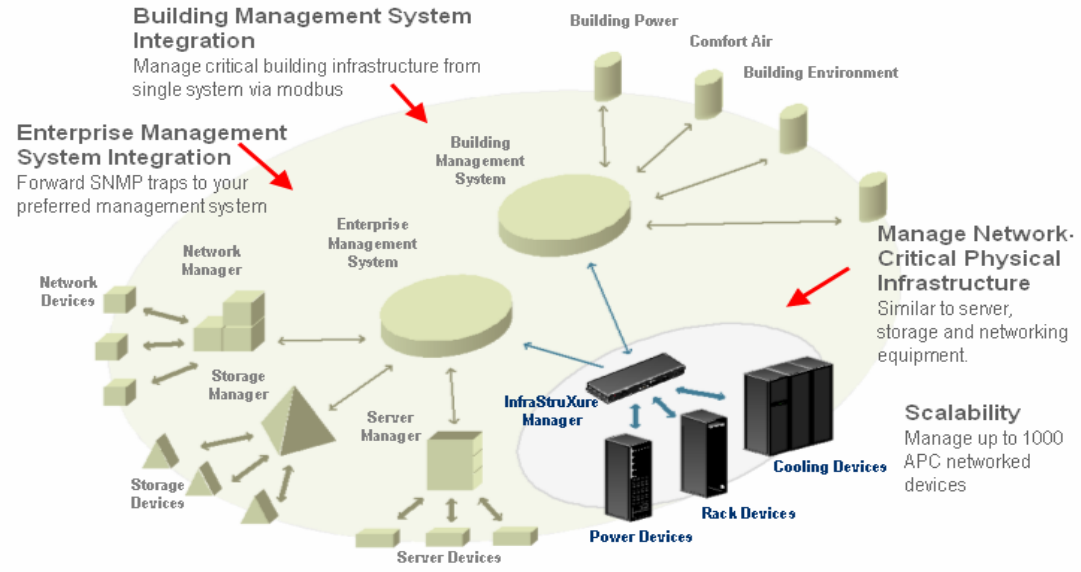

Figure 9, InfraStruXure Manager Integration Map

Conclusion

The evolution of blade centers present many challenges to NPCI. Traditional NCPI design and solutions fail to provide ample power, cooling and ventilation required to support this technology. In order to achieve the availability levels that blade servers demand careful consideration and planning needs to be taken into account. APC has spent and continues to spend significant R&D dollars to address all aspects of blade center environmental demands and in doing so has the only NCPI architecture that meets all the demands of a high density computing environment. A complete end to end solution using that includes all components of NCPI may be configured using the online APC InfraStruXure Build-Out Tool.