Server Technology, Inc.

Blade Server Power Solutions:

Cabinet Level Power Distribution Solutions for High Density Cabinets

February 28, 2006
White Paper STI-100-001

Server Technology, Inc.
1040 Sandhill Drive
Reno, NV 89521
+1 (775) 284-2000
www.servertech.com

Blade Server Power Solutions

INTRODUCTION

With the high cost of data center floor space and current advances in technology, new installations with denser cabinets that require more power and cooling continues to be the trend. Besides the challenges that new installations present, equipment cabinet upgrades can also be a problem as the existing power and cooling currently provided may not support the new cabinet configuration. Surveys show that Information Technology equipment is typically replaced every 2 to 5 years depending on the individual organization and its needs. Surveys also show (See Chart 1) that when asked about their top 3 concerns; Heat/Power Density is the number one concern of Data Center Management.

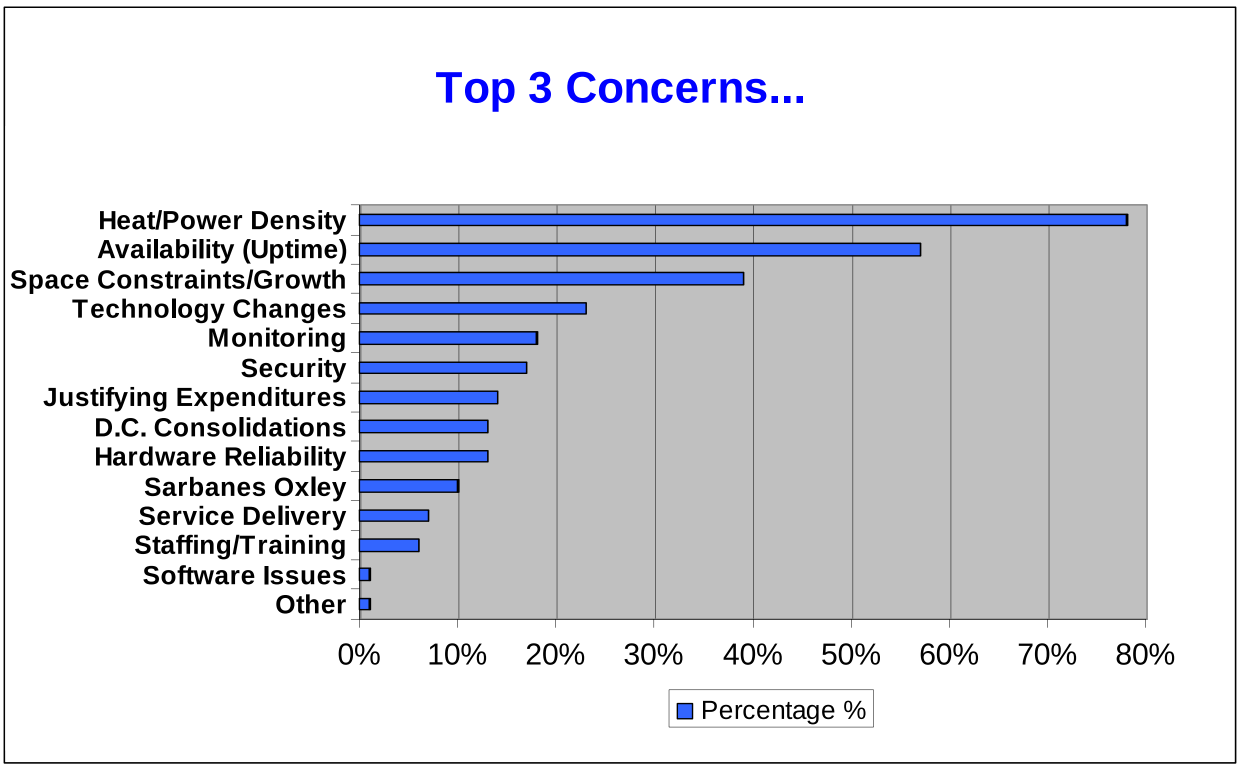

Chart 1

Note: From Spring 2005, Data Center User's Group Conference, The adaptive Data Center: Managing Dynamic Technologies

High density applications like cluster server configurations have in some cases pushed the kW power demands as high as 40 kW per cabinet. The required power depends on the equipment, how dense the cabinet is and whether redundancy is required. This has led to new and innovative solutions for providing cabinet level power utilizing CDU's (Cabinet Distribution Units).

The focus of this paper is to investigate blade server power trends at the cabinet level and offer power solutions depending on your application. This will make the critical infrastructure choice of providing power at the cabinet level easier to understand and manage.

POWER TRENDS IN THE DATA CENTER

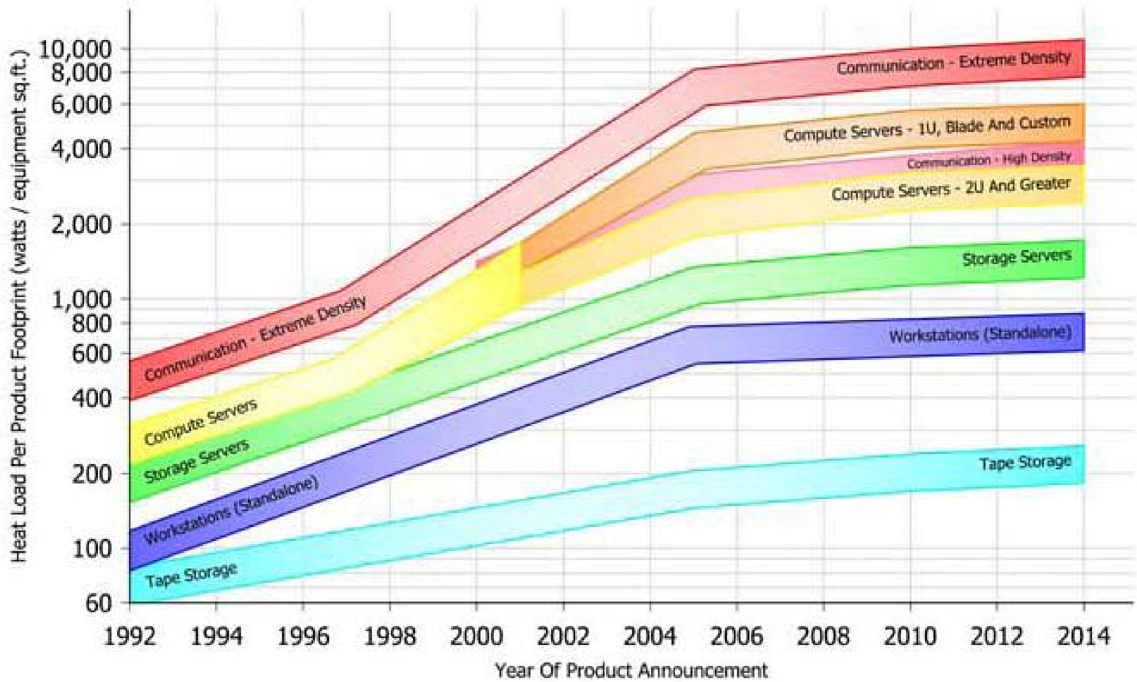

Graph 1

ASHRAE, Datacom Equipment Power Trends and Cooling Applications, Chapter 3, Fig 3.10 New ASHRAE updated and expanded power trend chart, © 2005 American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc., www.ashrae.org.

The overall trend is that the newer equipment in general generates more heat and requires more power for the same footprint. This is mostly due to more processing power per equipment footprint; as greater processing efficiencies are not being offset by the increased processing capacity for a given equipment enclosure. Power trend charts, like the one above, point to higher density loads that are greater than any loads that have been experienced in the field to this point. Due to these trends watts per Sq Ft no longer accurately defines the power requirements and the discussions have moved to a more precise watt per equipment sq. ft. metric. For the purposes of this paper the overall power requirements are the total kW required by the rack not factoring in the equipments Sq Ft.

Graphs like the one above (See Graph 1) are powerful tools that should be used to plan for current projects as well as for expansions and should also be considered when designing new grass roots datacenters. The need for CDU's to adapt to the ever increasing power requirements is crucial to supporting today's high density cabinets. Not only must today's high power demands be met to support Blade server and Cluster Server for server farm applications, but IT professionals must also be prepared for future upgrades and expansions.

BLADE SERVER OVERVIEW

It is now estimated that after just three years on the market, blade servers account for seven percent of all server shipments sold; with estimates showing that by 2008 they will make up thirty percent of all servers sold. It is expected that they will be the fastest growing server form factor through 2009.

There are many factors to consider when deciding if blade servers are right for your organization. These include but are not limited to costs, work loads and how predictable the work load is, performance and power consumption, uptime and repair, proprietary designs and disk storage considerations. With 3rd generation blade servers on the market the jury is still out on whether they will truly fulfill their promise of resiliency, repair efficiency, cost efficiency and dynamic load handling. Blades must continue to evolve toward this vision while still offering improvements in modularity, back-plane performance and the aggregation of blades into a single virtual-server image.

Currently five manufacturers make up more than 75% of all the blade servers sold in the market. These manufacturers are HP, IBM, Sun, Fujitsu/ Siemens and Dell. With the list of target applications growing for blade servers and the number of manufacturers increasing the future for this technology certainly looks bright.

POWER SELECTION AND DISTRIBUTION FOR HIGH DENSITY CABINETS

There are many different circuit combinations that can be used to provide power to today's densly populated cabinets. Today's total power requirements for rack mounted devices, per cabinet, can be as low as a 2-3 kW's but may exceed 40kW depending on the application, whether redundant power is required and of course, the devices that are installed within the cabinet.

Many servers today are designed to take a large range of input voltage from 120V to 208V. Typically smaller servers that may or may not be rack mountable require 120 V. Most of today's larger high density servers, such as blade servers require 208V power. While the floor space required to achieve a constant level of computing and storage capacity has shrunk, the energy efficiency of the equipment has not increased at the same rate. As a result, the power consumed and the heat dissipated has increased significantly within the cabinet foot print.

With increased power consumption (a trend that is likely to continue) comes important decisions that must be made with regards to cabinet level power requirements and selecting the proper CDU (Cabinet Distribution Unit). Tough choices like what power is brought down to the cabinet level such as 120V 15 or 20A, 208V 20 or 30A, 208V 30A 3-Phase or even 208V 60A 3-Phase are important decisions that effect costs, performance and the ability to provide for future expansion.

The number of power drops that must be run to each cabinet greatly affects the cost and whether or not there is sufficient power available for current needs and expansion. This is especially true in critical applications where redundant power is necessary. Each power drop run to a cabinet costs $500 to $1,500 or more; as well as having other consequences such as the number of CDU's required or whether or not cable management is a problem. This cost varies depending on several factors including panel space, distance, and what area the receptacles will be located in.

120 V For Server Applications

120V is a common power source for many server applications. Though only practical for the smallest installations 120V plugs are commonly available in most buildings. The problem is that most 120V outlets are 15A, though it is becoming more common to also find 120V 20A outlets as well. If you examine 120V at 15A it only provides 1.8 kW of power, and when derated per Underwriters Laboratory requirements, for safety purposes, this turns out to be only 1.44 kW of available power. (**Note:** Derated power is calculated as the total power times a .8 safety factor. **From this point on in the paper all power ratings listed will be derated.**) This does not take into account any losses within the Servers power supply (see your specific server power supply ratings for more information). Most devices today have internal power factor correction giving a 1 to 1 correlation between Watts consumed and the power supplied to the device making the power factor irrelevant when doing our calculations.

Therefore, even with the best conditions having a 120V 20A circuit the maximum available power per cabinet would be 1.92 kW. With many individual servers today requiring much more power this is not a viable solution for most cabinets with one or more servers or other devices installed. For example a single blade server chassis can require as little as 2 kW of power to more than 15 kW (fully populated with maximum load conditions). With 6 devices in a blade server rack (84 total servers) the total power required can be over 40 kW or far more than the amount of power than a single 120 V 20A power drop can supply. Running multiple drops for this application is far too costly and impractical especially when you consider how many drops would be required for a single cabinet.

208 V For Server Applications

Most servers today are capable of accepting 208V power, making it a viable choice for today's server applications. Depending on the available on-site power the 208 V drop may be either 20A or 30A. This gives us 3.33 kW and 5 kW of power respectively. Note: CDU's that accept input current of 30A or greater come with an integral power cord or multiple power cords depending on the device per UL standards.

208 V power provides a natural progression for server applications that require more power than a 120 V 20A (1.92 kW) drop can provide. 208 V power requires fewer drops to be run to the cabinet and the overall power supplied to the cabinet is increased. For certain data center environments a single drop might be all that is required to power a whole cabinet of equipment (two if redundancy is required), providing a large savings in the number of drops that must be run to the cabinet. Multiplied over a large number of cabinets this will represent a significant cost savings.

For blade servers 208 V 30A drop could be required for each server chassis within the cabinet depending on the server and how much of a safety factor will be used. Giving 208 V power an advantage over 120 V power in certain applications but still not optimizing either the number of power drops or the number of CDU's required for each cabinet.

208 V 30A 3 Phase For Server Applications (see Appendix A)

208 V 3-Phase has all of the advantages of 208 V power and more. The first advantage is that 208 V single phase power is easily derived from 3-Phase power within the CDU. This makes 3-Phase power very viable for sites that have the ability to supply it at the cabinet level. Also many blade cluster applications are demanding this increase in power for operation. Server clusters are installed with multiple chassis in a single cabinet where there might be anywhere from 6 to over 12 server chassis depending on their size in a single cabinet. New techniques like virtualization are also making these installations more popular as they can reduce HW costs by maximizing the overall processing capacities. Virtualization accomplishes this by allowing IT Managers to move applications from one system to the next either for

balancing demand or consolidating workloads. A typical 208 V 30A 3 Phase drop provides 8.65 kW of available power per drop. Blade Server clusters today can require 21 to over 40 kW per cabinet making multiple 30A drops required to supply the required amount of power.

Fortunately, 120 V single phase is also easily derived from 208 V 3-Phase Wye power which allows the CDU to have 120 V utility outlets if needed for peripheral devices or support equipment power when doing service or maintenance.

With more overall power being delivered at the cabinet level the number of drops required to be run to the cabinet is greatly reduced saving time, effort money, and installation complexity.

The last advantage is the ability to have enough power on hand within the cabinet to easily allow for expansion of the data center. If this is a consideration then a 208 V 30 A 3-Phase power solution just may be the answer without having to rerun additional drops or re-cable any of the devices in the cabinet. This is especially true today as some cabinets are setup knowing that though their full capacity is not required initially it will certainly be required for future growth. With the overall goal being that the IT infrastructure is a tool for growth and expansion not an inhibitor of future growth.

3-Phase power can allow a blade server application to provide as little as two power drops per cabinet in some applications reducing cost and cable management issues. This of course depends on several factors including the number of servers, whether redundancy is required, cooling capabilities, etc.

208 V 60A 3 Phase For Server Applications (see Appendix A)

Today's high density cabinet designs have increased power demand so radically that to meet the power needs 208 V 60A 3 Phase circuits are being utilized. Again these installations have all of the advantages of the 208 V 30 A 3 Phase applications with even more deliverable power at the cabinet level. This power configuration delivers over 17 kW of power at the CDU distribution level per drop.

Many of the blade applications like IBM's® Blade Center, HP's® ProLiant and Sun's Fire Blade Server applications are requiring this increased demand in power. Depending on the application the difference between a 30 A and 60 A 208 V 3-Phase solution may only be one or two additional drops per cabinet. However, with a large installation the additional cost of adding 1 or two drops per cabinet becomes very significant. These additional drops must also be considered for retrofitted cabinets where power availability or panel changes are required to provide the required power. Additional savings may also result in the reduction of CDU's required to be installed in a cabinet. This is roughly based on the number of drops running to a particular cabinet. So the choice between 30 A and 60 A drops often comes down to user preference, the application, the available power within the building, and if future expansion is being planned for. Therefore, all factors must be reviewed and in certain applications the savings between using 30A or 60A drops may not be as pronounced.

For example a high density cabinet might require 43 kW of power. This requires only (3) 3-Phase 208 V 60 A drops per cabinet but would require (5) five 3-Phase 208 V 30A drops. The difference being only two power drops per cabinet which may or may not result in a significant savings depending on the number of cabinets and the installation.

60 A 208 V 3-Phase installations are becoming more prevalent as power redundancy is becoming more common. CDU's can supply redundant power incase there is a power loss on one of the feeds coming into the cabinet. Many servers today have redundant power supplies just in case a power supply should fail. Most servers today come with multiple input cords just for these purposes.

Note: These high density applications may require liquid cooling. Having the available cooling will often determine the cabinet density and therefore the amount of power that is required for each cabinet. Due to high density cabinets "hot spots" are often introduced to the data center making the old term of kW per sq/foot irrelevant and forcing data center design engineers to be much more interested in the actual kW or power required for a particular cabinet when determining the overall cooling capacity that is required within the data center.

Table 1 provides a summary of some of the different power options that are available.

Table 1
Summary of Power Options

	Delivered Power in kW (with .8 SF)	# of Branch Circuits	Branch Circuit Protection	120 V Support	208 V Support
120 V 15 A	1.44	1	N	Y	N
120 V 20 A	1.92	1	N	Y	N
208 V 20 A	3.33	1	N	N	Y
208 V 30 A	5	2	Y	N	Y
208 V 3 Phase 30A	8.65	3	Y	Y (with WYE)	Y
208 V 3 Phase 60A	17.3	6	Y	Y (with WYE)	Y

Additional Things to Consider when Specifying a CDU

When specifying a CDU there are a long list of additional items that need to be considered. Some but not all of the key items that require consideration include: receptacle control, current metering, redundancy and remote management and monitoring.

Receptacle Control

Once the cabinet has been populated and your systems are up and running is often when disaster strikes. In some cases personnel will get into the cabinet for repair or maintenance and make cable changes on the fly. These changes can often bring down a particular branch circuit if they are not careful, forcing the shut down of a portion or even all of the devices within the cabinet. One way to avoid this problem is to install a CDU that has individual outlet control so that unused outlets may be turned off by the systems administrator ensuring that unwanted changes cannot be made to the cabinet's cable configuration unless they are authorized.

These problems can also occur during maintenance, cleaning or some other function within the datacenter where a server or another device gets plugged into a CDU located in an adjacent cabinet. Usually this is done because they needed power and it was the closest available receptacle. Situations like this can often exceed the current rating of that branch circuit causing it to fail. This has recently become more prevalent as IT personnel are trying to maximize their cabinet space by operating considerably closer to the maximum available power levels within the cabinet.

Power up sequencing is another receptacle control feature that is critical to a safe restart after power has been lost. Sequencing limits the inrush current on the power feed to one server ensuring that both power supplies of a dual-supply server receive power simultaneously. This avoids load imbalances and averts cumulative inrush that could otherwise trip the upstream circuit breaker. Adjustable sequencing parameters and unique programmable "Wake-Up" states such as "ON", "OFF or "Last State" also ensure proper server operation after a power loss or during start up.

Current Metering

Utilizing a CDU that provides the current reading for each protected branch provides several valuable features. The first is that the operator always knows how close they are to exceeding the overall current rating. This allows the operator to run the cabinet with maximum efficiency and ensure that they are best utilizing their assets within their expensive cabinet space. Example: Without branch circuit current monitoring the user has no idea how much current each particular branch circuit is drawing and therefore how many more devices can be added to a particular branch or cabinet. These current meters in 3-Phase power applications can display the individual phase currents and allow for balancing loads across the phases.

The standard safety rating for derating power capacity is typically 80%. Without a CDU that provides the current reading per branch circuit many facilities have no idea how close or far away they are from actually exceeding the ratings on the power drop and risk overloading their branch circuits.

Redundancy

Both IBM and HP blade servers today come with multiple power input feeds. These provide both redundant power to the server and the ability to compensate should one of the power supplies within the server fail. These are all part of the safety measures designed to ensure proper up time.

Many of the CDU's available today will also have dual power infeeds either for power redundancy or to meet the high power demands of today's server environments. In some cases as many as 42 outlets per power drop are supplied for a total of 84 outlet receptacles within a single, dual power input CDU.

WHAT DOES THE FUTURE HOLD?

Only time will tell if future capacity demands will continue to increase power consumption and cooling levels to higher and higher levels. Current ASHRAE projections certainly predict that they will. Other technologies such as virtualization, dual core processors, and new more efficient processor designs along with current and future software solutions will ultimately determine the outcome.

CONCLUSION

It is clear that by close review of the application and it's power requirements decisions can be made that will allow for future expansion, denser cabinets, greatly reduced costs (in many cases), power redundancy (if required) and will optimize cable and CDU requirements. It is only through making the power requirements a priority at the beginning of the project rather than an after thought that these goals are achieved. This has become especially important with today's higher power demands and the ever increasing push for lower costs, greater efficiencies and improved uptime availability.

APPENDIX A:

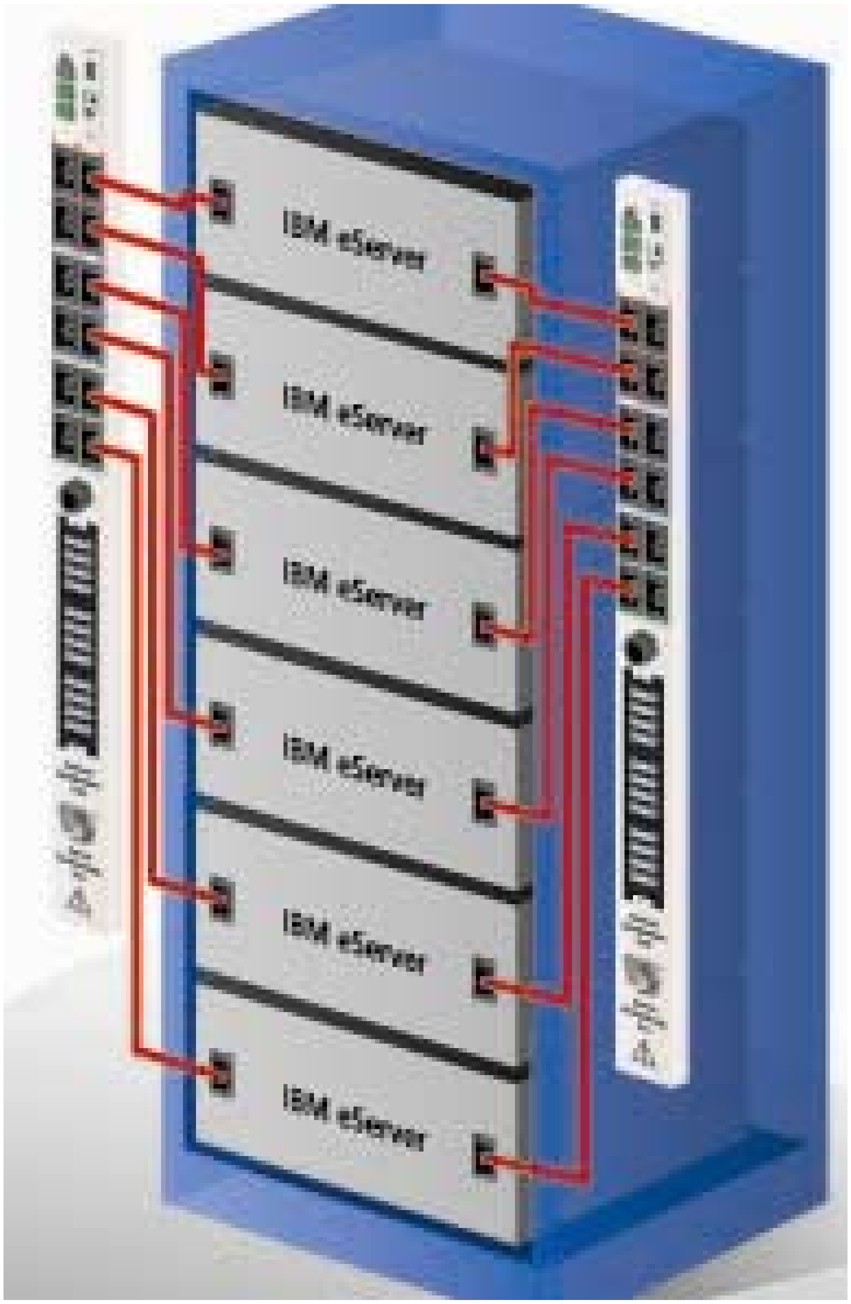

High Density Cabinet 208 V 3-Phase 60 A Infeed

- 6 BladeCenters per Cabinet
- Two 3-Phase Power Whips per Cabinet (208 V/60A)
- Two Switched 3-Phase CDU's per Cabinet
- 6 Branch circuits, two outlets per branch
- Load Monitoring
- Temperature and Humidity Monitoring
- Trending data for Temp/Humidity and Load
- Create groups of associated outlets and reboot with one point and click

Figure 1

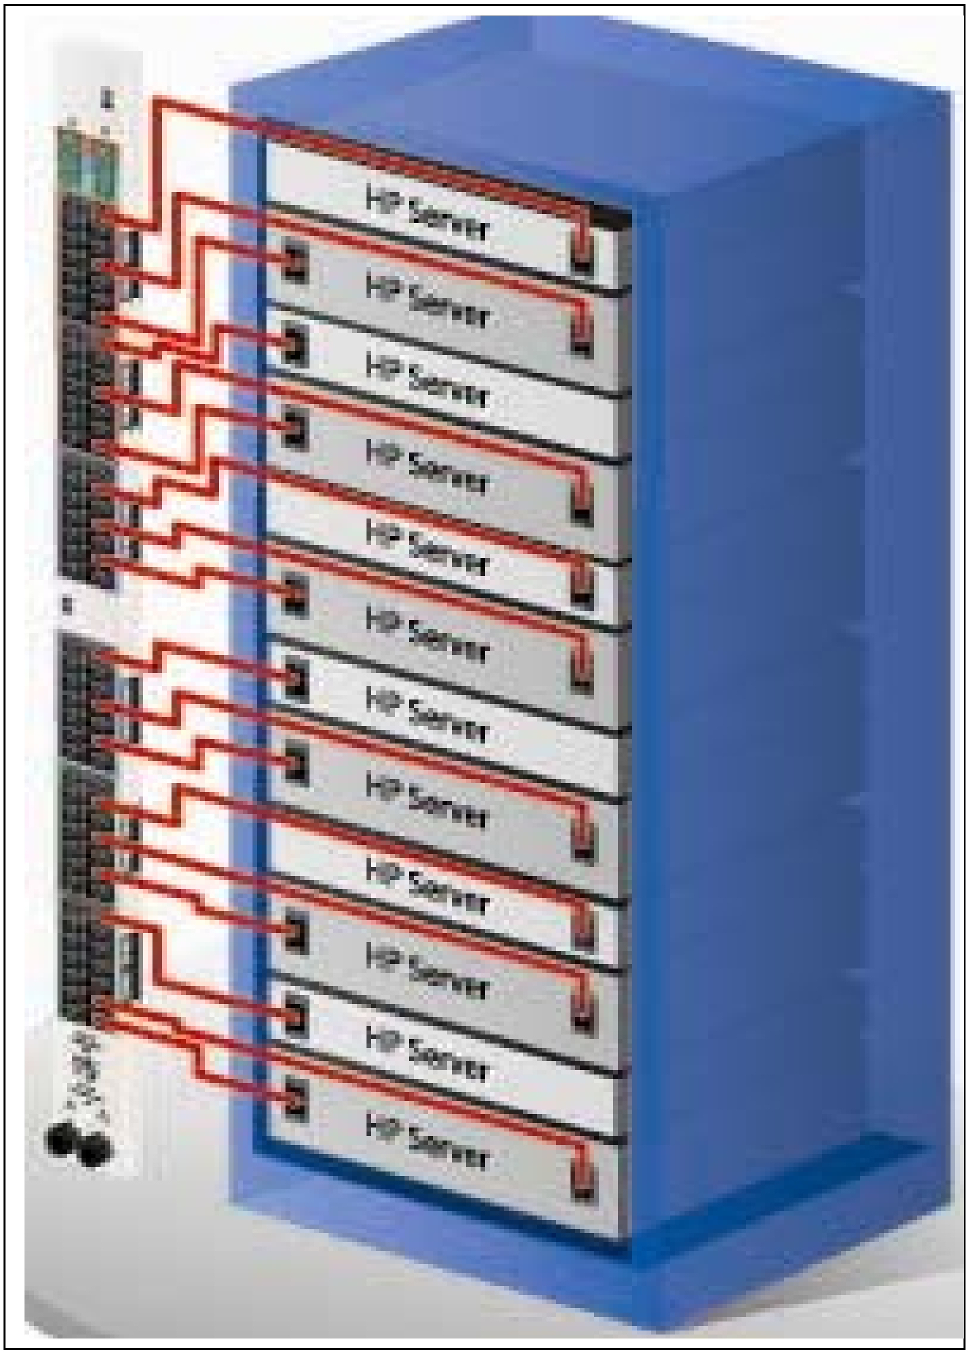

High Density Cabinet 208 V 3-Phase 30 A Infeed

- 12 HP Servers per Cabinet
- Two 3-Phase Power Whips per Cabinet (208 V/30A)
- One dual feed Smart 3-Phase CDU's per Cabinet
- 6 Branch circuits per CDU
- Load Monitoring
- Temperature and Humidity Monitoring
- Trending data for Temp/Humidity and Load
- Create groups of associated outlets and reboot with one point and click

Figure 2

@server is a registered traded mark of IBM.
HP is a registered trade mark of Hewlett-Packard.