Cisco Unified Computing System Blade Server Chassis Power Modes

White Paper

October 2010

Contents

Introduction

The Cisco UCS 5108 Blade Server Chassis is powered by 208 V (220 V outside the US) 50-60 Hz AC. Up to four power supply units (PSUs) may be installed, each with a single C20 connector for input power connection to the UCS 5108 power supply. This document applies to all versions of the Cisco Unified Computing System Manager. Where specific versions are required, that is highlighted.

The total number of power supplies required for operation depends on the total system power consumption and the level of redundancy required by the end user.

Cisco UCS Redundancy Modes

Cisco Unified Computing System (UCS) enables end users to select the power mode that fits their redundancy needs through the Cisco Unified Computing System Manager (UCSM). Within the Cisco UCS Manager, the power redundancy mode is selected via the global policy tab, as shown below.

Figure 1. Configuring the Global Power Policy

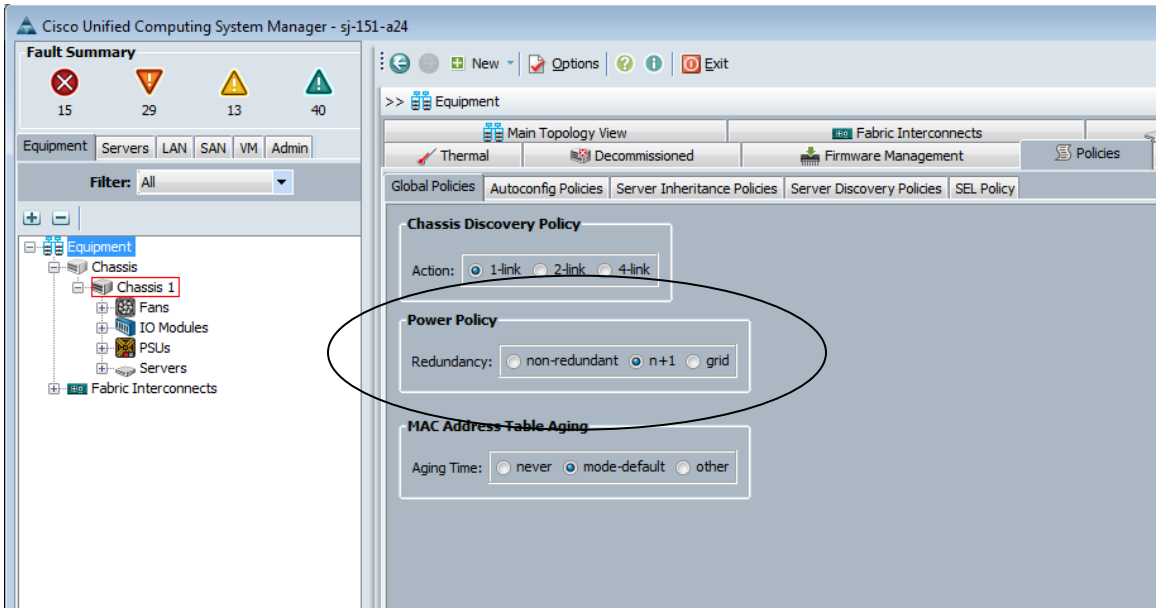

Figure 2. Enlarged view of Power Policy Options

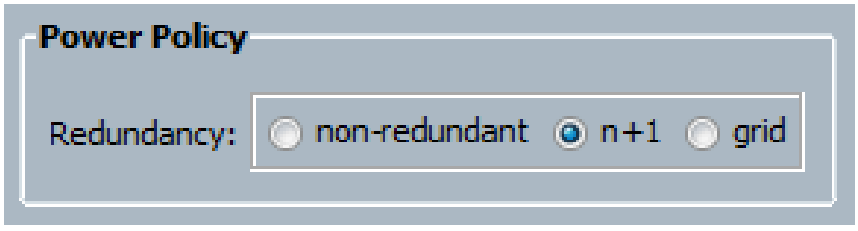

Since this power policy option is global, it is inherited by all the chassis managed by the Cisco UCS Manager instance. Non-redundant means that uptime cannot be guaranteed in the event of a failure. N+1 means that the system can tolerate the failure of one supply. Finally, grid redundancy (N+N) means that the system, if wired correctly into dual independent AC feeds, can tolerate the loss of one of those grids or half of the power supplies.

The Value of N for Cisco Unified Computing System

When describing power supply configurations, “N” refers to the minimum number of supplies required to deliver power without any redundancy. In other words, N+1 describes a power supply configuration that has one extra power supply installed and active.

Other than special scenarios, N=2 for UCS. N=2 is referred to as “standard mode” for the power supplies in the Cisco UCS 5108 Blade Server Chassis through the rest of this document.

Beginning with Cisco UCS Manager version 1.4.x or later, Cisco Unified Computing System has been enhanced to support especially high density configurations. Any following discussion around this new high-density mode applies only to Cisco UCS Manager v1.4.x or later. Specifically, in high density mode, three power supplies are made active (N=3). The high-density Mode (N=3), is enabled for only the two following configurations:

1. With the global power supply policy set to non-redundant, you must have at least three active supplies (installed, functioning and connected to AC)
2. With global power supply mode set to N+1, you must have four active supplies installed.

If these conditions are met, the system is automatically placed into high-density mode. As explained later, the benefit of high-density mode is that server blades are allowed to potentially consume more power and maximize their performance.

When to Use High-Density Mode

The high-density mode of operation is useful for chassis configurations which exceed the standard mode capacity (5000W DC total power). To make this determination, the UCS Power Calculator (http://www.cisco.com/assets/cdc_content_elements/flash/dataCenter/cisco_ucs_power_calculator/) can be used to estimate chassis power consumption for any given configuration. As a rule of thumb, if the maximum AC power estimate from the Cisco UCS power calculator exceeds the 5000W limit, then configuring the system so that it enters high-density mode should be considered for optimal performance. It is important to point out that the power calculator makes maximum power estimates based on worst case utilization. Power calculator maximum numbers are only achievable by high-performance computing applications. Enterprise applications typically use much lower power and 50% utilization is a reasonable figure to use.

Transitions Between Standard and High-Density Mode

When the appropriate conditions are met as outlined above, **all qualifying chassis are automatically and dynamically switched to high-density mode.** For example, if Cisco UCS Manager is configured to use an N+1 global power redundancy policy,, all chassis that have four functional power supplies will be in N=3, 3+1 mode. If a chassis with three supplies has a fourth supply installed and connected, that chassis is automatically and dynamically switched from standard-density to high-density mode.

If the event of a failure, chassis in high-density mode are not dynamically switch back to standard density. For example, if the power supply is in high-density N+1 mode (N=3) and a supply fails, the system will report a fault and continue to actively use all three power supplies. If, on the other hand, the IO modules are rebooted or if the chassis loses power with only three functioning supplies, the system will switch to standard density (2+1) mode. In this example, chassis belonging to a UCS domain with a global N+1 redundancy policy will re-enter high-density mode if a fourth supply is installed.

Blade Level Power Allocations

The Cisco UCS blade server chassis has automated internal mechanisms to prevent the power consumption from exceeding the allowable power. How much power is available depends on whether the standard mode or high-density mode of operating is in effect.

To guarantee that power use stays below the limit, the chassis implicitly caps the overall power consumption of the servers. This is done via the same mechanism used to implement chassis power capping (as part of the UCS group power capping feature). To learn more, please consult the group power capping feature white papers and documentation.

For chassis that are not part of a power group, the following blade level power caps are automatically set:

- Standard Mode (N=2)

  Blade level power consumption is implicitly capped to 550W for half-slot blades or 1100W for full slot blades.

- High Density Mode (N=3)

  Blade level power consumption is implicitly capped to 600W for half-slot blades of 1200W for full width blades.

It is important to clarify that it is very difficult to hit even the 550W limit for a half width blade. Doing so requires the blade to be:

- populated with maximum memory
- using the highest power CPUs and
- running extremely demanding applications that stress CPU and memory
    - High Performance Computing is one example of such an application

It is not expected that any enterprise application will come close to this limit.

Redundancy Mode Details

Choosing which redundancy mode to use is a decision for the end user and should be based on their specific availability requirements. Most users will want to select N+1 or grid redundant (N+N) mode.

Non-Redundant (N): In Non-redundant mode, the system may go down with the loss of any supply or power grid associated with any particular chassis. To operate in non-redundant mode, each chassis should have at least two power supplies installed. Supplies not used by the system are placed into standby. Which supplies are placed into standby depends on installation order (not on slot number). Load is balanced across active power supplies, not including any supplies in standby.

Standby power supplies are activated in the event of a failure. In the event of an active power supply failure, Cisco Unified Computing System will enable a standby supply to avoid downtime. If the load is low enough, no service interruption will occur. However, in the event a failure does happen, it is possible that failover will not occur fast enough to avoid downtime.

Best practices are for all deployments to be run in redundant mode. It is never recommended that users run any Cisco UCS chassis without some level of redundancy.

Figure 3 shows the power connection schemes for both standard mode and high density mode. With the global power policy set to non-redundant (Figure 3), any chassis with three or more functional supplies will switch to high-density (N=3) mode automatically.

Figure 3. Suggested Wiring for Non- Grid Redundant Mode

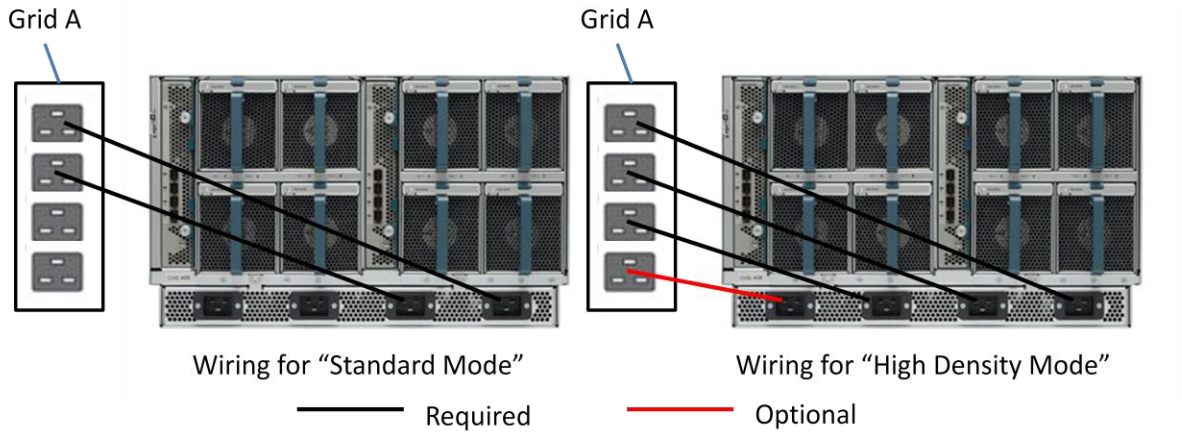

Figure 4 shows connection schemes for the power supplies. The key point is that the correct number of power supplies should be installed and plugged into AC power.

Redundant (N+1): In N+1 mode, the chassis tolerates the failure of any single supply without any interruption of service. In this mode, the chassis should have at least three supplies installed. To run any given chassis in redundant high-density (N=3) mode, four supplies must be installed and connected.

Figure 4. Suggested Wiring for N+1 Redundant Mode

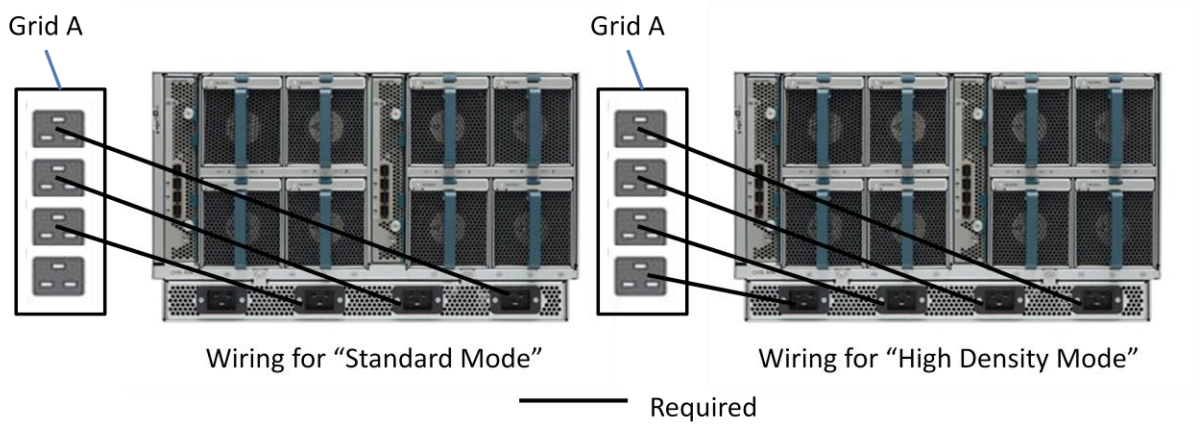

Grid Redundant (N+N) - below: The purpose of the grid redundant mode is to enable a configuration that can tolerate the loss of either a power supply or a input power circuit. In grid-redundant mode the system can withstand the loss of any two power supplies. The Figure 5 shows the required wiring Scheme for Grid Redundant power where N=2.

For grid redundant operation, it is important to connect the supplies to the independent grids as shown in Figure 5. Not connecting the grids in a supported configuration will cause a grid (circuit) loss to appear as if multiple power supplies have failed. High-density mode is not supported for grid redundant operation because grid redundancy requires a N+N implementation where N=2.

Figure 5. Required Wiring for Grid Redundant (N+N) Mode

Critical Loss of Power

As for any system, even with redundancy, there always exists some catastrophic sequence of events that will result in down time.

Some examples of events which would result in critical loss of power for the Cisco UCS Blade Server Chassis are:

- Losing a single supply in non-redundant mode
- Losing two supplies in N+1 Mode
- Losing three supplies or (one supply + one grid) in grid redundant mode

If the system is lightly loaded, there is a good chance that all the blades will continue operation even if the system has nominally insufficient power. In the worst case, if the system is truly deficient and there is no way to carry out a graceful shutdown, the blades will shut down in slot number order until there is enough capacity to power the remainder.

Summary

The Cisco UCS power configuration approach enables users to meet their redundancy and power density needs. The approach is simple yet flexible enough to adapt to the varying needs for different applications within the data center.

The global power policy provides users a single interface for configuring their desired power redundancy state. Most users will desire some amount of power redundancy to maximize server uptime. The global power policy is shared among all the chassis managed by a given UCSM instance.

The new density mode offerings available in Cisco UCS Manager version 1.4.x or greater provide support for highly dense configurations running the most demanding applications. The selection of which density mode is used (standard or high-density) is made automatically and dynamically by the system on a chassis-by-chassis basis. This selection is made exclusively based on power supply unit population. This allows some of the chassis within a Cisco UCS domain to operate in standard-density mode while others operate in high-density mode. Only chassis running the densest blades with the most demanding applications will benefit by running in high-density mode. Any chassis running less than 5000W will not benefit from operating in high density mode.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA

Document number: UCS-TR100020 10/10