Enterprise Environmental Planning

Power Configurator User Guide
v2.1.3

CALCULATE INDIVIDUAL POWER CONSUMPTIONS

Use the Power Configurator to obtain three groups of environmental information from uniquely configured systems. These include idle, maximum and load factor power consumptions. Scale the load factor of your configuration from 0-100% utilization for an accurate understanding of your environmental requirements.

GENERATE POWER REPORT

The Power Configurator tool allows you to import certain hardware configurations to determine if your configuration has sufficient power for redundancy. Once imported, generate a report for additional information on your configuration.

ENVIRONMENTAL PLANNING

Use the Power Configurator tool to help determine the best fit PSU capacity to install in your server. Also use the tool to help determine the power capacity requirements for PDU and UPS installations.

Author: Rani Doughty
rdoughty@lenovo.com

Questions / Comments: power@lenovo.com
Data Center Services, Enterprise Business Group

Revision History

1.0.0 – November, 2012	First major release.
2.0.0 – May, 2014	Second major release. New document creation.
2.0.1 – August, 2014	Amended heat loss picture. Added Questions/Comments email link on first page. Amended some items in the 'Using the System x Power Configurator' section. Added in hyperlinks for easy navigation. Added in configuration rules 'All Chassis'. Changed 'Power Supply Ratings' page to include 80 Plus Titanium.
2.0..2 – October, 2014	Updates to contacts and emails
2.1.0 – October 2014	Updates to pictures, functions and format. Added Turbo Boost section.
2.1.1 – March 2015	Added section on Power Configurator import function for CFR files
2.1.2 – April 2015	Update to template
2.1.3 – October 2015	Update to installation instructions

Reviewers:

Matthew Archibald – Data Center Services, Enterprise Business Group
marchibald@lenovo.com

Jerrod Buterbaugh – Data Center Services, Enterprise Business Group
jbuterbaugh@lenovo.com

Table of Contents

Introduction

The System x Power Configurator is a software tool designed to assist with calculating System x Systems environmental information. The data in the System x Power Configurator tool is derived from running real-world workloads across a number of configurations to properly characterize component power consumption under various conditions. The current workloads are a combination of Floating Point and Small FFT Processor workloads as well as running configuration-tuned versions of HPL (Portable Implementation of the High-Performance Linpack Benchmark) and Stream to exercise multiple sub-systems of the IT systems. The two highest power consuming sub-systems in a non-HPC (High Performance Computing) IT system are processor and memory, so focus is given to exercising those sub-systems to correctly model power consumption under traditional and virtual workloads. For HPC type IT systems, some configurations contain GPGPU (General Purpose Graphical Processing Unit) or MIC (Many Integrated Core) I/O adapters. These adapters are characterized by running highly parallel graphic rendering workloads for complete characterization. All tests are conducted using default uEFI/BIOS settings.

The System x Power Configurator provides three groups of environmental information. The first represents Idle or minimum power consumption, the second is Maximum power consumption, and the third is Load Factor. Load Factor is a scale factor between Idle and Max that can represent any configuration's total aggregate system utilization for a specific workload.

The data reported by System x Power Configurator can be used in certain cases to determine electrical wiring and levels of redundancy. The data reported by System x Power Configurator represents a worst-case power consumption value under normal operating conditions and may not model power consumption under component failure conditions. Final determinations should be made by persons skilled in the art or by contacting power@lenovo.com for assistance.

You can obtain the latest copy of this document and the configurator tool from this web page:

System x Power Configurator tool:
http://www.ibm.com/systems/bladecenter/resources/powerconfig.html

About This Guide & System x Power Configurator

This guide provides explanations and examples on how to use the System x Power Configurator tool.

If you are using this guide and tool to assist you in sales, keep in mind the following information.

There is no standard for vendors testing and publishing power consumption. Vendors may publish power consumption numbers and are not required to disclose the workload capacity the hardware is operating under. This will alter the power values that are published from vendor to vendor. Under the same workloads each vendor should not vary any more than approximately 10% in their power consumption. System x runs a specific version of Linpack with a CPU workload (mPrime). This may give a higher maximum power value compared to other vendors however it very accurately represents the system power consumption under load and is safe to use for planning and budgeting your power circuits.

System Requirements and Support

The following are the minimum system requirements to use the System x Power Configurator tool.

Minimum System Requirements

In order to successfully open and use the System x Power Configurator Tool on an individual workstation or virtual machine, your system must meet the below minimum specifications:

At least 256MB available memory
20MB available hard disk space
Monitor resolution of 1024 by 768 pixels
Operating System: Linux, Mac OS, Microsoft, Solaris
Minimum Java Runtime Environment (JRE) v7 Update 45

If you do not have at least JRE 7 Update 45 you will need to download and install it for the program to run. You can obtain JRE 7 and above via the below link:
http://www.java.com/en/download/manual.jsp

Tool Release Dates

In general, the tool has rolling updates that coincide with major System x hardware announcements. Minor updates are normally rolled up and released with major updates. The tool is released internally and takes approximately 1 week to be available on the web externally. If you require a pre-release version from the internal site, please email power@lenovo.com.

On occasion bug fixes and minor updates can also mean a new version of the tool is released to the web. Ensure you check the download page at the below link to ensure you are running the latest version.

http://www-03.ibm.com/systems/bladecenter/resources/powerconfig.html

Installation and Planning Support

The System x power and thermal team can be contacted via the following address: power@lenovo.com

If you have any questions concerning installation and planning for power, thermal, or mechanical, please contact us at this address. Support is provided around System x, BladeCenter, Flex System, iDataPlex, and NeXtScale Systems.

Downloading the System x Power Configurator

The System x Power Configurator (program) is released to the Internet for general availability to the public. The System x Power Configurator can be obtained via the below link:

http://www-03.ibm.com/systems/bladecenter/resources/powerconfig.html

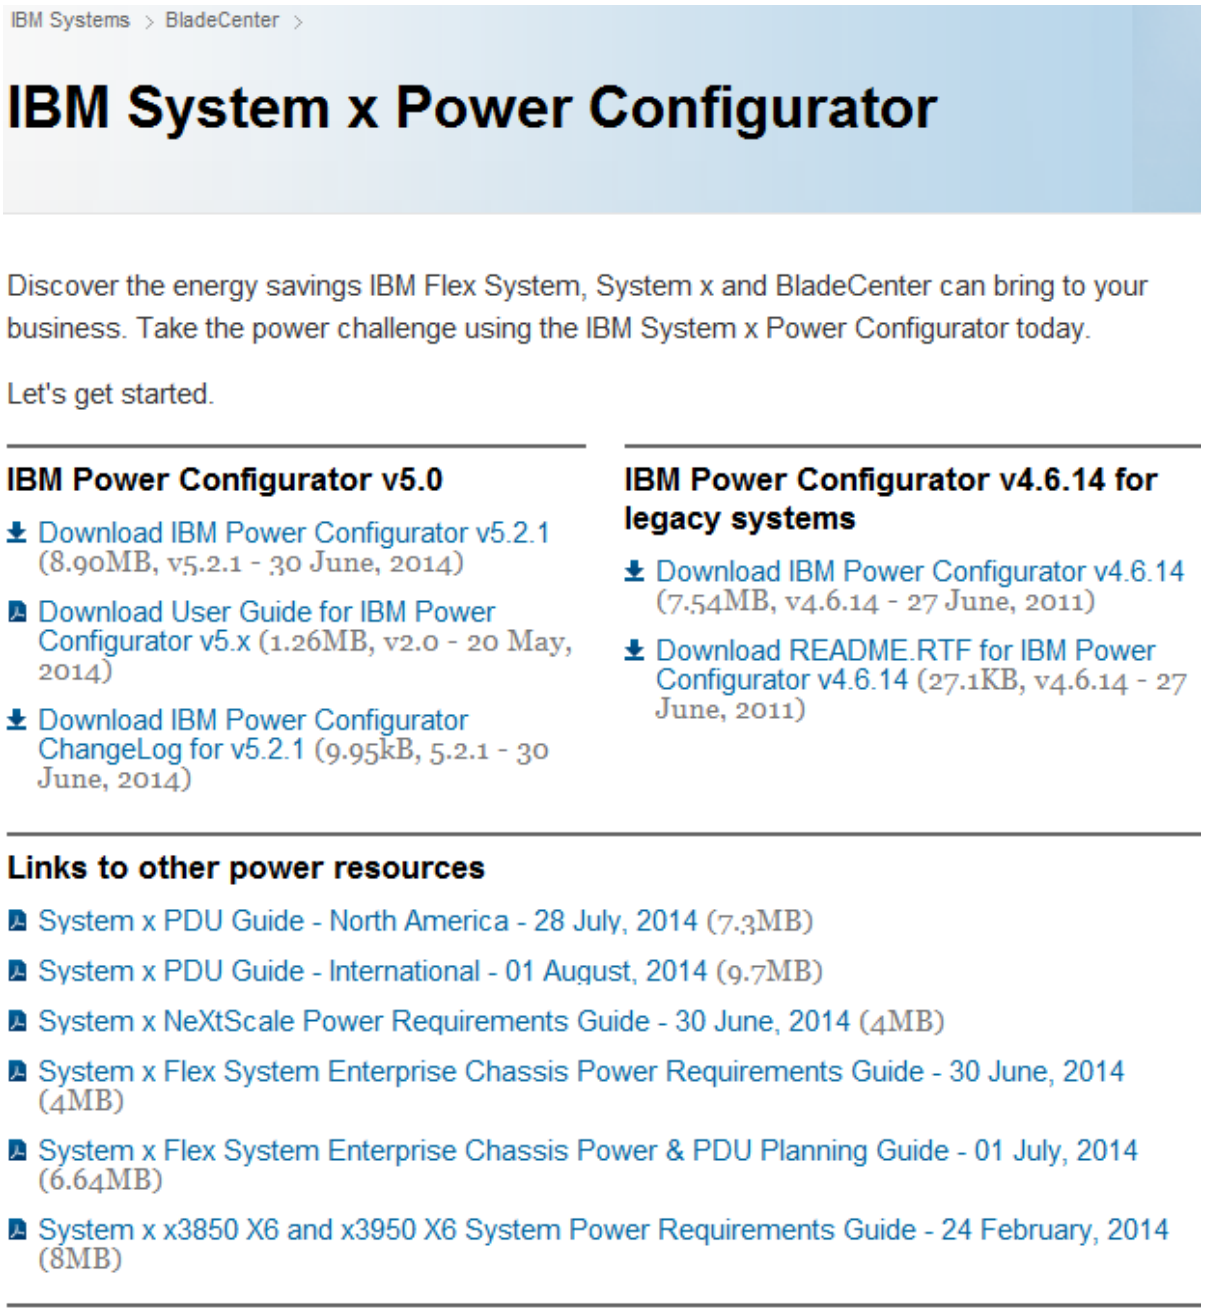

Figure 1: System x Power Configurator and Power Guides - download page

Click on the "Download System x Power Configurator v5.x" link (circled above) and save the compressed file to your system.

Launching the System x Power Configurator Tool

The System x Power Configurator is downloaded as a compressed file. Once downloaded to your system you will need to extract the contents of the compressed file in order to run the program.

In order to extract the contents of the file you will need to have appropriate software to allow you to decompress the folders contents. Follow the below instructions to decompress and launch the Power Configurator:

- 1. Locate the PowerConfigurator file you just downloaded (this will be in the form of .zip for windows, se for Linux, or .app for Macintosh).
- 2. Double click on the compressed file to invoke your decompressing software.
- 3. Extract the contacts of the compressed file to a dedicated folder on your system.
- 4. Once extracted, launch the tool by double clicking on the below file:

  PowerConfigurator.jar

NOTE: If you do not decompress and extract the files to your system the program may hang at the splash window and will not run. It may also cause the program to only load partial sections of the program and render it unusable.

Once the System x Power Configurator program has been invoked and loaded, the below window will appear where you can begin configuring your settings and hardware. See Figure 2 on page 10.

There are some functions in the tool that are currently being worked on. For a progress report on the functions and features of the tool, please email power@lenovo.com.

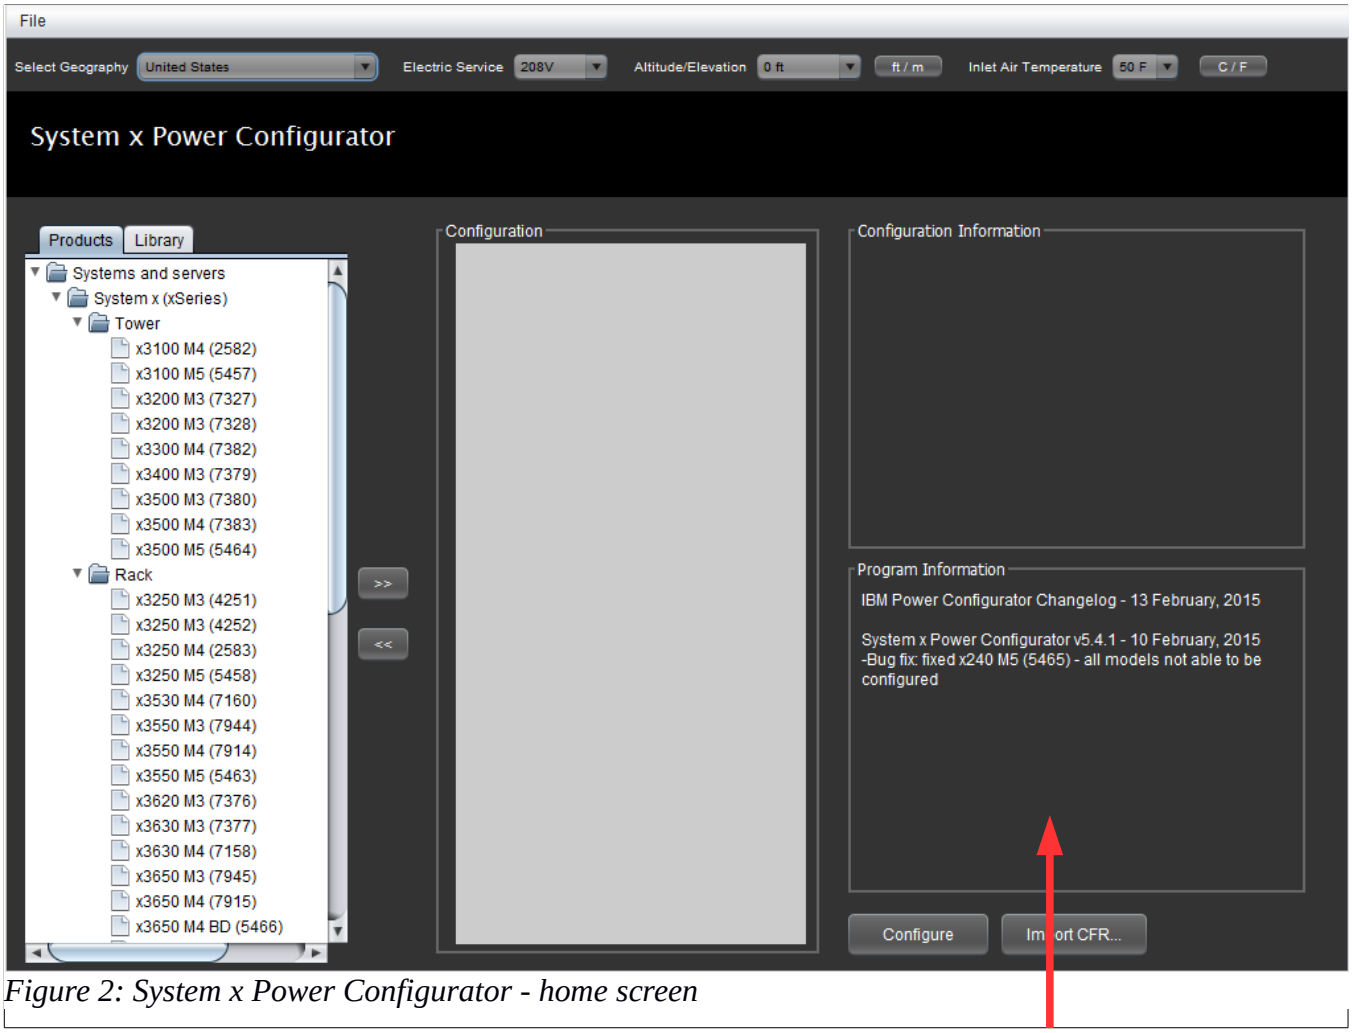

Figure 2: System x Power Configurator - home screen

From the home screen you can change your settings, select your hardware to configure and view the change log. The Change Log as seen in Figure 2 displays the date of release and the changes since the last version.

Configuring your Settings

Once launched, you will have the option to set your Geography, Electronic Service, Altitude/Elevation and Inlet Air Temperature for your power configuration. These options are selectable in the tool bar located at the top of the System x Power Configurator program. See Figure 3.

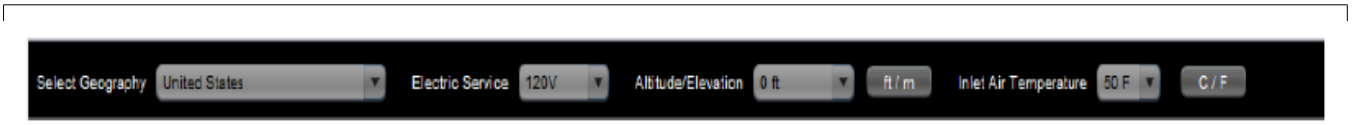

Figure 3: System x Power Configurator - tool bar settings

Selecting Geography

The Geography refers to the country that your hardware will be used in.
To set the geography of the hardware, select the drop down box from the below menu and click your location. Alternatively the tool will allow you to type ahead your country in the drop down menu for a quick selection. The default geography selection is the United States. See Figure 4.

Figure 4: Setting Geography

Selecting Electric Service

The Electric Service refers to the voltage available to you to power your hardware. The voltage selection in the menu will change depending on the geography selection. Power numbers will vary depending on your input voltage, so it is important you select the appropriate country and voltage for your configuration.

Use the drop down menu to change the voltage. You can also type ahead the number in the drop down menu for a quick selection. See Figure 5.

Figure 5: Setting Voltage

Selecting Altitude/Elevation (ft / m)

The Altitude/Elevation refers to the height of your hardware above sea level. The Altitude/Elevation of your hardware can effect the operating activity of your servers fans and PSUs which may slightly change your power results.

To set the Altitude/Elevation for your power configuration select the drop down box from the Altitude/Elevation menu and click the appropriate altitude/elevation. Alternatively you can select the drop-down menu and type the number for a quick selection. You can swap between feet and meter by clicking the ft / m button. See Figure 6.

Figure 6: Setting height

Selecting Inlet Air Temperature (C / F)

The Inlet Air Temperature refers to the environmental temperature that your hardware is operating in. To set the Inlet Air Temperature of your power configuration select the drop down box from the Inlet Air Temperature menu and click your temperature. Alternatively you can select the drop-down menu and type the number for a quick selection. You can swap between Celsius and Fahrenheit by clicking the C / F button. See Figure 7.

Figure 7: Setting Elevation

Using the System x Power Configurator Tool

The following section provides an example of configuring a rack/tower server and an example of configuring a chassis/node system.

Configuring System x Rack/Tower Servers

To configure the power for your System x Rack or Tower server, click to highlight your server from the left hand menu under the 'Products' tab menu. Once selected, use the >> key to add the server to the 'Configuration' panel. Hit configure to configure the servers hardware. See Figure 8.

This will launch the 'Server Builder' window where you can configure the internal options for your Rack/Tower server selection. See Figure 9 on page 14 for an image of the Server Builder page.

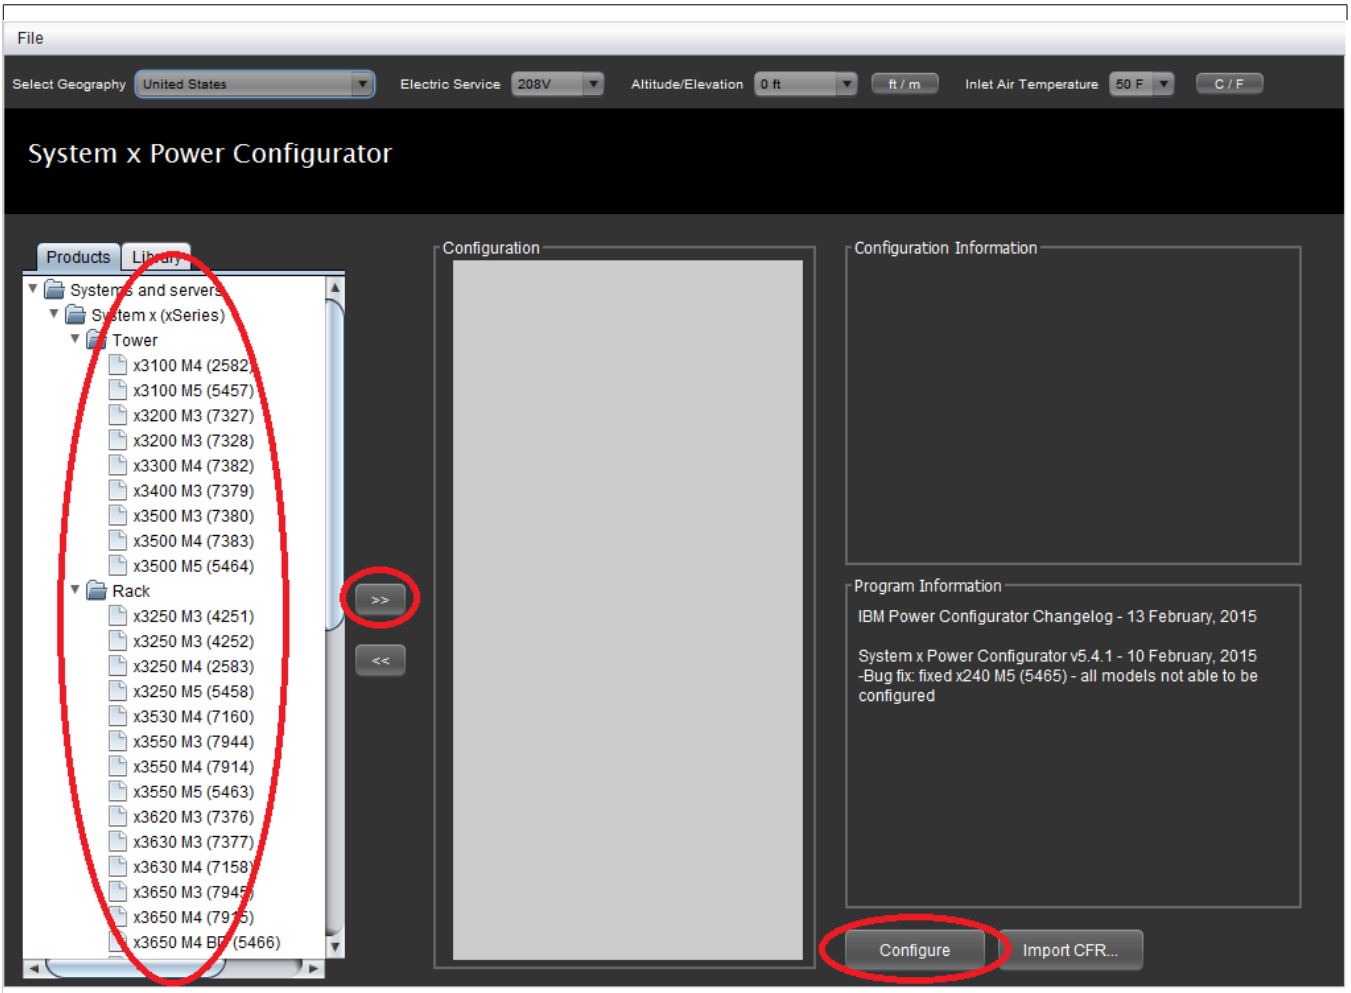

Figure 8: System x Power Configurator - home screen selection

Use the drop down menus to select processors, DIMMs, HDDs, PCI Adapters, power supplies, optical drive, and tape.

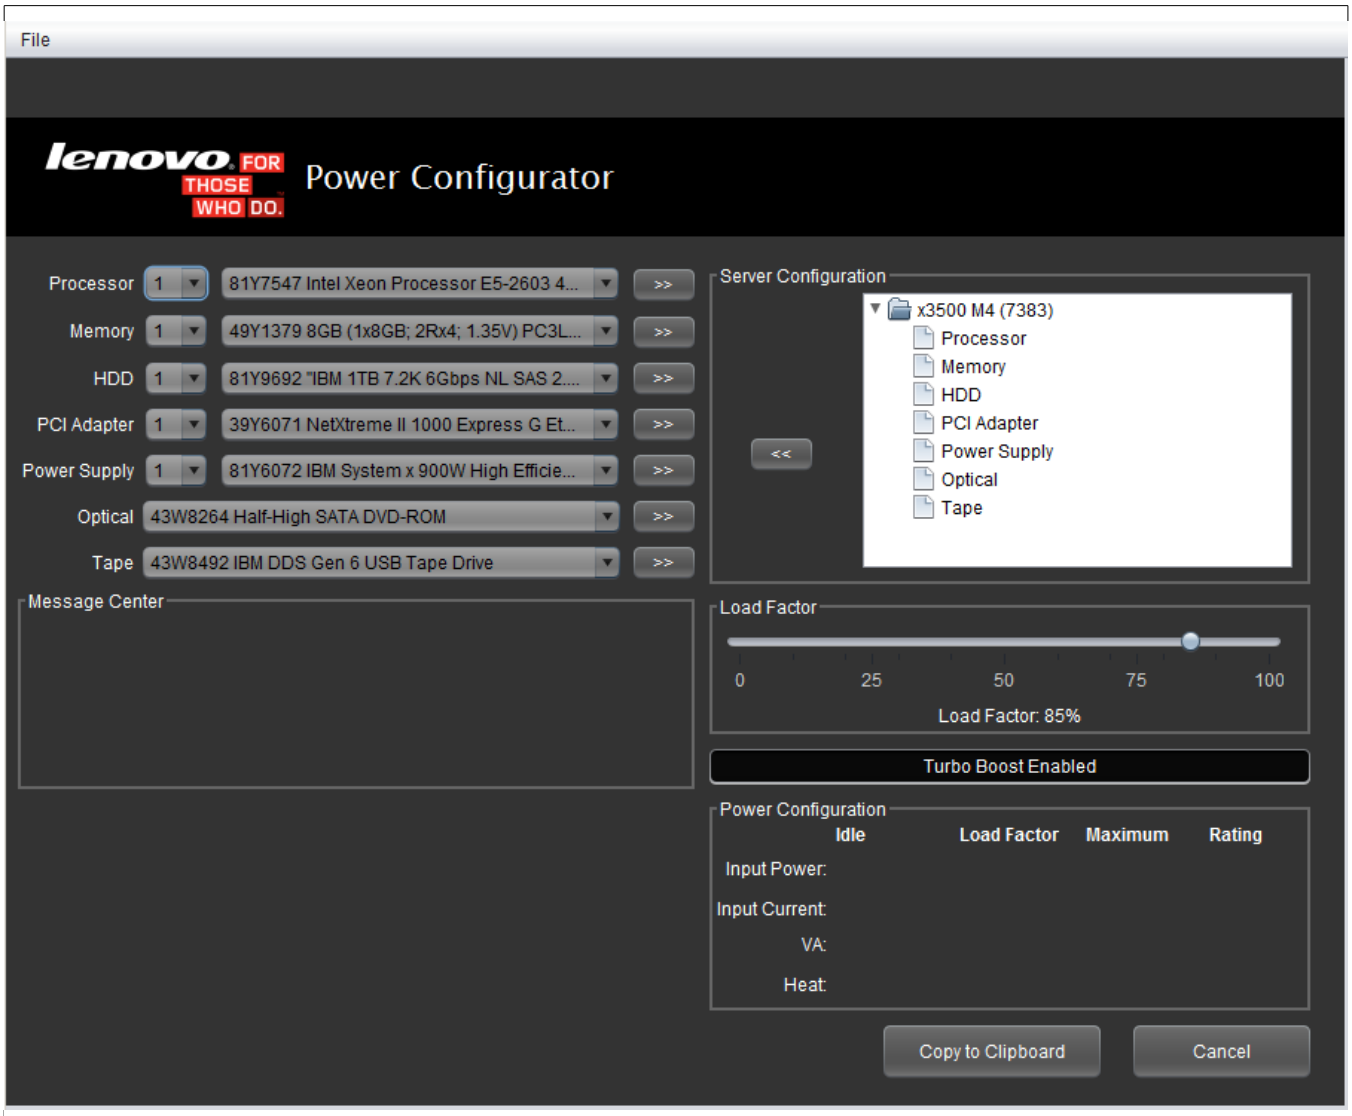

Figure 9: System x Power Configurator - Server Builder screen

Once you select your option use the >> button to add your selected hardware to the 'Server Configuration' panel. See Figure 10

You can specify the total number of individual items to add to the server using the drop down numbers in front of the descriptions.
If the drop down menu turns grey, this indicates the maximum number of items you can install has been reached.

Use << the button to remove items from the Server Configuration panel if required. You can remove all items under a heading in this panel by highlighting for example, the 'Processor' or 'Memory' heading and using the remove button.

Once all of your hardware has been added, your power consumption will be displayed in the 'Power Configuration' panel. See Figure 10.

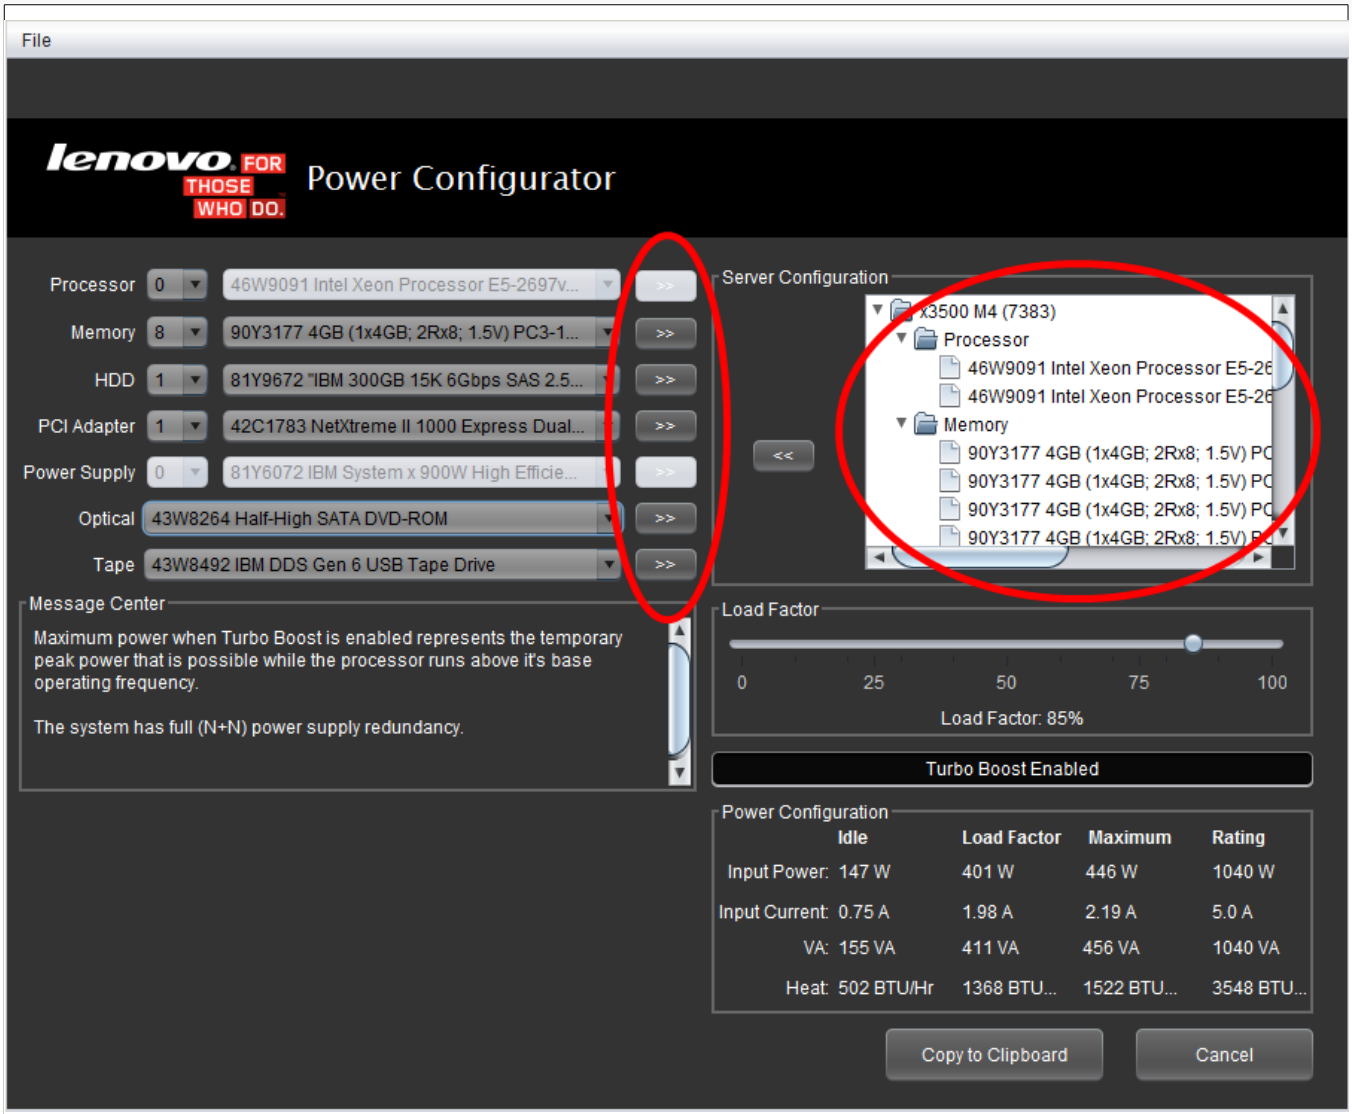

Figure 10: System x Power Configurator - Server Builder screen configuration example

Note: Your power values will not appear until you have selected at minimum processor(s), DIMM(s), and the power supply unit(s).

If you wish to start over or begin a new configuration, close the tool and relaunch the PowerConfigurator.jar file from the folder to begin a new configuration.

Configuring System x Chassis/Nodes

To configure the power for your System x Chassis and Nodes, click to highlight your Chassis selection from the left hand menu under the 'Products' tab menu. Once selected, use the ›› key to add the server to the Configuration panel. Hit configure to configure the chassis and node(s) hardware. See Figure 11.

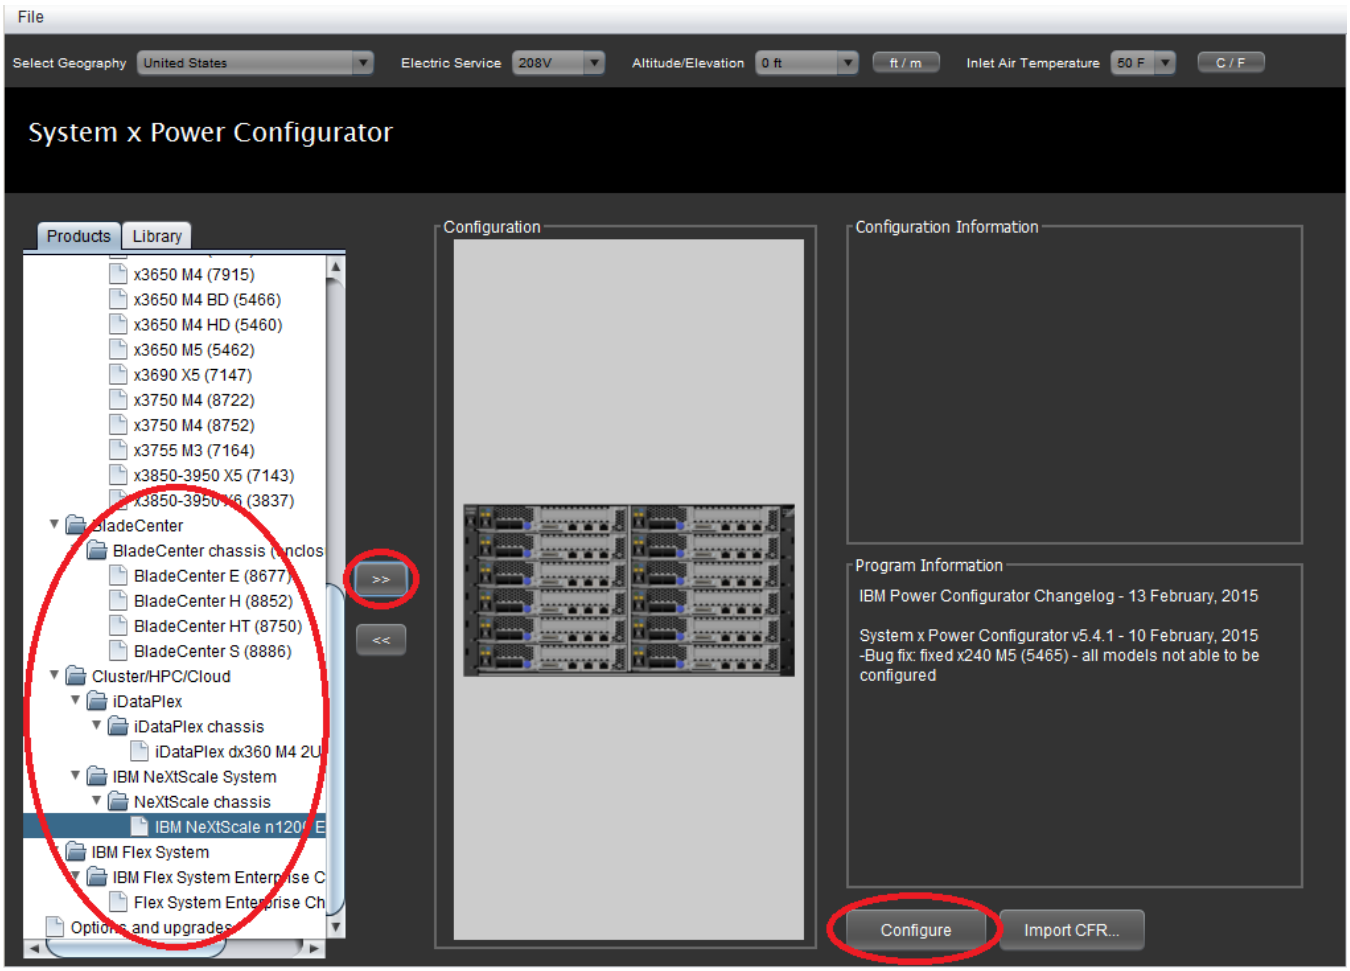

Figure 11: System x Power Configurator - home screen selection

This will launch the 'Chassis Builder' window where you can configure the internal hardware for your chassis including nodes, switches, management modules(s), power supply units (PSU), and fans for your chassis selection. See Figure 12 on page 17 for an image of the Chassis Builder page.

Configure the chassis' internal switches, PSU's and fans before you configure the nodes. To configure the chassis, click the 'Configure Chassis' button.

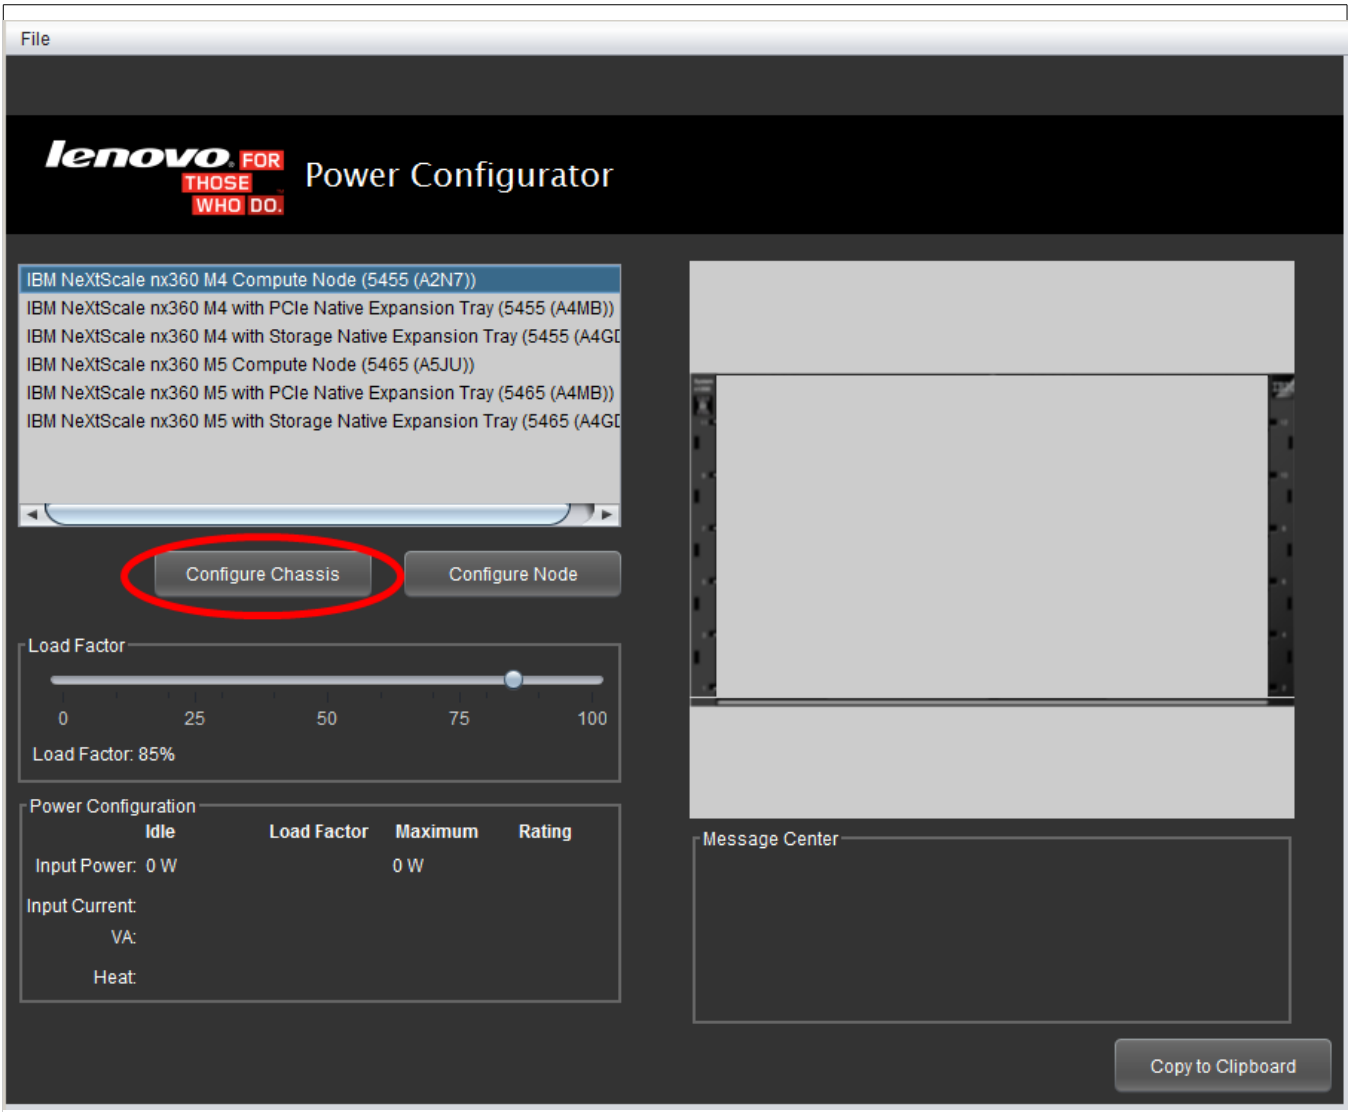

Figure 12: System x Power Configurator - Chassis Builder screen configuration example

Note: The Power Configuration panel will not display any power values until you have configured the chassis with PSUs first.

The 'Chassis Configurator' page is shown in Figure 13 on page 18.

Use the drop down menus to select power supply units (PSUs), and fans, and where applicable, I/O switch modules, and management module(s).

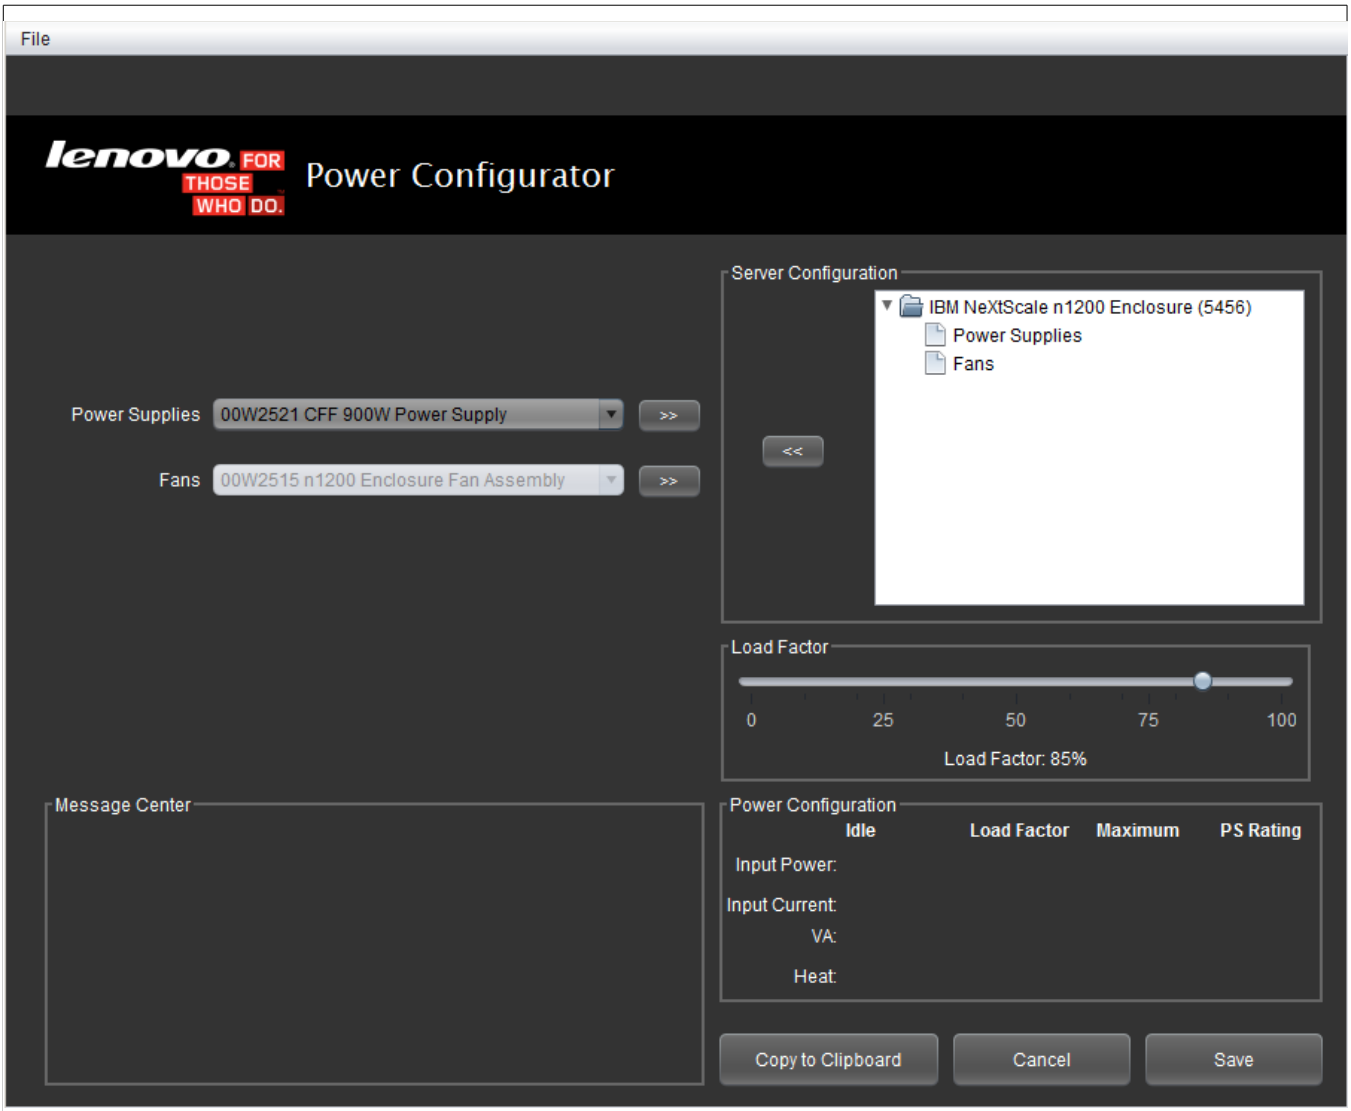

Figure 13: System x Power Configurator - Chassis Configurator screen example

Once you select your option 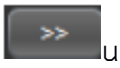 use the button to add your selected hardware to the 'Server Configuration' panel. See Figure 14 on page 19.

If the drop down menu turns grey, this indicates the maximum number of items you can install has been reached.

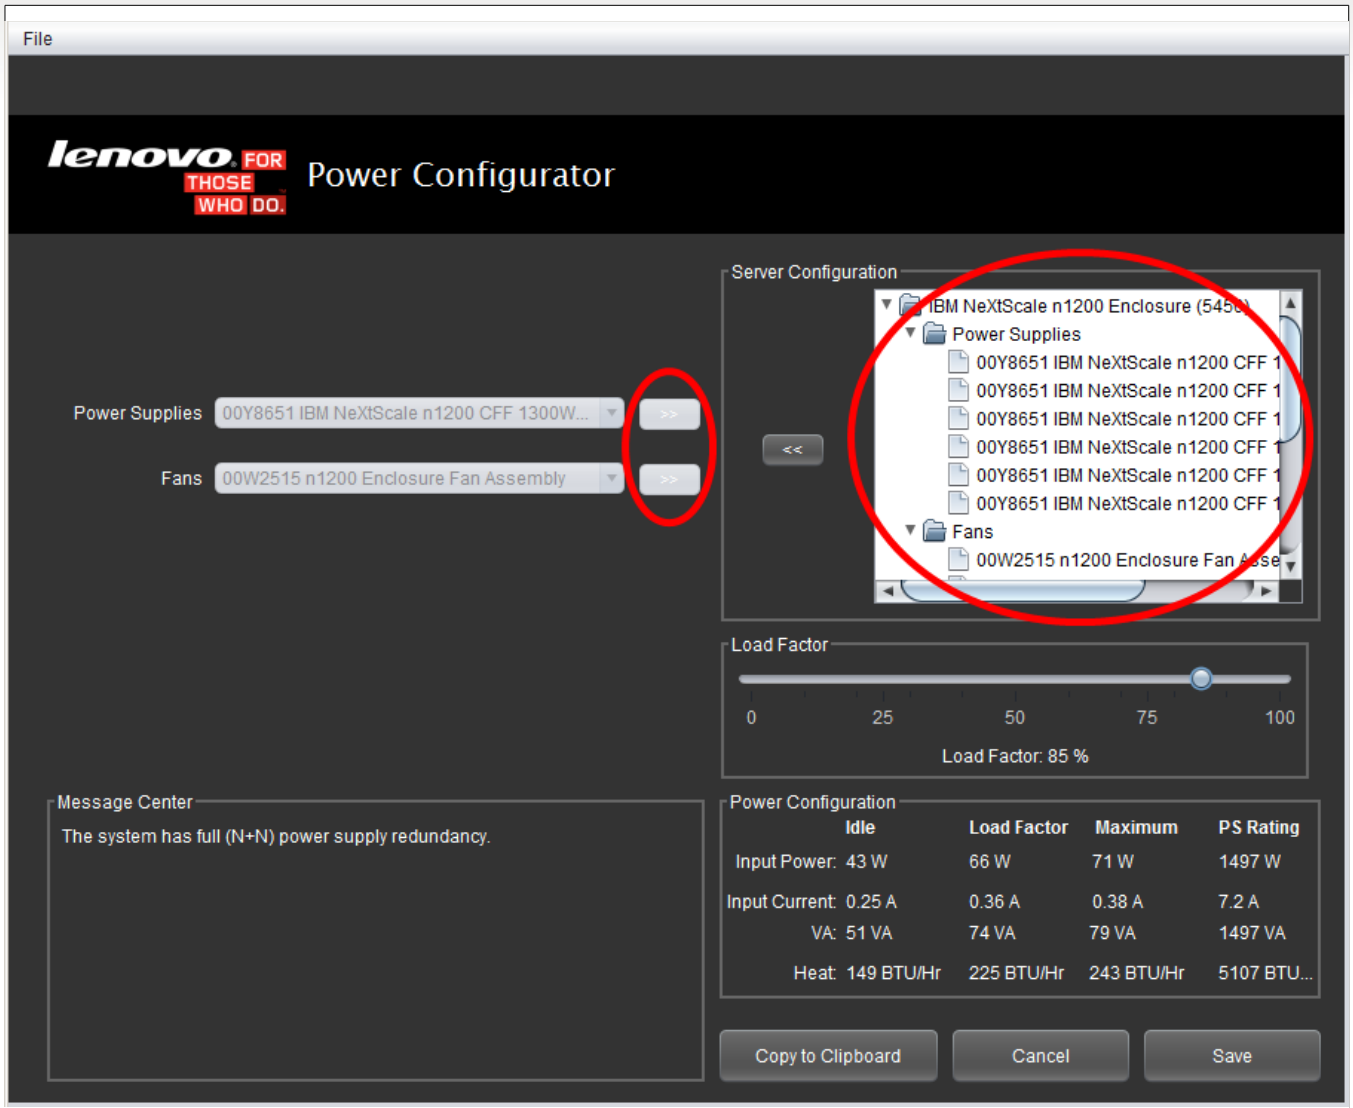

Figure 14: System x Power Configurator - Chassis configuration screen example

Once all hardware has been selected. Hit the 'Save' button. This will save the chassis hardware to the Chassis Builder page as seen in Figure 15 on page 20.

The power numbers for the chassis will be displayed in the 'Power Configuration' window.

Once the chassis has been configured you can begin to configure the nodes. Click on the 'Configure Nodes' button, circled in Figure 15.

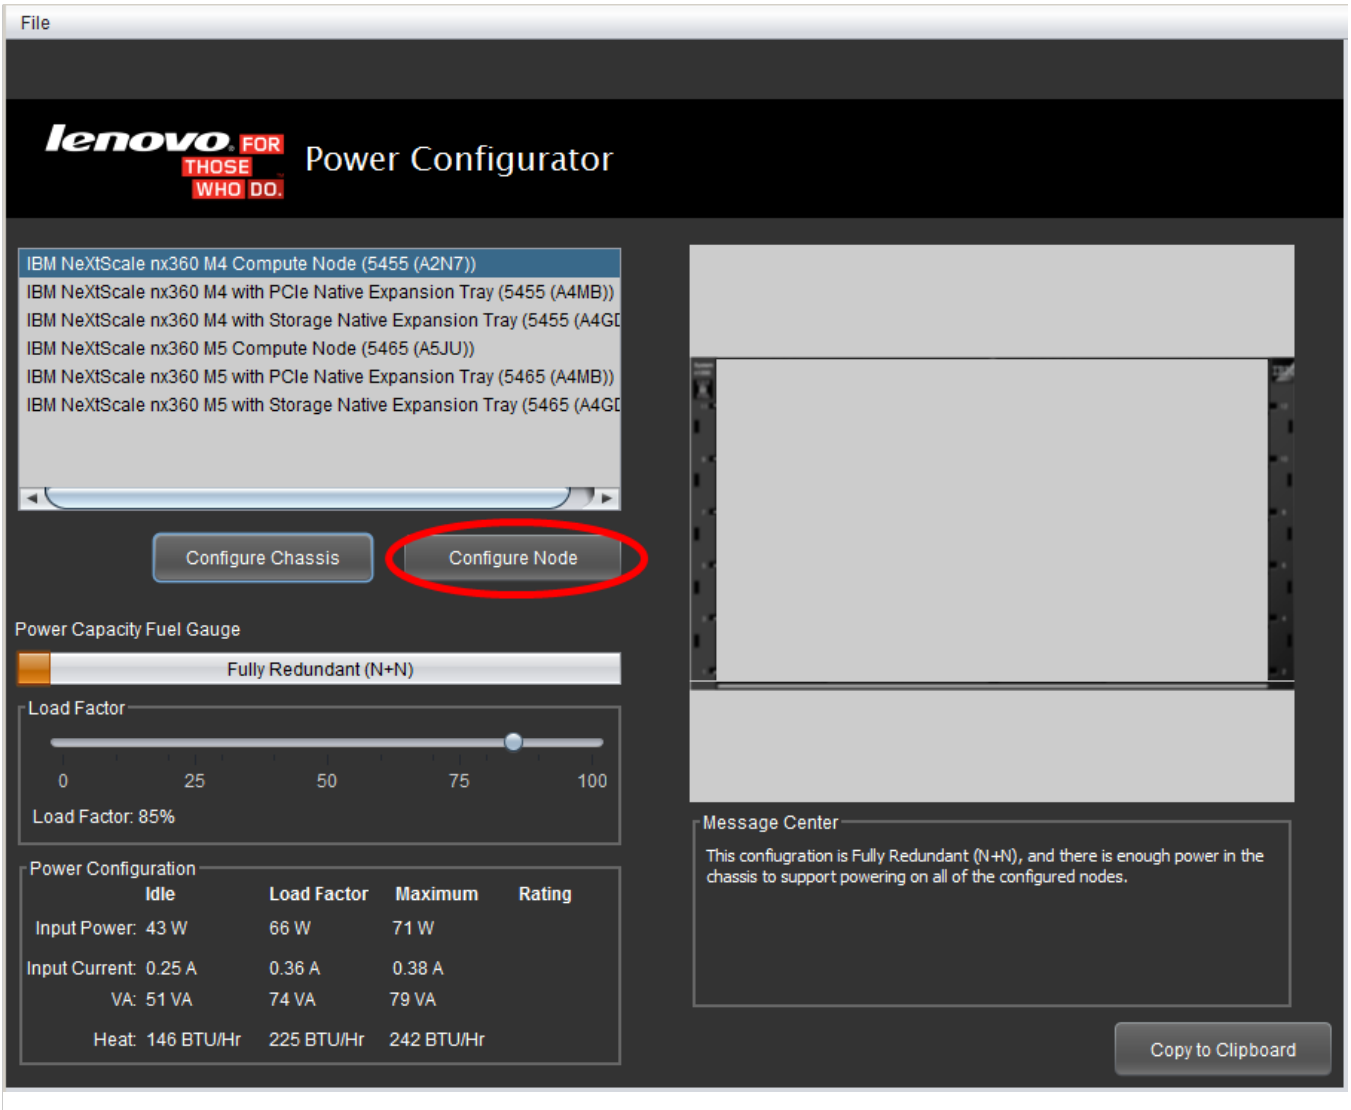

Figure 15: System x Power Configurator - Chassis Builder screen configuration example

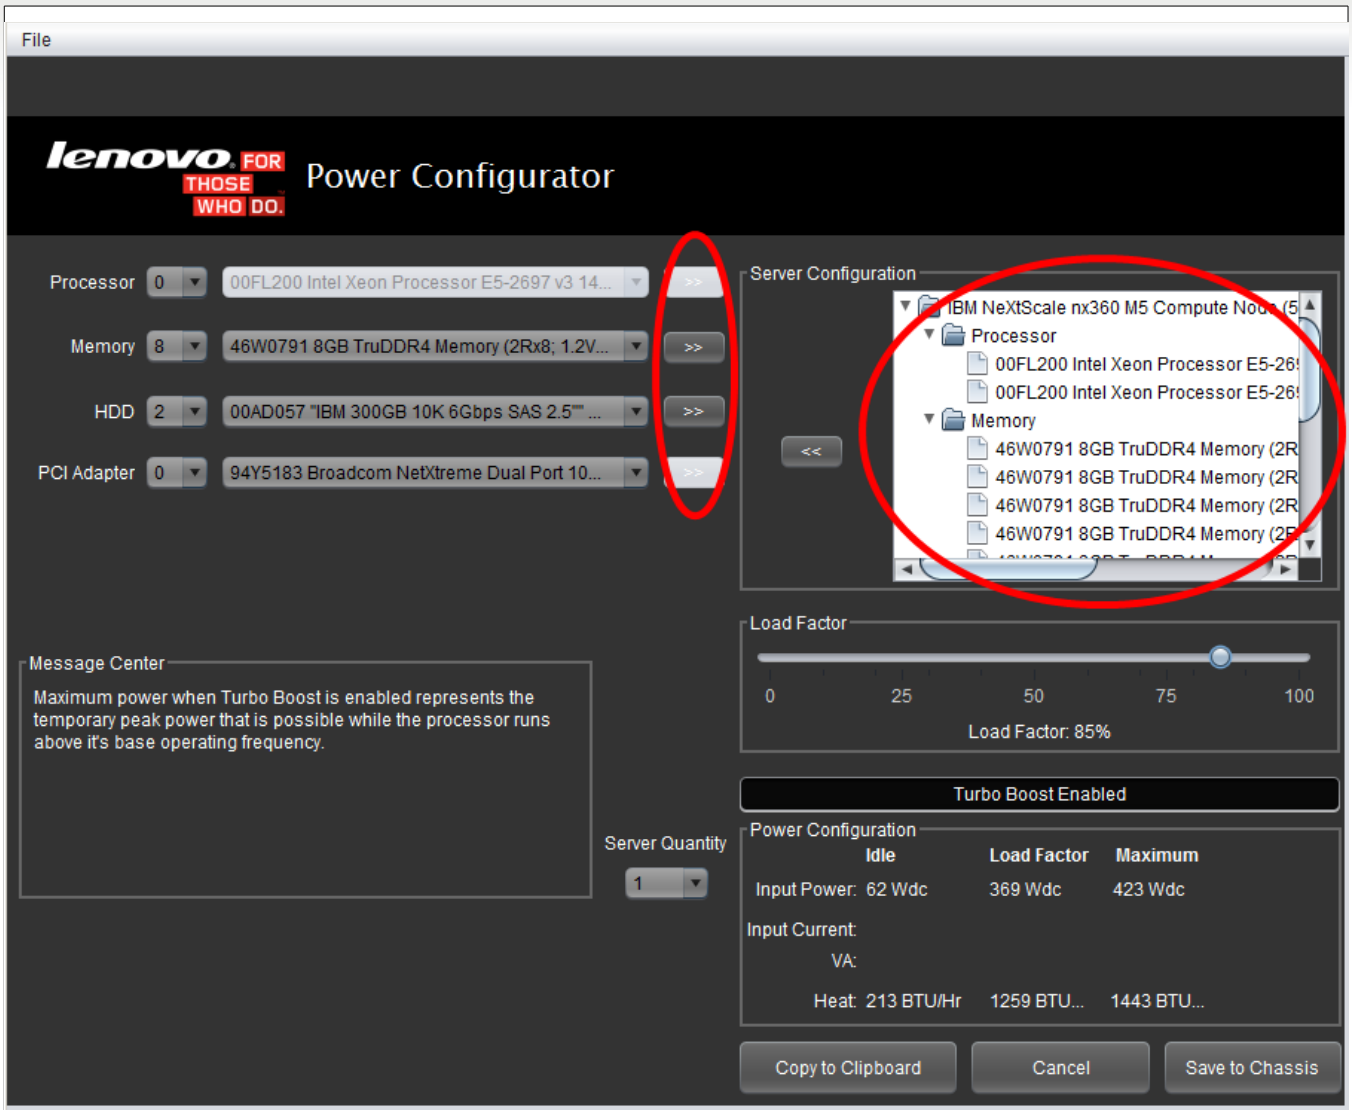

Figure 16: System x Power Configurator - Node Builder screen configuration example

The 'Server' Builder' page will be displayed as seen in Figure 17 on page 22.

You will also notice a 'Power Capacity/Fuel Guage' bar indicator appear. This function is discussed in the Power Capacity/Fuel Gauge Indicator section on page 28.

Note: Not all configurations will have a power capacity bar indicator. This only applies to specific hardware such as NeXtScale.

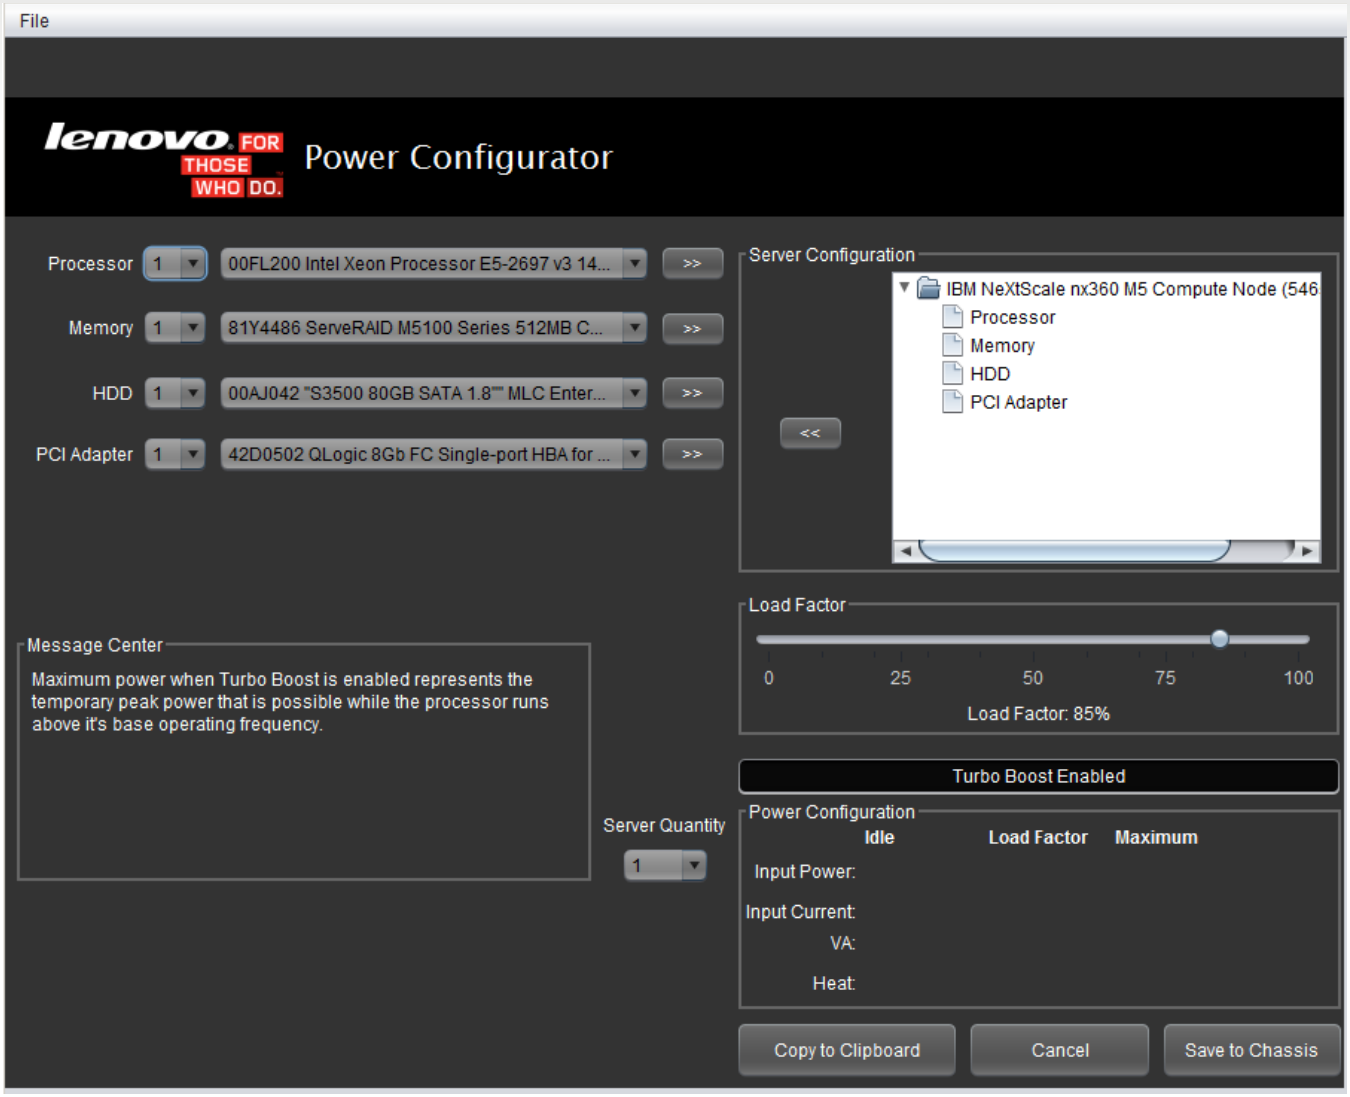

Figure 17: System x Power Configurator - Node Builder screen configuration example

Use the drop down menus to select processors, DIMMs, HDDs, and PCI Adapters.

Once you select your option 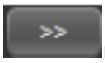

use the button to add your selected hardware to the 'Server Configuration' panel. See Figure 16 on page 21.

You can specify the total number of individual items to add to the server using the drop down numbers in front of the descriptions.

If the drop down menu turns grey, this indicates the maximum number of items you can install has been reached.

If you are installing multiple of the same node change the 'Server Quantity' at the bottom of the page.

The power values will not display until you have configured processor(s) and memory. The power results for the individual node will appear in the 'Power Configuration' panel, as seen in Figure 16. Once the node has been configured click the 'Save to Chassis' button.

Note: The configuration will be saved and the total value of the power numbers calculated for the node(s) and chassis will be displayed on the 'Chassis Builder' page. See Figure 18 on page 23.

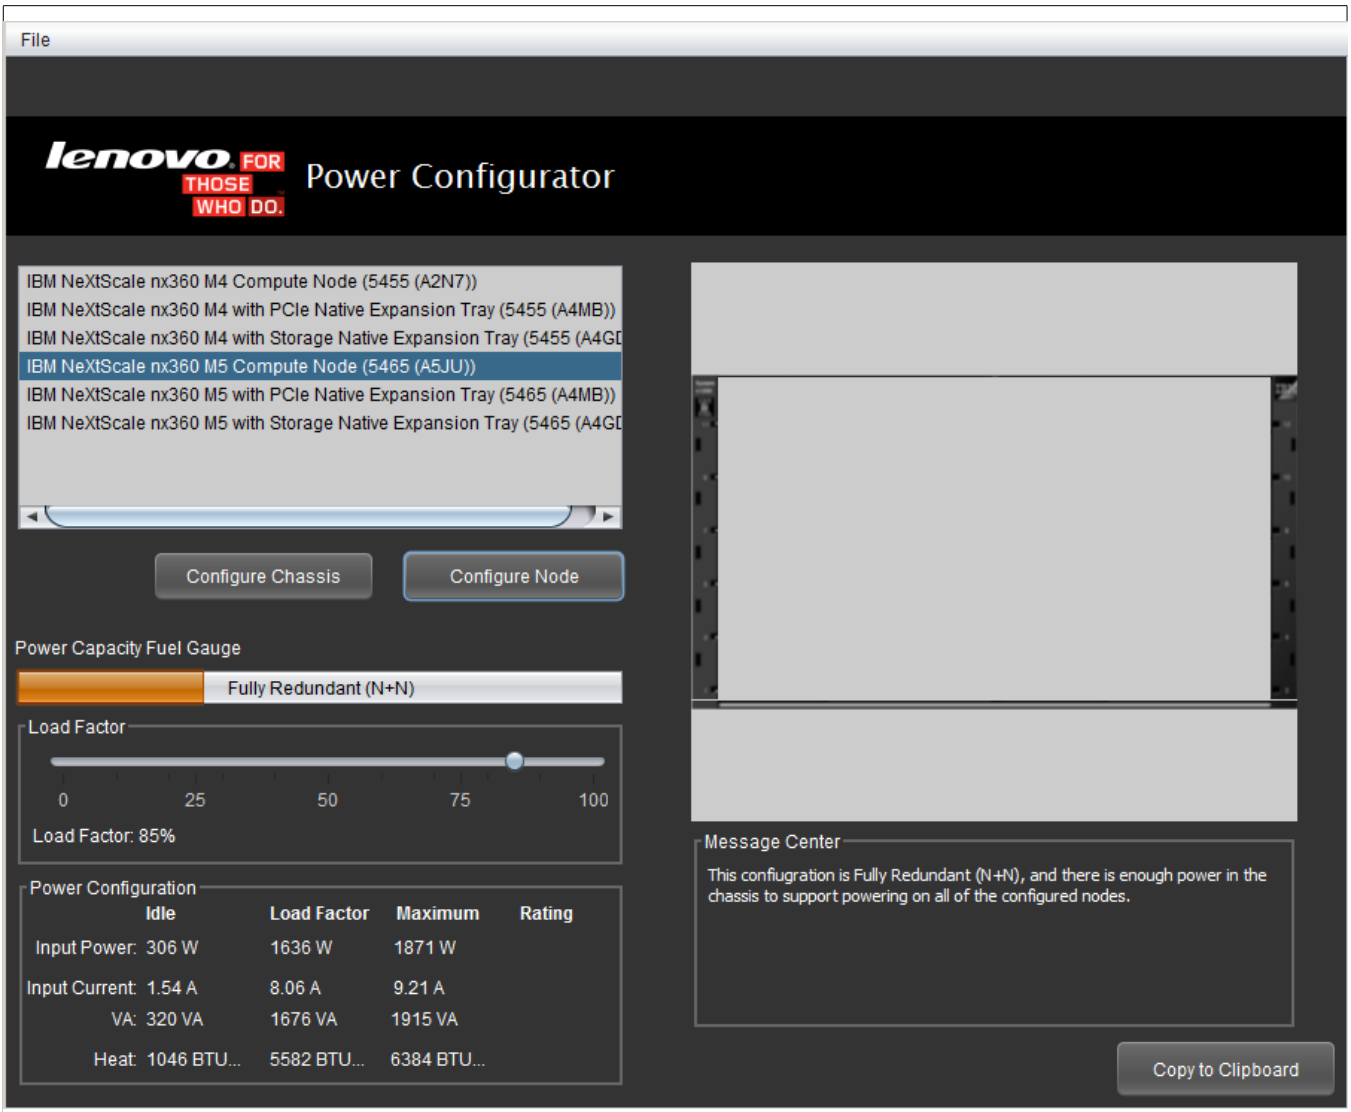

Figure 18: System x Power Configurator - Chassis Builder screen configuration example

Configure additional nodes as needed and save to the chassis. The power numbers will update on the Chassis Builder page with each new node added.

Note: The picture will not display the nodes added. This is a feature enhancement that will be enabled in an upcoming release of the tool.

Functions and Features of the System x Power Configurator

This section discusses the functions and features of Power Configurator. For additional information on features supported please contact power@lenovo.com.

Load Factor

By default the 'Load Factor' for every configuration is set to 85%.

You can adjust the load factor of the server by sliding the dial located in the 'Load Factor' panel. See Figure 19. You should adjust this setting to suit your hardware's work load.

As you change the load factor of the server the power numbers in the 'Power Configuration' panel will automatically update to reflect the new load.

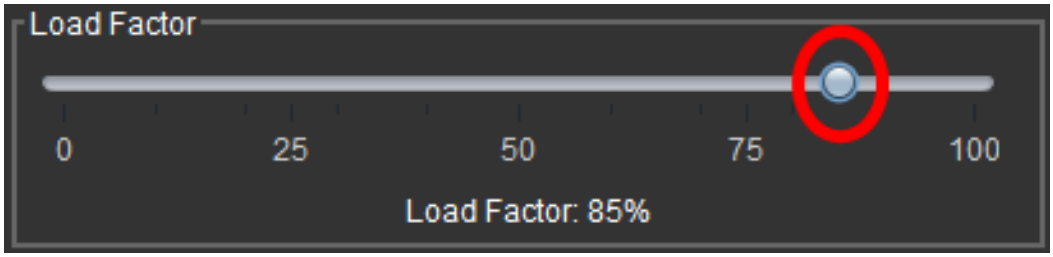

Figure 19: System x Power Configurator - Load Factor

Turbo Boost

Depending on current workload, Turbo Boost technology allows a dynamic increase in the clock speed of the active cores to gain a performance boost. Turbo Boost can be engaged with any number of cores enabled and active. While it will increase the performance of both multi-threaded and single-threaded workloads, the system will also draw more power due to the increased use of the processor cores.

The Power Configurator allows you to configure the power of your system with and without Turbo Boost enabled and adjust the power draw appropriately. Figure 20 displays the Turbo Boost button.

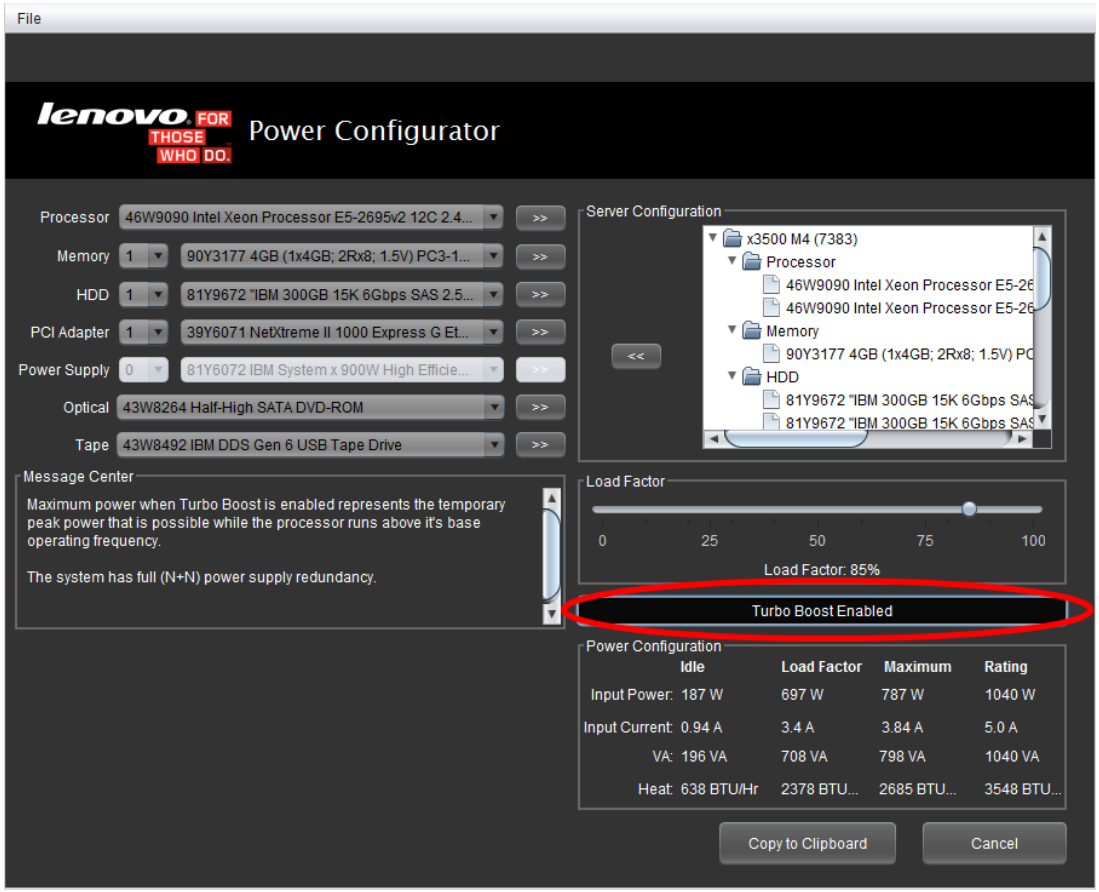

Figure 20: System x Power Configurator - Turbo Boost button

By default Turbo Boost is **enabled** once you start configuring your system in Power Configurator. Ensure you click the button to **disabled** it if you do not intend to use Turbo Boost.

Note: When temperature, power, or current exceeds factory-configured limits and the processor is running above the base operating frequency, the processor automatically steps the core frequency back down to reduce temperature, power, and current. The processor then monitors temperature, power, and current, and re evaluates.

The Message Center

When configuring hardware in the Power Configurator the 'Message Center' panel will update with information on your systems power supply redundancy based on the hardware and number of PSUs installed. An example of where to find the 'Message Center' is seen below in Figure 21.

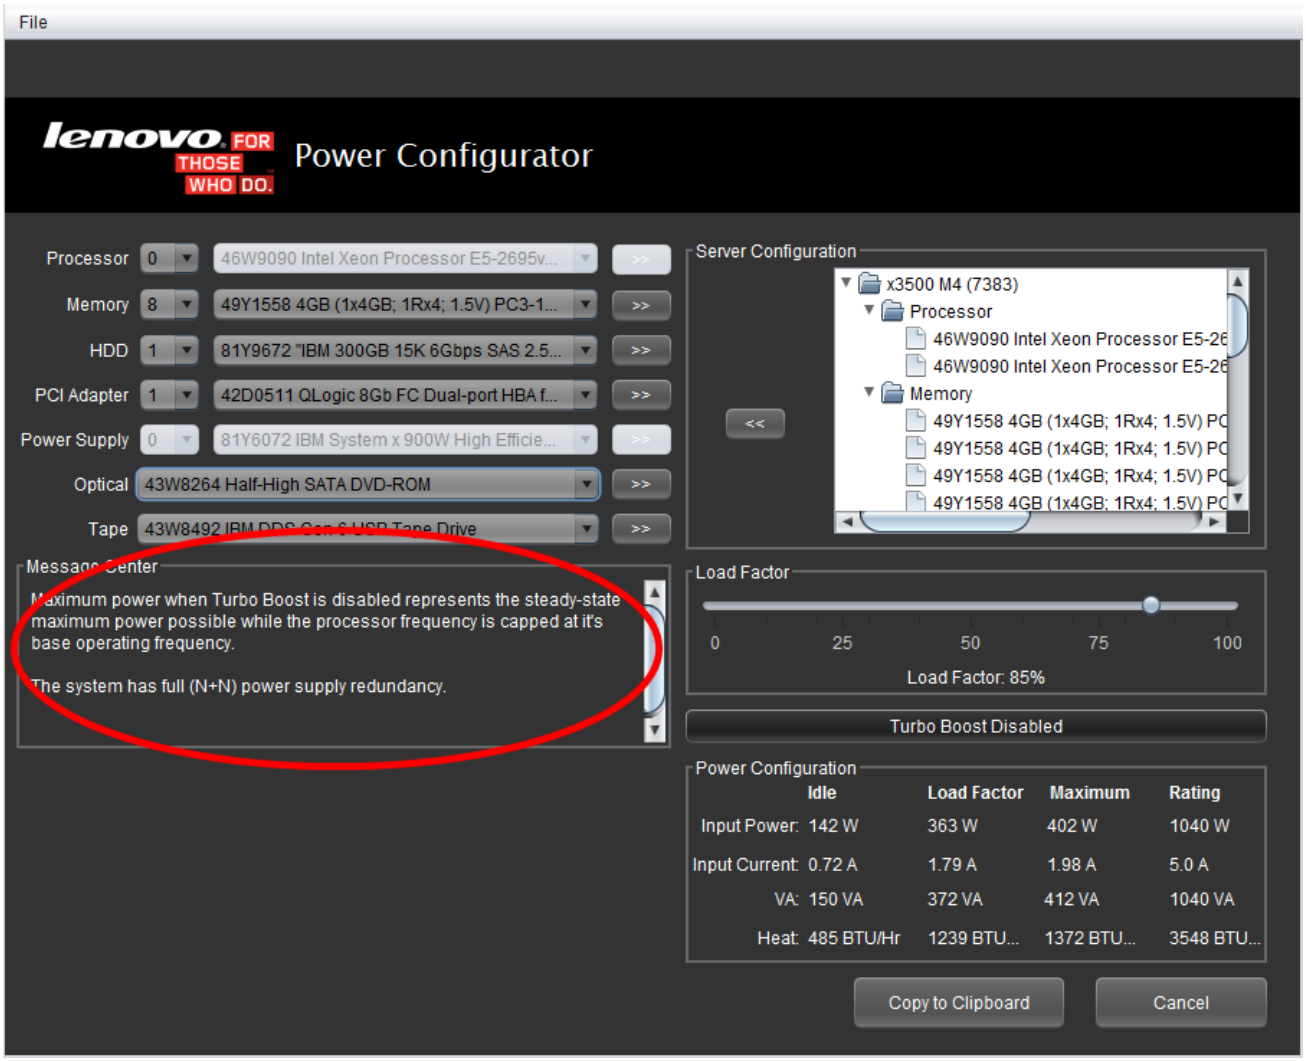

Figure 21: System x Power Configurator - Message Center

Copying Results to Clipboard

Once you have configured your hardware in the tool, you have the option to copy the power results to our clipboard. The following figures is an example of copying the power results of a tower server to clipboard and the resulting output pasted to a spreadsheet. See Figures 22 and 23.

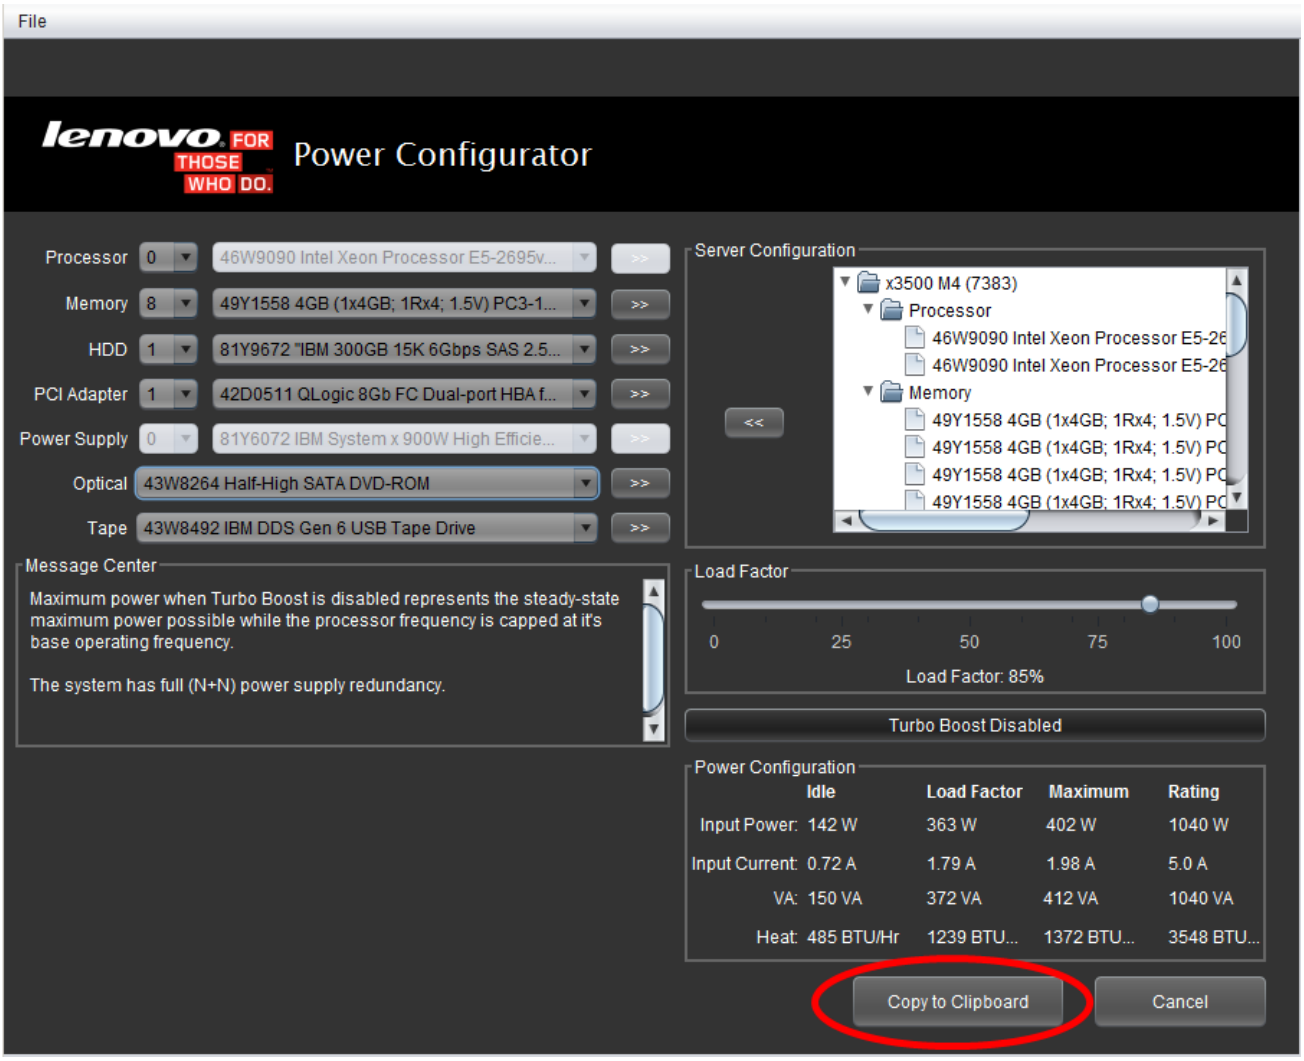

Figure 22: System x Power Configurator - copy to clipboard

Country:	United States			
Operating Voltage:	120 V			
Ambient Temperature:	50 F			
Elevation:	0 ft			
Configured Machine Type:	7383			
	Idle	Load Factor	Maximum	Rating
Input Power (W):	77 W	249 W	280 W	1224 W
Input Current (A):	0.72 A	2.15 A	2.42 A	10.2 A
Apparent Power (VA):	86 VA	258 VA	290 VA	1224 VA
Heat Generation (BTU/Hr):	264 BTU/Hr	849 BTU/Hr	956 BTU/Hr	4176 BTU/Hr

Figure 23: System x Power Configurator - copy to clipboard results

Power Capacity/Fuel Gauge Indicator

Certain chassis/nodes will have a *'Power Capacity/Fuel Gauge'* bar as an indicator for the redundancy level that should be employed in the IMM2.1.

NOTE:

The Power Capacity Fuel Gauge and the actual power calculated from Power Configurator are not linked and do not represent each other.

The Power Capacity Fuel Gauge represents the *Power Budget* number based on worse case scenarios of the hardware installed. This number is determined at system start up by the Node Manager Power and Thermal Utility Tool. (Refer to the Node Manager Power and Thermal Utility Tool section for additional details).

The Fuel Gauge is only shown in Power Configurator to help you determine the minimum policy that needs to be set in place to ensure the system will boot. Refer to the individual system Power Planning Guide for additional details. The Power Planning Guides can be downloaded from:
http://www-03.ibm.com/systems/bladecenter/resources/powerconfig.html

The power numbers displayed in Power Configurator represents the estimated power that can be drawn from the specific hardware installed in the system.

The bar indicator will advise of the level of redundancy needed to be employed at the IMM level for your particular configuration. The bar indicator will update as you add more hardware/nodes to the enclosure. It will notify if the configuration is in one of the below states:

- **N+N** – fully redundant,
- **N+1** – 1 spare PSU,
- **N+1 with OVS –** over-subscription of power may occur,
- **N** – no-redundancy, or
- **Not Supported** – not enough PSUs/power to power on all nodes.

Or:

- **Redundant without Throttling –** Throttling will not be allowed
- **Redundant with Throttling –** Throttling will be allowed
- **Non-Redundant –** System will not be power redundant

The 'Message Center' will also update as hardware is added, advising on ways to make the configuration fully redundant (example: "Remove N # of nodes to make the configuration fully redundant").

See Figure 24 on page 29 for an example of the fuel gauge.

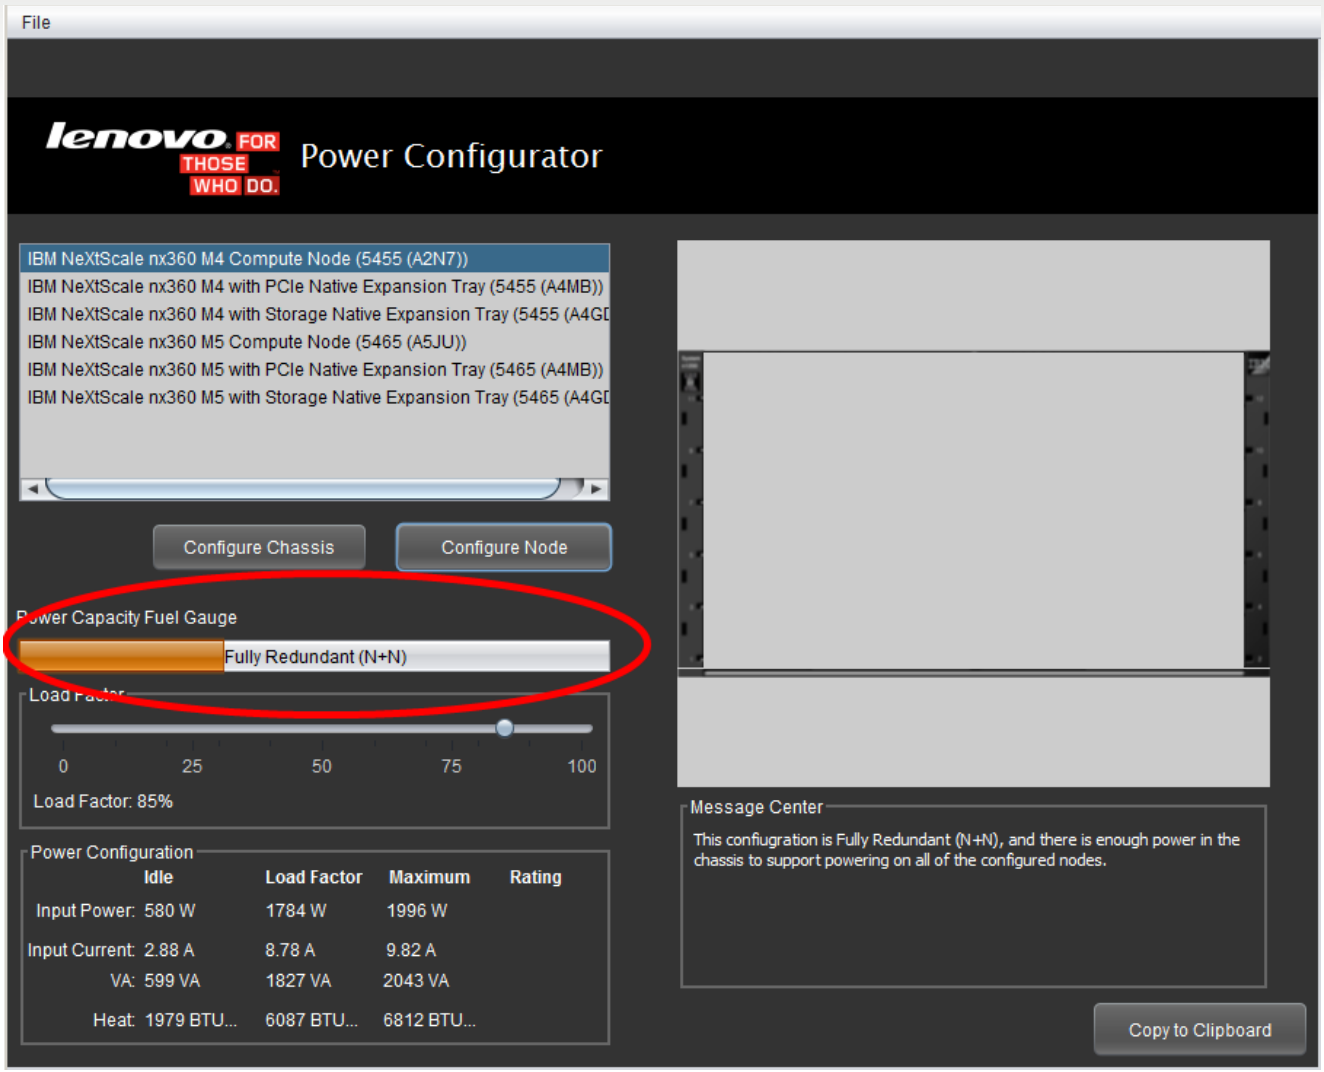

Figure 24: Power Capacity/Fuel Gauge indicator

Import CFR file

The Power Configurator allows the import of Blue Horizon and x-Config CFR files for power validation. The purpose of importing a CFR into the Power configurator tool is to execute power evaluations on the configuration. This evaluation will determine if the power supplies configured in each system within the CFR file is a supported configuration or if there may be potential power problems.

To import a CFR file click either File › import or the import button found on the bottom right of the home screen, as seen below.

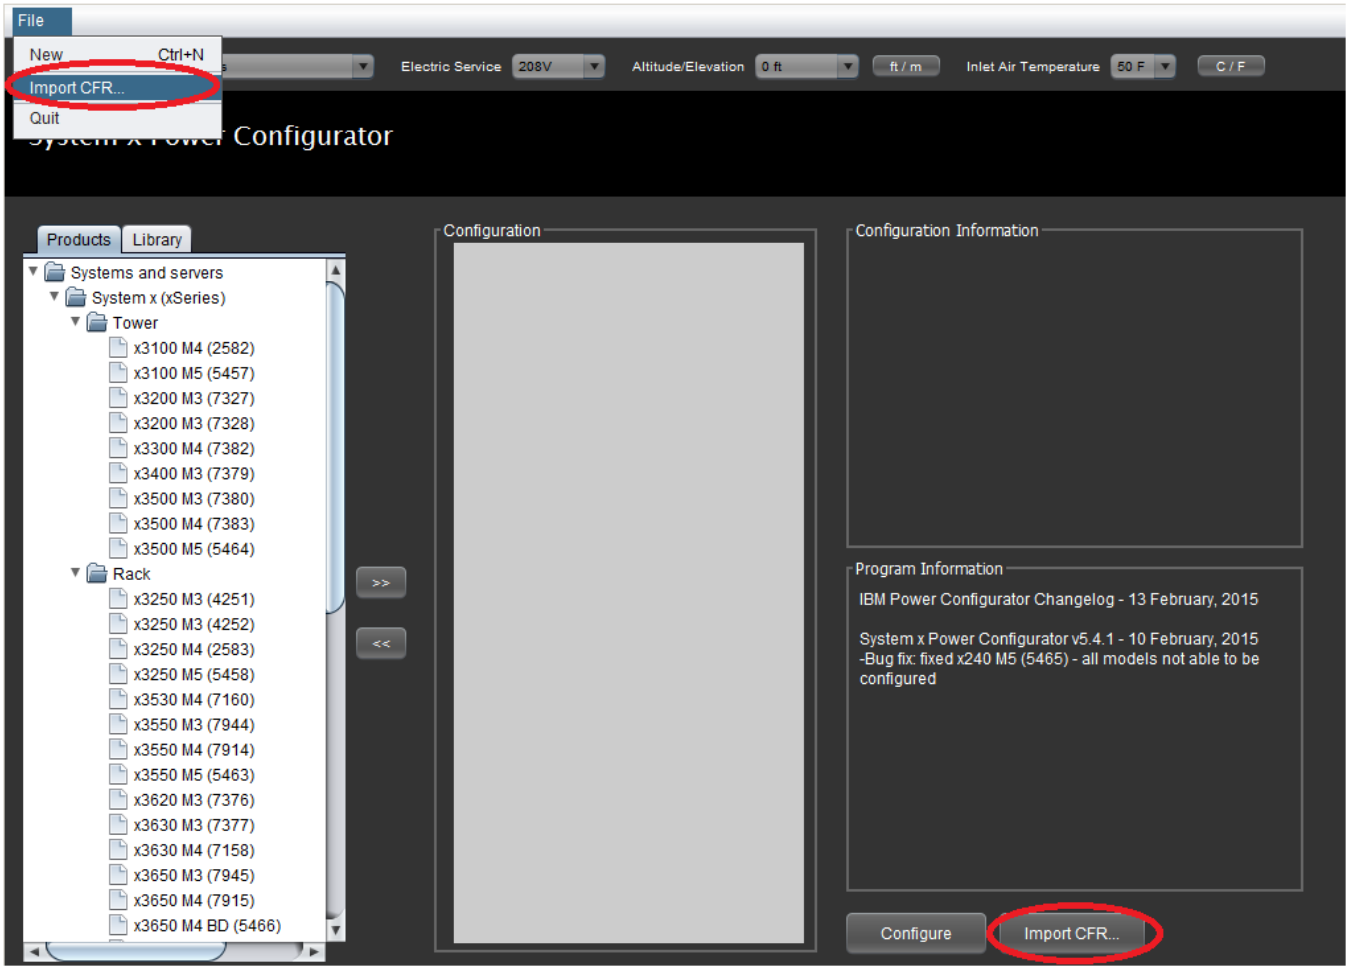

Figure 25: Import CFR file buttons

Once import option has been selected, locate the CFR file you require power validation on and select open.

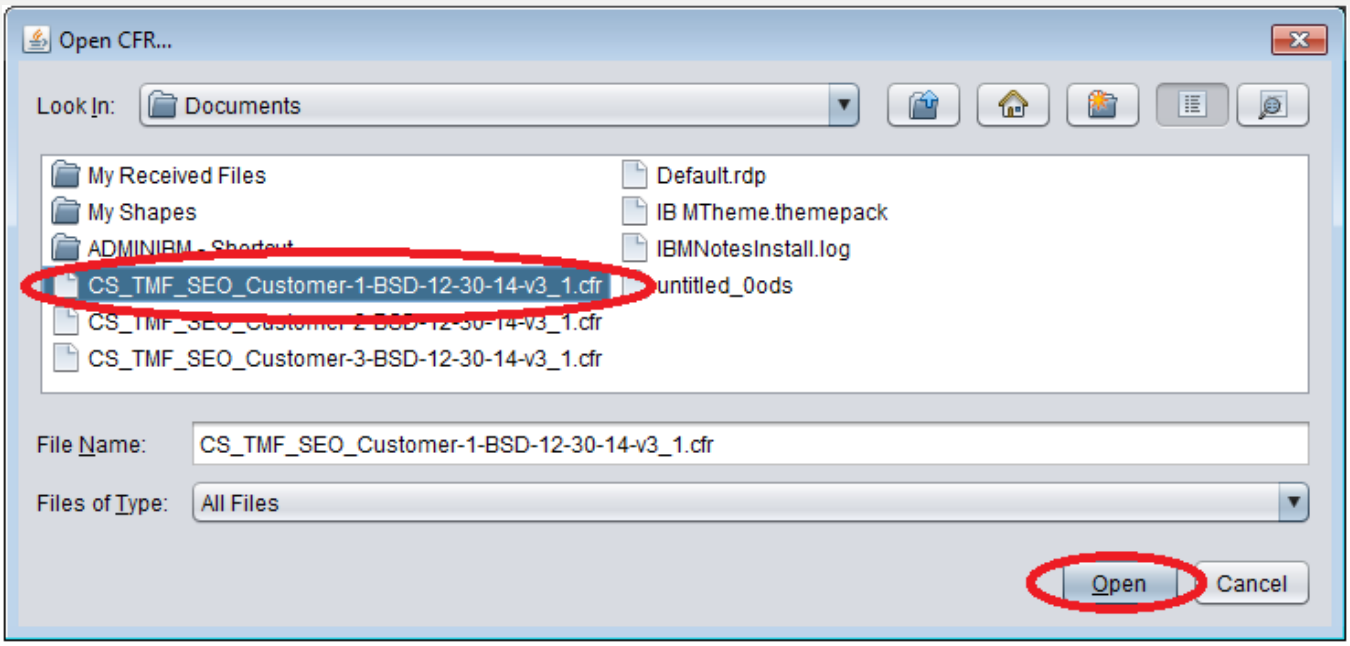

Figure 26: CFR open screen

The file will be imported into the tool and automatically analyzed. The status and any errors or warnings will be displayed in the Configuration Information window, as seen here.

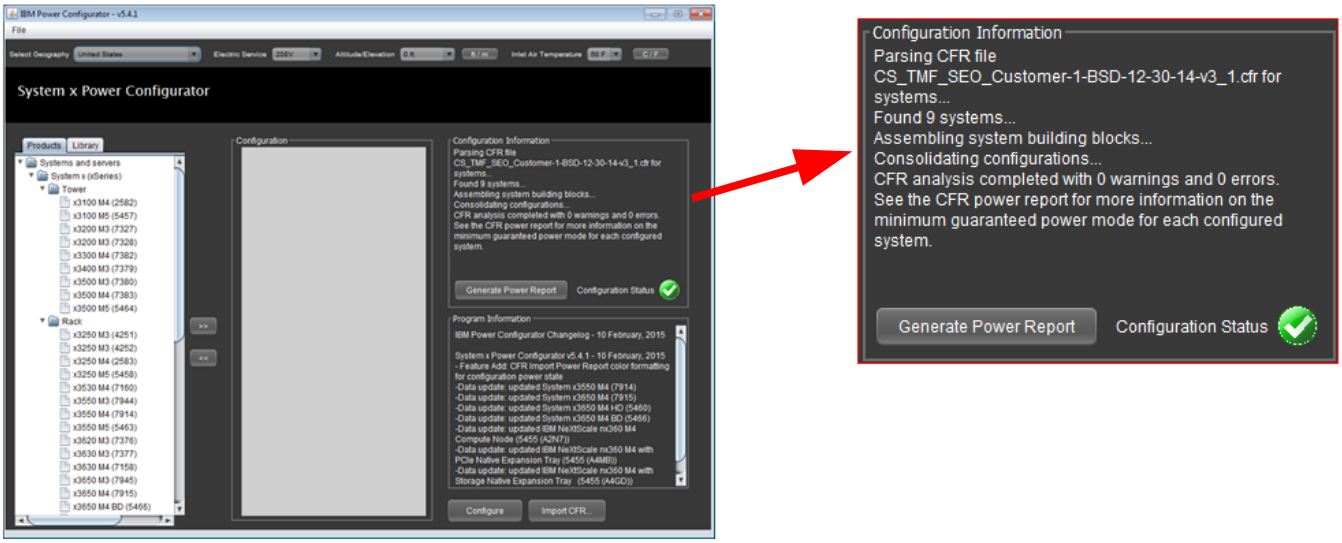

Figure 27: Configuration Information screen with imported CFR file

Note that this function supports multiple configurations in the one CFR file. The tool will analyze each configuration separately and report the status of each system in the power report.

There are 3 possible status results the tool will produce. These include: A green tick, yellow exclamation mark, or a red cross.

- A green tick indicates the configuration has enough power supplies or sufficient capacity to power the system.
- A yellow exclamation mark indicates a warning that the configuration is not redundantly powered or only has one power supply installed. This means that there is sufficient power to operate the system, but the system may shut down if a power event occurs. The system may also operate in a power managed state if the load exceeds the power supply oversubscription level.
- A red cross indicates an error that there is not enough power with the installed power supplies to support operating the system in a non-power managed state. There may also not be enough power capacity for the system to power on.

Figure 27 on page 31 shows the Power Configurator Configuration Information window with a CFR import with a green tick, indicating no errors.

Generating a Power Report

To view additional information about the configuration, click the Generate Report button. An example of a report with no errors is shown below.

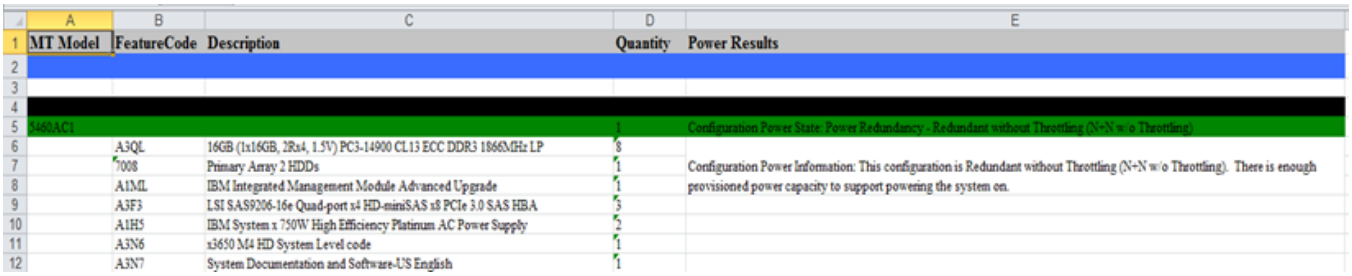

	A	B	C	D	E
1	MT Model	FeatureCode	Description	Quantity	Power Results
2					
3					
4					
5	5460AC1			1	Configuration Power State: Power Redundancy - Redundant without Throttling (N+N w/o Throttling)
6		A3QL	16GB (1x16GB, 2Rx4, 1.5V) PC3-14900 CL13 ECC DDR3 1866MHz LP	8	
7		7008	Primary Array 2 HDDs	1	Configuration Power Information: This configuration is Redundant without Throttling (N+N w/o Throttling). There is enough
8		A1ML	IBM Integrated Management Module Advanced Upgrade	1	provisioned power capacity to support powering the system on.
9		A3F3	LSI SAS9206-16e Quad-port x4 HD-miniSAS x8 PCIe 3.0 SAS HBA	3	
10		A1H5	IBM System x 750W High Efficiency Platinum AC Power Supply	2	
11		A3N6	x3650 M4 HD System Level code	1	
12		A3N7	System Documentation and Software-US English	1	

Figure 28: Generated report from an imported CFR file with no errors

If there are multiple systems configured in the CFR file, each system in the power validation report will be separated by a black bar. Directly under the black bar will be a color coded bar to indicate the power validation status (green, yellow or red) as seen in figure 28 above.

The minimum supported power mode is also listed in the right hand column of the power validation report. For instance, in the figure 28 example the minimum power mode is N+N without throttling.

Supported Hardware for Importing

The power configurator import function supports the following hardware.

System x3500 M5 (5464)
System x3550 M5 (5463)
System x3650 M5 (5462)
NeXtScale n1200 Enclosure Chassis (5456)
NeXtScale nx360 M4 Compute Node (5455)
NeXtScale nx360 M5 Compute Node (5465)

System x3550 M4 (7914)
System x3550 M4 (7914 E-2600 v2)
System x3650 M4 (7915)
System x3650 M4 (7915 E-2600 v2)
System x3650 M4 HD (5460)
System x3650 M4 BD (5466)

As of March 2015, the power configurator import function only supports CFR files from **Blue Horizon** or **x-Config**, however please check back at the power configurator download page for updates to the tool as support for other CFR files will be made available soon.

Power Configurator download page:
http://www.ibm.com/systems/bladecenter/resources/powerconfig.html

Node Manager Power and Thermal Utility Tool

The Node Manager Power and Thermal Utility tool is an integrated software tool for determining the as-configured total power budget for all new System x, Pure Flex, iDataPlex, and NeXtScale systems. This technology takes a more granular approach of determining system and chassis power budget than using look-up tables in system management devices. The benefits of this allow power policies to be set based on actual component power consumption under any supported operating condition or workload. Power policies are able to be more accurately maintained without unnecessary over-budgeting to ensure as much available power is provisioned by the system as the policies allow. This prevents resiliency and performance impacts such as unexpected throttling and the system powering off unexpectedly.

The Node Manager Power and Thermal Utility Tool functions by running separate, sub-system specific workloads and then calculates a total worst-case power consumption estimate. The result of the Node Manager Power and Thermal Utility Tool is reported to the respective management interface for determining power-on support and redundancy policy of the supported systems. The Node Manager Power and Thermal Utility Tool result is not directly reported to any user interface, but the power policies are managed by this configuration specific power budget. This means that as the configurations change, the enclosure will automatically manage the provisioned power according to the power policy set by the end user.

Notes and Rules when Configuring

The following rules apply when configuring the below systems when using the System x Power Configurator. The Power Configurator is based on complex algorithms to determine the power measurements. Failure to adhere to the rules may result in incorrect algorithm calculations and incorrect power results displaying.

Flex System

- Configure the chassis first. Configure the FSM second if one is required. Configure the nodes last.

BladeCenter

- Configure the chassis first. Configure the nodes last.

iDataPlex

- Configure the chassis first. Configure the nodes last.

NeXtScale

- Configure the chassis first. Configure the nodes last.

All Chassis

- The final power results will only be displayed once the chassis is configured with PSU(s).
- Power displayed when configuring the nodes is DC power consumption of the node and is noted as Wdc. This number does not include chassis infrastructure, cooling, or power supply efficiency.
- If a mistake is made in configuring a chassis and you need to remove an item, restart the tool to ensure accuracy in the final power calculation.

All nodes

- Nodes require at minimum CPU and memory for power data to be calculated for the node. The final power number for all nodes requires PSUs to be added to the chassis.

All Rack/Tower

- Power results will only display if you have configured at minimum: CPU, memory and PSU(s).

General Rules

- Relaunch the tool if you need to start a new configuration.
- Relaunch the tool if you exit a window and re-enter a window to make additional changes.

Power Definitions Explained

The following panel is the 'Power Configuration' panel that displays the power results. Each heading is described below.

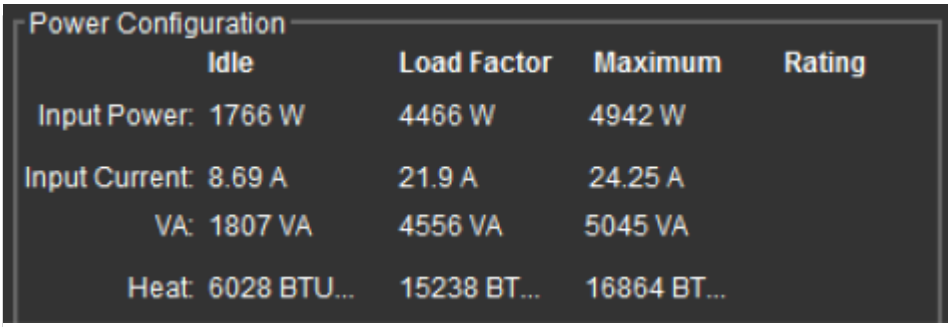

Figure 29: Power Configuration Panel example

Input Power: Refers to real power or active power of your configuration based on your configured hardware. This is measured in watts (W).

Input Current: Refers to apparent power of your configuration. It is reported as single-phase current draw. The current reported does not represent the individual phase currents of a three-phase circuit. It is measured in amps (A).

VA: Refers to the measurement of power of your configuration from direct current (DC) electrical circuit. It is measured in volt-ampere (VA).

Heat: Refers to the heat dissipation of your configured hardware. It is measured in British thermal unit (BTU).

Idle: Refers to power drawn with the machine(s) logged into the OS, and no other applications running. Memory, HDDs, optical drives, & tape drives are not being utilized.

Load Factor: Refers to the load work your configuration will encounter based on the Load Factor percentage (%) you configure on the hardware.

Maximum: Refers to the maximum measured power drawn by the machine with the CPUs 100% utilised & memory is utilized 10% more than base utilization.

Rating: Refers to the label rating for a fully configured server or node. E.g. all DIMM slots populated, highest power CPUs configured, all drives populated, etc. The rating is further explained in the next section.

Note: Some systems do not display a rating because the rating is determined by configuration. These systems should use the power supply rating listed on the power supply or in the product publications to determine the overall chassis rating.

The Rating Label Explained

The "Rating" is the label rating of the product. It is the absolute worst case power consumption for a fully configured system or chassis. This means all DIMM slots, CPU sockets, PCI slots, HDD slots etc. are fully populated and are running at maximum capacity assuming the worst case power load across the entire system/chassis. It is calculated on assuming the highest power consuming hardware is installed, for example: 130W CPUs, quad rank DIMMs, 15k HDD spindles etc. It also assumes highest fan speeds.

For these reasons, if you are running a system with 95W CPUs for example, it is not realistic to provide the rating number as an accurate guide for your systems power load. Use the System x Power Configurator tool to better define your hardware's power draw. The rating label is circled below in Figure 30.

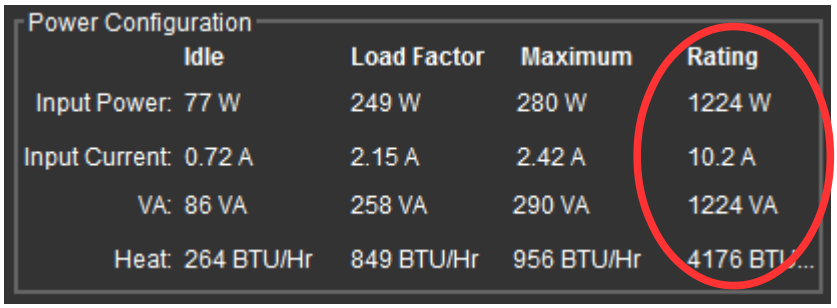

Figure 30: System x Power Configurator - label rating

80 PLUS Explained

80 PLUS is a performance specification for power supplies used within servers and computers. To meet the 80 PLUS standard, the power supply must have an efficiency of 80% or greater, at 20%, 50%, and 100% of rated load with PF of 0.9 or greater. The standard has several grades, such as Bronze, Silver, Gold, Platinum, and Titanium.

The power supplies used in System x high-end Systems are hot-swap high efficiency 80 PLUS Platinum or Titanium power supplies. The efficiency varies by load.

For additional information on the 80 PLUS and ratings, refer to the Power Supply Ratings section.

Power Supply Efficiency, Capacity and Ratings Explained

There can often be confusion when talking about a Power Supplies efficiency, its capacity, and its rating. The following section provides a brief understanding of each of these when talking in terms of a Power Supply.

Power Supply Efficiency and Capacity

Power Supply efficiency is a ratio of how much power is wasted when power enters the PSU from a power source (such as a PDU or wall socket) versus how much is outputted and used by the hardware being powered.

To provide a simple example – if a servers PSU has a *rating* of 80% efficiency and is capable of pulling 600W *capacity* from a power source, then the usable power will be 80% of the 600W capacity being pulled (which equates to 480W usable). The remaining power is converted in to heat and lost. See Figure 31 for an example.

For the customer, this means paying for 600W's of power but only being able to use 480W. Because of this, it becomes important to provision PSUs with high efficiency ratings, particularly in large Data Centers with multiple racks of equipment.

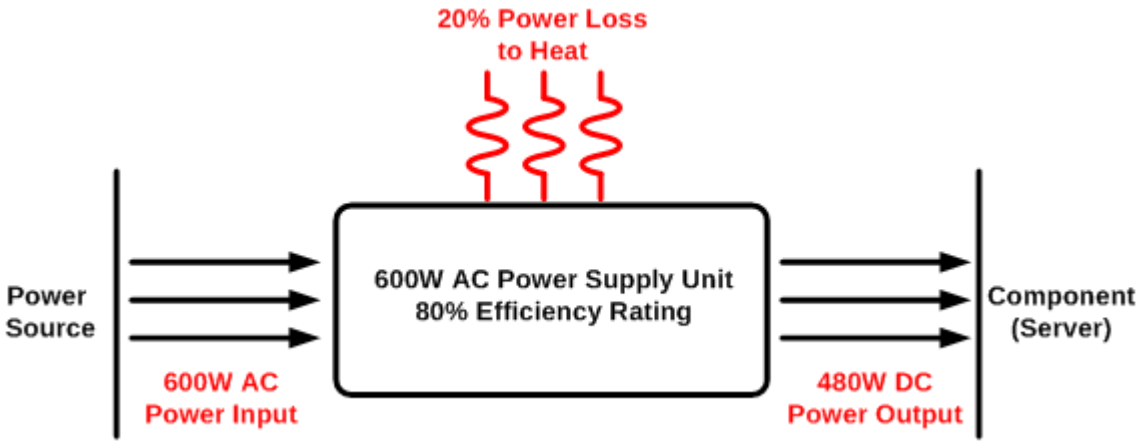

Figure 31: Simple example of Power Supply efficiency

This is why the 80 PLUS certification program was created. This certification helps to identify PSU with the best ratings. The more efficient the power supply, the better rating the PSU will have, and this will mean less power being drawn from the power source, decreasing the total cost of ownership (TCO). The 80 Plus certification program and its ratings are discussed further in the next section.

Power Supply Ratings

A PSU efficiency (as discussed in the previous section) will determine its rating. The better the efficiency of the PSU, the better the rating. The 80 PLUS certification program has a number of ratings. These are labelled – 80 Plus, Bronze, Silver, Gold, Platinum, and Titanium, of which at 230V internal redundant the efficiencies are rated between:

80 Plus Bronze: 81% - 85%
80 Plus Silver: 85% - 89%
80 Plus Gold: 88% - 92%
80 Plus Platinum: 90% - 94%
80 Plus Titanium: 90% - 96%

There are exceptions to the effectiveness of the rating and the efficiency of the PSU. The numbers will vary depending on the amount of power required from the server which is directly related to the load of the server. The load of the server refers to the capacity at which the hardware runs.

Typically, the least efficient will be when a server is running at low load requiring less power, the most efficient will be at approximately half to three quarter load, and expect slightly less efficiency at maximum load.

For example:

	20% Load	50% Load	100% Load
80 PLUS Platinum Standard	90.00%	94.00%	91.00%

Example – 80 PLUS report: http://www.plugloadsolutions.com/psu_reports/System x_7001700-XXXX_900W_SO-571_Report.pdf

System x high-end servers are certified for 80 Plus Platinum or 80 Plus Titanium, and will operate between 90% - 94% efficiency for Platinum or between 90%-96% efficiency for Titanium, depending on the load.

Refer to the 80 PLUS Explained section on page 39 for additional information on System x 's 80 PLUS rating specifications.

System x Power Configurator Link

From the below link you can download the most up to date version of the tool.

System x Power Configurator Tool, User Guide, Legacy Tool, Power and PDU guides:

http://www.ibm.com/systems/bladecenter/resources/powerconfig.html

Helpful Links

The below links are other resources related to power and cooling.

Hubbell - Twist Lock Plug/Outlet Catalog (Includes NEMA Outlets)

http://www.hubbellcatalog.com/wiring/catalogpages/section-b.pdf

Hubbell - Pin & Sleeve Plug/Outlet Catalog (Includes IEC309 Outlets)

http://www.hubbellcatalog.com/wiring/catalogpages/section-E.pdf

System x Configuration and Options Guide

http://www.ibm.com/systems/xbc/cog/

System x and BladeCenter Reference Sheets

http://www.redbooks.ibm.com/Redbooks.nsf/pages/xref?Open


lenovo
FOR THOSE WHO DO.
Data Center Services
Rack, Power & Cooling
Designed to work seamlessly together
Do you have questions about Rack, Power, Thermal, or Mechanical?
power@lenovo.com is your source for answers!
Download the Power and Cooling Guides from:
http://www.ibm.com/systems/bladecenter/resources/powerconfig.html