Let's go

The bridge to possible

X-Series Power and Cooling

Be blown away by our shocking revelations!

Scott Garee, UCS Platform TME

Agenda

- Introduction
- Chassis power architecture
- Chassis cooling architecture
- Power/Cooling policies
- Power measurements
- UCS Power Calculator
- Tuning for power efficiency
- Monitoring power/cooling
- Troubleshooting info

Power architecture

cisco Live!

Chassis power architecture

- Power entry modules (PEM) provide cable sockets for up to 6 power supplies (PSU)
- Power and thermals are monitored by the CMC on the IFMs
- PEMs also group input connections and PSUs per grid source for N+N redundancy
- Populate PSUs evenly across PEMs

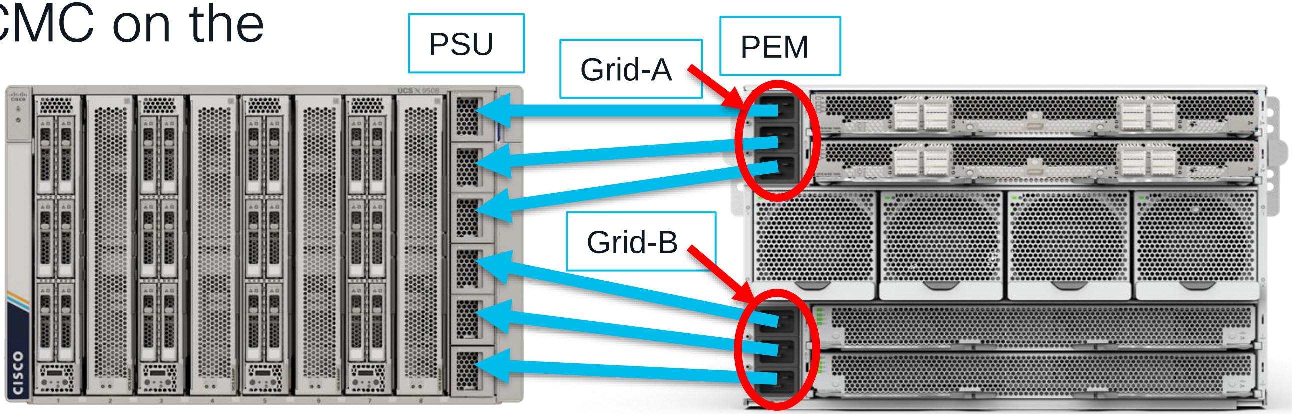

Chassis power supplies

- 2800W @ 54V DC output
  - High line 208-240V AC
  - 1400W low line 100-115V AC
- Titanium rated
- 92% efficiency at maximum output
  - 2800W output
  - ~3044W (15A) input power @ 208VAC
  - PSU efficiency loss is due to heat and internal fan power.

Note: All power info in this document is based on high-line (200-230V AC).
For low-line cases recompute based on low-line voltage and 1400W max PSU Output.

X-Series 54V Power Distribution
What's the Big Deal?

Distribute same power with 80% less current
Less space and materials

Reduces heat generation and lower cooling burden
54V fans 25% more efficient than 12V

Improves AC-DC conversion efficiency 1.5-2% over 12V

Less copper required to carry current
Smaller server power distribution components

=

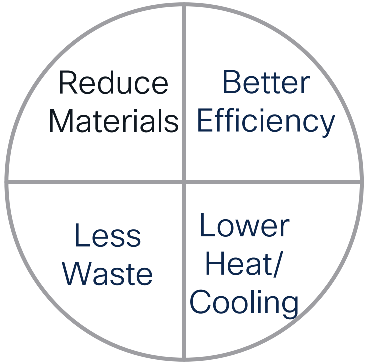

Power output to single 54V DC bus

- Total load is always shared across all active PSUs, independent of redundancy mode.
- Power supplies in standby mode (output off) do not share load.
- Power available for allocation is limited by redundancy and power limit settings.

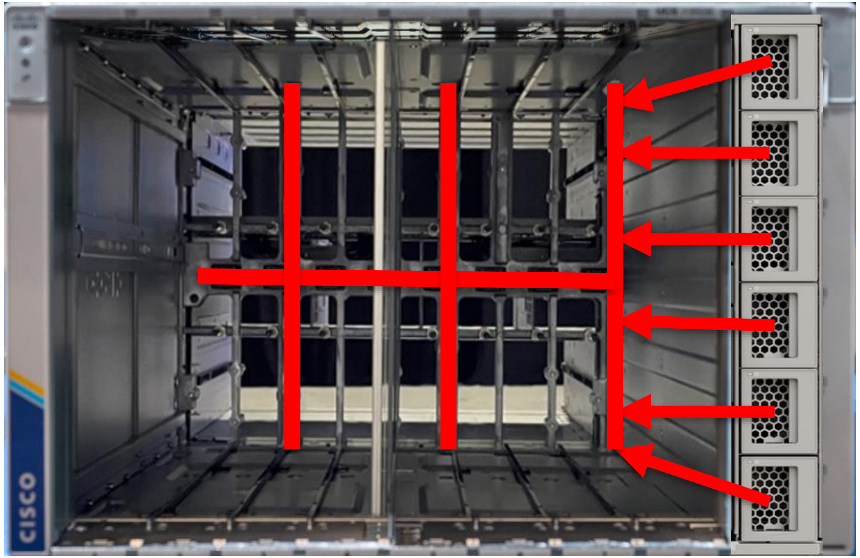

Cooling architecture

Lots of fans to move lots of air - efficiently

- Maximum chassis airflow is approximately 1100 CFM
- Approximately 110 CFM per node, 40 CFM per fabric module and 15 CFM per PSU
- Hot-swappable N+1 redundant
- Large higher voltage system fans provide more airflow at lower speeds, that lowers fan noise and reduces power consumption

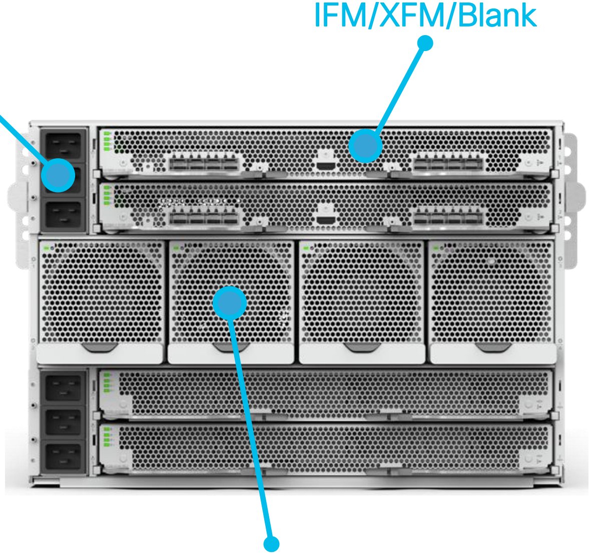

Zoned Based Cooling: Dealing with 350W CPUs

- Vertical and horizontal zones
- Intelligent, Independent fan speed control based on sensors participating in each zone
- Reduces fan power consumption by right-sizing air flow on a per-zone basis, and increases overall power efficiency

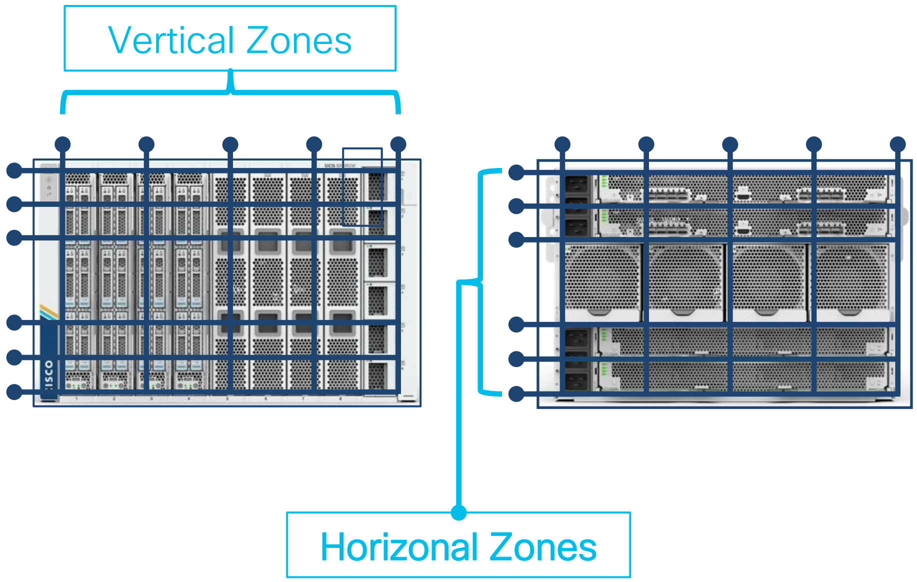

Fan Module Redundancy

- Hot-swappable, dual rotor, N+1 redundant
  - If one rotor/motor fails, the other will keep running to maintain cooling.
  - In failure conditions (or fan removal)
    - Remaining fans increase speed to max value to compensate
    - Throttling will be used to maintain thermals if necessary (uncommon.)

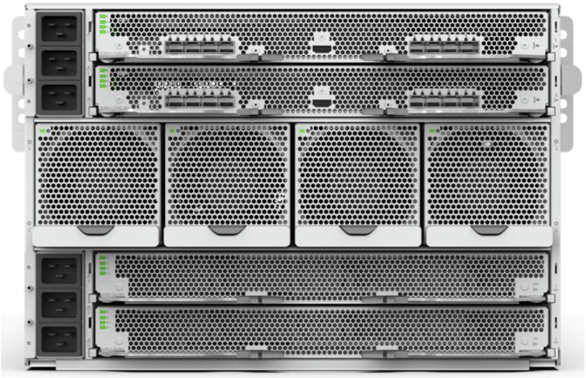

Infra FW 4.2(2d) or 4.2(3c) or later recommended

Cooling and Air Flow Optimizations

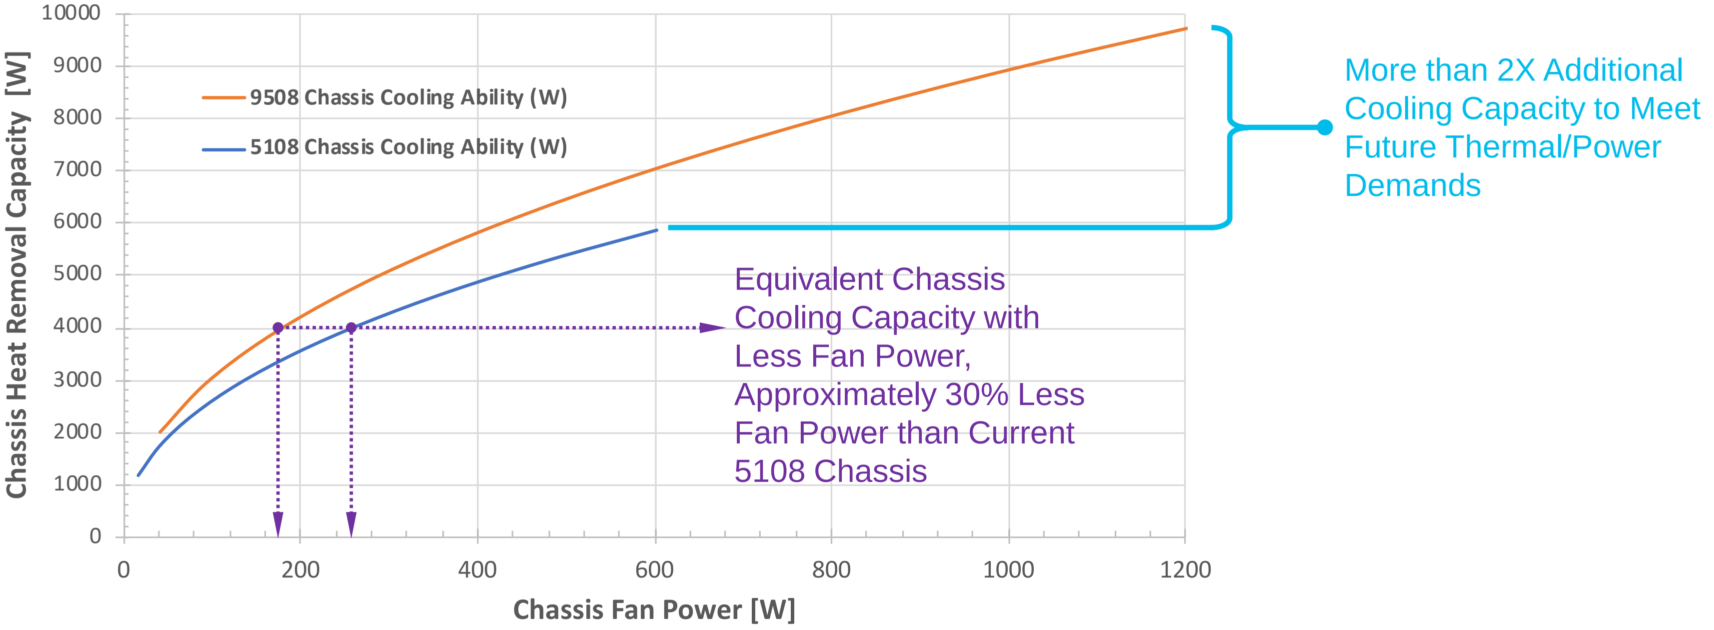

Fan Speed Sensor: Data Collection

- CMC collects temperature sensor data every 5s and processor power data every 1s

- CMC evaluates temperature and power sensor values to (1) maintain current operating speed (2) incremental increase/decrease or (3) maximum fan speed increase

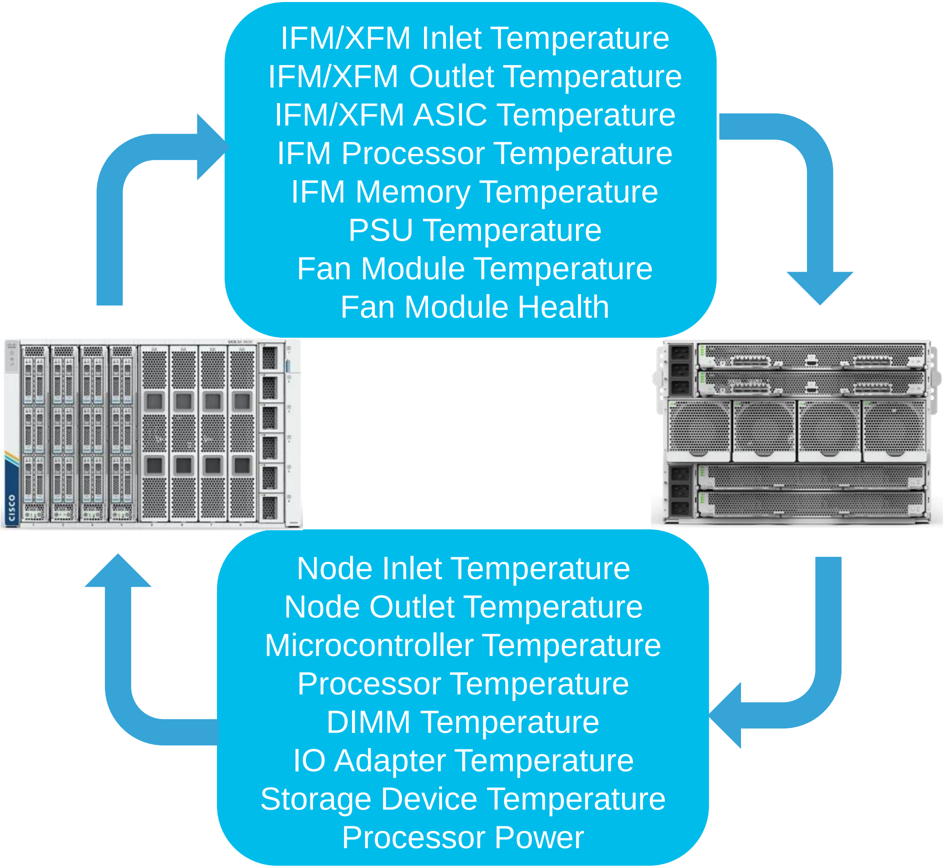

Fan Speed Sensors: Chassis

- Chassis inlet is a pseudo sensor determined by the lowest node inlet temperature
- Chassis inlet sensor assigned a single dynamic maximum speed increase threshold, e.g. 35C for ASHRAE Class A2
- Fan health has single maximum speed increase threshold (80% <= 27C, 90% 28C, 100% >= 29C)

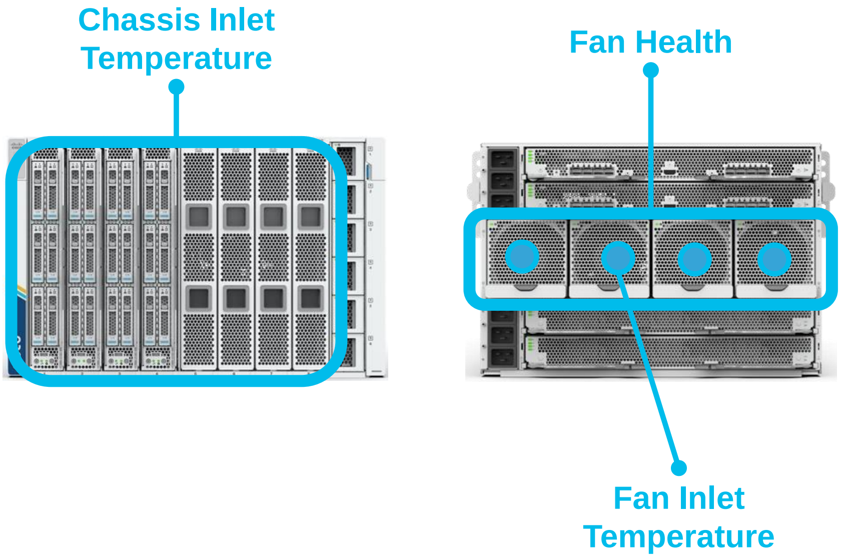

Airflow Paths

- Airflow path is front to back
- Low impedance (no midplane) chassis design reduces fan power consumption and improves overall power efficiency
- Cooling aligned with compute thermal demand

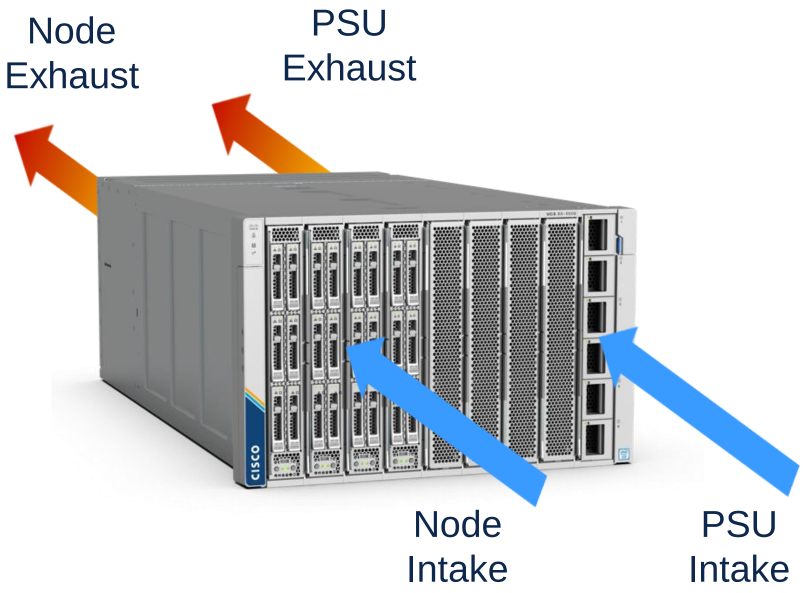

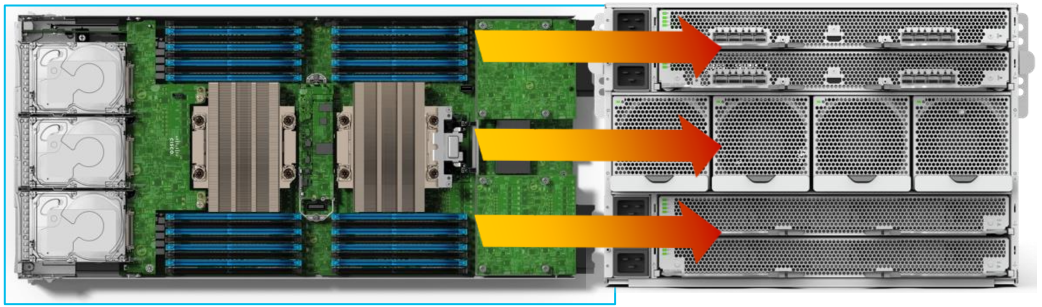

Power features

Power redundancy

- The PSU population and redundancy modes determine how much total power can be consumed without power faults due to PSU output loss. Power allocation will have implicit cap at max value.
- Power is reserved to meet redundancy requirement.
- **Total power is still shared across all populated PSUs.**

Total and per-PSU maximum power by population and redundancy mode, with none in standby						
Mode	1 PSU	2 PSU	3 PSU	4 PSU	5 PSU	6 PSU
N	2800/2800	6087/2800	8400/2800	11200/2800	14000/2800	16800/2800
N+1	N/A	2800/1400	5600/1867	8400/2100	11200/2240	14000/2333
N+2	N/A	N/A	2800/933	5600/1400	8400/1680	11200/1867
N+N (Grid)	N/A	N/A	N/A	5600/1400	N/A	8400/1400

Power Save Mode

- Allows unneeded power supplies to be turned off to improve efficiency
- Time to power up prevents use for redundancy
- Supplies needed to meet redundancy requirements are not placed into standby mode.
  - Example: Six PSU in N+2 with 4500W load. Two PSU required for load, plus two redundant. Remaining two can be put into standby mode.

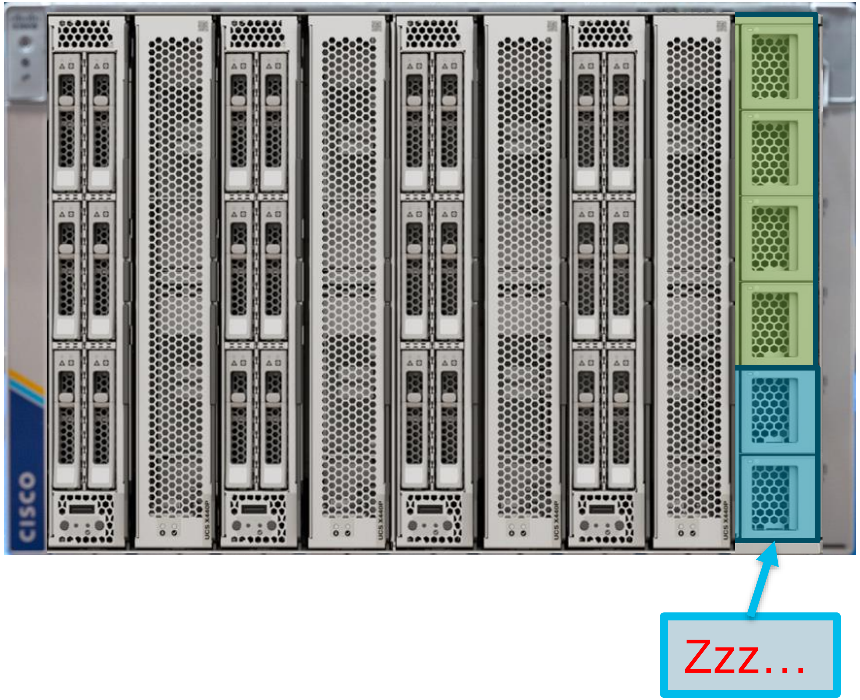

Extended power mode

- When operating with redundant supplies the total power allocation/consumption can be increased by 15%.
- Since load is always shared across all active PSUs (including redundant supplies) the per-PSU load is always less than the maximum when redundancy is intact. E.g. 50% for all PSUs at 100% load in N+N mode.
- If a PSU fails, or a grid fails when in N+N, the system will immediately reduce the allocation back to 100% of the non-redundant load.
- If load was greater than 100% blades will be power capped to bring the load to 100% or less.
- Examples
  - Six PSU in N+N: Normal limit is 8400W. Extended power will allow 9660W load.
    - 9660W / 6 PSU = 1610W/PSU = 4830W/Grid
  - 6 PSU in N+2: Normal limit is 11,200W. Extended power will allow 12,880W load.
    - 12880W / 6 PSU = 2147W/PSU. If one PSU fails the limit will be reduced to 8400 * 1.15 = 9660W

Extended power illustration

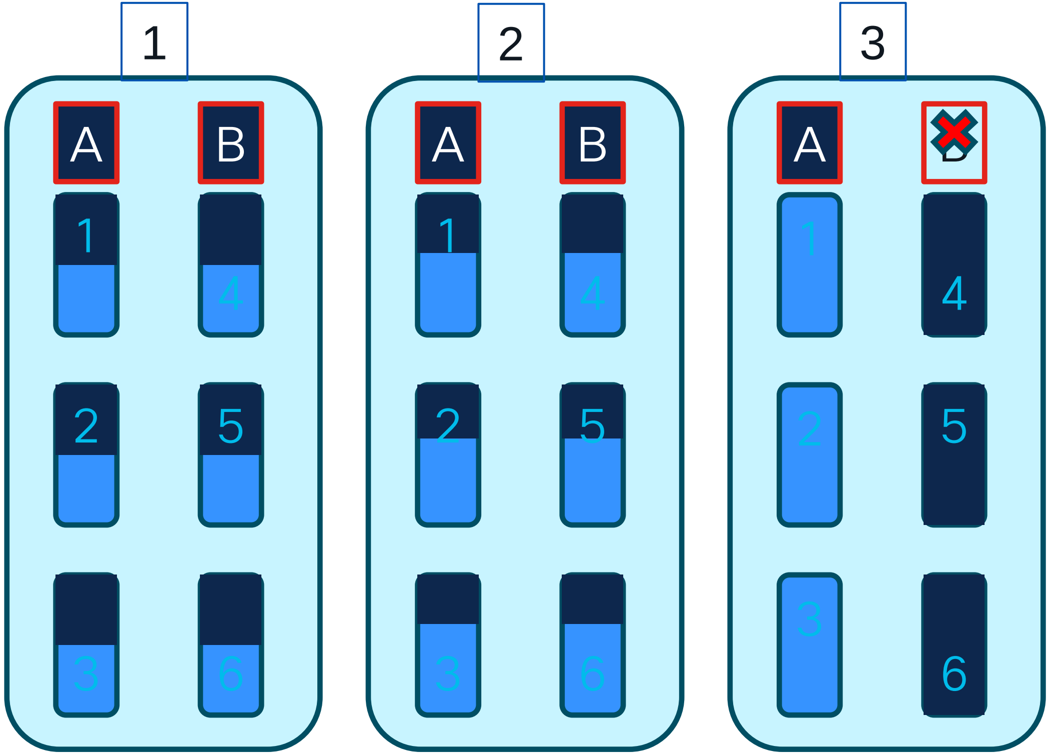

Example workload consumes maximum power

1: Grid mode, not extended
Each PSU at 50% (1400W)
Workload capped at 8400W

2: Grid mode, extended
Each PSU at 57.5% (1610W)
Workload capped at 9660W

3: Grid mode, grid failed
Each PSU at 100% (2800W)
Workload capped at 8400W

Failure state does not change with extended power

Power profiling

- Calculates minimum and maximum consumption using a tool integrated in the system BIOS.
- At initial blade discovery, allocation will be for blade SKU max
  - E.g. 1610W for X210c M7
- During inventory discovery (if not profiling) the catalog estimate is determined from the blade inventory
- If profiling is enabled the server will run a characterization workload to determine min and max required power

Power limit

- Power can be capped at the chassis input
- Actual power consumption is limited by throttling servers
- Minimum cap prevents limit beyond ability to throttle
- Can prioritize servers to order throttling
- Throttling is very granular, using P-States
- Extreme emergency throttling (e-brake) possible for PSU overcurrent and thermal issues.

Power allocation

- Power is allocated in priority order from the server Power policy, from highest to lowest, until all power is allocated.
  - Base allocation for chassis
  - Minimum power for each server in each priority
  - Maximum power for each server in each priority
- The total power to allocate is the lowest of the configured limit in the chassis Power policy or the limit based on the PSU redundancy.
- Allocations can adjust if Dynamic Power Rebalancing is enabled.
- Chassis power allocation can also be rebalanced.

Emergency brake

- Hardware based hard throttling mechanism controlled by the Chassis Management Controller (CMC) on the IFM.
- Triggered when Grid power is lost and load is in the extended power range, on PSU over-current, and thermal overtemp.
- Once load is below available power the brake is released and normal throttling is used to maintain power below new limit.

Power and Thermal Policies

Power policy

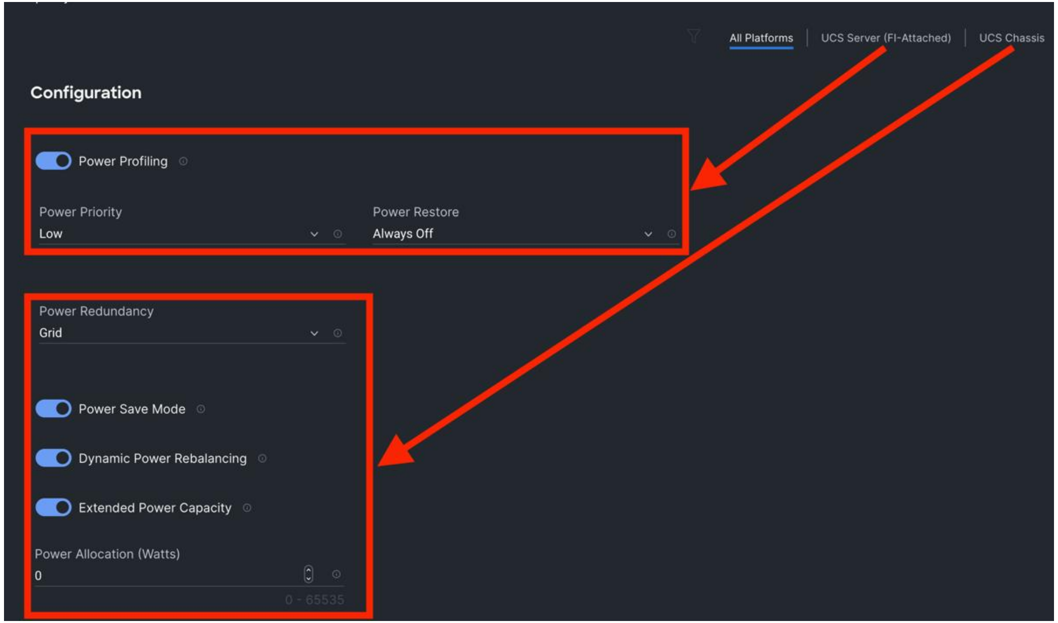

- Single policy for both chassis and server
- No overlapping settings
- Recommend separate policies to avoid confusion

Power policy - Server

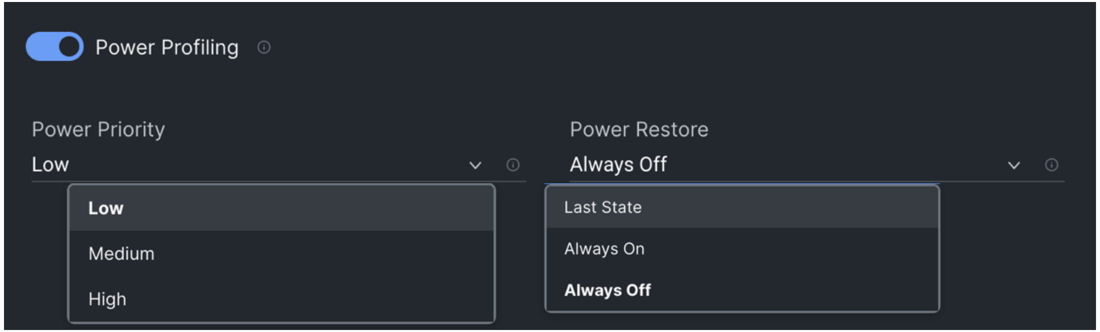

- Recommend profiling on all servers. PCIe node included with server.
- Highest priority servers are allocated any power available above minimum.
- Once higher priority servers have their max allocation, lower priority servers are allocated what is left.
- Power restore refers to chassis power restoration. Last State recommended.

Power policy - Chassis

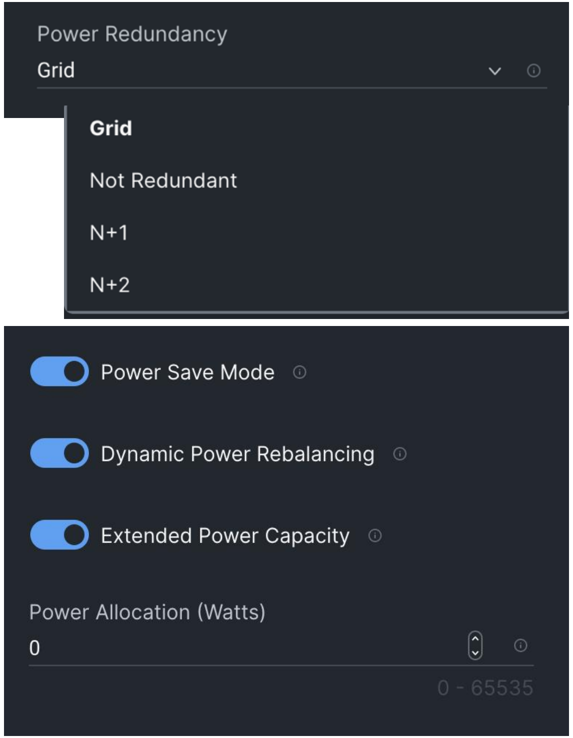

- Redundancy per customer requirements
- Default options recommended for others
- Dynamic power rebalancing will re-allocate power between servers if competing for power in power constrained environments.
- Power allocation sets an input power limit. Throttling is used to stay within the limit.

Thermal policy

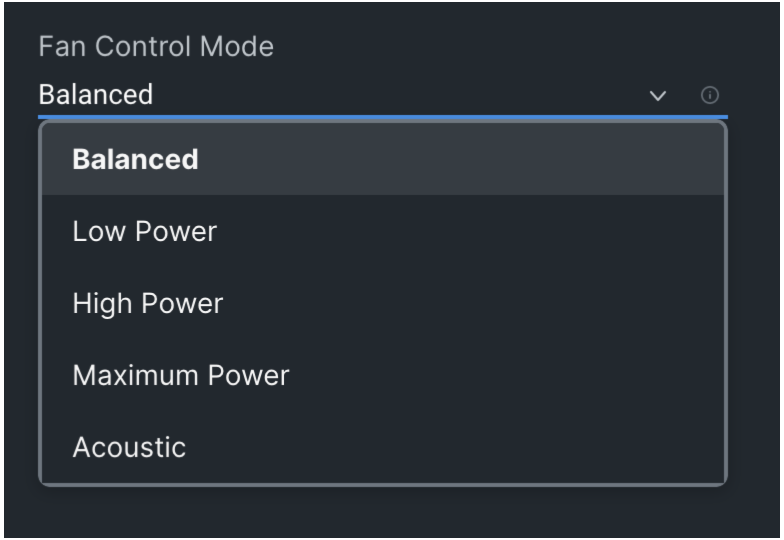

- Applies to chassis or rack mount servers
- Sets the base fan speed
- Workloads with frequent transitions from idle to max may throttle while fan speeds ramp from lower levels.
- Increase policy to prevent throttling.

Order from lowest to highest: Acoustic, Low, Balanced, High, Maximum

Chassis inventory power information

- Configured settings
- Chassis totals
  - Required input power based on node inventory
  - Available output power (per redundancy mode, PSU count)
- Chassis power required
  - Static default for fans, IFM, XFM
  - Dynamically adjusted as servers power up
- Server power required
  - Blade SKU based (No knowledge of what’s in it)
  - Catalog/inventory based
  - Profile based

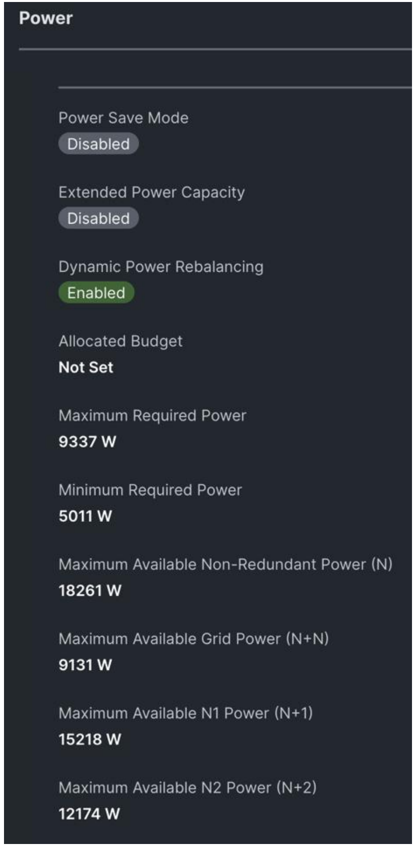

Measuring power

Points of measurement

- Input power
  - The power consumed at device input
  - At the PSU input for chassis measurements
  - At the hotswap controller input for chassis components
- Output power
  - The power available downstream of a power source
  - The PSU output (or commonly the sum of all PSU outputs for the chassis
  - The device power converter output for components
  - Device power – actual device consumption available from some devices

Input power calculations

- Input power (from the power source) will always be greater than measured consumption at the components or PSU output, due to PSU and power conversion efficiency losses.
- PSU loses power as heat and due to fan power used to remove heat.
- Blade loses power as heat due to voltage conversion at the blade input.

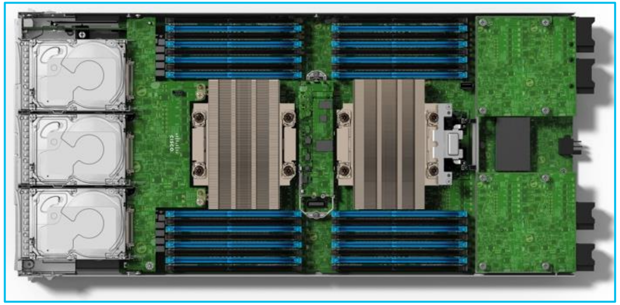

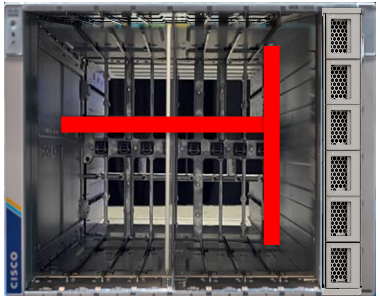

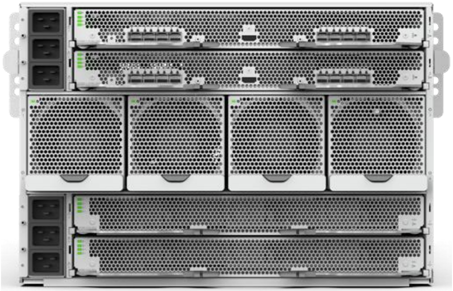

Monitoring power - Server Statistics tab

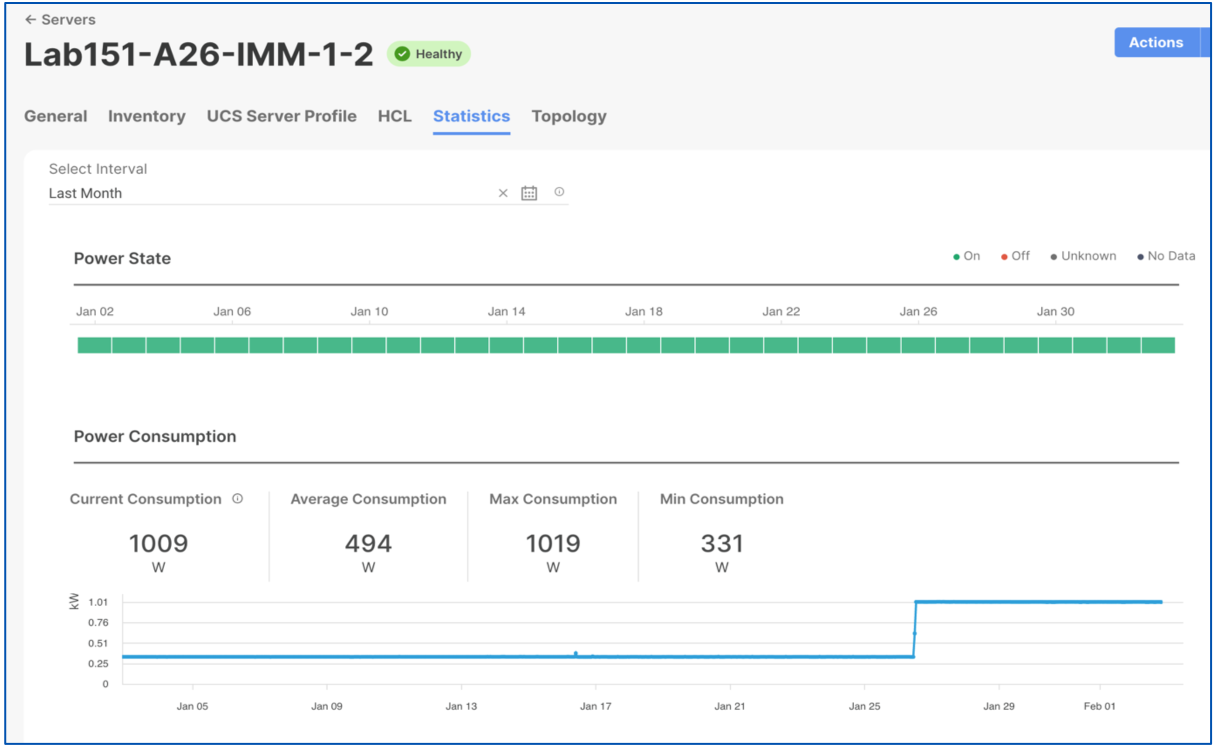

Input power consumed by the blade

Monitoring power – Metrics explorer

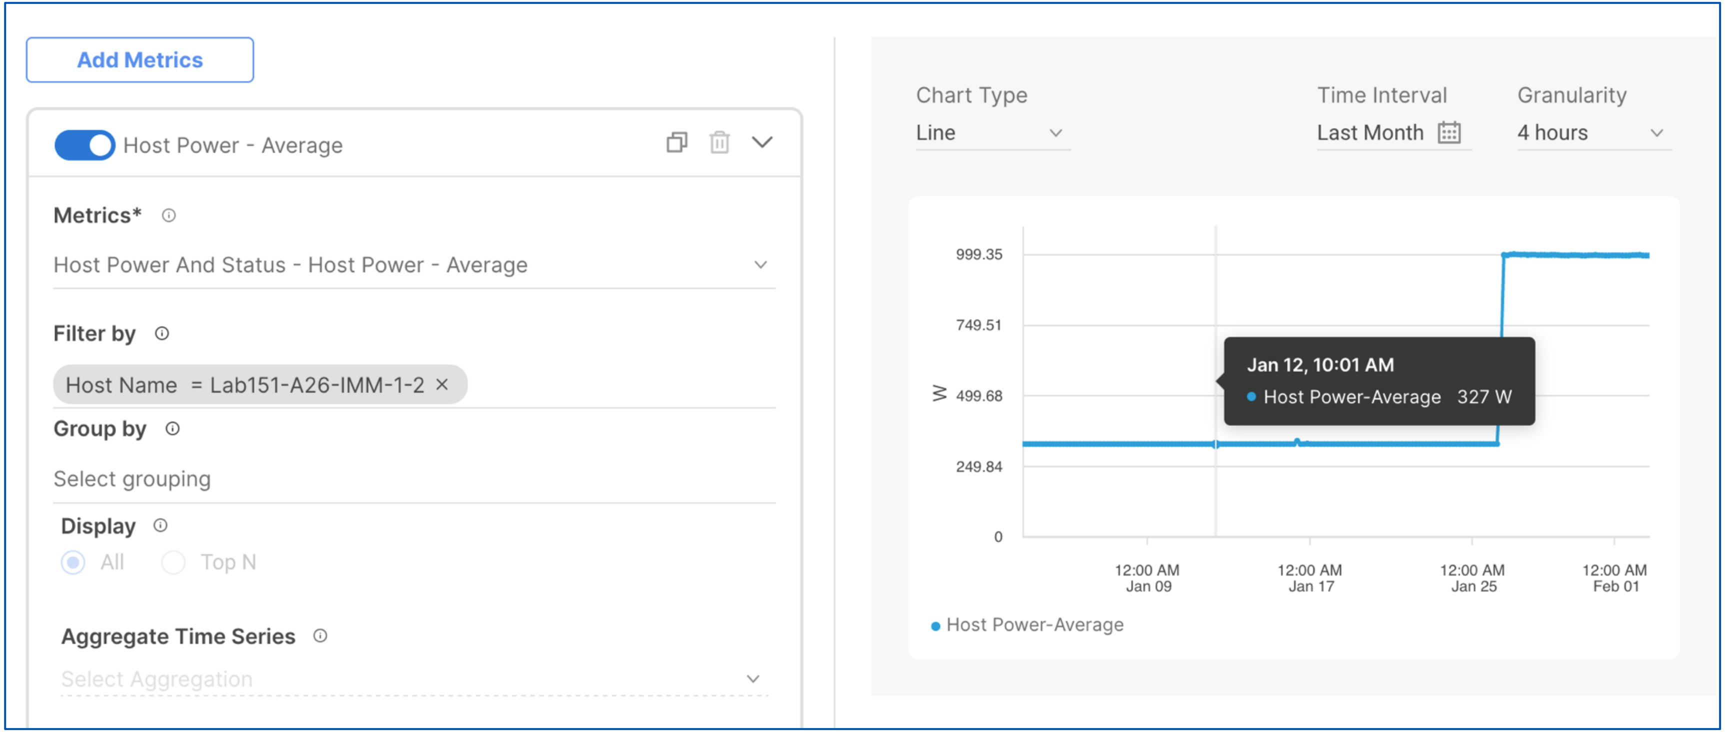

Input power consumed by the blade

Input power calculations – PSU efficiency

- At maximum output the PSU operates at ~92% efficiency.
- Efficiency is highest at 40-60% load (~96% efficiency for titanium.)
- To calculate input power from PSU output power or chassis power measurements or estimates, use the efficiency.
  - E.g. 5100W chassis measured power / .94 = 5526W estimated input power
  - 5526W / 208V = 27A input current

Input power calculations – Maximums

- 2800W (maximum output) * PSU_count
  - Six PSU in Grid mode, with a failed grid = 8400W / 0.92 = 9130W per PEM max
  - 9130W / 208VAC = ~44A per PEM
- In Grid extended mode, with no failure, max total power = 9660W
  - This is spread across both grids, so 4830W per PEM = ~24A per PEM/PDU
  - On grid failure this number does not double to 48A, because power will be limited to 8400W out/9130W in.
- The majority of workloads do not consume maximum power
  - Workloads should be understood, so power can be better estimated to avoid unnecessary over-provisioning

Power Calculator Demo

UCS Power calculator

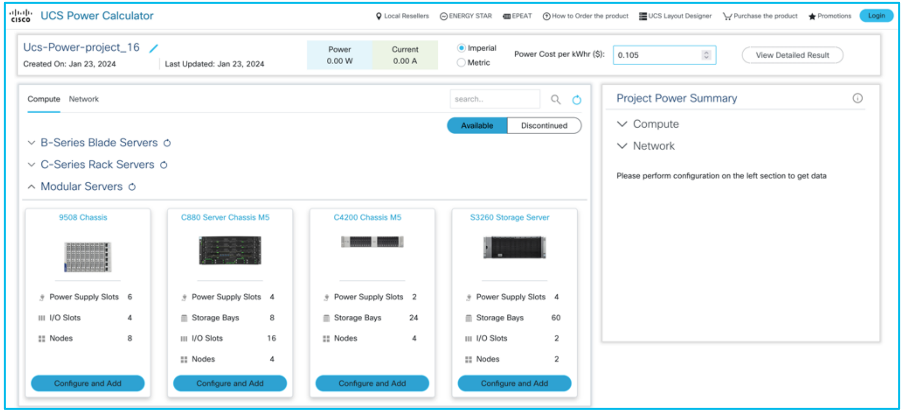

- Actual consumption of real workloads in customer environments can be estimated using the UCS Power Calculator
  - http://ucspowercalc.cisco.com
- For most workloads, significantly less than maximum power
  - Contention at maximum consumption would result in undesirable application performance for most applications, including VMs

Power/Performance Tuning

Intel SST-PP - Speed Select Performance profiles

- Multiple options in 4th and 5th Gen Xeon Scalable
- Maximize efficiency with a single CPU SKU, adapt to changes
- Options to flex cores, clocks, or both to adjust power demand

Model	Description	CPUSpeed	Cores	Rel perf	Cache (MB)	Cache/Core	Socket TDP	Core TDP
Flex clock 6548N	Intel(R) Xeon(R) Gold 6548N CPU @ 2.8GHz	2.8	32	0.93	60	1.9	250	7.8
	Intel(R) Xeon(R) Gold 6548N CPU @ 2.6GHz	2.6	32	0.87	60	1.9	225	7.0
	Intel(R) Xeon(R) Gold 6548N CPU @ 2.3GHz	2.3	32	0.77	60	1.9	205	6.4
Flex cores 8558U	Intel(R) Xeon(R) Platinum 8558U CPU @ 2.0GHz	2.0	48	1.00	260	5.4	300	6.3
	Intel(R) Xeon(R) Platinum 8558U CPU @ 2.0GHz	2.0	40	0.83	260	6.5	270	6.8
	Intel(R) Xeon(R) Platinum 8558U CPU @ 2.1GHz	2.1	32	0.70	260	8.1	250	7.8
Flex both 6548Y+	Intel(R) Xeon(R) Gold 6548Y+ CPU @ 2.5GHz	2.5	32	0.83	60	1.9	250	7.8
	Intel(R) Xeon(R) Gold 6548Y+ CPU @ 2.6GHz	2.6	24	0.65	60	2.5	225	9.4
	Intel(R) Xeon(R) Gold 6548Y+ CPU @ 3.0GHz	3.0	16	0.50	60	3.8	205	12.8

OPM – The power efficiency easy button

- Optimized Power Mode
  - Up to 20% less power with <5% less perf[1]
- New simple enable/disable BIOS setting
  - 1.0 with 4th Gen (Sapphire Rapids)
  - 2.0 with 5th Gen (Emerald Rapids)
- Configures multiple power/performance settings
  - C-States: Fast C1E, Package C6
  - Active Idle Mode: Decoupled core and un-core clocks
  - Most effective at medium CPU utilization (40-60%)
  - Still effective at higher levels (~70%)

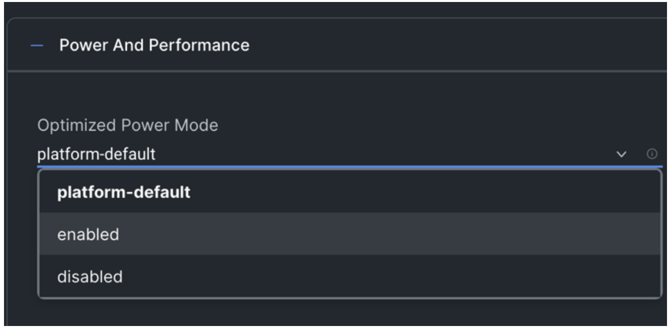

Intel Optimized Power Mode

- Intel OPM is a new BIOS-level capability configurable on UCS M7 servers
- Intel data shows that best results occur at lower CPU utilization
  - (low to medium utilization levels have the highest power savings)
- Easy button within UCSM and Intersight

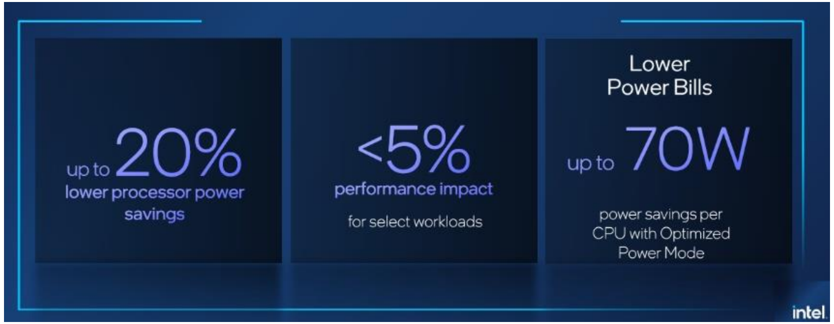

Intel: Power and Performance curves

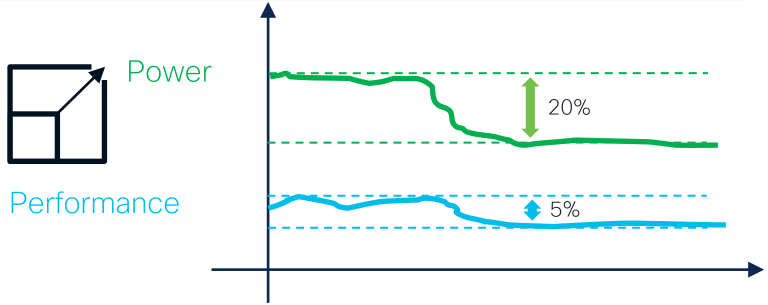

Source: E6 at intel.com/processorclaims

Select workloads: SpecJBB, SPECINT and NGINX key handshake

Cisco Policy with Intel Optimized Power Mode

UCS Manager 4.3.2

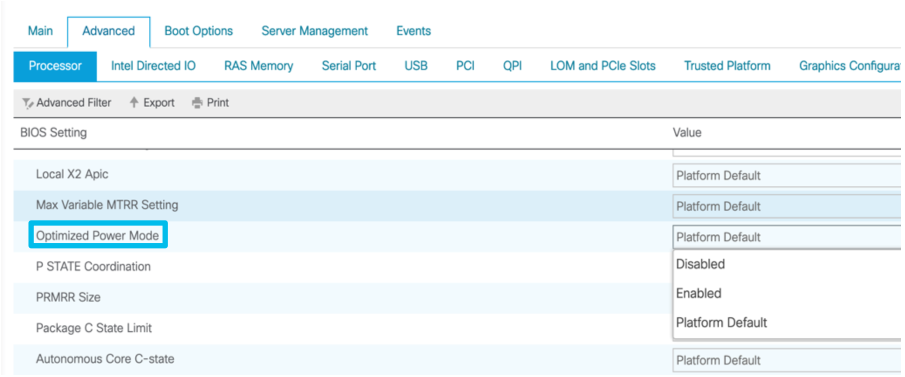

Intersight

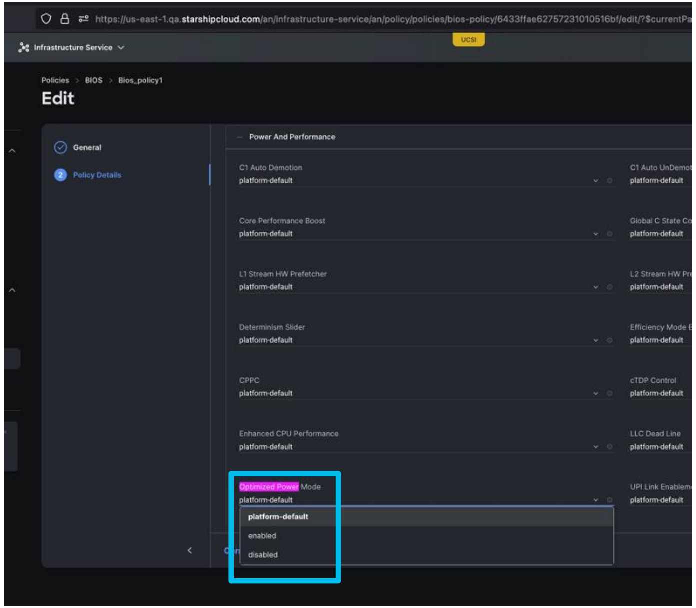

OPM and SpecINT Results

- Increased work (threads), lower the power savings % but smaller performance difference
- Highest power savings at lower to midrange utilization

			OPM Disabled				OPM Enabled					
	BIOS settings	Fan Policy	Threads	Int Rate (OPM Disabled)	Average Power Consumed W (OPM Disabled)	CPU Usage	Threads	Int Rate (OPM Enabled)	Average Power Consumed W (OPM Enabled)	CPU Usage	Power Saving	Performance diff
8462Y+/32 Cores Idle Power: 256 Watts CPU Power (cTDP) - 300 Watts	Default	Acoustic	8	80.9	414	6.3%	8	77	347	6.4%	19.3%	-5.1%
	Default	Acoustic	16	156	519	12.4%	16	149	445	12.5%	16.6%	-4.7%
	Default	Acoustic	32	297	696	25.1%	32	286	614	25.1%	13.4%	-3.8%
	Default	Acoustic	48	415	815	37.5%	48	404	742	37.5%	9.8%	-2.7%
	Default	Acoustic	64	504	845	50.1%	64	496	789	50.1%	7.1%	-1.6%
	Default	Acoustic	90	521	852	70.3%	90	514	808	70.3%	5.4%	-1.4%
	Default	Acoustic	128	589	880	100%	128	587	863	100%	2.0%	-0.3%
	Performance BIOS	Max Power	8	80.4	801	6.3%						

Range most modern workloads operate in

Power optimization – VDI example

- Validated during Cisco UCS VDI CVD development
- Power consumption estimated using UCS Power Calculator
- Testing executed with and without OPM enabled

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_m7_vmware_vsp_8.html
http://ucspowercalc.cisco.com

VDI Workload Performance Without Energy Conservation

- Balanced fan policy, Intel OPM, BIOS and ESXi OS power settings *Disabled*
- CVD publishable VSI Score of 819

OPM-Disabled-ESX-HighPowerf-8Hr-02

Successfully completed Login VSI test with **318** **knowledgeworker-win11** sessions. VSImax (system saturation) was not

Test result review

320 sessions were configured to be launched in **2880** seconds.

In total **2** sessions failed during the test:

- **0** sessions was/were not successfully launched
- **2** launched sessions failed to become active
- **318** sessions were active during the test
- **0** sessions got stuck during the test (before VSImax threshold)

High Perf with fans set to Balanced

With **318** sessions the maximum capacity VSImax (v4.1) **knowledgeworker-win11** was not reached with a Login VSI baseline perfomance score of **819**

Login VSI index average score is **754** lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of **819** is: **Good**

PASS

Login VSI Base line indicator

VSI Baseline	Performance
0-799	Very Good
800-1299	Good
1200-1999	Reasonable
2000-9999	Bad

VDI Test Power Results

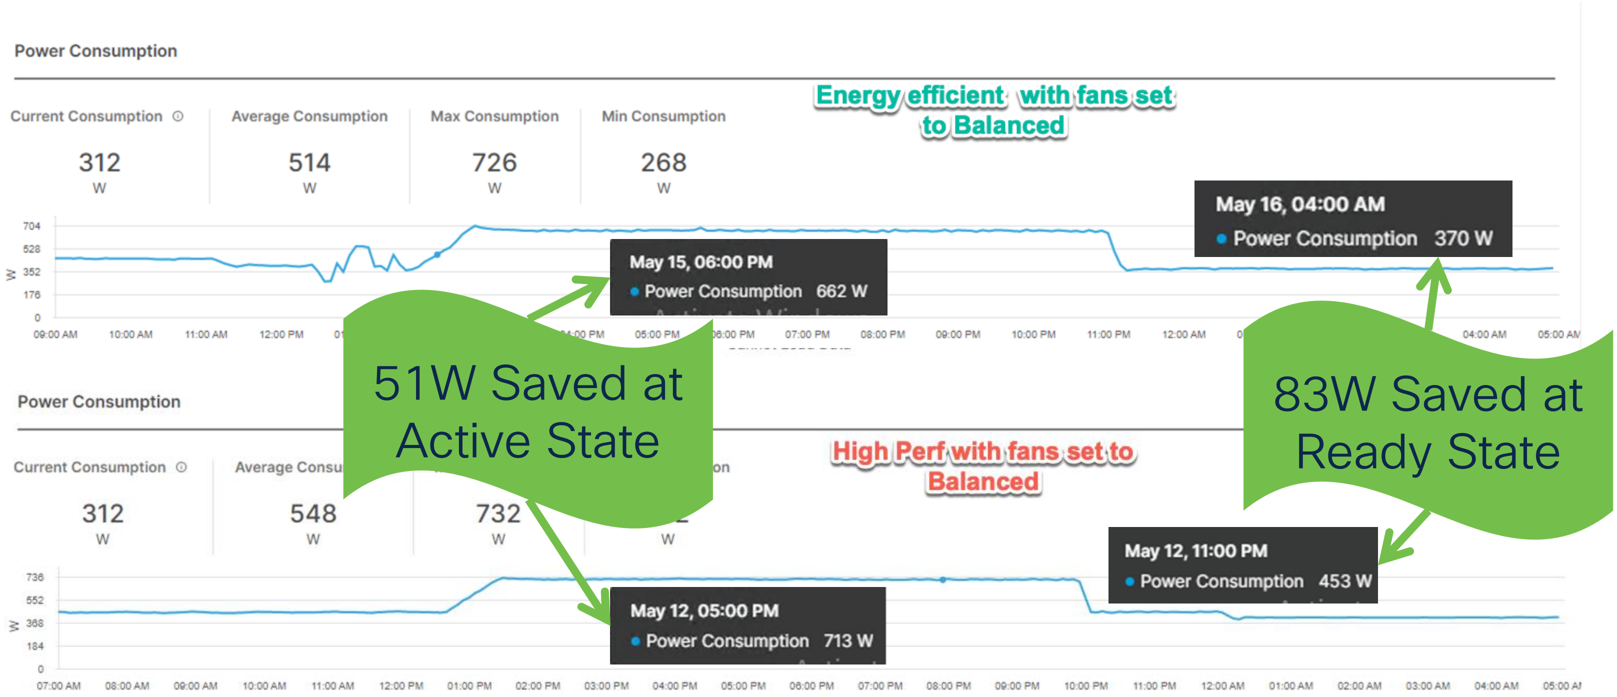

VDI Workload Performance With Energy Conservation

- Balanced fan policy, Intel OPM, BIOS and ESXi OS power settings *Enabled*
- Respectable VSI Score of 914 vs 819
- Performance difference **11.6%**
- UX will not suffer with this difference

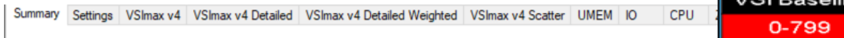

OPM-Enabled-ESX-LowPowerf-10

Successfully completed Login VSI test with **318** **knowledgeworker-win11** sessions. VSImax (system saturation) was not reached.

Test result review

320 sessions were configured to be launched in **2880** seconds.

In total **2** sessions failed during the test:

- **0** sessions was/were not successfully launched
- **2** launched sessions failed to become active
- **318** sessions were active during the test
- **0** sessions got stuck during the test (before VSImax threshold)

With **318** sessions the maximum capacity VSImax (v4.1) **knowledgeworker-win11** was not reached with a Login VSI baseline performance score of **914**

Login VSI index average score is **835** lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of **914** is: **Good**

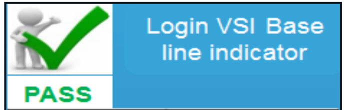

VSI Baseline	Performance
0-799	Very Good
800-1299	Good
1200-1999	Reasonable
2000-9999	Bad

Estimated per node power reduction: 51W at load, 16 servers saves 55,644 kWh/Year

Additional tuning options

- UCS Tuning Guides include extensive details on tuning for performance or energy efficiency
- Updated for each generation
- UCS M7 Performance Tuning Guide
- Other guides in our White papers repository

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/ucs-m7-platforms-wp.html
https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/white-paper-listing.html

Troubleshooting

Power and thermal goodies – Blade power

- From SSH session on the FI
  - connect cimc <chassis#/slot#>
  - power

```
OP:[ status ]
Power-State:                [ on ]
Master-State:               [ Master ]
VDD-Power-Good:             [ active ]
Power-On-Fail:              [ inactive ]
Power-Ctrl-Lock:            [ unlocked ]
Power-System-Status:        [ Good ]
Front-Panel Power Button:   [ Disabled ]
Front-Panel Reset Button:   [ Disabled ]
```

Power and thermal goodies – Slot inventory

- From SSH session on the FI
  - connect iom <chassis#> (FI-A will connect to IFM-1 and FI-B to IFM-2)
  - show platform software cmcctrl blade inventory

```
Blade,Status,BladeClass,BmcSlot,Occupied,Make,Product,SerialNo
B1,DISCOVERED,0,B1,occupied,"Cisco Systems Inc","UCSX-210C-M7","FCH270978H1"
B2,UNKNOWN,0,B2,empty,"","",""
B3,UNKNOWN,0,B3,empty,"","",""
B4,DISCOVERED,0,B4,occupied,"Cisco Systems Inc","UCSX-210C-M7","FCH270978FY"
B5,DISCOVERED,0,B5,occupied,"Cisco Systems Inc","UCSX-210C-M7","FCH270978H7"
B6,UNKNOWN,0,B6,empty,"","",""
B7,UNKNOWN,0,B7,empty,"","",""
B8,UNKNOWN,0,B8,empty,"","",""
```

Power and thermal goodies – Thermal status

- From IFM: show platform software cmcctrl thermal status

```
...
master:        1              # MASTER (otherwise start from the other FI)
...
fan policy:
 active:       5              # ACOUSTIC
 configured:   5              # ACOUSTIC
state:         2              # OK
sound power:   87             # dBA
airflow:       497            # CFM
ambient temp:  23             # Degrees Celsius
...
fan[main/4].fault/read/req:    0/45/55 # OK (fan status, 55 PWM request, 45 current)
...
blade[1].present/state: 2/2    # PRESENT/OK (blade healthy and requesting cooling)
blade[2].present/state: 2/1    # PRESENT/COOL (healthy, no additional cooling needed)
...
zone[4].pwm/ hot_cpu / hot_non_cpu:    58/ blade1:P2_TEMP_SENS.69.91 / N/A
```

Power and thermal goodies – Power info

- For the even more adventurous…
  - Look in a chassis techsupport for the file
    techsupport_detailed_iocard1/cmc/log/pwrmgrcli.log (or iocard2)
  - No interface contract! Contents can change without notice.
- PSU inventory/details/statistics
- Power settings/redundancy status
- Power consumption/estimates/allocation
- Power issue messages

Power info – PSU inventory

- Presence, enable, output status, input status

```
...
PSs present          : 1 2 3 4 5 6
PSs RMT on            : 1 2 3 4 5 6
PS DC ok             : 1 2 3 4 5 6
PS AC ok             : 1 2 3 4 5 6
```

Power info – PSU details

- Input/output measurements, thermals, high/low line

```
Power supply: 1
        Input Voltage   : 206250 mV
        Input Current   : 3968 mA
        Input Power     : 816000 mW
        Output Voltage  : 53992 mV
        Output Current  : 14500 mA
        Output Power    : 783000 mW
        Max Power       : 2800 W
        Temperature 0   : 23.98 C
        Temperature 1   : 36.98 C
        Temperature 2   : 41.98 C
        Fan Speed 0     : 7800 rpm
        Fan Speed 1     : 5968 rpm
        Input source    : HIGHLINE
```

Power info – PSU statistics

PSU statistics:

	PSU1	PSU2	PSU3	PSU4	PSU5	PSU6	Total
Presence	Present	Present	Present	Present	Present	Present	
AC/DC Input	Present	Present	Present	Present	Present	Present	
DC Output	Present	Present	Present	Present	Present	Present	
Input Voltage	206.250 V	205.500 V	205.750 V	205.000 V	204.750 V	205.000 V	
Input Current	3.968 A	3.980 A	4.000 A	4.070 A	4.007 A	4.000 A	
Input Power	816.000 W	816.000 W	820.000 W	830.000 W	820.000 W	817.000 W	4919.000 W
Output Voltage	53.992 V	54.015 V	53.992 V	53.953 V	54.007 V	54.015 V	
Output Current	14.500 A	14.390 A	14.250 A	14.765 A	14.265 A	14.578 A	
Output Power	783.000 W	777.000 W	770.000 W	797.000 W	770.000 W	788.000 W	4685.000 W
Max Power	2800 W	2800 W	2800 W	2800 W	2800 W	2800 W	16800 W

Power info – Settings (per policy)

```
...

power_save_mode        : Disabled
power_extended_mode    : Enabled
power_rebalance        : Enabled
power policy           : Grid
chassis powercap       : Enabled
chassis powerlimit     : 0W
blade1 priority        : 3
blade2 priority        : 3
blade3 priority        : 3
blade4 priority        : 3
blade5 priority        : 3
blade6 priority        : 3
blade7 priority        : 3
blade8 priority        : 3
```

Power info – Redundacy status

```
Redundancy status:
================================
...
Cluster master          : no
Policy                  : Grid
State                   : ok
Total power available   : 16800
Policy powerlimit       : 8400
Min PSU num req:        : 3
PSU input line          : High
Power save mode         : Disabled (All PSUs active)
-----------
Grid                    : 1
        Active PS       : 1 2 3
        Spare PS        :
        Unavailable PS  :
-----------
-----------
Grid                    : 2
        Active PS       : 4 5 6
        Spare PS        :
        Unavailable PS  :
```

Power info – Actual consumption

```
Chassis power consumption report:
--------------------------------------------------------------------
       |presence | cur     | avg     | min     | max     | fault   |
--------------------------------------------------------------------
Blade1  | Present |    530W |    553W |    241W |    938W | No      |
--------------------------------------------------------------------
Blade2  | Present |   1010W |    429W |    203W |   1097W | No      |
...
FAN4    |         |    283W |         |         |         |         |
--------------------------------------------------------------------
FEM1    |         |     43W |         |         |         |         |      (FEM = XFM)
--------------------------------------------------------------------
FEM2    |         |     45W |         |         |         |         |
--------------------------------------------------------------------
IFM1    |         |    165W |         |         |         |         |
--------------------------------------------------------------------
IFM2    |         |    168W |         |         |         |         |
--------------------------------------------------------------------
Chassis |         |   4689W |   3241W |   1970W |   7006W |         |
```

Power info – Consumption limit estimates

```
Chassis power info report:
(BladeX prof_min, prof_max, min_pwr, max_pwr, pwr_limit, host_pwr are collected from BMC)
-------------------------------------------------------------------------------------
        |presence | prof_min| prof_max|  min_pwr|  max_pwr|pwr_limit| fault   | host_pwr|
-------------------------------------------------------------------------------------
Blade1  | Present |   585W |  1151W |   968W |  1608W |   980W | No     | On     |
-------------------------------------------------------------------------------------
Blade2  | Present |   539W |  1120W |   968W |  1608W |  1139W | No     | On     |
-------------------------------------------------------------------------------------
...
Blade8  | Present |   569W |  1157W |   968W |  1608W |   786W | No     | On     |
-------------------------------------------------------------------------------------
SubTotal|         |  4562W |  9144W |  7744W | 12864W |  7762W |        |        |
-------------------------------------------------------------------------------------
...
IFM2    |         |        |        |        |   300W |        |        |        |
-------------------------------------------------------------------------------------
SubTotal|         |        |        |        |  2084W |        |        |        |
-------------------------------------------------------------------------------------
Chassis |         |  6646W | 11228W |  9828W | 14948W |        |        |        |
-------------------------------------------------------------------------------------
Chassis Required Power Range       |  6646W | 14948W |        |        |        |
```

Power info – Allocation

```
Chassis power allocation report:
--------------------------------------------------------------------------------
Chassis powerlimit: 9660                                                       |
--------------------------------------------------------------------------------
Chassis blade available budget: 8683                                           |
--------------------------------------------------------------------------------
Chassis blade allocated budget: 7762                                           |
--------------------------------------------------------------------------------
ebrake_min_req_psu: 0 | prereq_ebrake =            N/A                         |
--------------------------------------------------------------------------------
       |presence | init_pwr| dyn_pwr | ok_to_pwr_on | ok_to_pwr_prof | fault   |
--------------------------------------------------------------------------------
Blade1  | Present |   980W |   980W | Yes          | No             | No      |
--------------------------------------------------------------------------------
Blade2  | Present |   944W |  1139W | Yes          | No             | No      |
--------------------------------------------------------------------------------
...
Blade8  | Present |   981W |   786W | Yes          | No             | No      |
--------------------------------------------------------------------------------
SubTotal|         |  7762W |  7762W |
--------------------------------------
Ch Base |         |  2084W |   977W |
--------------------------------------
Total   |         |  9846W |  8739W |
```

The bridge to possible

Thank you

cisco Live!

Let's go