




# Inovação além da computação com UCS-X e Intersight

Quint@s Quinze - 4 de maio de 2023


Rafael Guerra  
Technical Solutions Architect - Cloud & Data Center Brazil




# Agenda




Challenges



Know the UCS-X



Use Cases




Demo

# Hardware Evolution in Cloud Era


# Compute Platform Directions 2022 and Beyond

## *Disruptive Eco-System Ingredients*



### Compute


High Watt CPU, GPU, DPU, IPU  
Persistent Memory, Liquid Cooling,  
SmartNIC, FPGA Offload



### Protocols


RoCEv2, SMBDirect, NVMe-oF,  
GPUDirect

Session ID



### Fabric

400G PCIe Gen4/Gen5/Gen6 CXL  
1.1/2.0/3.0



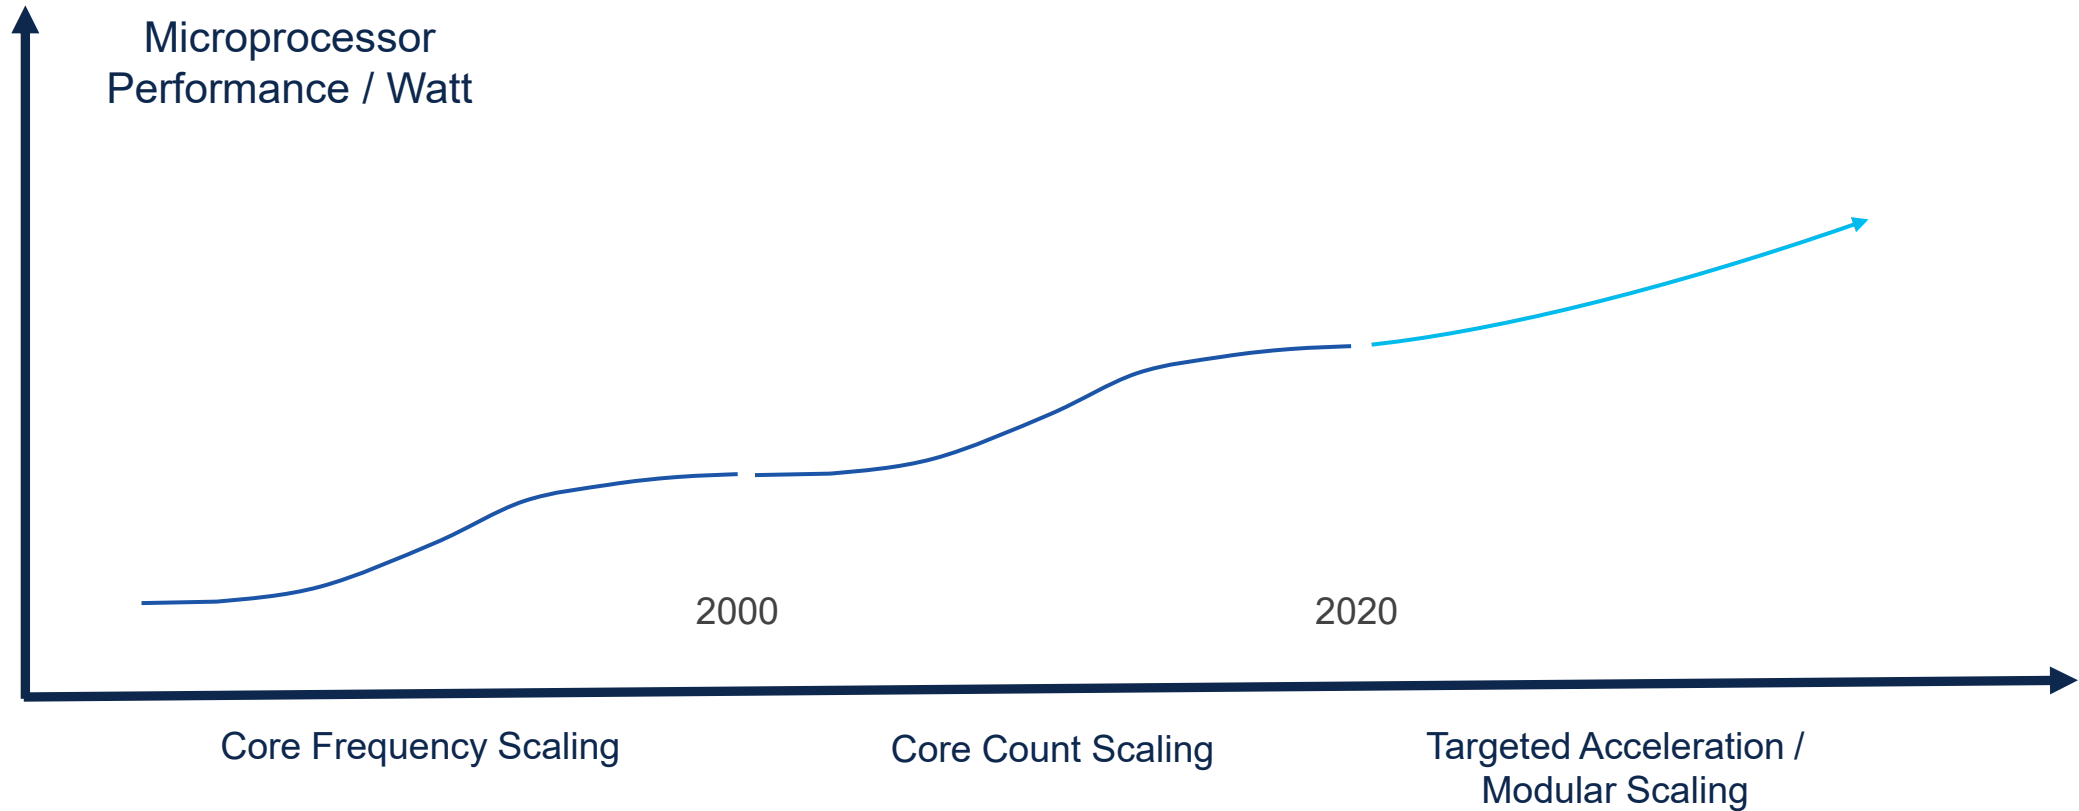
### Storage

32/64TB NVMe Drive, EDSFF, Low  
Latency NVMe

# Silicon Directions



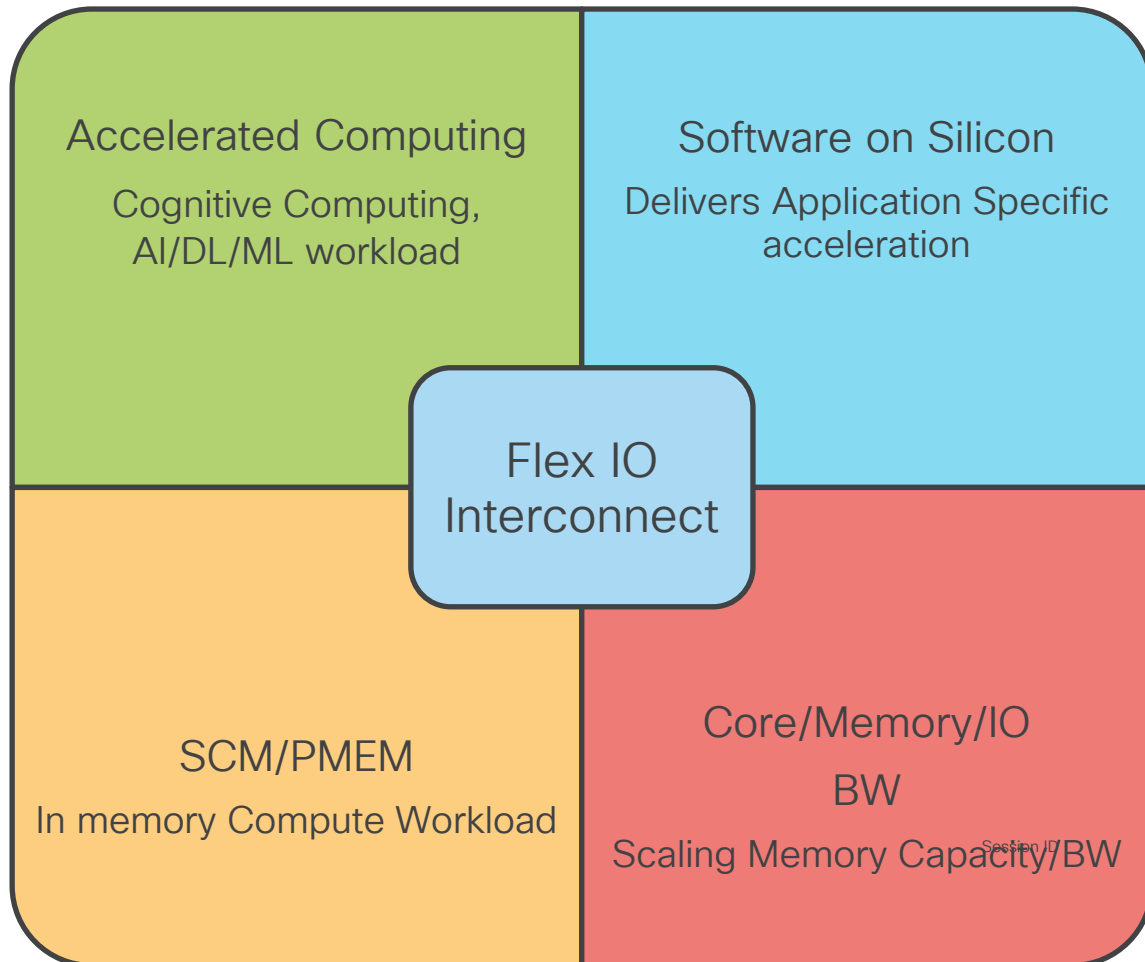
# Performance Efficiency – Server Market



# Hardware Disaggregation



# IO Interconnect



How do interconnects meet new requirements ?

- Latency Guarantees
- Coherency (Load / Store)

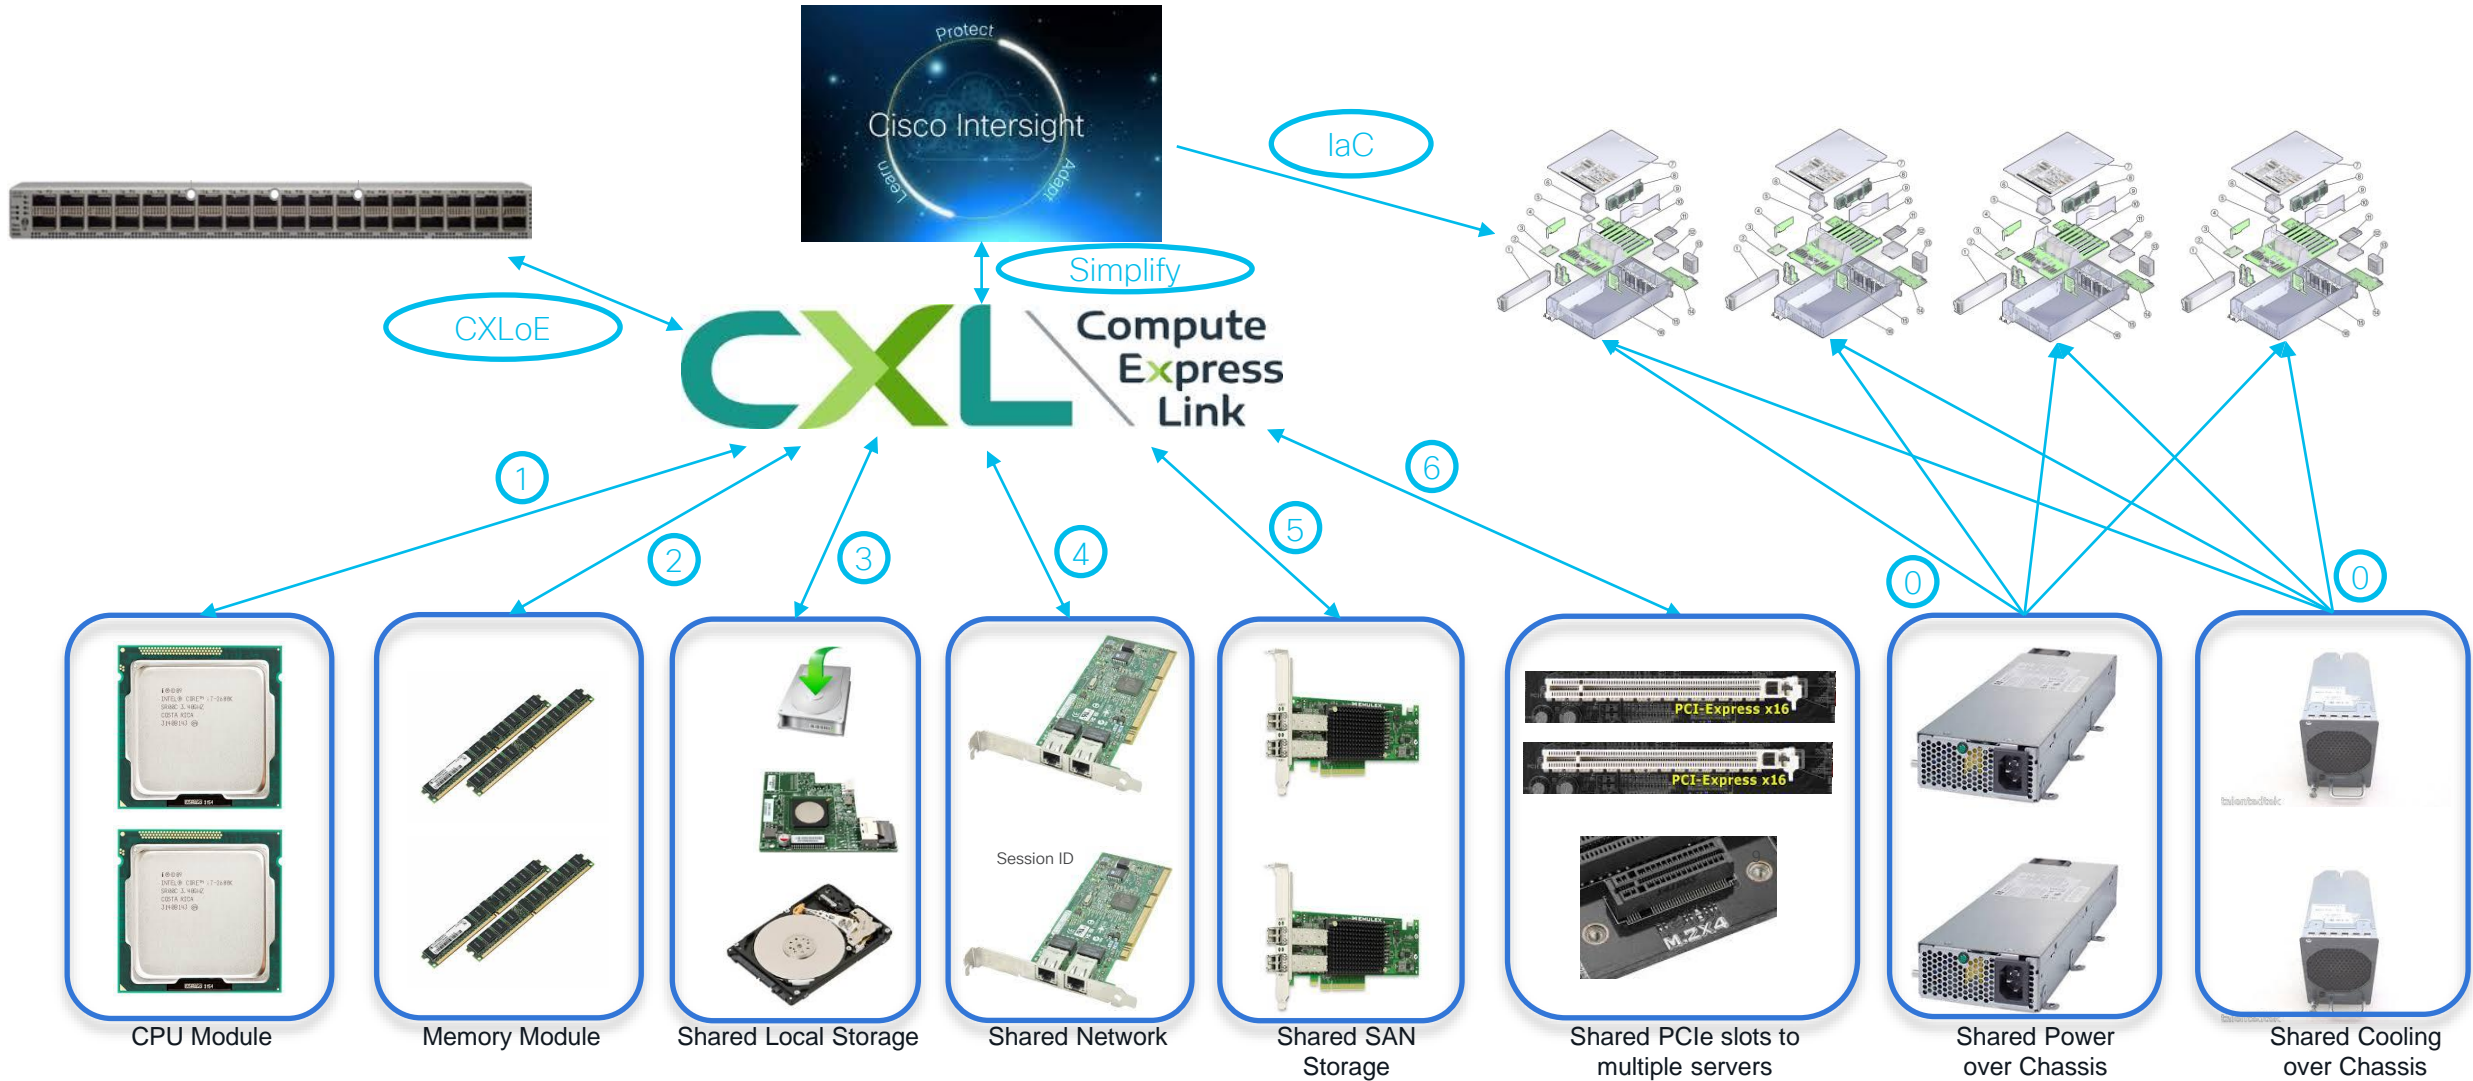
Options:

- Ethernet
- PCIe Gen4,5,6,X
- CCIX, OpenCAPI
- Gen-Z
- Infiniband
- NVLink
- **CXL** (Cisco is one of the Promoters/Founding Members)

Industry is investing in multiple solutions, CXL is gaining most traction, need a flexible architecture to support multiple approaches.




# Server Dis-Aggregation: Driven by Cost & Sustainability



# Blade vs. Rack? That Separation is also Fading...

- New compute innovations are blurring long-standing assumptions
- These innovations have impacts to networking, security, storage, etc.
- A robust orchestration layer is needed to keep simplicity



Item	Rack	Blade	Future
More Drives	24-48	2-4	6+N
More PCIe	2-12	2-4	2+N
Standard PCIe	Open	Closed	Open
Higher Power Peripherals / Cards	1-6	1	1+N
Direct to ToR switch	Always	IO Modules	Either
Unified IO	Separate	Default	Default
Shared Power & Cooling	Minimal	Default	Default
Unified Management	Minimal – no blade	Default – no rack	Default – any form factor
Compute Fabric (CXL, GenZ, etc.)	Add on	Add on	Integrated


Session ID

# Liquid Cooling




# Power & Cooling

- CPU, GPU and Switch ASIC power requirements moving from ~205W TDP today to 350W+ in the coming years
- Traditional fan cooling consumes lot of power and less efficient as system power increases
- Passive cooling is approaching its limitation
- Liquid cooling technology to address future cooling requirement with significantly better cooling efficiency & reduced noise levels
- Closed loop liquid cooling provides a retrofit solution
- Future Data Center designs will need to provision for Rack level liquid cooling infrastructure (with external Cooling Distribution Unit - CDU)




# From Air to Liquid Cooling

## Air-Cooled Traditional




- Embedded Heat Pipe Heat sinks carrying water droplets
- Fan primary mover of hot air
- Maxes @350 W TDP
- About 10% Cooling Power
- Exponentials increase in cooling power beyond 350W TDP

## Closed Loop-Liquid Cooling




- Close Loop Cold Plates liquid cooling
- Single Phase –Leakage concern
- Pump/Radiator circuitry
- Radiator is air-cooled
- @500 W TDP for 2U
- @15% Cooling Power for 500 W
- Lower Cooling Power for 350W TDP
  - Better for NetZero
- Blades stay at 350 W

## Liquid Cooling Rack-level close loop



- Open Loop Cold Plates
- No coolant connection outside of the rack
- Single Phase(water) or 2 Phase (refrigerant)
- 500 W TDP
- Common for rackmounts and blades
- About 15% Cooling Power when paired with Liquid-Air CDU
  - Minimal advantage for NetZero

## Liquid Cooling Rack-level open loop



- Open Loop Cold Plates
- Coolant collection tied to facilities
- Single Phase ( water) 2 Phase (refrigerant)
- 500 W TDP and beyond
- Works for rackmounts and blades
- Common for rackmount & blades
- About 5% Cooling Power when paired with Liquid-Liquid CDU
  - Better for NetZero
  - About 1 kW of cooling power savings for a loaded UCS-X

# Persistent challenges for IT

Architectural silos  
spawn islands of  
infra and operations



Reduced visibility,  
increased cost  
and risk

Application diversity  
drives infrastructure  
diversity




Specialized systems  
increase operational  
complexity

Fragmented  
technologies impede  
time to value



Valuable IT energy  
wasted on platform  
integration

# Intel® Xeon® Roadmap Update



2017 - Current

April 2021

Early 2023

2024+

<p><b>Cascade Lake</b> 2<sup>nd</sup> Gen Intel® Xeon® SP Shipping</p>	<p><b>Ice Lake</b> 3<sup>rd</sup> Gen Intel® Xeon® SP Shipping</p>	<p><b>Sapphire Rapids</b> 4<sup>th</sup> Gen Intel® Xeon® SP Q1 2023</p>		<p><b>Next Rapids</b> Future Gen Intel® Xeon® SP 2024+</p>
<p>28 cores, Intel 14nm 205W, 6CH DDR4-2933 48L PCIe 3.0 3x UPI 1.0 (10.4 GT/s) Intel PMem™ 100 Series</p>	<p>40 cores, Intel 10nm 270W, 8CH DDR4-3200 64L PCIe 4.0 3x UPI 1.0 (11.2 GT/s) Intel PMem™ 200 Series</p>	<p>60 cores, Intel 7 (10ESF) 350W, 8CH DDR5-4800 80L PCIe 5.0, CXL 1.1 4x UPI 2.0 (16 GT/s) Built-In Accelerators</p>	<p>64 cores, Intel 7 (10ESF) 350W, 8CH DDR5-5600 80L PCIe 5.0, CXL 1.1 4x UPI 2.0 (20 GT/s) Improved Perf. by ~15%</p>	<p>&gt;64 cores 8CH DDR5-6400 PCIe 5.0, CXL 2.0 Advanced UPI speed Modular Architecture</p>
<p><b>Purley Platform</b> (2S/4S/8S)</p>	<p><b>Whitley Platform</b> (2S)</p>	<p><b>Eagle Stream Platform</b> (2S/4S/8S)</p>		<p><b>Next Stream Platform</b> (2S/4S/8S)</p>

4<sup>th</sup> Gen Xeon SP (Sapphire Rapids) Brings a Massive Jump in Platform Capabilities and Performances over 3<sup>rd</sup> Gen Ice lake

# DDR5 vs DDR4 – DDR5 Benefits



M7 Memory DIMM server technologies	DDR4	DDR5	DDR5 Benefits
Speed	1.6 to 3.2Gbps data rate; Max 3200MT/s	4.8 to 8.4Gbps data rate; First generation: 4800MT/s	DDR5 debuts at 4800MT/s, 50% increase over DDR4-3200
Operational voltage	1.2 Volts	1.1 Volts	Reduce power; at 1.1V, DDR5 consumes ~20% less power than DDR4
Power Management	On Mother board	On DIMM PMIC	Better power efficiency, better scalability; Improves signal integrity, and reduce noise
DRAM DIMM Type	RDIMM (16GB-64GB) LRDIMM (128GB, 256GB)	RDIMM	Registered DIMM only
Channel architecture	72-bit data channel (64data + 8 ECC) 1 channel per DIMM	40-bit data channel (32 data + 8 ECC) 2 channels per DIMM	Higher memory efficiency and lower latencies of data access for the memory controller
On Die ECC (Error Correction Code)	No	Yes	Correct bit errors within the DRAM chip; Increase reliability, reduce defect rates
Max DRAM die density	Up to 16Gb SDP --> 64GB DIMM	Up to 64Gb SDP --> 256GB DIMM First gen: 16Gb SDP, 64GB DIMM	Higher Single Die per Package DIMM, eventually saving cost


DDR5 is faster, more efficient, and will scale to higher densities/speeds.

On M7, with Sapphire Rapids, we will support DDR5-4800, 50% BW increase over M6.



# Cisco UCS

It's not a server.




It's a system.

# UCS Fabric Simplified


Simplicity

Resiliency

TCO Reduction



TRADITIONAL RACK



Ad Hoc and Inconsistent

TRADITIONAL BLADE



Structured, but Siloed and Complicated

CISCO UCS



Simplified, Optimized and Automated

# 7<sup>th</sup> Gen UCS Servers and GPU support

## X-Series Compute Nodes



UCS X210C M7  
(2-socket)

UCS X410c M7  
(4-socket)



Run any app with  
adaptable system

## PCIe Nodes with Intel or nVidia GPU + UCS X-Fabric Technology



Accelerate  
modern apps

## UCS C-Series Servers




UCS C220 M7 Rack Server



UCS C240 M7 Rack Server

Run more apps  
in a familiar form factor


# Meet the UN/BOX




## UCS X-Series powered by Cisco Intersight



Cloud  
Operated



Unparalleled  
Flexibility



Future-ready



# Radically simplified hybrid cloud infrastructure


Simplify with  
cloud-operated infrastructure



Simplify with a system designed  
for modern applications



Simplify with a system  
engineered for the future




UCS X-Series

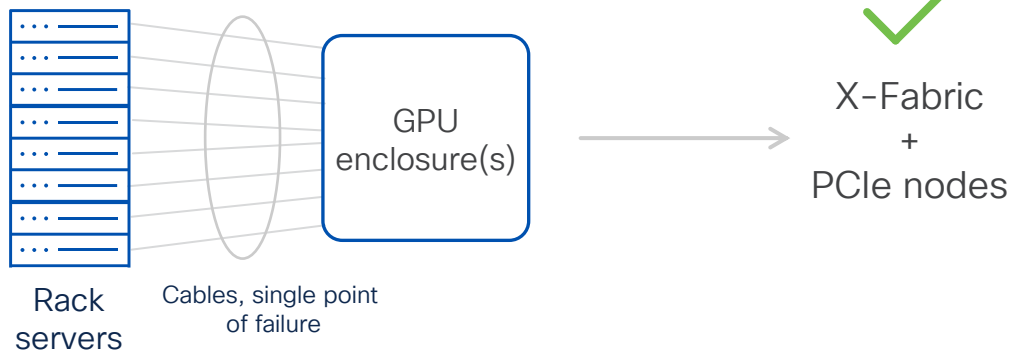
# Industry-leading simplicity

## Conventional approaches

1 | Silos of multiple ethernet and SAN fabrics and adapters



2 | Complex PCIe connectivity to external accelerators




## Cisco solution

UCS X-Series




Cisco InterSight



# UCS X9508 System chassis

## Chassis

7RU IO direct connect  
8 flexible slots  
Optical ready  
Liquid cooling ready




## Power and cooling

6x 2800W PSU  
54V power distribution  
4x 100mm Dual rotor fan

## Ethernet fabric

Two Ethernet Modular fabrics  
2 TB/s throughput




## X-Fabric modules

Two flexible IO modules  
For storage and GPU nodes connectivity




7 RU

### Chassis



8 Nodes




Compute Node  
With 6xSFF, 2xM.2

500+ Cores,  
~100TB Mem, >1PB  
Drives,

PCIe Nodes  
Drive Nodes\*


Intelligent Fabric Modules  
Fabric for Converged Traffic  
(Ethernet, Mgmt, FC, NVMeOF etc)



Replace Fabric for  
higher bandwidth

Replace Fabric  
when PCIe/CXL Gen  
changes

Backplane less design For Fabric  
(Simpler Upgrade, No Cap on bandwidth)



Flexible X Fabric  
(Fabric for PCIe, CXL Traffic)



Power and Cooling  
Innovation!

Power Distribution, Fan Speed Sensors, Fan  
Control Policy, Air Flow, Others

High Energy Efficiency and  
Power Envelope For Future



## Consolidate Rack Workloads



AI/ML



Accelerated VDI




Big Data, SDS, Containers



Traditional Blade Workloads

## UCS X-Series with X-Fabric



Up to 960

Cores  
per Chassis  
(M6 or M7)

24

GPUs  
per Chassis



200G

Bandwidth to  
compute node

1 PB

of storage

# UCS X210c M7 2S Compute Node

## Key Features

### CPU

Up to 2 Intel 4<sup>th</sup> Gen Xeon CPUs

### X-Fabric Mezz

Connect to PCIe Node

In future, to Memory/Drive Nodes

### VIC

Up to 2 VICs

100G Cisco mLoM

25G Cisco mLoM

25G Cisco VIC Mezz

Agg NW 200G per node

### Memory

Up to 32x DDR5 4800 MT/s DIMMs

### GPU

Up to 2 x PCIe GPU

+

Up to 2 x NVMe PCIe Drives


### Drives

2 x M.2 SATA HW Raid1/ 2 x M.2 NVMe

Up to 6 x SAS/SATA/NVMe

- RAID Controller for SAS/SATA

- Up to 90TB (15TB x 6) per Blade



# UCS X410c M7 4S Compute Node

## Key Features

### Form Factor

Occupies two slots in chassis

### CPU

Up to 4 Intel 4<sup>th</sup> Gen Xeon CPUs


### X-Fabric Mezz

Connect up to 2x PCIe Nodes

In future, to Memory/Drive Nodes

### Memory

Up to 64x DDR5 4800 MT/s DIMMs



### VIC

200G Aggregate/ 100G per Fabric  
100G 5<sup>th</sup> Gen VIC mLOM

or

25G 5<sup>th</sup> Gen VIC mLOM  
25G 5<sup>th</sup> Gen VIC Mezz

### Drives/GPU

#### Front Mezz options


1. 6 SAS/SATA with HW RAID
2. 6 NVMe PCIe Gen4 x4
3. 2 NVMe PCIe Gen4 x4, plus 2 PCIe GPUs

#### Internal

1. 2x M.2 SATA HW RAID1
2. 2x M.2 NVMe

# GPU on X210C

- 2x T4 (with modified heat sink) in the front mezz slot
- Supports 2x NVME drives
- High GPU density solution with GPU node (24xT4 in 7RU with 4 nodes)





# UCS X-Fabric Technology and PCIe Nodes with GPU

Open, modular design enables compute and accelerator node connectivity

PCIe node supports up to

- 4x Intel Data Center GPU Flex 140\*
- 2x Intel Data Center GPU Flex 170\*
- 2x Nvidia A16
- 2x Nvidia A40
- 4x Nvidia T4
- 2x Nvidia A100



UCS X-Fabric Technology


- ✓ Based on native PCIe Gen. 4
- ✓ Provides GPU acceleration to enterprise application
- ✓ No backplane or cables = Easily upgrades



# UCS Fabric Topologies :

## Chassis Bandwidth Options for X-series

- Wire once for Connectivity and Discovery
- Add cables for increased bandwidth



**2X 1 LINK**

200 Gbps per Chassis

**2X 2 LINK**

400 Gbps per Chassis


**2X 4 LINK**

800 Gbps per Chassis

**2X 8 LINK**

1600 Gbps per Chassis


# Use Cases



Remote Work (VDI)



Cloud Native



Critical Apps

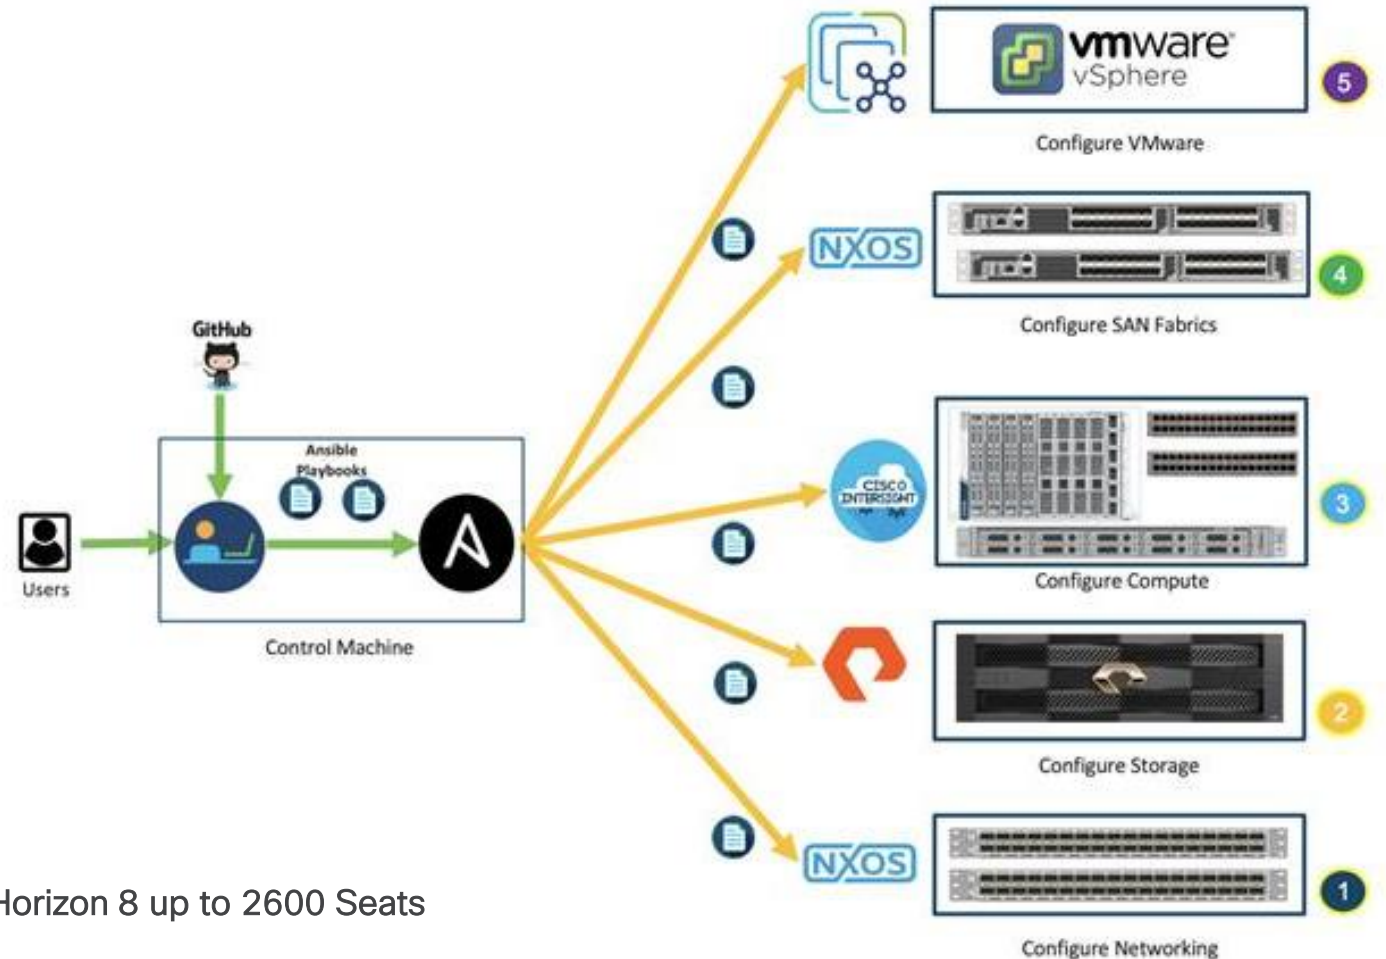


# Remote Work

## Use Cases

FlashStack VDI using Cisco UCS X-Series M6 with VMware Horizon 8 up to 2600 Seats

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flashstack\\_ucs\\_x\\_series\\_vmware\\_vsphere.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flashstack_ucs_x_series_vmware_vsphere.html)



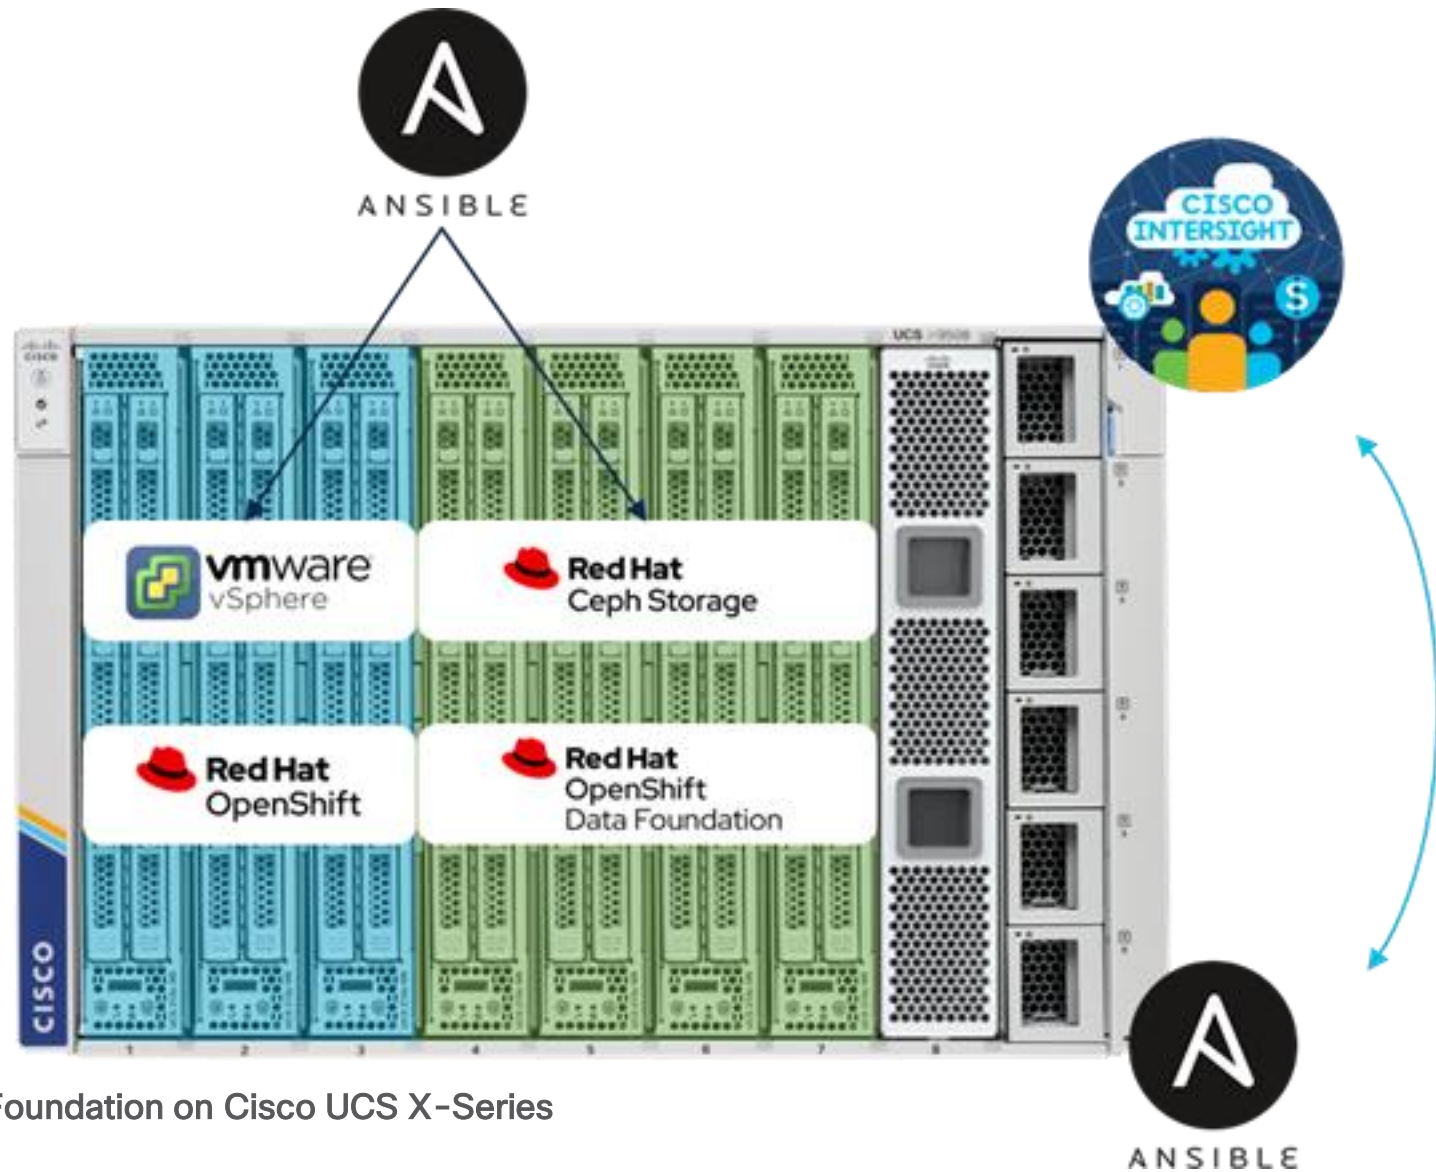


# Cloud Native

## Use Cases

Red Hat OpenShift Container Platform with OpenShift Data Foundation on Cisco UCS X-Series

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/ucs\\_xseries\\_ocp\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/ucs_xseries_ocp_ci.html)



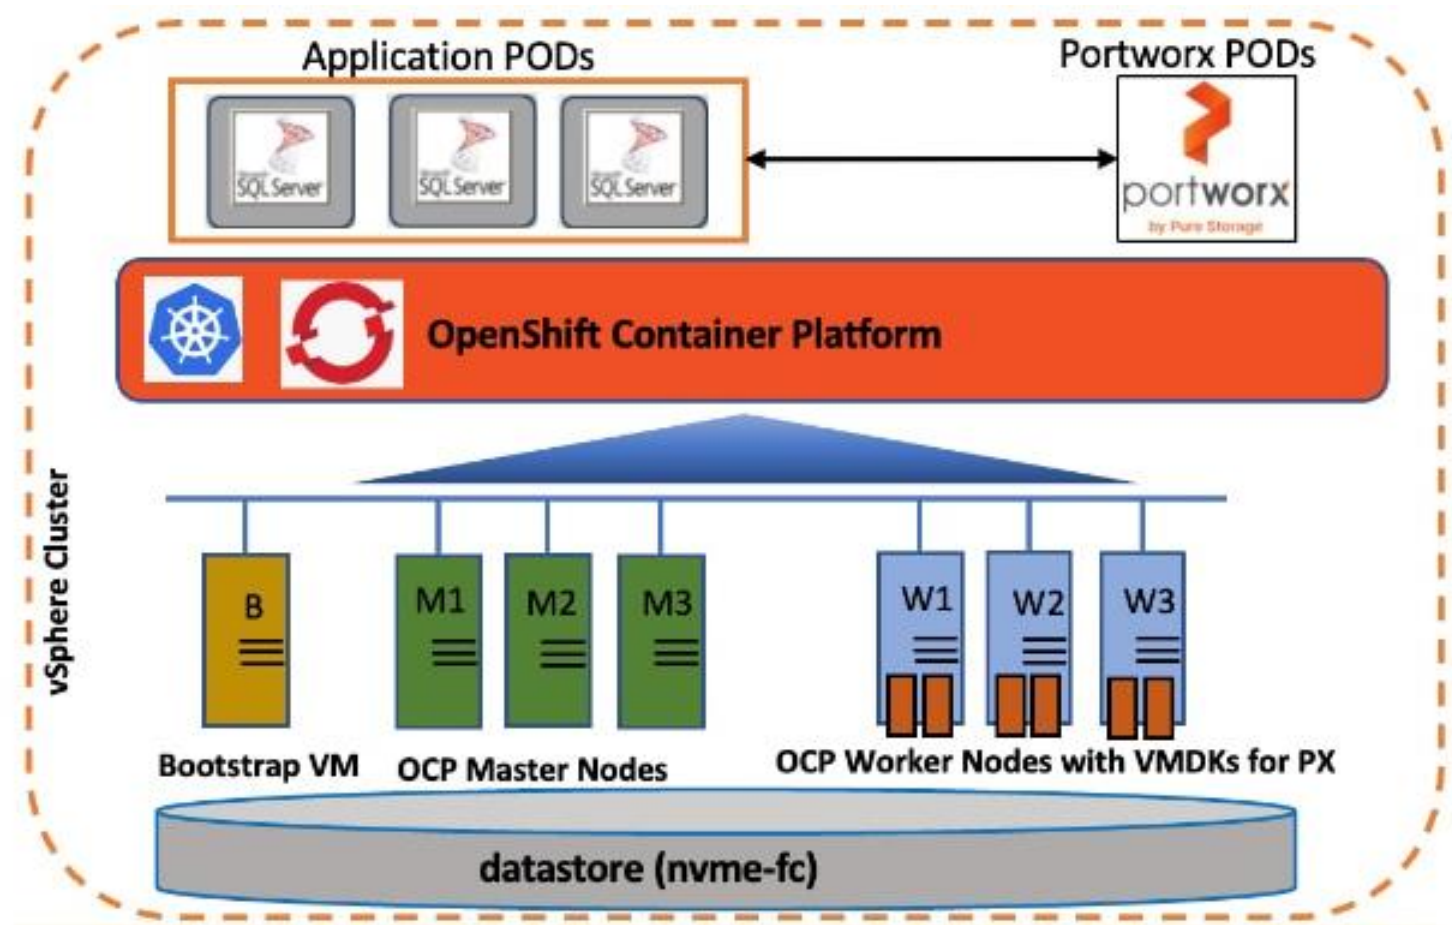


# Critical Apps

Use Cases

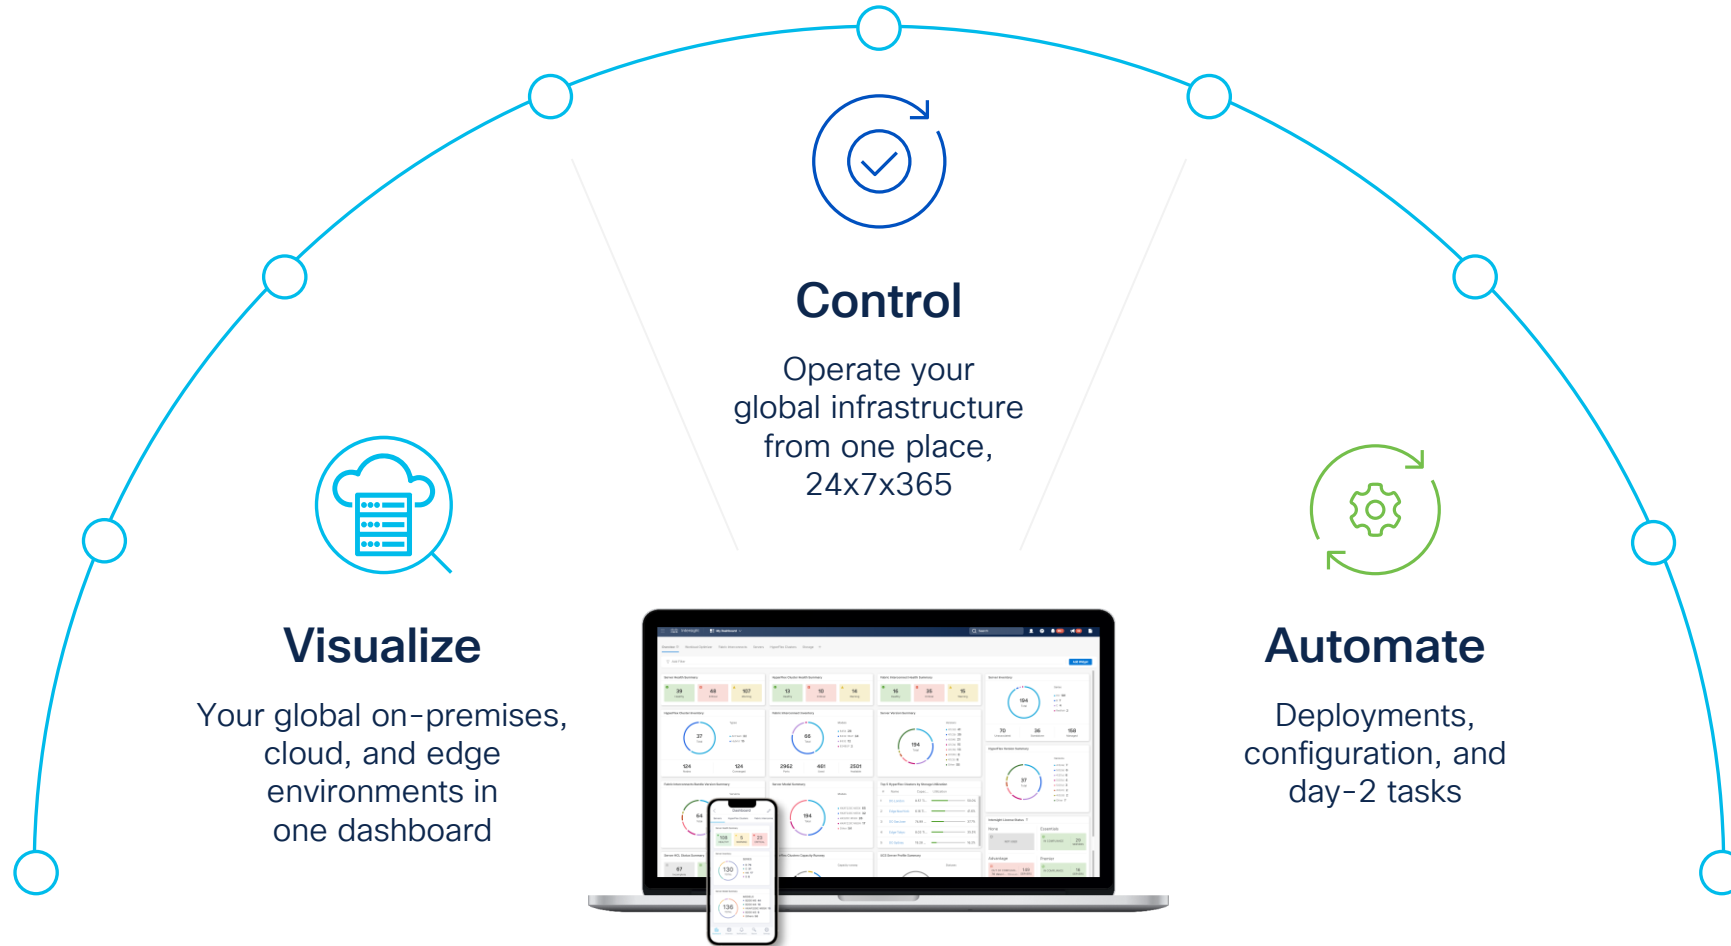
FlashStack for SQL Server 2019 with Cisco UCS X-Series and Pure Storage FlashArray//XL170

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/ucs\\_xseries\\_flashstack\\_sql\\_openshift.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/ucs_xseries_flashstack_sql_openshift.html)



# Modernize infrastructure management

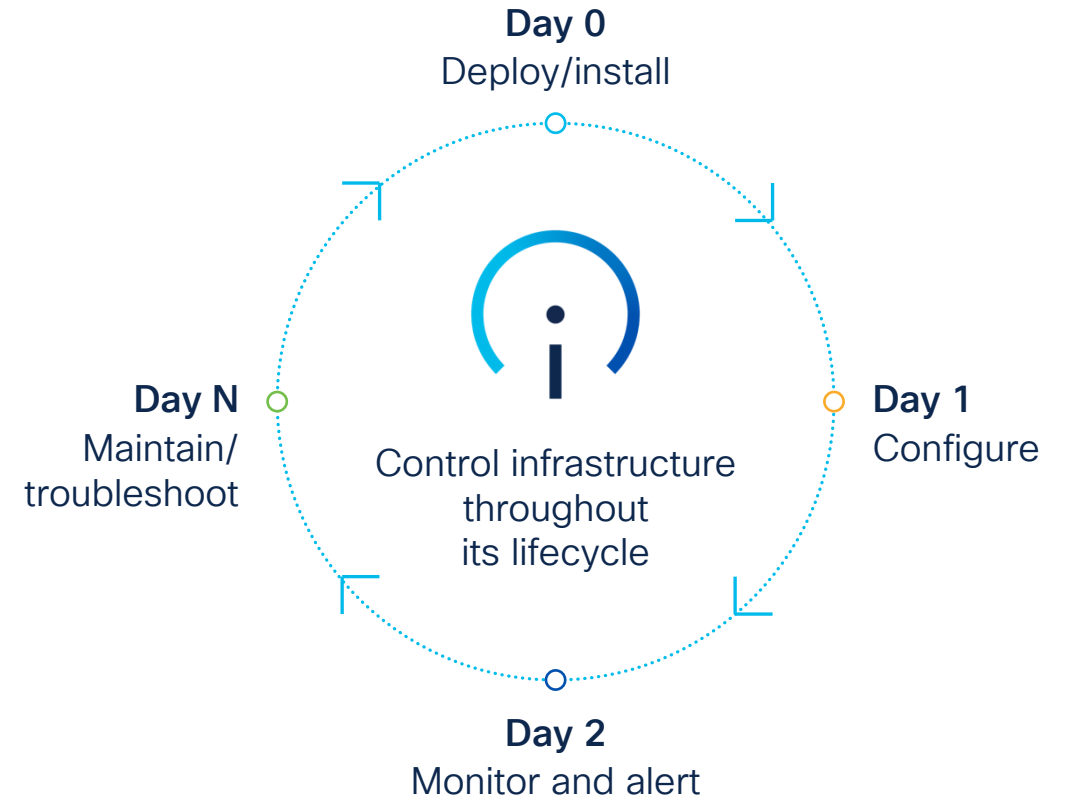
Cisco Intersight Infrastructure Service



# Control

## Infrastructure lifecycle management

Deploy, configure, and operate servers, VMs, storage, and networking throughout their lifecycle, from one place... anytime, anywhere



“Intersight has reduced the weight of routine deployment and monitoring tasks so we can focus on the work that helps our business.”

Director of IT and Communications, regional medical center

# Prevent and resolve

**Infrastructure issues with custom insights and automated support**

- Security advisories and field notices
- HCL baselining and notifications
- Proactive RMAs
- Connected support (TAC)
- Firmware upgrades
- Compliance

“You know you are being proactive when an engineer replaces a hard drive that was displaying anomalies before it fails and causes problems downstream.”

Distinguished Engineer, global technology conglomerate

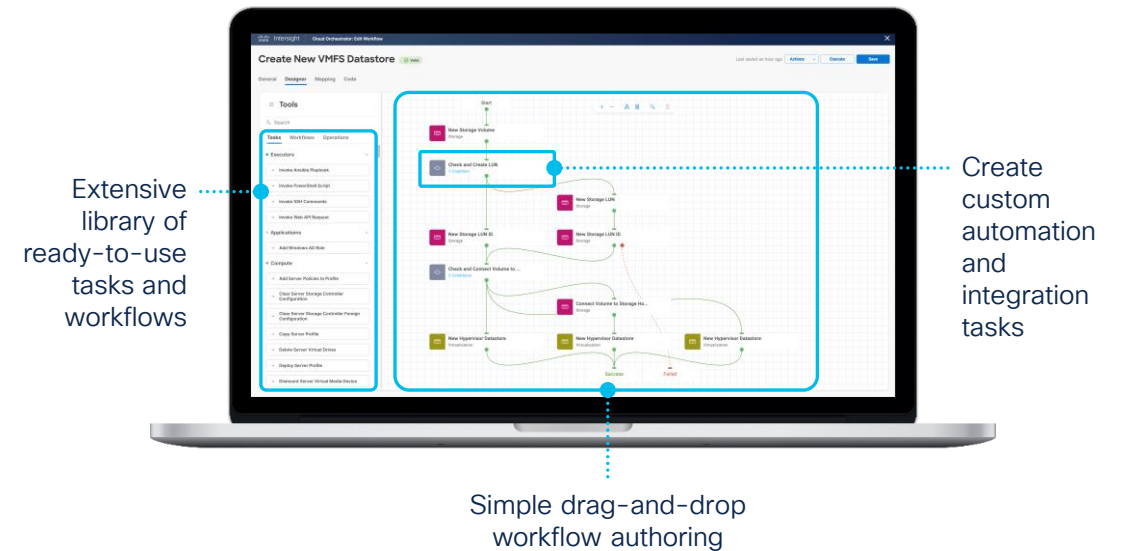
# Automate

## Create automated workflows to...

- Accelerate delivery of apps and infrastructure
- Reduce risk with standardization
- Simplify cross-domain orchestration

## Two options

- Workflow designer  
Drag-and-drop workflow authoring
- APIs/SDKs  
For using your existing automation tools (Ansible, Terraform, etc.)



“With consistent and repeatable workflows offered by Cisco Intersight, IT teams can save time and decrease operational costs.”

# Integrate

## Cisco Converged Infrastructure Solutions with UCS X-Series



Simplified cloud-based management of solution components

Monitoring, orchestration, and workload optimization for different layers of solutions

## Storage plugins



See storage capacity, controllers, configuration, snapshots, and replication



Automate storage management and orchestration



Use storage task library to create and execute workflows

## ServiceNow plugin



Incident management

Basic configuration

Inventory module

## Third-party server support



Dell PowerEdge/HPE ProLiant



See inventory and health

Add to automated workflows



# Deployment options



CISCO INTERSIGHT

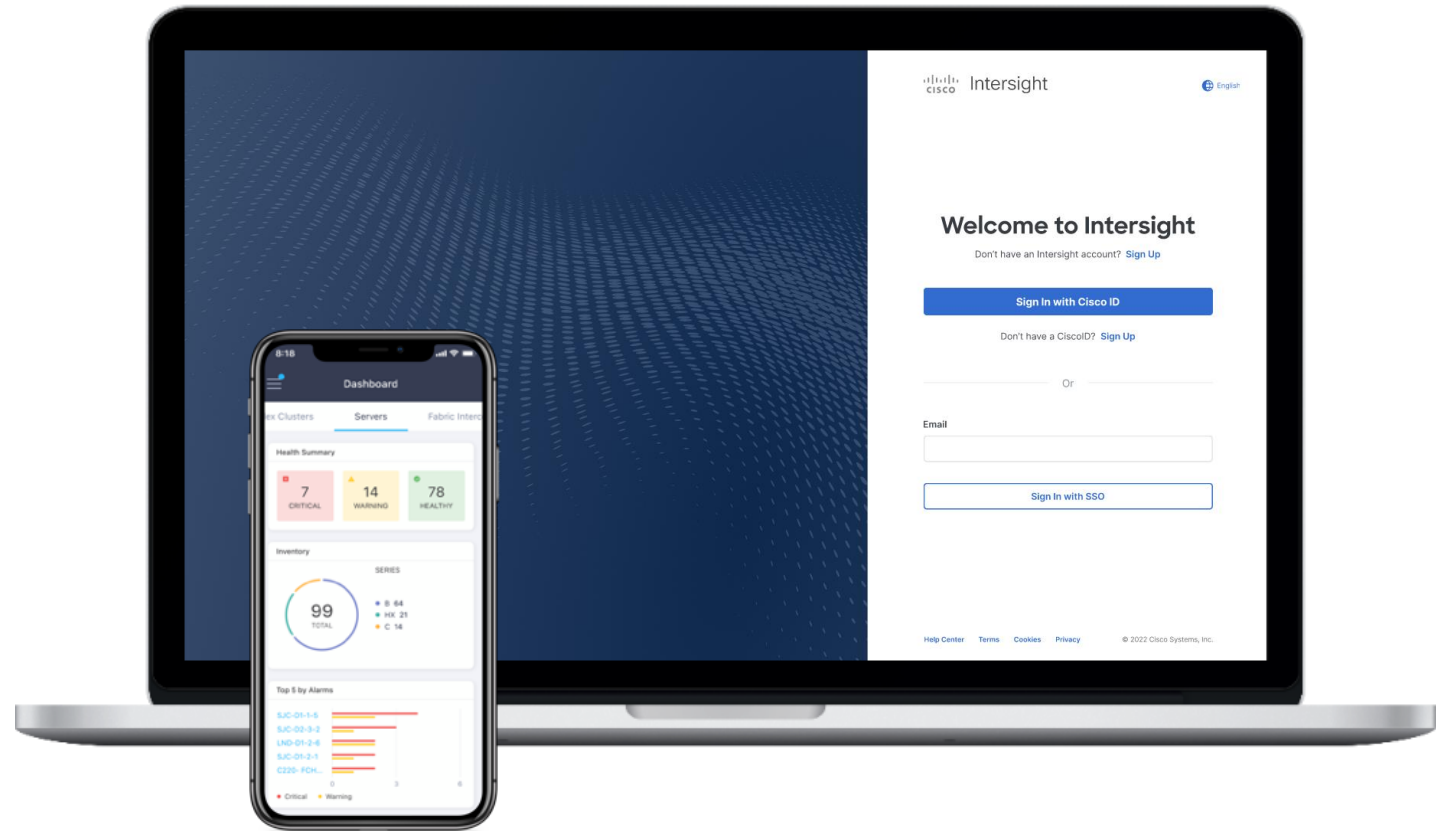
## Intersight.com

Cisco-hosted SaaS



## Virtual appliance

Connected  
Private (air-gapped)



# When it comes to efficiency, rack servers fall short

## X-Series advantages over rack servers



### Modular

Produces roughly 50% lower consumption of raw materials over three generations than rack servers



### Energy-efficient

Reduces the number of necessary inverters, power conversions, and by design runs more efficiently lowering total energy consumption



### Usage Aware

Uses constant monitoring and cooling algorithms to dynamically adjust power usage for optimal efficiency

---

## Sustainable by design



### Sustainable packaging

Offers multipack options to reduce packaging for high-volume orders



### Recycled materials

Plastic parts use post-consumer recycled resin



### Efficient Components

Utilizes efficient components like high-efficiency, Titanium-rated power supplies

---

## Embedding sustainability through the product lifecycle



### Supply Chain Sustainability


Cisco works to uphold human rights, promote worker well-being and minimize negative environmental impacts in our supply chain.



### Product Takeback and Reuse

Return end-of-use gear for free in a simple, secure, and sustainable way. Certified remanufactured products and available through Cisco Refresh.

# X-Series industry-leading sustainability by design



## Architectural Power & Lifecycle Optimization

### Standard out of the box

- Unified IO and 100G for all servers reduces adapter, cabling, network and storage port needs
- Efficient 54V power distribution to minimize transmission and conversion loss within all system elements
- Input power monitoring at system level allows a quicker reaction based on real usage to realize 3-4% overall savings
- Titanium rated PSU

### Optional add-on configuration

- X fabric allows for lifecycle upgrades of server modules while extending the life of many existing components (PSU, Fan, Case, Risers, and more in future)

## Cooling & Fan Optimization

- Zone-based cooling to intelligently operate fans where they are needed and not running them at common largest need
- Large real estate footprint on server blade modules allows for simpler airflow efficient fan power usage
- Midplane free design allows for future technologies such as liquid cooling to change the math on PUE in your DC without forklift upgrade of entire racks

# Workloads Trending Toward Liquid Cooling

Using GPUs, FPGA,  
Accelerators and  
Frequency  
Optimized  
Processors



# Intersight Infrastructure Services Simplifies Sustainable Operations



## Power policies

- Global setting for Chassis Power behavior
- Redundancy, Power Save modes

## Power profiling policy

- Node, Blade and Rack server
- Tests and determines actual min and max power draw

## Power capping policy

- Node, Blade and Rack server
- Individual and group power capping
- Group level or individual manual

## Dynamic power rebalancing

- Competitive advantage
- Chassis and Rack group level redistribution of available power

## Power group policy

- Set of chassis that draw power from the same PDU and maintain under the advisable limit

## Fan speed control policy

- Global or individual non disruptive server setting
- Controls power usage and airflow noise levels

## Power allocation policy

- Global setting to specify policy-based power capping or manual node power cap is used

## Server power control policy

- Chassis and Rack level
- Priority set to calculate initial power allocation and discretionary use

## BIOS and OS power policy

- Use BIOS tokens to adjust power based on app needs
- OS level power recommendations

# Intersight Services simplifies power management operations-example

This screenshot shows the 'Create' page for a BIOS policy in the Intersight Infrastructure Service. The left sidebar lists various system components, with 'Configure' selected. The main area is titled 'Policies > BIOS' and contains a 'General' tab and a 'Policy Details' tab. The 'Policy Details' tab is active, showing a table of configuration options:

Power And Performance	
C1 Auto Demotion platform-default	C1 Auto UnDemotion platform-default
Core Performance Boost platform-default	Global C State Control platform-default
L1 Stream HW Prefetcher platform-default	L2 Stream HW Prefetcher platform-default
Determinism Slider platform-default	Efficiency Mode Enable platform-default
CPPC platform-default	cTDP Control platform-default
Enhanced CPU Performance platform-default	LLC Dead Line platform-default
UPI Link Enablement	UPI Power Management

A 'Cancel' button is visible at the bottom left of the main content area.

This screenshot shows the 'Create' page for a Power policy in the Intersight Infrastructure Service. The left sidebar is the same as in the previous screenshot. The main area is titled 'Policies > Power' and contains a 'General' tab and a 'Policy Details' tab. The 'Policy Details' tab is active, showing configuration options:

**Policy Details**  
Add policy details

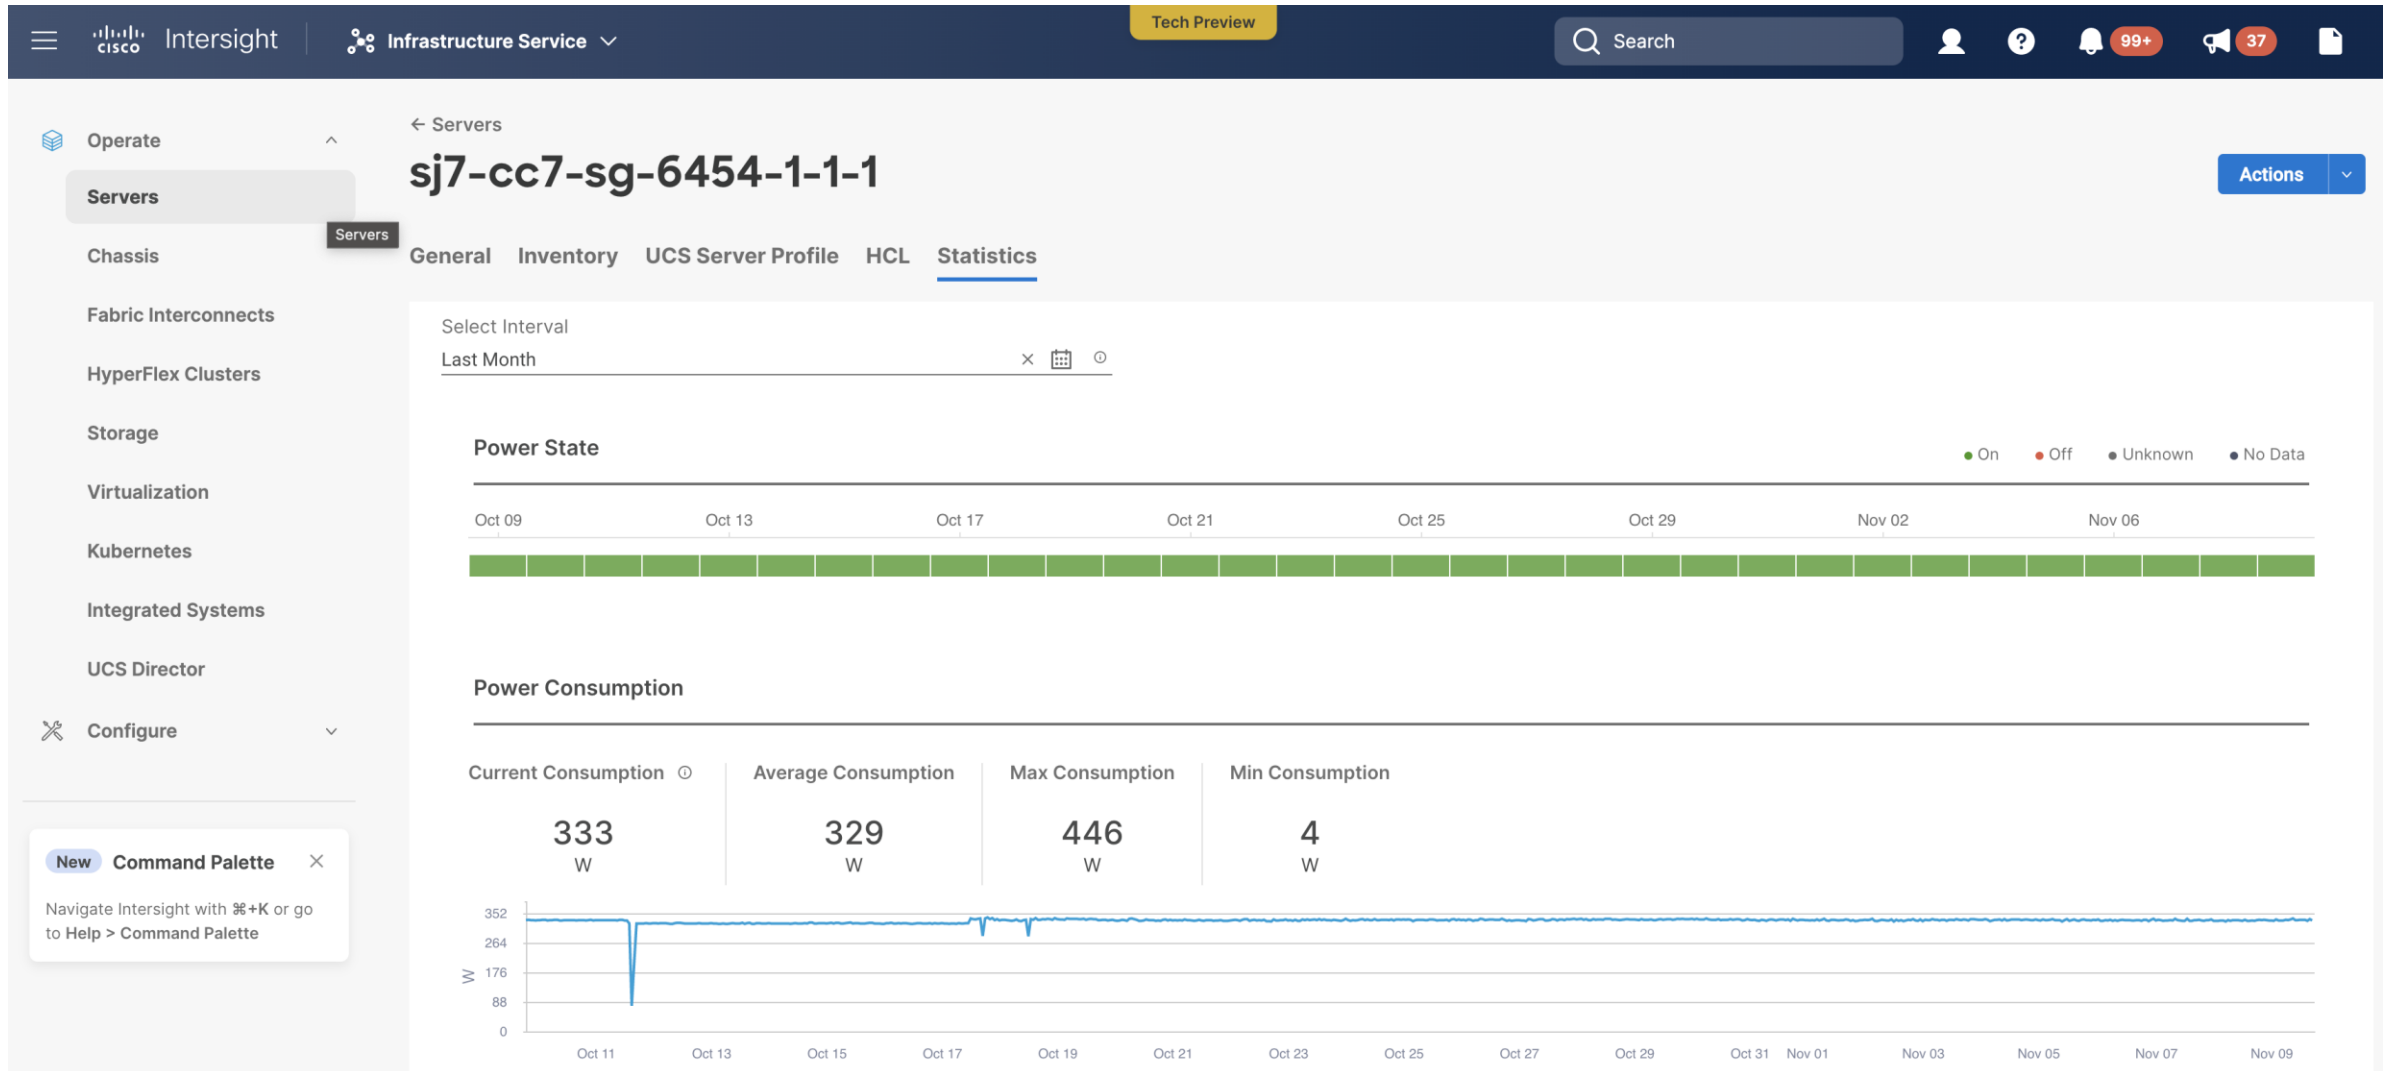
Platform filters: All Platforms | UCS Server (FI-Attached) | UCS Chassis

**Configuration**

- Power Profiling
- Power Priority: Low
- Power Restore: Always Off
- Power Redundancy: Grid
- Power Save Mode
- Dynamic Power Rebalancing
- Extended Power Capacity
- Power Allocation (Watts): 0

'Cancel', 'Back', and 'Create' buttons are visible at the bottom.

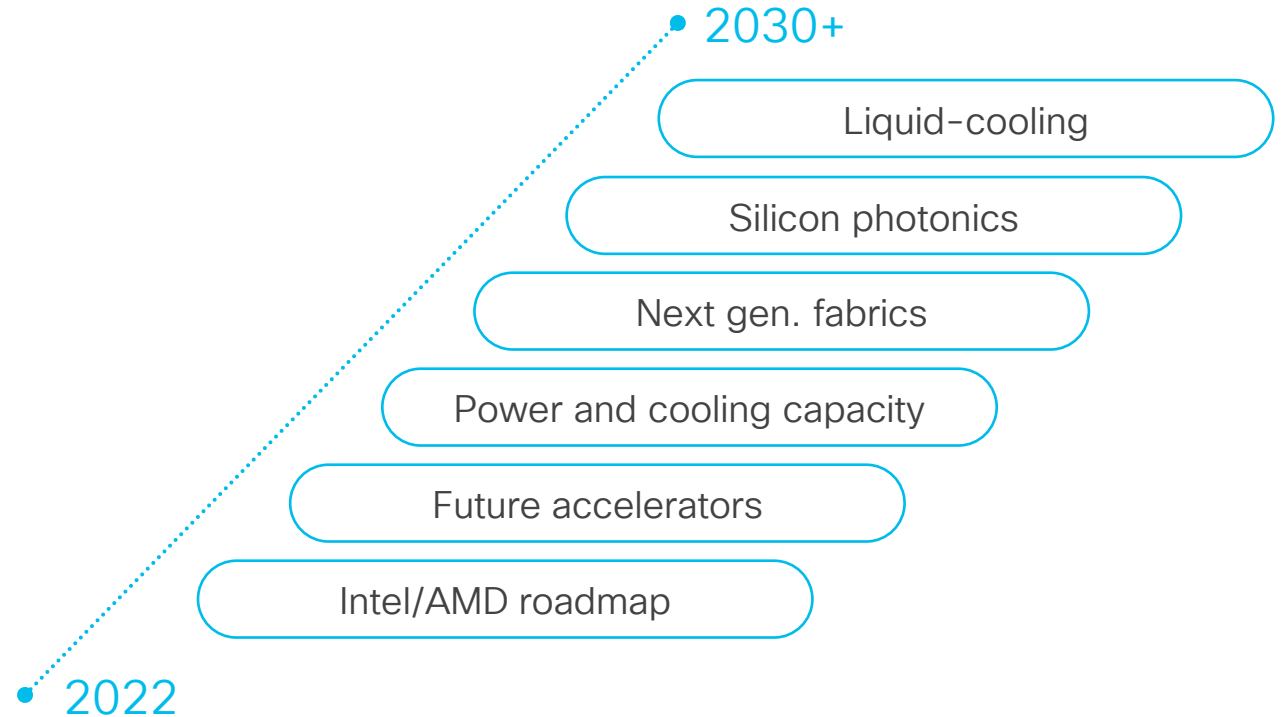
# Visibility & Trending Server Power Utilization



# Future-ready design

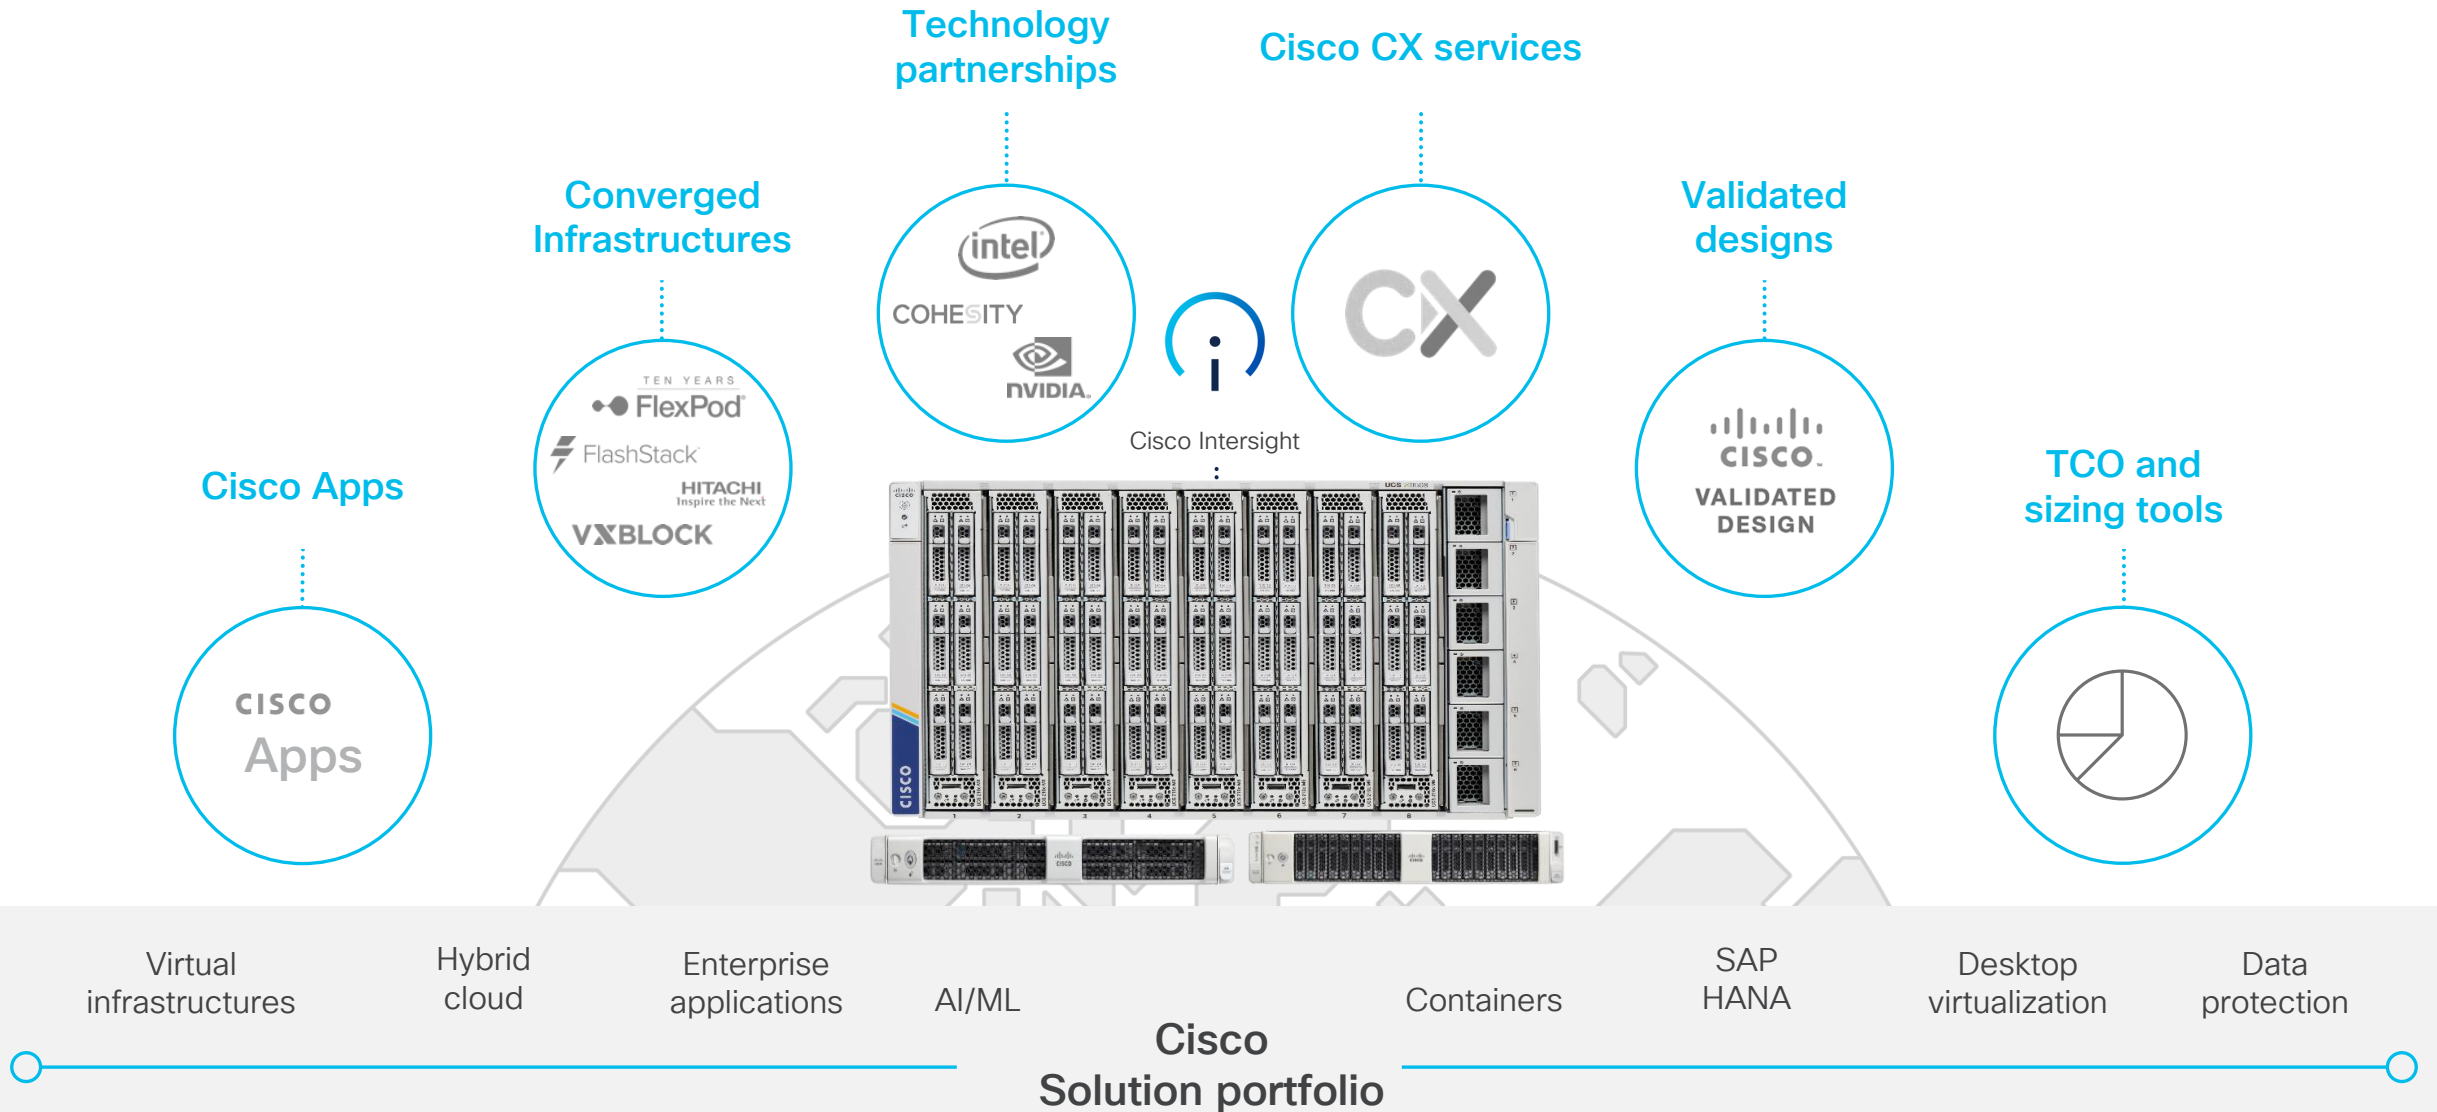
Engineered for the next decade

Protect investments  
and focus on innovation





# Computing for the next decade



# Resources



## Data Center Design Guides:

<https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/data-center-design-guides-all.html>

## Cisco UCS X-Series Modular System:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-x-series-modular-system/index.html>



## Cisco Intersight help center:

<https://intersight.com/help/saas>

## UCS Solutions Team GitHub Repository:

<https://github.com/ucs-compute-solutions>





The bridge to possible